Paper Interfaces: an HCI Approach to Geometry Education

THÈSE N^{O} 5579 (2012)

PRÉSENTÉE LE 21 DÉCEMBRE 2012
À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS
CRAFT - GESTION
PROGRAMME DOCTORAL EN INFORMATIQUE, COMMUNICATIONS ET INFORMATION

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Quentin BONNARD

acceptée sur proposition du jury:

Prof. P. Ienne, président du jury
Prof. P. Dillenbourg, Prof. F. Kaplan, directeurs de thèse
Prof. M. S. Horn, rapporteur
Prof. B. Signer, rapporteur
Prof. A. Wegmann, rapporteur

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2012

À ma famille, mes racines et mes ailes, mon exemple et ma force.
J’espère que ce travail vous rendra aussi fier de moi que je le suis de vous.

À Huong-Ly, de loin la plus importante et la plus belle découverte dans ma vie.

Acknowledgements

Cette thèse est le fruit de beaucoup d'efforts, mais pas seulement des miens. Je profite donc de la publication de ce document pour remercier tous les participants. Je le fais en français, la langue dans laquelle je pense et ressens, en gage de sincérité.

Tout d'abord, ce travail a été validé par un jury d'experts, que le Professeur Ienne m'a fait l'honneur de présider. Les professeurs Beat Signer, Micheal Horn et Alain Wegmann ont lu avec attention ce document, ont proposé de pertinentes améliorations, et m'ont permis de la défendre en personne.

Le jury était aussi composé de mes directeurs de thèse, les professeurs Pierre Dillenbourg et Frédéric Kaplan, qui m'ont guidé tout au long de ce périple. Pierre m'a fourni tout ce dont j'avais besoin pour ce travail, m'a guidé, m'a poussé à dépasser mes limites, et à aller un peu plus loin encore. Frédéric m'a fait bénéficier de ses lumières et de ses horizons, passionnés et passionnants, tout en gardant les pieds sur terre et la main à la pâte.

La valeur de cette thèse vient de son ancrage dans la réalité de l'école, et donc de la collaboration de qualité qui nous a été offerte par messieurs Jean-Claude Marguet et Alain Ramelet et leur service de l'enseignement obligatoire de Neuchâtel. Il est difficile d'exprimer à sa juste valeur la gratitude que j'ai pour les enseignants Olivia Angeli, Fabio Fébo, Denis Trachsel, Raphaël Brissat et Elisabeth Rolli qui m'ont donné sans compter de leur expertise, temps, énergie, et enthousiasme. Leurs élèves, mais aussi ceux venus pour les événements organisés par le bureau de l'égalité des chances de l'EPFL, m'ont aussi offert de leur patience et de leur créativité.

Ma thèse s'est appuyée sur la TinkerLamp, un système déjà existant et développé par Guillaume Zufferey, Patrick Jermann, Olivier Guédat, Son Do-Lenh, Aurélien Lucchi, et David Bréchet. Ils nous ont laissé en héritage un système de qualité et un important travail sur le logiciel. Ce logiciel a servi de base à une fructueuse collaboration avec Sébastien Cuendet, Nan Li, Nikos Maris et Andrea Mazzei. Andrea a d'ailleurs été plus qu'un collaborateur, partageant avec moi son expertise de la vision par ordinateur. De même, Patrick m'a tout appris des statistiques.

D'autres compagnons m'ont aidé dans l'exploration des interfaces papier pour la géométrie. Chia-Jung Chan Fardel a illuminé le tout début du voyage d'une activité complète et inspiratrice. Christophe Schild et Grégory Del Colle ont défriché l'hostile algorithme de reconnaissance utilisé dans le Chapitre 8 – Andrea l'a transformé en terre fertile. Carlos Sanchez-Witt est parti en éclaireur pour le Chapitre 5. Michael Chablais est à l'origine de l'architecture logicielle de l'activité du Chapitre 6. Anna Geiduschek a continué l'exploration alors que j'étais retenu

par la rédaction de cette dissertation, en créant une impressionnante activité originale.

Amanda Legge mérite bien un paragraphe à elle toute seule pour sa contribution. Elle a irradié de son enthousiasme, de sa méticulosité, et de sa créativité la moitié du Chapitre 7 et le Chapitre 8 qui, s'il avait existé sans elle, aurait été bien amoindri. Ne s'arrêtant pas là, elle a eu l'incommensurable patience de relire la quasi-totalité de ce document.

Mes collègues Himanshu Verma, Son Do-Lenh, Kshitij, Julia Fink, Andrea Mazzei, Nan Li, Sébastien Cuendet et Hamed Alavi ont aussi partagé de bonne volonté cette pénible tâche de relecture. Ingrid Le Duc et Siara Isaac ont quant à elles significativement amélioré la présentation orale de ce travail.

Tout le Craft (celui d'hier, d'aujourd'hui et de demain) a d'ailleurs participé à ce travail en fournissant un cadre de travail stimulant, et un lieu de vie agréable, sous la bienveillance universelle et parfois sucrée de Florence Colomb.

Enfin mes amis, au Craft et ailleurs, et bien sûr ma famille et Huong-Ly ont soutenu cette thèse comme ils m'ont toujours soutenu, en me donnant ma force chacun à leur manière.

Lausanne, le 30 novembre 2012

Quentin

Abstract

Paper interfaces are an alternative for controlling a computer. Typically, users interact with pieces of paper which are detected by a camera and augmented with relevant information by a projector. The development of paper interfaces, historically, aimed at merging digital and physical versions of documents, enabling to transparently work on one or the other. Furthermore, in the new era of natural interaction techniques, the special affordances of paper can be of great value as a basis for tangible interaction and Augmented Reality: digital objects, linked to paper artefacts, can be manipulated by folding, cutting, orienting, etc.

In the context of classroom technologies, paper interfaces are especially appropriated, because paper is integrated and ubiquitous in the school environment and learning processes. Students and teachers are familiar with its properties and know how to interact with it. The goal of this dissertation is to explore the possibilities of using paper as a support for the learning content and more importantly as a tangible body. We focus on geometry education at primary school, because it is a domain where these two aspects of paper can be extensively exploited: pupils can write formulas or draw figures, and they can also move cardboard shapes, or fold along symmetry axes.

We designed five series of pedagogical activities: classifying quadrilaterals, mastering the protractor, communicating angles, exploring symmetries, and describing physical transformations with geometrical concepts. These activities are experimented in increasingly valid settings, such that the last series took place in regular classrooms. We also developed methods to analyse the learning activity happening during these experimentations.

Our studies revealed important insights on paper interfaces and their application in classroom education. The variety of collaborative scripts that could be created shows the flexibility of paper as a material for building user interfaces that support pedagogical designs. Such flexibility can be further exploited to enable teachers to create their own pedagogical Augmented Reality applications.

We also observed many examples of uses that were not intended in the original design, which we refer to as “creative appropriation”. As a result, our paper interfaces were integrated in the everyday conditions of the classroom, used intuitively by the pupils, and managed autonomously by the teachers.

Keywords: Paper Interfaces, Augmented Reality, Tangible User Interfaces, Tabletop Computing, Geometry, Computer Supported Collaborative Learning.

Résumé

Les interfaces papier sont une alternative pour contrôler un ordinateur. Un utilisateur interagit avec des morceaux de papiers qui sont détectés par une caméra et un projecteur les augmente avec les informations qui s'y rapportent. Historiquement, les interfaces papier avaient pour but de réunir les versions papier et électronique de documents, pour permettre de travailler de manière transparente sur l'une ou sur l'autre. Par ailleurs, selon la tendance actuelle à rendre les interactions plus naturelles, les particularités du papier sont un point de départ prometteur pour créer des interfaces tangibles et en Réalité Augmentée : des objets virtuels, liés à des artefacts en papier, peuvent être manipulés en les pliant, en les coupant, en les pivotant, etc. Dans le contexte des technologies en classe, les interfaces papier sont spécialement pertinentes, parce que le papier est intégré et omniprésent à l'école. Les élèves et les enseignants sont familiers avec ses propriétés, et savent comment l'utiliser. Le but de cette thèse est d'explorer les possibilités du papier comme support de contenu et comme objet tangible. Nous nous concentrons sur la géométrie à l'école primaire parce que c'est un domaine où ces deux facettes du papier peuvent être exploitées : les élèves peuvent écrire des formules ou dessiner des figures, mais il peuvent aussi manipuler des formes en carton ou plier une feuille selon un axe de symétrie.

Nous avons développé cinq séries d'activités pédagogiques utilisant ce système : la classification des quadrilatères, l'utilisation du rapporteur, la description d'angles, l'exploration de la symétrie, et la description de transformations physiques avec des objets géométriques. Ces activités ont été testées dans des conditions de plus en plus réelles, de sorte que les expériences finales ont eu lieu dans des classes ordinaires. Nous avons aussi développé des méthodes pour analyser l'activité d'apprentissage ayant eu lieu lors de ces expérimentations.

Les études nous ont beaucoup appris sur l'utilisation d'interfaces papier en classe. La variété des activités qui ont pu être créées montre la flexibilité du papier comme support d'interfaces à but pédagogique. Une telle flexibilité peut être davantage exploitée pour donner aux enseignants la possibilité de créer leur propre activité en Réalité Augmentée.

Nous avons aussi observé plusieurs exemples d'utilisations qui n'étaient pas prévues au moment de la conception, ce que nous appelons “appropriation créative”. Il en résulte que nos interfaces papiers ont été intégrées dans l'environnement de la classe, utilisées intuitivement par les élèves, et gérées de manière autonome par les enseignants.

Mots-clés : Interfaces Papier, Réalité Augmentée, Interfaces Utilisateur Tangibles, Surfaces Interactives, Géométrie, Apprentissage Collaboratif Assisté par Ordinateur.

Contents

Acknowledgements vii

Abstract/Résumé ix

List of figures xix

List of tables xxvii

Introduction 1

Position . . . 2
Overview . . . 3

Bookmark 5

1 Paper Interfaces 7

1.1 Class I – Augmented Paper Documents . . . 9
1.1.1 The DigitalDesk . . . 9
1.1.2 Synchronizing Paper and Digital Documents . . . 9
1.1.3 Digital Pens . . . 9
1.1.4 Mobility of the Feedback . . . 11
1.1.5 Augmented Writing . . . 11
1.1.6 Augmented Active Reading . . . 11
1.1.7 Augmented Reading . . . 12
1.1.8 Augmented Collaboration . . . 12
1.1.9 Borderline: Augmented Notebooks . . . 13
1.2 Class II – Structured Sets of Augmented Paper Documents . . . 13
1.2.1 Augmented Books . . . 13
1.2.2 Augmented Education . . . 13
1.2.3 Multimedia Scrapbooks . . . 14
1.2.4 Augmented Prototyping . . . 14
1.2.5 Augmented Collaboration . . . 14
1.2.6 Augmented Presentation . . . 15
1.2.7 Printout Tracking . . . 15
1.3 Class III – Tangible Paper . . . 16
1.3.1 Cards . . . 16

1.3.2 Augmented Reality Support 17
1.3.3 Non-Augmented Reality Uses of Tangible Paper 17
1.3.4 Scaffolding Interaction 18
1.4 Class IV – Tangible Paper Controllers 18
1.4.1 3D Surface 18
1.4.2 Paper Cubes 19
1.4.3 Paper Solids 19
1.4.4 Sticking Tags 20
1.4.5 Other 3D Paper Shapes 20
1.5 Map: our Activities in the Continuum of Paper Interfaces 20
1.6 Tangible User Interfaces 21
1.6.1 Examples of TUIs 21
1.6.2 Epistemology 23
1.7 Open Issues 24
1.7.1 Deeper Exploration of Paper as an Interface 24
1.7.2 Structuring the Design of Paper Interfaces 24

2 Computer-Supported Geometry Education 27
2.1 Computer Interfaces for Learning Geometry 27
2.1.1 Interacting with the Logo Microworld via a Keyboard 27
2.1.2 Dynamic Geometry Software: Manipulating Geometry with the Mouse . 29
2.1.3 Tangible User Interface: Geometry in the Shared Physical Space 31
2.1.4 Digital Pens 33
2.1.5 Paper Interfaces 34
2.2 Addressing the Challenges of Learning Geometry with Paper Interfaces 35
2.2.1 Language and Diagrams 35
2.2.2 Figure vs. Drawing 36
2.2.3 Microsurface and Macrospace 37
2.2.4 Empirical and Deduction 37
2.3 Elements of Teaching 38
2.3.1 Collaborative Scripts: Orienting Toward Productive Interactions 38
2.3.2 The Three Circles of Interaction with DGS 39
2.3.3 Orchestration: the Real Conditions 40
2.3.4 Integration 40
2.3.5 Indicators 41
2.3.6 Map: Teaching Geometry with our Activities 42
2.4 Conclusion 43

3 Research Questions and Method 45
3.1 Research Goals 45
3.2 Framework of the Design Space 46
3.2.1 Computational Model for Paper Input Devices 46
3.2.2 Temporal Structure of Paper Interfaces 48

3.2.3 A Map of Paper Interaction . . . 49
3.3 Method . . . 49
3.3.1 Design Based Research . . . 50
3.3.2 Participants . . . 51
3.3.3 Maps: the Iterations of our Activities . . . 52
3.3.4 Material . . . 54
3.4 Metroscope . . . 54
3.4.1 Metroscope: A Camera-Projector System . . . 55
3.4.2 Comparison with Other Augmenting Technologies . . . 56
3.4.3 Data Analysis . . . 58

4 Classifying Quadrilaterals 63
4.1 Description of the Activity . . . 64
4.1.1 Interface Elements . . . 64
4.1.2 Scenario . . . 66
4.1.3 Design Goals . . . 67
4.1.4 Position in the Framework . . . 68
4.2 Pilot Study . . . 68
4.2.1 Objectives . . . 69
4.2.2 Procedure . . . 70
4.2.3 Results . . . 70
4.2.4 Lessons for the Following Iterations . . . 73
4.3 In-Situ Study . . . 73
4.3.1 Objectives . . . 74
4.3.2 Procedure . . . 74
4.3.3 Results . . . 74
4.4 Conclusions . . . 82

5 Discovering the Protractor 87
5.1 Description of the Activity . . . 87
5.1.1 Interface Elements . . . 87
5.1.2 Scenario . . . 88
5.1.3 Design Goals . . . 89
5.1.4 Position in the Framework . . . 90
5.2 Controlled Study . . . 91
5.2.1 Objectives . . . 91
5.2.2 Procedure . . . 91
5.2.3 Results . . . 93
5.3 Complementary Study . . . 95
5.3.1 Objectives . . . 95
5.3.2 Procedure . . . 96
5.3.3 Results . . . 96
5.4 Conclusion . . . 100

6 Describing Angles 103
6.1 Description of the Activity . . . 104
6.1.1 Interface Elements . . . 104
6.1.2 Scenario . . . 106
6.1.3 Design Objectives . . . 107
6.1.4 Position in the Framework . . . 107
6.2 Preliminary Study . . . 107
6.3 Informal Study . . . 109
6.3.1 Objectives . . . 109
6.3.2 Procedure . . . 110
6.3.3 Results . . . 110
6.4 Formal Study . . . 115
6.4.1 Objectives . . . 115
6.4.2 Procedure . . . 116
6.4.3 Results . . . 117
6.5 Complementary Study . . . 128
6.5.1 Description of the changes . . . 128
6.5.2 Procedure . . . 130
6.5.3 Results . . . 131
6.6 Conclusion . . . 133

7 Exploring Symmetries 135
7.1 Description of the Demonstration Activity . . . 135
7.1.1 Generic Sheet . . . 135
7.1.2 Example Sheets . . . 136
7.1.3 Design Objectives . . . 138
7.1.4 Position in the Framework . . . 138
7.2 First Teacher Design . . . 138
7.2.1 Objectives . . . 139
7.2.2 Procedure . . . 139
7.2.3 Results . . . 139
7.3 Second Teacher Design . . . 146
7.3.1 Objectives . . . 146
7.3.2 Procedure . . . 147
7.3.3 Results . . . 148
7.4 Controlled Study . . . 156
7.4.1 Objectives . . . 157
7.4.2 Procedure . . . 157
7.4.3 Results . . . 159
7.5 Creative Activity . . . 168
7.6 Conclusion . . . 170

8 Cutting and Folding 173
8.1 StarrySheets 173
8.1.1 Principle 173
8.1.2 Implementation 174
8.1.3 Feedback 175
8.1.4 Application 176
8.1.5 Online Recording 176
8.1.6 Design Goals and Context 177
8.2 AngleHunt 179
8.2.1 Description of the Activity 179
8.2.2 Study 180
8.3 Sympliage 184
8.3.1 Description of the Activity 184
8.3.2 Study 185
8.4 Triangram 190
8.4.1 Description of the Activity 190
8.4.2 Study 190
8.5 Messangles 194
8.5.1 Description of the Activity 194
8.5.2 Study 195
8.6 Discussion 201
8.6.1 Validation from the Metroscope 202
8.6.2 Validation without the Metroscope 202
8.6.3 Mobility 203
8.6.4 Appropriations 203
8.6.5 Intuitiveness 204
8.6.6 Summary 204
8.7 Conclusion 204

9 Findings 207
9.1 How do Pupils Use a Paper Interface? 207
9.1.1 Presence 207
9.1.2 Side 209
9.1.3 Position 210
9.1.4 Rotation 212
9.1.5 Edges 213
9.1.6 Folds 214
9.1.7 Ink 214
9.1.8 Ephemeral and Persistent Properties 215
9.1.9 Creative Appropriation 216
9.2 How do Pupils Learn with a Paper Interface? 217
9.2.1 Collaboration 217

9.2.2 Exploration 221
9.2.3 Splitting as Pedagogical Design Principle 223
9.3 How do Teachers Use a Paper Interface? 224
9.3.1 Appropriation of the Functionalities 224
9.3.2 Instantiation of the Activity 226
9.3.3 Just In Time Adaptations 227
9.3.4 Orchestration 228
9.3.5 Integration 229
9.3.6 Continuum of Appropriation 229
9.4 Design Principles 230
9.5 Conclusion 232

10 Conclusion 233
10.1 Overview 233
10.2 Limitations 234
10.2.1 Experimental Methods 234
10.2.2 Generalisation 236
10.3 Perspectives 238

A Plan d’Étude Romand 239
A.1 Second Cycle 239
A.2 Third Cycle 240

B Planar Calibration 243

Bibliography 266

Curriculum Vitae 267

List of Figures

1 Presence of paper and computer in the classroom. 2

1.1 The continuum of paper interfaces link the use of paper as a support for documents to the use of paper for its tangible properties. 8

1.2 Digital pens (a) embed a camera which allows the pens to see the microscopic Anoto pattern printed on regular sheets of paper (b). The position is coded using the small displacement of dots relative to a regular grid. Source: Anoto Group AB 10

1.3 Liao et al. (2008) defined selection gestures (highlighted in red): crop marks (1), lasso (2), underline (3), margin bar (4). The stitching mark (5) goes across two different documents. The “Z” eraser gesture (6) deletes an unwanted crop mark. Source: (Liao et al., 2008) 12

1.4 Printers hanging on the walls of hospital, outputting patients’ constants. 15

1.5 Three documents are placed on the interactive surface. Circles denote their area of influence, and squares in the intersections of circles show how many other documents have the same keywords as the corresponding documents. Source: (Koike et al., 2000) 16

1.6 A flexible, paper-like surface containing infra-red markers, enabling the detection of how they are stacked (a), leafed through (b), or folded (c). Source: (Holman and Vertegaal, 2008) 18

1.7 A cube made of augmented paper, which allows a digital pen to outline regions to remove, which are then reflected on the digital model. Source: (Song et al., 2007) 19

1.8 An illustration of three categories of TUIs: interactive surface, token+constraint, and constructive assembly (from left to right). Source: (Ullmer et al., 2005) . . . 21

2.1 3 beans can not be arranged into several equal rows, but 10 can. This means that 3 is prime, and 10 is not prime. 28

2.2 The drawing (left) resulting from the commands (right) given to the *turtle*. Source: http://en.wikipedia.org/wiki/File:Ucblogo.png 28

2.3 Clairaut explaining in 1741: “Assume, for example, that BC turning around point B moves away from AB to get closer to BE; it is clear that while BC would turn, angle B would continually open; & in contrast angle C would become more and more narrow; which would first let us think that, in this case, the decreased size of angle C would be equal to the increase of angle B, & thus the sum of the three angles A, B, C, would remain the same, whatever the inclination of lines AC and BC, on the line AE.” Source: (Laborde, 2000) 29
2.4 Examples of DGS 30
2.5 The trace tool amplifying a phenomenon. Source: (Laborde, 2000) 30
2.6 Examples of TUI for geometry. 32
2.7 An augmented book to develop spatial abilities. 34
2.8 Augmented tabletop to develop spatial abilities. 35
2.9 The “proof without words” of the Pythagoras theorem. The diagram is simple, but its meaning is not obvious. Explaining this figure would defeat its purpose, and thus, is left as a riddle to the reader. 36
2.10 A square (left), when printed “on its edge” (middle) can be mistaken for a rhombus (right). 36
2.11 The three levels of the ArgueGraph script (Dillenbourg and Hong, 2008): (1) Students take a quiz. (2) Students are paired according to the quiz and discuss. (3) Pairs write a common questionnaire. (4) The teacher reviews and formalizes the answer of the whole class. (5) Each student writes a structured synthesis. . 39

3.1 The life time of our three categories of paper interaction. 48
3.2 The Metroscope, our camera-projector system, along with various types of objects which can be augmented on a table: sheets, cards, tools and wooden blocks. 55
3.3 Artefacts resulting from capturing the picture of a projection using a rolling shutter camera. The Metroscope uses a global shutter which avoids these artefacts. 56
3.4 A decision tree showing the technological alternatives to augment paper. 57
3.5 The Metroscope in less than ideal (lighting and humidity) conditions. Note the tissue on the top to absorb the water dripping from the roof of the tent. 58
3.6 Pictures of the experimental set-up at the front of a classroom (left) and the back (right). 60
3.7 Screenshot of the tool, with which a single time line (top left) or user defined bookmarks (top right) allow the user to synchronize the display of several multimedia data. In this screenshot, from top to bottom, and left to right: a series of snapshots taken from the camera of the Metroscope, a screencast of the projection by the Metroscope, a video from the left side of the room, the panoramic video from below the lamp, and a video from the right side of the room. 61

4.1 The components of Quads. (Left) Five cardboard quadrilaterals are classified into two groups on the instruction sheet, a card displays the measure of the angles of a rectangle. (Right) The feedback card displays the validation text. . . 63

4.2 A QR Code of the url http://short.epfl.ch/quads which points to a video demonstration of Quads. 64
4.3 The numbered cardboard quadrilaterals used in the study. 64
4.4 The English translation of the three tool cards (left, with a green outline) and the feedback card (right, with an orange outline) 65
4.5 An example of the information displayed by each of the cards. 65
4.6 The lay-out of pages 1 (left), 6 (middle) and 7 (right) of the leaflet. 66
4.7 A square (left), when printed "on its edge" (middle) can be mistaken for a rhombus (right). 68
4.8 The five quadrilaterals used in the pilot study. 70
4.9 The classifications proposed by the pupils out of the 16 possible classifications. Each node represents one possible classification. Their colour indicates how often they were proposed by a group. Each classification is linked to the classifications that are possible after changing one element. 71
4.10 The single display of the cards shown in Figure 4.5 can be combined into a "test bench" (right), where cardboard shapes are placed to display all of their characteristics at the same time. 72
4.11 The physical set-up of the experiment: three pupils sitting under the Metroscope, in a spare room of their school. 73
4.12 Expected classifications for the first six pages of the study. The criteria for each classification is shown on the right. The classification of page 6 is not a partition: a square is in both the rectangle and rhombus classes. 75
4.13 Time spent by each group on each activity. 76
4.14 The traces of the various pieces of paper. 78
4.15 Average speeds of the different kinds of paper artefacts. 79
4.16 The y-coordinate of the position of the various elements of the paper interface as a function of time, for one group and one exercise. This allows us to visualize the position of the cardboard shapes (blue and orange, depending on the expected classification) relative to the areas on the exercise sheet (delimited in black). . . 80
4.17 Number of vertical transitions per exercise sheet. 81
4.18 Number of vertical transitions per exercise sheet. 81
4.19 The positions of the pupils. 82
4.20 The x-coordinate of the position of the various elements of the paper interface as a function of time, rotated horizontally. This shows the position of the pieces of paper relative to the first exercise sheet (see Figure 4.19b) 83
4.21 Time spent on each section of the collaboration surface by the pieces of paper. 84

5.1 The various elements of Angoli: the two control cards and two of the angle measure cards. The one on the left is flipped and shows the measure of the angle constructed with the blue control card (70°). 87
5.2 A QR Code of the url http://short.epfl.ch/angoli which points to a video demonstration of Angoli. 88

5.3 The lines building 138° clockwise and 42° counter-clockwise overlap. This relation correspond to the fact that supplementary angles sum to 180°. 90
5.4 One exercise representing a protractor used on the pre- and post-test. 92
5.5 Effect of the condition on the learning gain. 93
5.6 Visualization of the feedback given to one of the groups in the continuous condition. The value of the angles built, together with the angle expected and their supplements, are plotted against time. The various values are plotted only when a feedback was shown. 94
5.7 Translated transcript of an episode of conflict between pupils A,B,C and D in Angoli. 99

6.1 The components of the interface of SpaceJunk. 103
6.2 The observer team (left) measures the orientation to give to the laser, and communicates it to the controller team (right). 104
6.3 A QR Code of the url http://short.epfl.ch/spacejunk which points to a video demonstration of SpaceJunk. 104
6.4 Laser Sheets. 105
6.5 Satellite View Sheets. 106
6.6 Two used ammunitions on the left, and a stack of remaining ammunitions (the first of which has already been annotated on its back). The green augmentations mean that the shot was a hit, and the red augmentation means that the shot was a miss. 106
6.7 Typical group dynamics depending on the size of the group. 109
6.8 The first class of the study listening to the explanation by one experimenter. Photo Alain Herzog 111
6.9 Pupils sticking their feedback note on the whiteboard. Photo Alain Herzog . . . 112
6.10 A sample of the various formalisms invented to describe an angle. 113
6.11 Histogram of the final scores of the 36 groups. 115
6.12 The sub-teams on the left are switching their roles, and the group on the right has just hit another satellite. 116
6.13 Evolution of the scores. Each marker is a shot. The line shows the score (i.e. the number of hit targets) as a function of time (in minutes). Markers in the middle of a segment correspond to missed shots; markers on a corner correspond to hits. 118
6.14 The angles shot by each group as a function of the expected target angle. Correctness can be estimated by the closeness of the line $y = x$ (light green, from bottom left to top right). If a shot is close to the line of equation $y = 180 - x$ (light grey, from top left to bottom right), it means that the angle is the supplement of the expected angle. An expected angle outside the $[0, 180]$ range means that the satellite was not attainable with the laser used for the shot. A vertical line with a lot of points signifies a lot of errors on a specific target. The dashed line shows in which order the shots were made (two consecutive shots are linked by a dashed line). 119

6.15 Examples of a correctly (left) and incorrectly (right) aligned protractor. The angle is outlined in red. 120
6.16 Close-up of a satellite view where multiple lines towards the satellite on top show that the origin of the angle was not immediately understood. 120
6.17 Positions of the laser control relative to the laser sheet (represented as a grey rectangle) for each group. The arc gives the scale of the protractor. The colours of the rainbow are used as an indication of time (from red to purple). 124
6.18 Exploitation of the traces of manipulation. 125
6.19 Distribution of the radius of the ellipses approximating the movements of the pupils with the control cards relative to the laser sheet. 126
6.20 A visualization of the model of the trajectories followed by the control card. The protractor and the sheet (in gray) are drawn for reference. 127
6.21 Three zones defined to refine the model of the trajectories made by the control card. 128
6.22 An improvement of the interface for the activity shown in Figure 6.1. The laser (1) is at a fixed place on the desktop. It is triggered by flipping a card (2), indicating its readiness (3). The selection of the laser happens by rotating another card (4). The two sub-teams communicate with a full size sheet (5). 129
6.23 New elements of the interface. 130
6.24 Evolution of the scores. Each marker is a shot. The line shows the score (i.e. the number of hit targets) as a function of time (in minutes). Markers in the middle of a segment correspond to missed shots; markers on a corner correspond to hits. 132
6.25 Distribution of the radius of the ellipses approximating the movements of the pupils with the control cards relative to the laser sheet. 133

7.1 A QR Code of the url http://short.epfl.ch/kaleidoscope which points to a video demonstration of Kaleidoscope. 136
7.2 A view from the top of the Kaleidoscope sheets, with a schema indicating the axes reflecting the input area. The greyer sides of the sheets are actually the projection of the other side of the axes. Note that anything in the input area is reprojected, such as fingers, not only the ink on the paper. 136
7.3 From left to right: a vertical axis, a horizontal axis, an oblique axis, plus-axes, cross-axes, and star-axes. The grey k's illustrate how the Metroscope would augment the black k's on each sheet. 137
7.4 The sheets of the sample activity. 137
7.5 An overview of pupils manipulating outlines placed on paper cut-outs. The pupils align the symmetry axis they find on the outlines with the symmetry axis of the symmetry sheets. If the projection of the outline matches their outline, their axis is correct. 142
7.6 Figures to complete by reflecting the outlines across the axes. The Metroscope projects the symmetric projections of the pre-printed figure and the pupils' answer. 143
7.7 The teacher explaining the exercise to the class. 143

7.8 The teacher using the +axis to show a vertical symmetry axis on a pentagon. . . 144

7.9 A pupil checking the symmetry axis of his lunch during the break. 146

7.10 The lamp in the back of the classroom. 147

7.11 The sheets of the booklet for the fourth graders. 149

7.12 The sheets of the booklet for the third graders. 150

7.13 Left: the input for the generation of a mandala, i.e. a drawing in a quadrant of a Kaleidoscope sheet. Right: the output of the drawing after extraction with a scanner, which could be sent back to the teacher to distribute to the class. . . . 151

7.14 From left to right: the output of one group of fourth graders for steps 1, 2, and 3, and the resulting mandala rendered from a scan of the third sheet. 156

7.15 Set of symmetry exercises, each with a different axis type. From left to right: vertical axis, horizontal axis, +axes, and ×axes 158

7.16 A pupil placing triangles by *taping* pre-cut, red, triangle paper-pieces. Note the band of reinforcement rings between the sheet and the pencil case. 158

7.17 Number of placed triangles based on the worksheet axis, for each session. The first row corresponds to the morning session, and the second row to the afternoon. For each type of exercise (vertical, horizontal, perpendicular, and diagonal axes), the histogram shows how many of the triangles were placed. The maximum number of triangles is 5 for the single axis, and 15 for the double axes. 160

7.18 Percentage of incorrectly placed triangles based on the worksheet axis. 161

7.19 A case of superfluous triangles added by the pupil. 162

7.20 Order effects based on ‘draw first’ or ‘tape first’ conditions. The left side of each graph represents the morning session and the right side represents the afternoon session, after switching modes. 163

7.21 Results of the questionnaire. 163

7.22 Set-up of the Metroscope in the classroom. 165

7.23 The ×axes exercise sheet of a pupil who approximated the Metroscope by folding.166

7.24 Average duration of time spent at the Metroscope per number of visits *per pupil* 167

7.25 (Left) Triangle placements depending on whether the Metroscope was used. (Right) *Correct* triangle placements depending on whether the Metroscope was used. 168

7.26 The two phases of augmented creation of paper cut-outs. 169

7.27 A cut-out in space. 169

7.28 A kaleidoscope sheet limited in time. The augmentation disappears when the bar (on the right) is empty. 170

8.1 Detecting the deformation of a set of blobs using StarrySheets. 174

8.2 Feedback on the relative position of two StarrySheets. Red lines correspond to errors. In this case, the line can be interpreted. The upper cluster looks like a whirl, corresponding to the rotation between the expected setting and the detected one. The lower cluster seems like attracted by the upper cluster, corresponding to the displacement between the expected relative position and the detected one. 176
8.3 Checking the cut of a StarrySheet. 177
8.4 Checking the merge of the edges of a StarrySheets. 177
8.5 Checking the overlap of StarrySheets. 178
8.6 Checking the fold of StarrySheets. 178
8.7 The map to be recomposed by pupils. 179
8.8 Points linking pieces of maps (left) according to the instruction sheet (right). The piece of map in the middle contains the points E and F. The piece of map on the bottom left contains the point G. The instruction sheets defines the angle FEG as 85° wide and gives a schema of it. 180
8.9 The set-up for AngleHunt. 181
8.10 Examples of multiple hands manipulations. Photo Himanshu Verma 182
8.11 Different group dynamics. 183
8.12 A teacher watching and intervening when necessary. Photo Alain Herzog 184
8.13 Front and back of the sheets used in Sympliage. 184
8.14 Pupil checking their folds. 185
8.15 The Metroscope, at the back of the classroom. 186
8.16 A common misunderstanding. A0 was intended to name a corner marked by a cross (see left), but pupils often thought that A0 was a cell of the grid (see right). 187
8.17 Examples of appropriation. 188
8.18 The expected solution (left) could be compared visually with a pupil's solution (partial, right). 189
8.19 Two steps of Triangram. 190
8.20 A pupil having difficulties bringing his assembly to the Metroscope. 191
8.21 A board helped carry the pieces of paper. 192
8.22 On the right: the range of feedback. From bottom to top: green dots (correct), red points (small imprecision), a long red line (detection artefact) and red lines (bigger imprecision). 193
8.23 Assembling pieces of StarrySheets in Messangles. 195
8.24 Creating a Messangles exercise. 195
8.25 Two strips of tape fixing the assembly of triangles at once, which is faster than individually taping each pair of neighbouring triangles. 197
8.26 Two pupils showing their StarrySheet at the same time, resulting in a falsely negative feedback. 198
8.27 Dynamics of the classroom. 199
8.28 A pupil measuring the angle of a paper triangle. 200

9.1 The triangles were placed assuming a vertical axis and the sheet was configured for a horizontal reflection. To fix this, Mia cut out the fiducial markers to align the expected and projected axes. 217
9.2 An overview of pupils manipulating outlines placed on paper cut-outs. They align the symmetry axis they find on the outlines with the symmetry axis of the symmetry sheets. If the projection of the outline matches their outline, their axis is correct. 225

B.1 The displacement between the expected position and the average of the detected positions, with the planar calibration (left) and the spatial calibration (right). The point on the tip of the arrow was expected at its base. For readability, the length of the arrows are scaled up by 10. 245
B.2 Distribution of the distance (in millimetre) between the expected position of the circles, and the actually detected position, for the spatial calibration (left) and the planar calibration (right). 246

List of Tables

3.1 Examples of actions covering the space defined by the framework. 49
3.2 Summary of the groups of subjects who tested our designs. 53

4.1 Position of Quads in the framework 69
4.2 Possible extensions of Quads in the framework 86

5.1 Position of Angoli in the framework 90
5.2 Order of the angles to be built. (CCW and CW respectively stand for counterclockwise and clockwise). 92
5.3 Order of the angles to be built. (CCW and CW respectively stand for counterclockwise and clockwise). 96

6.1 Position of SpaceJunk in the framework. 108
6.2 Statistical models of the influence of the sheet on the trajectories of the control card. A model defining three areas (in, on the border of, or outside of the sheet) explains best the fact that the elliptic trajectories of the controller card are less and less circular as the movements are wide. 127
6.3 Differences between the two versions of SpaceJunk. 128
6.4 Position of the new version of SpaceJunk in the framework 130

7.1 Position of Kaleidoscope in the Framework 138
7.2 Differences between resulting designs. 151
7.3 Position of Kaleidoscope in the framework 159

8.1 Position of activities based on StarrySheets in the framework 201
8.2 Overview of the comparison of the various AngleHunt, Sympliage, Triangram, and Messangles 204

9.1 Scattering of the interfaces for the various activities. 219

Introduction

There is a newcomer at school. The children like him already and they play with him on evenings and weekends. The teachers would love for the children to not only play but also work with the newcomer, because he has a lot of potential. He learns quickly and he is especially good at mathematics (his parents are great mathematicians). Unfortunately, he understands teachers and pupils very poorly. Teachers need to make a huge effort to have him do what they expect. The pupils seldom ask him for help on their homework, because he cannot read properly. As a result, in most classrooms, the newcomer spends most of his time ignored in the back of the classroom, or even worse: isolated in a special classroom with his fellows.

This is not sad – the newcomer is just a computer after all – but its potential to help teachers and pupils is underused. Most of the pupils' work involves paper, but it cannot be easily incorporated in activities involving computers. Conversely, most activities involving computers do not incorporate other tools. There is a very limited compatibility between the two technologies (paper and computers). However, they are not mutually exclusive. In fact, the seminal work of Wellner (1993) showed that paper and computers could complement each other by integrating the work on paper documents and on digital documents into one digital desk. With such a digital desk, paper, pen, and fingers act as direct interfaces to control the computer, without alien devices like keyboards or mice.

The paper interfaces we will discuss in this dissertation rely on a camera to detect the physical state and the content of pieces of paper, which will be used as input to a computer, which can output information directly on or next to the elements of the interface, via a projector.

Paper interfaces are a promising vector for the integration of computers in the classroom, where paper is ubiquitous (see Figure 1). There is paper hung on almost every vertical surface of the class (walls, closets, etc) or even on a thread suspended in the air. The pupils read textbooks, and complete exercise sheets printed by the teachers. They take their notebooks to do their homework and carry home their drawings to show their parents. Pietrzak et al. (2010) showed how to integrate multimedia content into this work flow without disrupting it. They used a digital pen to link the annotations of the students to digital objects and vice versa.

However, there is more to paper than just reading and writing. Many have predicted that

Figure 1: Presence of paper and computer in the classroom.

digital technologies would lead to the disappearance of paper. While some types of paper documents are disappearing and publishers are struggling, we are still far from the once predicted “paperless office” (Sellen and Harper, 2003). Our interest for paper is not a form of nostalgia. We explore the pedagogical advantages of paper in terms of functionalities in a very pragmatic approach. For example, one of the advantages is that the thickness of a document helps navigate it. Another is that the spatial disposition supports the linking of several pieces of information. Furthermore, paper is a body with physical properties: it can be flipped, moved, cut, folded, etc. These properties have been exploited in the long history of the use of paper and have now crystalized into well established practices. Many of these practices are involved in school education, and even the subject of teaching; for example, pupils learn how to use an agenda, transparent paper, etc. Practices involving paper are taught so early that they can be considered a part of our every day life.

Position

This dissertation builds on the integration of paper in everyday life, because we aim at reusing the skills related to paper to augment it with valuable capacities from computers. It is easy to restrict the augmentation of paper to truly valuable ones, because interaction with paper by itself is already very rich. We already mentioned how paper facilitates reading and writing, but paper appears in many other contexts: painting canvases, presenting posters, revising flash cards, sending postcards, playing cards, solving puzzles, folding origami, paying with banknotes, wrapping packages, etc.

Our hypothesis is that these skills can be largely exploited in an interface without hurting its integration in an environment – in our case, the classroom. We want to investigate this exploitation by going beyond paper as a support for information, and move the focus towards paper as a physical, tangible object. In this perspective, geometry at primary schools provides a rich testing ground. Paper is a perfect fit for fixed 2D shapes to be used as tangibles. There is a direct match between the translation and rotation of a geometrical 2D shape and its paper representation. Axial symmetry can also be directly manipulated by folding or flipping the paper. The shapes can also be cut to recompose them without changing the area (which is perfect to introduce area formulas, for example). The shapes can also be superposed for comparisons.

Overview

This dissertation is thus an exploration of the design space offered by paper interfaces in the context of geometry education at primary schools. It is articulated around three parts. The main part of the exploration corresponds to the presentation of five series of activities. These activities are different from each other, in order to cover the most from the domain of paper interfaces for geometry education.

First, in chapters 1, 2, and 3, we situate the activities on the various topics that will be part of the explorations, by placing them on *maps*. Chapter 1 starts by reviewing previous explorations, i.e. the work related to paper interfaces. Chapter 2 frames the context of the exploration by covering concepts and previous work in geometry education. Chapter 3 plans the exploration by describing our method and the framework to guide the creation of paper interfaces.

Second, we report on the steps of our exploration. Chapters 4 to 8 each describe an area of the design space that we explored with activities and studies. Some are more successful than others, but all carry important lessons. The subjects of the activities are the following:

- Quads (Chapter 4) introduces the criteria to classify quadrilaterals.
- Angoli (Chapter 5) aims at using and mastering the protractor.
- SpaceJunk (Chapter 6) shows the pupils how to use a protractor to communicate angles.
- Kaleidoscope (Chapter 7) lets pupils explore symmetries.
- StarrySheets (Chapter 8) are actually the base component of four activities: AngleHunt, Messangles, Sympliage, and Triangram. The first two address angle measurements, the last two address symmetries.

All these exotic names are not easy to remember. The reader should have been provided a bookmark, which contains a quick, visual reference to the activities. If this is not the case, the reader can tear up one of the spare bookmarks on the next page.

Third, we collect the fruits of our exploration. In doing so, Chapter 9 reassembles the lessons learned from each area of the design space into a more global overview. The bookmark will be particularly useful for this chapter, which frequently mentions the many studies done with the different activities. Chapter 10 concludes our findings and outlines the limitations of our explorations.

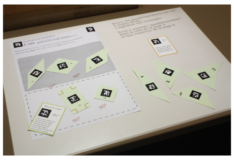

Ch. 4
Quads

Classifying
Quadrilateras

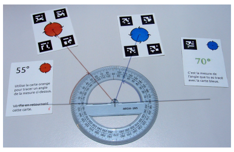

Ch. 5
Angoli

Discovering the
Protractor

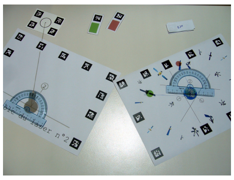

Ch. 6
SpaceJunk

Communicating
Angles

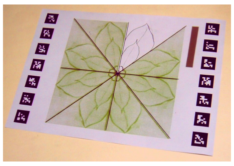

Ch. 7
Kaleidoscope

Exploring
Symmetries

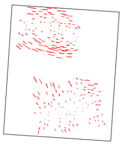

Ch. 8
StarrySheets

Folding and
Cutting

AngleHunt

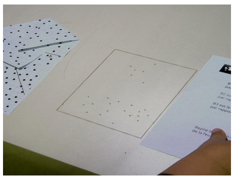

Sympliage

Triangram

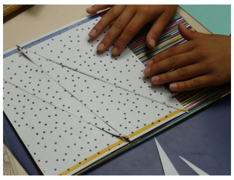

Messangles

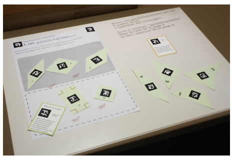

Ch. 4
Quads

Classifying
Quadrilateras

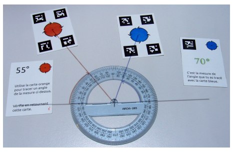

Ch. 5
Angoli

Discovering the
Protractor

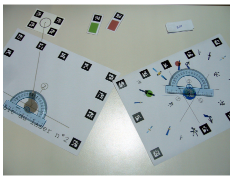

Ch. 6
SpaceJunk

Communicating
Angles

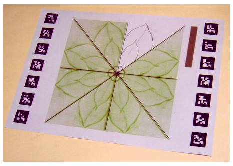

Ch. 7
Kaleidoscope

Exploring
Symmetries

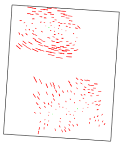

Ch. 8
StarrySheets

Folding and
Cutting

AngleHunt

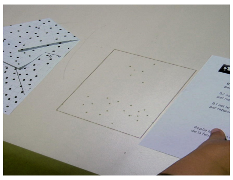

Sympliage

Triangram

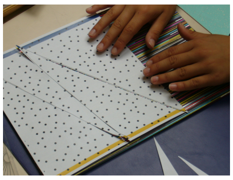

Messangles

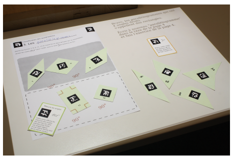

Ch. 4
Quads

Classifying
Quadrilateras

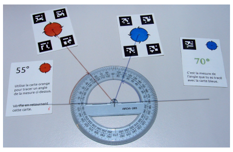

Ch. 5
Angoli

Discovering the
Protractor

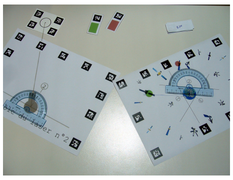

Ch. 6
SpaceJunk

Communicating
Angles

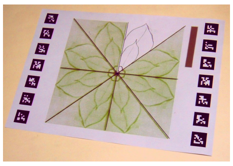

Ch. 7
Kaleidoscope

Exploring
Symmetries

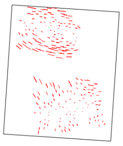

Ch. 8
StarrySheets

Folding and
Cutting

AngleHunt

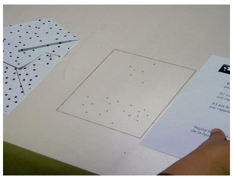

Sympliage

Triangram

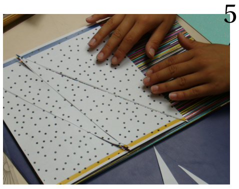

Messangles

1 Paper Interfaces

I am holding the morning newspaper in my hand, reading an interesting article. My eyes scan the page rapidly, I am focusing on the content of what I read. Paper offers a nice and handy interface to this reading activity but I am not paying much attention to these specific affordances of paper.

The day after, I find the old and now obsolete newspaper and decide to turn it into something else: a paper airplane. I fold it to make a perfect glider. What counts now is the material properties of the newspaper, its tangibility. I am not paying attention to its obsolete content anymore.

This dissertation deals with paper interfaces and their relevance for a particular domain, the learning of geometry in primary school. In this first chapter, we review former work on paper interfaces. Taken as a whole, this former research may look very heterogeneous and unstructured. What we argue in this first chapter is that this apparent heterogeneity comes from the fact that some work focuses primarily on paper interfaces as interesting content holders and others as tangible interfaces to control other systems.

A closer look will actually reveal that the research cannot simply be divided into two groups. A continuum of usages of paper interfaces exist from situations in which content is important to situations in which tangibility is the key.

This continuity is mainly due to the fact that paper documents cannot, generally, be considered in isolation. Rather, they often constitute systems, structured sets of physical objects. A book is a structured set of two dimensional sheets, organized sequentially in a volume. Its physical nature affords specific actions like flipping pages in particular ways. Content is still primary, but tangibility is more important than that of an isolated sheet of paper. Consider now a set of playing cards. The way they can be organized in space, stacked, and exchanged, determines the state of the game. The content of each card is atomic (much smaller than on a newspaper article). What counts is the ways they are manipulated and their position in the physical space. Finally, imagine a sheet of paper that could be folded in particular ways to control a computer simulation (e.g. a protein folding application). Paper, in this case, is not at all a content holder; it is its tangibility that is important.

Figure 1.1 gives an overview of this continuum, in which we propose to distinguish four classes:

- Class I – augmented paper documents. Only the content supported by paper is used. The actions applied on paper include reading, looking, writing, drawing, printing, skimming, filtering annotations, etc.

- Class II – structured sets of augmented paper documents. This class is one step toward taking advantage of the tangibility of paper and takes into account the fact that paper can be used to organize the content in space. The actions applied on paper include holding, placing, rotating, navigating, dispatching, assembling, turning pages, stacking, splitting, etc.

- Class III – tangible paper. This class is even one step closer to making use of the tangibility. The content these papers hold is less important (typically, it would just include an identifier). The actions applied on paper include flip, cut, fold, pinch, scratch.

- Class IV – tangible paper controllers. Paper is a material with which one can build tangible interfaces. The actions applied on paper include placing and orienting in space with precision, rubbing, hovering, shaking, grabbing, collecting, piling up, tearing up, etc.

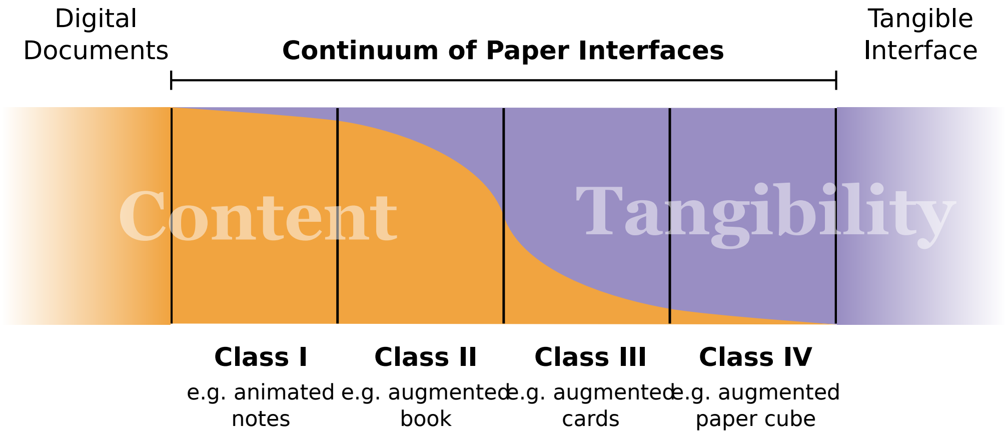

Figure 1.1: The continuum of paper interfaces link the use of paper as a support for documents to the use of paper for its tangible properties.

The rest of the chapter presents examples from each of these classes. Some interfaces stretch between more than one zone. In this case, we place the interface according to its most salient aspect. Interfaces on the border between two zones are used to amplify the details distinguishing the two zones.

1.1 Class I – Augmented Paper Documents

1.1.1 The DigitalDesk

Newman and Wellner (1992) and Wellner (1993) pioneered the use of paper as an interface to bridge the paper-digital divide. Based on the observation that the cultural form of the desk was a way to reduce the effort needed to learn a new interface, they reversed the logic by giving physical desktops electronic properties. This minimizes the learning curve of the interface, since it is already known.

The DigitalDesk used a camera to track documents and fingers, and a projector to augment them digitally from above. For example, DigitalDesk Calculator allows a user to input numbers from a sheet in a calculation without having to type them.

1.1.2 Synchronizing Paper and Digital Documents

Wellner (1993) described *paper pushing* and *pixel pushing* to distinguish the clerical work on physical desks and digital desks, which are often strongly related, yet isolated. One first step to prevent against this isolation is to allow the upload of paper documents to the digital world. Protofoil (Rao et al., 1994) aims at doing so with scanned and faxed documents, where a special cover sheet allows for the input of metadata about the document. This way, the digital version can be sorted in the system without another manual step after the scan.

The Protofoil coversheet was a *Paper User Interface* to the XAX document server (Johnson et al., 1993). The coversheet is a form filled out by the user to inform the server as to who is receiving the scan how to archive the document, e.g. category of filing, other recipients, or Optical Character Recognition (OCR) options. Ito et al. (1999) even proposed mounting scanners into desks, and scan *Memopads* where users would write their annotations to attach to the document.

1.1.3 Digital Pens

However, the DigitalDesk reduces one of the strong advantages of paper: its mobility. The augmentation can happen only on the DigitalDesk. Mackay et al. (1995) explains how Ariel, a system augmenting engineering drawing, became a portable version of the DigitalDesk: paper is marked with red dots, which allows a camera to situate a document and extract annotations made in the field. Mobility is critical for synchronizing annotations and data, such as information about location or photos in field biology research (Yeh et al., 2006).

Digital pens are another way of linking content in its digital and physical form. For example, Piper et al. (2013) used them to attach audio recording to photos. They consist of (slightly bigger than usual) pens, which embed a camera that sees the immediate neighbourhood of the pen tip. They are meant to be used with paper preprocessed with a printed microscopic pattern (see Figure 1.2), which allows the computer to precisely situate the position of the pen tip on the page. The pattern is almost invisible to the human eye and can be produced with regular printers.

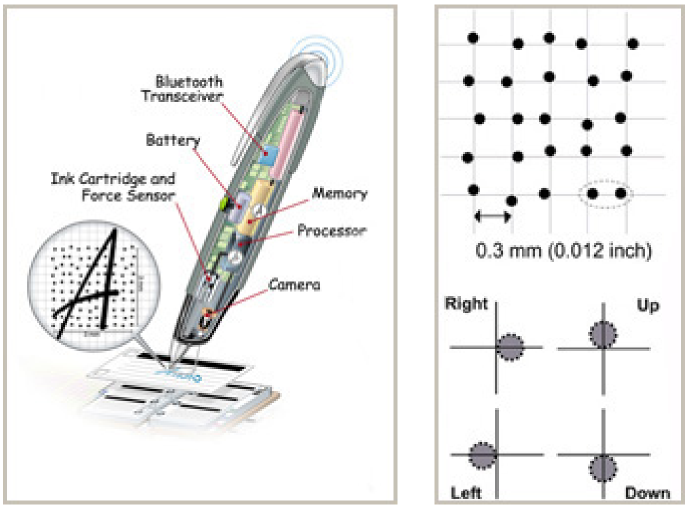

(a) A digital pen

(b) The Anoto pattern

Figure 1.2: Digital pens (a) embed a camera which allows the pens to see the microscopic Anoto pattern printed on regular sheets of paper (b). The position is coded using the small displacement of dots relative to a regular grid. Source: Anoto Group AB

Many approaches use such an augmented paper, and several frameworks have appeared already. PADD (Guimbretiere, 2003) is a cyclic process based on printers and digital pens that synchronizes annotations on digital and physical versions; digital documents are printed with the Anoto pattern when paper affordances are needed, and annotations made with a digital pen are added to the digital document. Steimle et al. (2009) proposed a similar framework to share annotations with other users.

Yeh et al. (2008) proposed an event-driven toolkit to develop paper applications based on digital pens. Letras (Heinrichs et al., 2010) is another development toolkit aimed at mobile devices to display augmentations. Norrie et al. (2006) described a framework to support the rapid development of paper applications, with models for documents and a server to synchronize various document-based services. It comes with an authoring tool (Weibel, 2009), and a framework to recognise pen gestures, i.e. strokes which correspond to commands (Signer et al., 2007). Steimle (2009) proposed six interaction techniques based on digital pens (inking and clicking) and augmented paper (moving, altering shapes, combining, and associating) to define three kinds of actions: annotating, linking, and tagging.

1.1.4 Mobility of the Feedback

It is also possible to replace the projector and screen of the DigitalDesk with something more mobile, allowing feedback in the field. For example, Liao et al. (2006) explored different types of feedback that could be embedded in the pen: an LED, vibrations, and voice and sounds. They studied which ones are adapted for three types of feedback (discovery, status, or task). Song et al. (2009) incorporated a miniature projector on a digital pen. Song et al. (2010) then separated the projection from the pen to a spatially aware projector, which allowed for bi-manual interactions with the augmentation.

1.1.5 Augmented Writing

Digital pens also call for augmented writing. For example, Guimbretiere (2003) addressed proof-editing. Weibel et al. (2008) showed how to use pen gestures to issue proof-editing commands (insert, delete, replace, move) to be reflected in a word processor. On a related note, Olberding and Steimle (2010) investigated how to add erasing capabilities into digital pens.

The applications of augmented writing are not at all limited to proof editing. Garcia et al. (2012) linked paper to a music creation software to let composers explore musical ideas using pen gestures. Tsandilas et al. (2009) let musicians define their own gestures to create their own musical language. Lee et al. (2008) made a purely pen-based tutoring system that scaffolded students in statics problems with preprinted sheets and audio feedback. Finally, the pen is not the only tool that can be augmented: Flagg and Rehg (2006) showed how painters can be helped by a computer color creation and painting subtask division.

1.1.6 Augmented Active Reading

Paper is a less distracting interface than a computer: Oviatt et al. (2006) showed that interfaces that depart from traditional tools induce an extraneous cognitive load and that the closer to traditional pen and paper the system is, the better the performance of students. In other words, digital stylus yielded better results than pen tablets, which in turn outperformed graphical tablets. Paper is thus an efficient choice for active reading. Arai et al. (1997) used the metaphor of a highlighter pen to mark items in the digital document from the paper version. This could be used to build a thesaurus, retrieve streets from a map, or access an encyclopedia.

XLibris (Crossen et al., 2001) aggregated and linked information from the digital world, e.g. for connecting readers of a same author. Chong and Kawsar (2010) identified the search function in books as the most missing feature. They used a smart phone recognizing the ISBN code of the book so that the search text entered by the user could be applied to the right document. Letondal and Mackay (2009) showed how the reading of scientific data could be augmented by the display of computations on the pen's LCD screen, and how annotations of interesting findings can be transferred back to the computer. Liao et al. (2008) used natural pen gestures for selection (circle, underline, vertical bar in the margin, crop marks), and other gestures for cross document actions (like copy, paste, delete, stitch, or hyperlink (see Figure 1.3). Erol et al. (2008) proposed a method for recognizing a portion of text of documents with smart phones, which allows the use of such devices instead of digital pens.

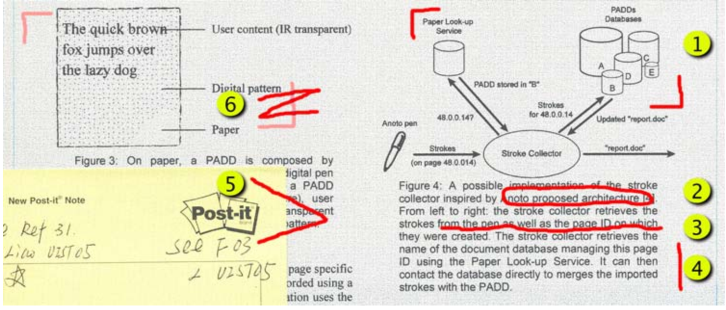

Figure 4: A possible implementation of the stroke collector inspired by Anoto proposed architecture [4]. From left to right: the stroke collector retrieves the strokes from the pen as well as the page ID on which they were created. The stroke collector retrieves the name of the document database managing this page ID using the Paper Look-up Service. It can then contact the database directly to merges the imported strokes with the PADD.

Figure 1.3: Liao et al. (2008) defined selection gestures (highlighted in red): crop marks (1), lasso (2), underline (3), margin bar (4). The stitching mark (5) goes across two different documents. The "Z" eraser gesture (6) deletes an unwanted crop mark. Source: (Liao et al., 2008)

1.1.7 Augmented Reading

More simply, reading on paper can be augmented with computer services (Dymetman and Copperman, 1998) and multimedia content, such as animated anaglyphs[1] (Chehimi, 2010).

Grasso et al. (2000) showed how to augment tourist information (guides, maps, flyers, and newspapers) by linking them to a website and large screen display in order to connect tourists with services, such as contextual information, ratings, and comments. Replacing the scanners with digital pens, Norrie and Signer (2005) proposed a similar, but real-time and mobile, solution for a large festival.

1.1.8 Augmented Collaboration

Since content on paper is easily seen (or even modified) by several people at the same time, paper supports collaborative tasks easily. Wellner (1993) proposed a DoubleDigitalDesk for remote collaboration, like Tang and Minneman (1991), who proposed a sketchbook augmented by a similar system, which projected the image of the interaction surface of one user on the interaction surface of the other user. Lange et al. (1998) described an immersive environment to create design requirements, including sticky notes and data reports, linked to augmented content. Haller et al. (2010) also supported meetings, integrating paper documents. Malmborg et al. (2007) used large sheets of paper collaboratively annotated as support for co-located design. Hong et al. (1999) imagined how cover sheets in printers could be used as bulletin boards.

---

[1] Anaglyph: stereoscopic 3D effect resulting from glasses where one lens is blue, and the other one red.

1.1.9 Borderline: Augmented Notebooks

Notebooks, as a support for augmented collaboration (Mackay et al., 2002; Tabard et al., 2008; Pietrzak et al., 2010; Malacria et al., 2011; Portocarrero et al., 2010), are on the border between the interfaces that see paper as a document, verses paper as document sheets. The latter classification would be justified by the fact that a notebook has several pages that can be flipped. However, notebooks often group pages for practical reasons, i.e. to handle a large number of pages, not because the aggregation or sequence has meaning. We contrast this with augmented books in the next section.

1.2 Class II – Structured Sets of Augmented Paper Documents

1.2.1 Augmented Books

Augmented books are on the other side of the border between Class I and Class II, compared to augmented notebooks. Flipping the page is indeed an important aspect of the interaction with a book. Books are designed to organize their content spatially. This can be contrasted with notebooks, where the content is simply accumulated on a contingency basis. Harrison et al. (1998) explored other affordances of paper for selecting a page: squeezing, holding, and tilting pages.

Back et al. (2001) focused on truly augmenting the book reading experience, i.e. on improving it without changing it. The set-up included a comfortable chair, high quality sounds, and invisible markers in a book, which played music and sound effects to accompany the reading experience. Ucelli et al. (2005) used a screen on the side to augment a book guiding children in the composition of colors. Scherrer et al. (2008) used a similar system, focusing on the precision of the augmentation. Koike et al. (2000) augmented instructional material with projections from above the desk. Asai et al. (2005) used a head mounted display.

Billinghurst et al. (2001b) showed how a book can be the common element of interfaces using various display technologies. They claimed that the book can be used as an anchor between reality, augmented reality, and virtual reality. For example, head mounted displays are more immersive than screens, but screens are more akin to collaboration in small groups, and a book can be used in both cases.

1.2.2 Augmented Education

Augmented books are easily applied to education (Billinghurst et al., 2001b; Ucelli et al., 2005; Martin-Gutiérrez et al., 2010; Koike et al., 2000; Asai et al., 2005), but they are not the only kind of document that can take advantage of tangibility. Zufferey et al. (2009) used paper forms to control pedagogical simulations of warehouses, which allows for the reorganization of the tabletop. Cuendet et al. (2011) illustrated how a sequence of sheets is mapped to a pedagogical sequence in exercises for carpenter training.

1.2.3 Multimedia Scrapbooks

Scrapbooks are another popular way to use the organisation of augmented paper in terms of pages. Sound can be attached to scrapbooks, which are combined with a playback device (Stifelman, 1996). Conversely, the scrapbook can be used as an interface to access attached sound and videos on a PDA (Klemmer et al., 2003), or to produce a virtual scrapbook on a website (West et al., 2007).

1.2.4 Augmented Prototyping

Paper sheets also have the property of being cheap and easily (re)producible, which can be used for prototyping. Mackay and Pagani (1994) split storyboard elements into individual pieces of paper that can be rearranged easily. Doering et al. (2009) found that paper is a good fit for the creative process consisting of collecting, relating, creating, and donating. This allowed to take advantage of the expressive power of digital content with the ease of use of the real objects.

Bähr et al. (2010) used augmented paper in order to design the look and flow of a mobile user interface. To do so, they print the image of a smart phone with an empty screen on several pieces of paper. These pieces of paper are tagged with a fiducial marker, which allows the computer system to identify each screen. This way, the designer of a smart phone interface can specify each step of the interaction on individual pieces of paper, and link them in a workflow.

Beyond prototyping, paper interfaces can be used to design paper interfaces. For example, the toolkit proposed by Piper et al. (2012) allowed end users to design their own paper applications, even if they had no programming knowledge.

1.2.5 Augmented Collaboration

In comparison to simple interfaces of Class I, those of Class II have additional benefits for augmented collaboration: their position relative to each other. This is especially powerful to organize ideas publicly, as is often done with sticky notes (Mistry et al., 2009; Klemmer et al., 2001) for physical presence notification (Moran et al., 1999), website design (Klemmer et al., 2008; Ljungstrand et al., 2000), organizing data in a grid (Jacob et al., 2002), commanding tactical operations (McGee et al., 2002), activating simple commands by folding (Probst et al., 2011), or bookmarking (Steimle et al., 2008). However, Do-Lenh et al. (2009) showed that this usability benefit does not necessarily have a positive influence on learning outcomes: it may be more productive to have to negotiate the access to the interface, because the argument helps in the understanding of the problem.

Mackay et al. (1998) investigated how air traffic control can be augmented: they used small pieces of paper representing information about the landing of flights, which are scheduled according to the position of these pieces of paper. They concluded that interfaces for air traffic control should use these small pieces of paper, because it is simpler to implement and more promising than a full replacement with keyboards and screens for example.

Tap (2004) went further using the spatial characteristic of paper, and proposed escaping the “gravity of the desk” on medical monitoring by putting the printers on the walls of hospitals (see Figure 1.4). The strips of data would attract medical professionals, and could be split to bring the collaboration where it belongs.

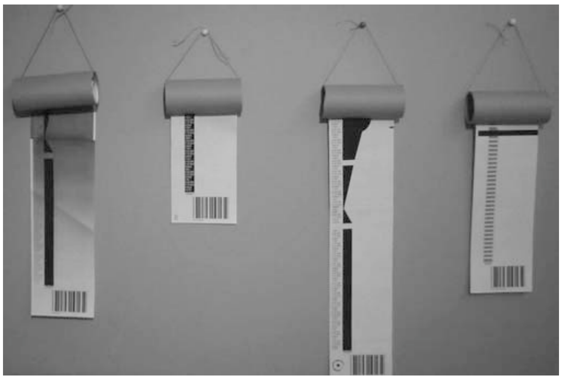

Figure 1.4: Printers hanging on the walls of hospital, outputting patients’ constants.

1.2.6 Augmented Presentation

Paper remains in presentations - in the form of notes - despite the ubiquity of computerized presentation tools. Nelson et al. (1999) showed that a deck of cards could be used to control a slide show. Signer and Norrie (2007) preferred an augmented paper representation of the slides, where a digital pen can select one slide to display and synchronously annotate it on the projection.

1.2.7 Printout Tracking

Paper sheets are cheap and easily produced, which has the side effect of creating barely manageable sets of documents. To address redundant printing, Smith et al. (2006) sent command intercepts to the printer in order to register the document. Then, a camera checked whether an already printed version of a document needed by the user was in the vicinity of the desk. Kim et al. (2004) also tracked documents on a desk, and allowed searches returning the location of a given document, even within a pile. On a related note, Steimle et al. (2010) allowed for the creation of mixed piles of physical and digital documents on a table top.

Tracking printouts brings us to the next border in the continuum of paper interfaces. Robinson and Robertson (2001) used paper sheets as anchors of digital information. Want et al. (1999) used them as links to their digital versions. Arai et al. (1995) used them as links to other digital objects. In these contexts, sheets are on the edge of becoming more significant in terms of their presence and position than by their content. For example, Koike et al. (2000) made circles appear around books, and used the intersections to formulate boolean queries involving the keywords associated with the books more naturally (see Figure 1.5).

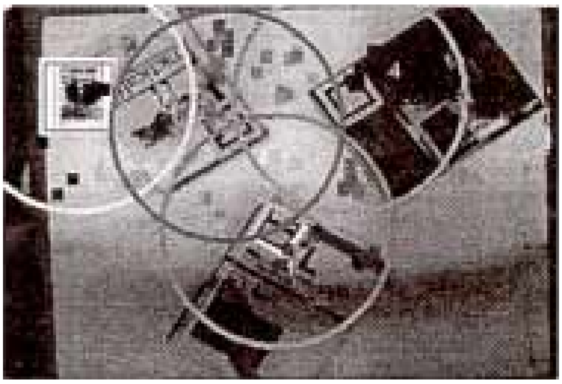

Figure 1.5: Three documents are placed on the interactive surface. Circles denote their area of influence, and squares in the intersections of circles show how many other documents have the same keywords as the corresponding documents. Source: (Koike et al., 2000)

1.3 Class III – Tangible Paper

1.3.1 Cards

We place cards on the tangible side of the continuum of paper interfaces. Their content is usually just an icon or short description of the object it represents. Cards are mostly used because they are easy to pick up, place, and move. Using cards as an interface dates back to Perlman (1976), where each card corresponded to a command of the programming language logo, and the rows holding them corresponded to procedures. Children can assemble programs simply by ordering the cards in the rows.

Cards are particularly popular in Augmented Reality games (Lam et al., 2006; Diaz et al., 2006; Tan et al., 2008). Cuendet et al. (2011) also showed how cards can be used as a tangible interface, where each card materializes an object to be manipulated, such as the endpoint of a segment, or an option to enable in the simulation.

Waldner et al. (2006) proposed card-like “tiles” to manipulate digital objects: the tiles represent actions and functions to be applied to the images. They compared this approach with a horizontal touchscreen and images simply printed on paper. Their hypothesis was confirmed: the paper condition offered the most efficient, familiar, and preferred mode of interaction and collaboration.

More simply, Cho et al. (2011) proposed attaching Augmented Reality videos to postcards, so that digital videos could be sent with traditional mail.

1.3.2 Augmented Reality Support

Billinghurst et al. (2001a) proposed the concept of Tangible Augmented Reality. It consists of interfaces attaching digital elements to physical ones, with the main goal of easing collaboration. As stated by Mackay and Fayard (1999), Augmented Reality provides a powerful alternative to the “keep it or replace it” model. Augmented Reality consists of displaying virtual elements among real ones. This requires an output device, most often a screen, a projector, or a head-mounted display. There are several ways to detect the real elements that will serve as input (Carter et al., 2010). Except for RFID, they rely on a digital camera to identify patterns printed on paper that are stuck to physical objects. These patterns can be the result of the position of words, or they can be specifically designed markers, such as 2D barcodes.

Van Krevelen and Poelman (2010) surveyed the work related to Augmented Reality. They concluded that these technologies still face critical issues preventing their adoption but expect this situation to change in less than a decade. The work related to Augmented Reality in education has been reviewed by Yuen et al. (2011), who concurred that these technologies will be increasingly adopted.

The work of Kaufmann and Dünser (2007) is particularly relevant to our work. Their Tangible Augmented Reality approach aimed at the collaborative learning of 3D geometry and spatial abilities using head mounted displays. They found that this interface was easy to learn (as opposed to CAD software), playful, and encouraged users to explore new functionalities. However, they also report that the Head Mounted Display had negative side effects on the subjects, such as eye strain or even “simulator sickness” (comparable to motion sickness).

Most Tangible Augmented Reality work is based on simple cardboard cut-outs of square 2D markers (Sidharta et al., 2006; Liu et al., 2007; Richard et al., 2007a; Freitas and Campos, 2008; Spikol and Eliasson, 2010), or even the raw sheet of paper containing a tag (Hsieh and Lee, 2008). Sticking markers on flat paddles is another solution (Kawashima et al., 2001; Dünser and Hornecker, 2007; Hornecker and Dünser, 2009).

1.3.3 Non-Augmented Reality Uses of Tangible Paper

Approaches using tangible paper for purposes besides Augmented Reality are more rare. Zhang et al. (2001) used a simple sheet as a background for a computer vision system to emulate a keyboard or a mouse with finger tracking. Blackwell et al. (2004) used cards with RFID tags as physical icons to collaboratively control an Information Retrieval System.

Ullmer et al. (2008) used tagged paper to specialize generic tangibles. For example, a generic tangible interface (called *core tangible*) could consist of 3 wheels. The reaction associated to the turning of these wheels depended on a parameter card. In a tangible interface for a radio, a card could map the rotation of a wheel to the volume, another card could map the rotation to the frequency, and yet another card could map the rotation to one of the preset stations.

1.3.4 Scaffolding Interaction

The difference between Class III and Class IV simply concerns the flatness of the paper interface. On the border, flat pieces of paper can be used to hold together elements that are not flat. More precisely, Pedersen et al. (2000); Millner and Resnick (2005); Hudson and Mankoff (2006); Block et al. (2008) used paper as a support to hold electronic parts together. This allowed them to assemble working prototypes, and they could draw on the paper to also draft the appearance of the system. Buechley et al. (2009) added the use of conductive paint to further integrate paper with the computational system. Indeed, paper is not only the physical support of the electronic components, but also the medium for the electrons that they exchange.

1.4 Class IV – Tangible Paper Controllers

1.4.1 3D Surface

Paper is perceived as flat, but it does exist in our 3-dimensional reality. Gupta and Jaynes (2006) tracked the orientation of pages in order to adapt projections and make the augmentations more seamless. Makino and Kakehi (2011) went even further and modelled the deformation of pages, which enables tangible interactions with pages, such as pinching. They also propose rolling the sheet into a half-cylinder, extending the interaction possibilities with actions like *pushing* or *blowing*. Holman and Vertegaal (2008) used infra-red markers to detect the shape of a flexible surface (see Figure 1.6), and invented related interaction techniques, such as sorting by shaking, browsing by leafing, or resizing by folding.

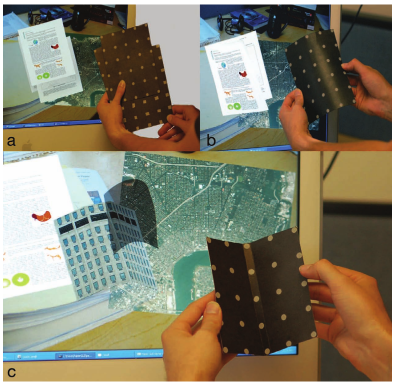

Figure 1.6: A flexible, paper-like surface containing infra-red markers, enabling the detection of how they are stacked (a), leafed through (b), or folded (c). Source: (Holman and Vertegaal, 2008)

Conversely, Koizumi et al. (2010) showed how to control the shape of a piece of paper by stimulating a Shape-Memory Alloy[2] embedded in the paper with sunlight, laser, or heat.

---

[2] A Shape-Memory Alloy is a metal that returns to a predefined shape when heated.

1.4.2 Paper Cubes

Paper also offers the possibility of being folded into a cube, which is a simple way to obtain a tangible Augmented Reality interface. For example, Lee et al. (2004) used cubes with Augmented Reality markers to author Augmented Reality animations, and Sin and Badioze Zaman (2009) used similar cubes to interact with the simulation of a solar system. Costanza et al. (2010) used paper cubes to control an electronic music creation software. They also provided a tutorial to assemble and set-up the rest of the system with a regular webcam, making the tangible interface *downloadable*. Moore and Regenbrecht (2005) associated cells of a map on a paper cube; by rotating the cube, it was possible to boundlessly navigate the map in any direction, even though the interface remained very small. Song et al. (2007) used a paper cube and a digital pen to combine the advantages of physical and digital modelling tools (See Figure 1.7).

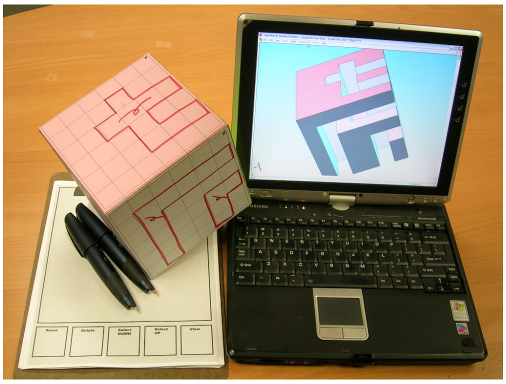

Figure 1.7: A cube made of augmented paper, which allows a digital pen to outline regions to remove, which are then reflected on the digital model. Source: (Song et al., 2007)

1.4.3 Paper Solids

Cubes are the simplest way to obtain a solid with paper, but not the only way. HyperGami (Eisenberg and Nishioka, 1996) and JavaGami(Eisenberg, 1999) are computer tools in which polyhedral forms can be created from a set of basic forms (e.g. platonic and archimedean solids), the faces can be decorated, and template printed. Then, the user can fold and assemble the solid, possibly sticking it to other solids to create more elaborate shapes. It is even possible to create complicated shapes (like a rabbit) from a single sheet, without having to cut it (Tachi, 2010). Song et al. (2006) used paper augmented with the Anoto pattern to produce such a solid, allowing a digital pen to control CAD software to modify the virtual model of the solid.

Saul et al. (2010) showed how paper solids can be augmented with electronic components to build interactive devices, such as speakers or robots. They also showed how genetic algorithms can evolve the shape of a solid to optimize it according to specific rules, e.g. a tall, stable lampshade that makes the most efficient use of paper.

1.4.4 Sticking Tags

More simply, tags printed on paper can be stuck on solids. For example, Cuendet et al. (2011) stuck tags on pieces of wood to help carpentry apprentices develop their spatial skills. Horn et al. (2011) marked blocks of wood, which represented pieces of programs, to let children control a robot. Jacucci et al. (2005) augmented an architectural mock-up in order to apply textures on it. Hornecker and Psik (2005) reported how a class of computer science students could develop a wide variety of interfaces by simply sticking tags on various physical objects. Jo and Kim (2011) even attached markers on humans to augment dramatic acting.

1.4.5 Other 3D Paper Shapes

Hendrix (2008) proposed a tool to help design pop-up books: users can place cuts and folds on a template, and the software gives a preview of the resulting shape after the template would be produced and folded. Eisenberg et al. (2003) gave further inspiration for paper realisations that could be augmented with *Sliceform Builder* and *Spectre*. *Sliceform Builder* helps design geometric models resulting from interlocking sets of planar pieces. *Spectre* allows for printing on several layers of transparent paper to visualize 3 dimensional objects by placing the layers regularly on top of each other.

1.5 Map: our Activities in the Continuum of Paper Interfaces

As a first reference point of our exploration of paper interfaces, we can already situate our activities into the four classes that we just defined:

- Kaleidoscope is a Class I interface: pupils draw on sheets, and the symmetric image of their drawing is projected. Only the content on the sheet is used.

- AngleHunt is a Class II interface: pupils have to reassemble several pieces of a document.

- Sympliage, Triangram and Messangles combine Class II and Class III characteristics: pupils have to transform pieces of paper and describe the process. The transformations exist both in a physical form, and in a content form.

- Quads and SpaceJunk are Class III interfaces: pupils manipulate various pieces of paper. The focus is on the manipulation, but the content still matters.

- Angoli is a Class IV interface: pupils manipulate flat, tangible controllers. The content of the pieces of paper is only used to identify them.

1.6 Tangible User Interfaces

The end of the continuum we just described overlaps with Tangible User Interfaces (TUIs), where paper is just one possible physical material that can be used to interact with a computer. We can illustrate the border between the continuum and TUIs with the two following papers. Terry et al. (2007) used wooden blocks and paper blueprints to query architectural documents: two wooden blocks are used to define the area of the blueprint to zoom into, another block filters displayed information, etc. Aliakseyeu et al. (2006) manipulated documents with non-paper tangibles: bricks are used to drag and drop projections of documents. Note that the two ends of the continuum of paper interfaces are not mutually exclusive: they correspond to different sets of possibilities that can very well be combined. Zufferey (2010) showed that the mix of tangibles and paper is particularly powerful in education, where manipulation is often key to experimentation, and documents help make reflection become persistent.

This section gives an overview of the work done on TUIs by first giving several examples, and then summarizing various taxonomies and frameworks related to TUIs.

1.6.1 Examples of TUIs

TUIs can be traced back to the *Slot Machine* of Perlman (1976). Wellner (1993) also mentioned *tactile interaction* as an added value of the DigitalDesk for electronic documents. Fitzmaurice (1996) really oriented his focus on *graspable interface*, more particularly on the usability aspect of haptic directness, i.e. the bimanual interaction and spatial reconfiguration of objects. The notion of *Tangible User Interfaces* has been developed by Ishii and Ullmer (1997), based on the idea of seamlessly linking digital data with physical objects.

We follow the categorization of Ullmer et al. (2005) to give a sample of tangible user interfaces: interactive surface, token+constraint, and constructive assembly (see Figure 1.8). Shaer and Hornecker (2009) reviewed far more examples in a wide variety of application domains.

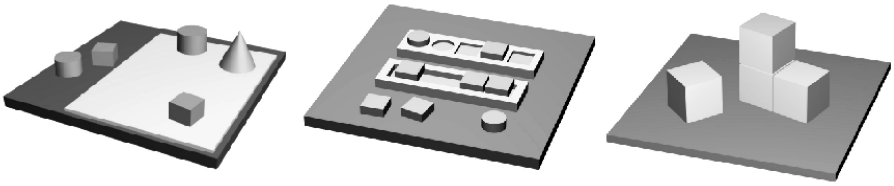

Figure 1.8: An illustration of three categories of TUIs: interactive surface, token+constraint, and constructive assembly (from left to right). Source: (Ullmer et al., 2005)

Interactive Surface

Ullmer et al. (2005) called Interactive Surface the category of TUIs where the spatial configurations of physical objects are used as input to a digital system, typically a simulation. The configuration of the objects is usually their identity, presence, position, and orientation. This information is often used in Augmented Reality approaches, where the augmentations come from a projector above the interactive surface. For example, Underkoffler and Ishii (1999) used this approach to simulate optics: the users can move the tangible representation of various optical equipment, and the effect of the representation on light is computed and projected.

Park and Woo (2006) showed an alternative technology: the projected augmentation originates from under the translucent table; cameras tracking elements of the interface are both above and below the table, and a large screen display is placed vertically. However, this complementary screen was left unused during the experiment, probably because it did not bring much advantage compared to the distraction resulting from having to look at two different areas at the same time.

Piper et al. (2002) illustrated a radical alternative: the configuration used as input is the surface itself, which is made out of clay and scanned with a laser. This setting is used for landscape analysis.

Constructive Assembly

The second category, Constructive Assembly, relates to interfaces consisting of modular elements which can be combined in order to build both the physical form of the object, and its meaning in the system. In its most common form, it embeds electronics to either sense or run the system. For example, in Topobo (Raffle et al., 2004), passive blocks can be assembled into a shape (e.g. an animal), and blocks containing a motor with kinetic memory repeat the assembly of movements to create the shape. Similarly, Zuckerman and Resnick (2003) proposed blocks with embedded electronics in order to teach dynamic systems to children. Each block has a function (either generating, transforming, or displaying values) and the blocks can be linked to each other. Children can see the evolution of the resulting system.

Building blocks can also be linked to an external display. For example, the ActiveCubes of Watanabe et al. (2004) embedded electronics to input and output information, infer how the blocks are assembled, and transmit this information to an external system. This external system will control the display of an object linked to the construction, e.g. the view of a 3D Model can be controlled with an assembly of cubes with gyroscopic sensors.

Token+Constraint

The last category is actually the contribution of Ullmer et al. (2005). It contains systems which take advantage of physical constraints of tangible elements (the tokens). The often cited historical example is the Marble Answering Machine developed by Durell Bishop during his Master thesis at the Royal College of Art in 1992. With this answering machine, the tokens are marbles, representing messages. The constraint is the place where the marbles are dropped: when a new message is recorded, a new marble pops up. The user can push it back to delete it, or place it in two different places to listen to the message, or call back.

The work of Blackwell et al. (2004) and Jacob et al. (2002) presented previously are other examples of Tangible+Constraint TUIs. The former used cards as tokens, and the holder as a constraint to build the organization of (the constraint on) the arguments (tokens). In the latter case, tokens (e.g. presentations in a conference) are organized into a schedule, which is enforced by the constraint of a grid.

Patten and Ishii (2007) used physical constraints to enforce logical ones (e.g. a strap around two tokens enforces a maximal distance). Interestingly, however, they do not classify their *mechanical constraints* as a Token+Constraint system. They explain this discrepancy by the fact that the system is not aware of these constraints, and does not interpret their meaning.

1.6.2 Epistemology

These three categories are not the only ways to taxonomize tangible interfaces. Researchers have been repeatedly trying to map the space of possibilities for TUIs. This space is huge, because it uses elements from the physical world, and users are already experts in this system. Mackay and Fayard (1999) made a similar comment about paper-based interfaces:

> “People go beyond officially-sanctioned uses, inventing new uses based on the situation at hand. A newspaper can be a child’s hat, a paper airplane or the lining for a bird cage; an envelope can be used as a scratch pad or a book mark.”

We review hereafter several efforts by TUI designers to describe or guide the design in this space of designs.

Holmquist et al. (1999) described a tangible interaction paradigm based on three kinds of tangible elements: tokens, containers, and tools. Tokens correspond to tangible elements that share a perceptual similarity with the represented digital object, in contrast with containers, which are generic physical objects that only temporarily link to digital objects. Tools represent functions or processes applied on the other objects.

Koleva et al. (2003) extended the granularity between token and containers, by defining six levels of coherence between a physical object and the associated virtual one. These levels range from the weak coherence of general purpose interface elements (like a mouse) to physical objects which closely represent the virtual one. To determine this level, Koleva et al. (2003) defined six properties of the link between physical and digital objects, such as its cardinality or its lifetime.

Fishkin (2004) proposed a simpler taxonomy which places TUIs in two dimensions. The first dimension is the degree of embodiment, which has four levels: full embodiment (input and output are merged in a unique object); nearby embodiment (around the object); environmental embodiment (around the user); and distant embodiment (the output is disconnected from the input). The second dimension is the degree of metaphor: none (no analogy); noun (the physical object looks the same as the digital one); verb (the actions on the physical object are similar to actions on the digital one); noun+verb (combining both); and full (physical and digital objects are indistinguishable for the user).

Finally, Hornecker and Buur (2006) broadened the scope of TUIs to *tangible interaction*, by describing more than the data-centred view of TUIs. They also add the *expressive-movement-centered* and the *space-centred* views of tangible interaction. The former focuses more on the interaction itself than the physical-digital mapping. The latter focuses on the position of the user in space.

1.7 Open Issues

The conclusion we draw from this literature review is two-fold: first, insufficient exploration of the possibilities offered by paper as an interface remains an open issue; second, the lessons to draw from this exploration have not been analyzed or outlined.

1.7.1 Deeper Exploration of Paper as an Interface

Placing the related work in our continuum revealed two kinds of approaches that have been less explored than augmented documents and tangible user interfaces. First, the approaches in the intermediate areas of the continuum offer more possibilities than the ends, because they combine characteristics of both. However, the literature is more shallow in these areas. Moreover, the continuum is not meant to be a one-to-one mapping: some approaches could very well stretch over several areas. For example, one can easily imagine a system that facilitates working on several documents at the same time, while also offering control over tools with a tangible interface.

Second, it seems that paper as a material has not been used to its potential. Among the tangible approaches, paper is mostly used to indicate its position, sometimes its spatial configuration, and rarely the shape of its surface. However, paper can be mechanically transformed (folding, cutting), and chemically transformed (burning, getting it wet), with effects that users can foresee. Paper itself even includes more than paper: cardboard is used, but there are many other alternatives to paper (transparent paper, filtering paper, paper-mâché, etc.). Finally, paper comes with a myriad of practices that can be exploited to build intuitive interfaces (e.g. playing cards, writing sticky notes, placing bookmarks, etc.).

1.7.2 Structuring the Design of Paper Interfaces

This exploration seems to be endless, which means that it would be beneficial to map it. The practice of working with documents and designing TUIs has been addressed much more, from an epistemological perspective, than paper interfaces. Klemmer and Landay (2009) proposed a toolkit to design TUIs, in which they distinguished between book and document, and integrated printers. However, as explained in the previous section, we feel that this is only scratching the surface of the possibilities of paper. In this perspective, a framework specific to paper would result in a deeper exploration.

Signer and Norrie (2010) set ambitious goals for such a framework. The framework should be independent from the augmentation technology, guide both the visual design of the paper interface and the interaction itself, and include authoring, publishing, and deployment mechanisms. A framework fulfilling these conditions could participate in the democratization of paper interfaces as much as Windows, Icons, Menus, and Pointer (WIMP) did for the personal computers.

In Chapter 3, we propose a framework inspired by the continuum of paper interfaces described here. It will emphasize the tangible side, as it appeared that the work related to paper interfaces was unbalanced in favor of the use of paper as a document.

2 Computer-Supported Geometry Education

Education is a vast research area, with a long history and many different approaches. This chapter aims at highlighting the prior work and concepts related to our application of paper interfaces for geometry education. We start by describing the existing types of interfaces used to learn geometry. We then we identify challenges for learning geometry, and show how paper interfaces could address them. Next, we review work related to teaching geometry.

2.1 Computer Interfaces for Learning Geometry

Computer-supported learning of geometry followed the general evolution of interfaces. At first, it made use of keyboard and screens for text-based interactions. The apparition of the mouse allowed for graphical user interfaces and more direct manipulations of geometrical objects. More recently, new interfaces, such as Tangible User Interfaces (TUI) and digital pens, extended the possibilities for interacting with a computer-supported geometry application. We first review each of these approaches, and then show how their features are not exclusive, as they be found in paper interfaces.

2.1.1 Interacting with the Logo Microworld via a Keyboard

Many computer-supported approaches to learn geometry are linked to constructivism, which we present very briefly hereafter. The main idea of constructivism is that knowledge is built upon an active interaction with an environment, which allows the construction of a representation of the environment. For example, Bruner and Kenney (1965) showed how children can construct knowledge about prime numbers by reorganizing beans into rows of equal length; a number of beans that can be arranged into only one row is prime (see Figure 2.1). In geometry, this construction can happen if the pupil can sense and act in a geometrical environment rather than in the physical world. Computers allow for the recreation of *microworlds* following different rules than the physical ones. Microworlds are used to construct meanings (Sutherland and Balacheff, 1999).

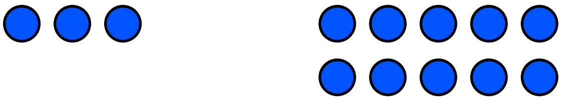

Figure 2.1: 3 beans can not be arranged into several equal rows, but 10 can. This means that 3 is prime, and 10 is not prime.

Historically, the concept of microworlds in geometry first appeared with the Logo programming language, as a reification of abstract objects (Minsky and Papert, 1972; Thompson, 1987). Logo is a functional programming language, to which Papert added *turtle graphics*, a set of commands to describe vector graphics (e.g. set pen down, move forward, turn right, etc.). Hölzl (1996) called these graphics intrinsic: commands are relative to the *turtle*. In the educational point of view, this makes the children map the movements of the turtle to those of their own body (see Figure 2.2). In other words, the movements children would do in the real world are mapped to commands in the Logo *microworld*. This microworld is a way to embody abstract objects (Minsky and Papert, 1972; Thompson, 1987) for learners to experiment with them.

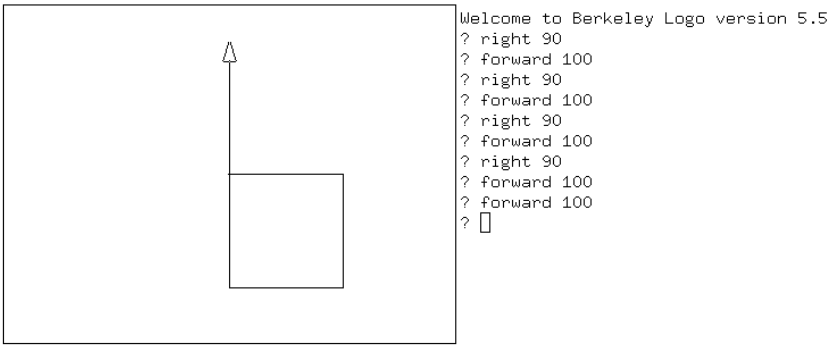

```
Welcome to Berkeley Logo version 5.5
? right 90
? forward 100
? right 90
? forward 100
? right 90
? forward 100
? right 90
? forward 100
? forward 100
? 
```

Figure 2.2: The drawing (left) resulting from the commands (right) given to the *turtle*. Source: http://en.wikipedia.org/wiki/File:Ucblogo.png

Pratt and Ainley (1996) noticed that Logo is an unusually stimulating environment for children, and cited many examples of children working creatively with Logo from Papert (1980); Ainley and Goldstein (1988); Blythe (1990). Papert (1972) explained: “When mathematizing familiar processes is a fluent, natural and enjoyable activity, then comes the time to talk about mathematizing mathematical structures, as in a good pure course on modern algebra.” Papert developed the constructionist learning theory to frame his work. Constructionism builds on the constructivist idea that knowledge is built by the learners, but insists that it is most effective by producing public objects, such as tangible objects or shareable code.

Logo is thus a geometry microworld with which pupils can interact through commands typed on a keyboard. However, the introduction of computer mice allows for a more direct manipulation. The potential benefits of such manipulation are illustrated hereafter.

2.1.2 Dynamic Geometry Software: Manipulating Geometry with the Mouse

Dynamic Geometry Software (DGS) are the most successful approach to integrate computers in geometry education. It is another kind of interface in the line of constructivist theories. It implements a variable figure, which is the result of a sequence of geometrical operations simulating ruler and compass constructions. The components of the figure can be directly manipulated with a mouse pointer, and the response of the figure follows geometrical rules rather than the strict transformation of the user. For example, if a point is defined as belonging to a line, the point will follow the projection of the mouse cursor on the line rather than the raw position of the cursor.

As pointed out by Laborde (2000), dynamic geometry is not new: in 1741, Clairaut introduced the constant sum of the angles of a triangle by explaining that the variation of one angle is compensated by the variation of another angle, as shown in Figure 2.3. At the beginning of the twentieth century, Méray proposed the teaching of geometry based on movement: translations for parallelism, rotations for perpendicularity, etc. He also proposed the idea of a variable figure, such as the spires of a spring which allow for the observation of the movements of points according to a constraint.

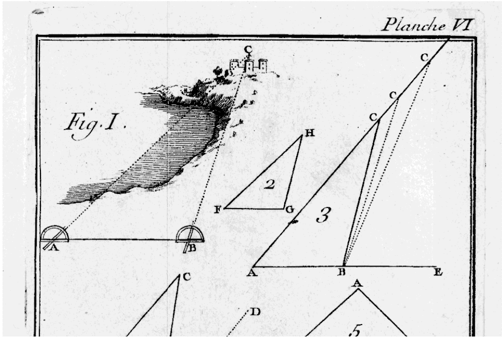

Figure 2.3: Clairaut explaining in 1741: "Assume, for example, that BC turning around point B moves away from AB to get closer to BE; it is clear that while BC would turn, angle B would continually open; & in contrast angle C would become more and more narrow; which would first let us think that, in this case, the decreased size of angle C would be equal to the increase of angle B, & thus the sum of the three angles A, B, C, would remain the same, whatever the inclination of lines AC and BC, on the line AE." Source: (Laborde, 2000)

Many examples of DGS have been developed: Cabri-géomètre[1] (Baulac, 1990), Geometer's sketchpad (Jackiw, 1995), Geometry Inventor, Euklid (Mechling, 1994), GEOLOG (Holland, 1993), Thales (Kadunz and Kautschitsch, 1993), The Geometric superSupposer, and GeoGebra[2]. Most of this software is designed for a personal computer, but there is also of version of Cabri for graphing calculators. See Figure 2.4 for screen shots. Platforms, such as the European project Intergeo[3], allow users to share digital content developed with DGS for teaching.

---

[1] CAhier de BRouillon Interactif (CA.BR.I.) - interactive sketchbook

[2] http://www.raumgeometrie.de

[3] http://i2geo.net/

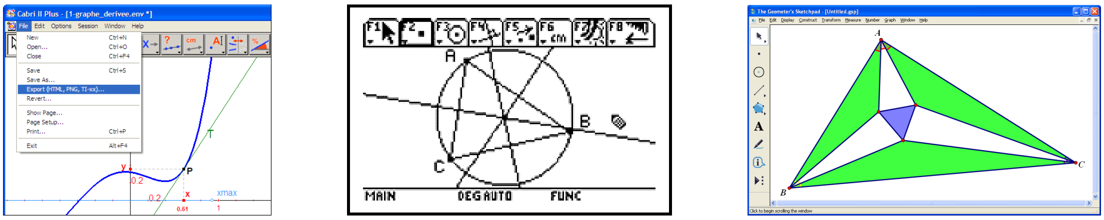

(a) Cabri-géomètre. Source: Cabrilog SAS

(b) Cabri-géomètre on a graphing calculator. Source: Texas Instrument, Inc.

(c) Geometer's sketchpad. Source: Wikipedia

Figure 2.4: Examples of DGS

In DGS, a diagram is constructed with a sequence of primitives, i.e. an assembly of technical symbols aiming at a cognitive symbol. As Mariotti et al. (1995) stated: the temporal sequence of the constructions' steps represents the counterpart of the logical hierarchy between the geometric components of a figure. More than an illustration of a cognitive symbol, the sequence elicits an interpretation. DGS forces the making of explicit properties needed for a construction (Clarou and Laborde, 2000), and its resulting behaviour requires the construction of an interpretation (Laborde, 2000).

DGS looks like an interactive version of geometry diagrams, but it is actually very different. As noted by Pratt and Ainley (1996), they allow for the drawing of objects on a geometrical basis rather than on a perceptual one. He further notes that this is not obvious for young children: they do not immediately see the difference with drawing software.

DGS can reproduce and amplify phenomena with the manipulation of a dynamic object. Laborde (2000) gives the example illustrated in Figure 2.5. The middle of the sum of two vectors with the same start point is constant for any start point. By allowing the tracing of this sum, DGS amplifies this phenomenon of the geometric world to make it appear.

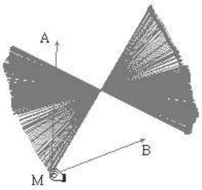

Figure 2.5: The trace tool amplifying a phenomenon. Source: (Laborde, 2000)

The benefits of DGS have been listed by Laborde and her coworkers (Laborde, 2002; Clarou and Laborde, 2000; Laborde, 2000). DGS helps the understanding of "what moves/what does not move" by enabling the dragging of certain points of the figure, which will react according to the rules given as its definition. For example, if two lines are defined as being perpendicular to one another, they can be moved up and down with relation to each other, but if one is rotated, the other is rigidly rotated in the same way.

Laborde (2003) noted how the intellectual curiosity is stimulated by causing a conflict between prediction and what actually happens. Pupils can be so surprised to discover that a construction results in a rhombus that they become eager to prove why.

Hölzl (1996) gave two reasons as to why users prefer DGS to Euclidean geometry (in the sense that only compass and straight edge can be used to define objects). He claims that pupils dislike Euclidean geometry for its tricky problems, and seemingly superfluous proofs. He also claims that teachers dislike Euclidean geometry for its absence of reliable and easily executable algorithms.

2.1.3 Tangible User Interface: Geometry in the Shared Physical Space

Although Tangible User Interfaces (TUI's) became popular with the work of Ishii and Ullmer (1997), tangible interfaces had been used before. Fitzmaurice et al. (1995) introduced the Graspable User Interface, an interface made of small physical artifacts that allow the direct manipulation of virtual objects. In an evaluation of the graspable interfaces, Fitzmaurice and Buxton (1997) argued that the mouse is a general, all-purpose, weak device that can be advantageously replaced by strong specific devices for a specific and limited task.

Indeed, attaching digital contents to physical tokens allow for rich, two-handed, and parallel interaction. Marshall et al. (2007) have categorized the possible benefits brought to learning by tangibles in four categories: collaboration, accessibility, novelty of links, and playful learning. Tangibles can provide external representations of a problem or an object, which play a key role in problem solving and learning (Ainsworth, 1999; Larkin and Simon, 1987) by helping the learner make inferences or by freeing up the cognitive load to allow the learner to focus on the core of his task. Another important dimension for learning is the coupling between cognition and physical experience (O'Malley and Fraser, 2004). This coupling can provide external representations of a problem or an object, which is instrumental in problem solving and learning (Ainsworth, 1999).

Do Lenh (2012) also listed the benefits of the many tangible user interfaces developed for education: engagement is increased (Horn et al., 2009b; Fjeld et al., 2007), multiple modes of communication (Evans et al., 2009) and representations (Ainsworth, 2006) are possible, and learning is more active. Furthermore, body movements help awareness of intentions (Hornecker et al., 2008), sensori-motor experience is beneficial for learning (Price and Rogers, 2004), small groups are more comfortable around an interactive tabletop than in front of a vertical display (Rogers and Lindley, 2004), and simultaneous inputs promote more equitable interactions (Stanton et al., 2002).

Hornecker and Buur (2006) broadened the scope of TUI to include *tangible interaction* by describing more than the *data-centred* view of TUI, which consisted of attaching virtual objects to physical ones. They described the *expressive-movement-centred* and the *space-centred* views of tangible interaction. The former focuses more on the interaction itself than the physical-digital mapping. The latter focuses on the position of the user in space. This is interesting if we draw parallels in the pedagogical context. TUI's involve physical behaviors that can be effective for learning, such as gesturing, physical movement, and embodiment (Goldin-Meadow, 2003; O'Malley and Fraser, 2004). The *expressive-movement-centred* view of tangible interaction can be seen as a way for pupils to enact a concept to learn. The *space-centred* deals with aspects like the division of labour, which can result naturally from the spatial disposition of tangible resources and orient the roles of learners in a group, as shown by Jermann et al. (2009).

Numerous instances of tangibles in learning can be found in (Shaer and Hornecker, 2009; Van Krevelen and Poelman, 2010; Yuen et al., 2011). Several TUIs have been implemented for learning geometry. Underkoffler and Ishii (1999) and Price (2008) used tangible tokens on an interactive surface to simulate optics configurations. Patten and Ishii (2007) used physical constraints to materialize spatial constraints, e.g. a disk to implement a minimum distance and a strap for a maximum distance. Kaufman used head mounted devices to allow the development of spatial skills in a collaborative way by attaching virtual 3D objects to physical tokens. Girouard et al. (2007) embedded RFID tags into cubes and connectors, which pupils used to complete exercises printed on cards, such as constructing an object with a given area, and then the computer could give feedback. See Figure 2.6 illustrations corresponding to the work described in this paragraph.

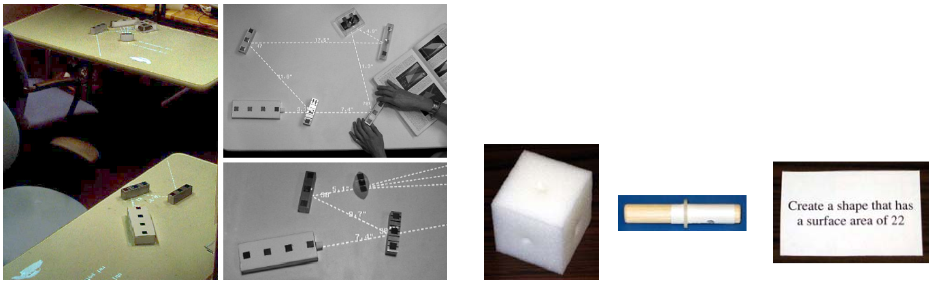

(a) Illuminating Light. Source: Underkoffler and Ishii (1999)

(b) Smart Blocks. Source: Girouard et al. (2007)

(c) Rubber band as contact constraint. Source: Patten and Ishii (2007)

(d) Strap and disk as max-min distance constraints. Source: Patten and Ishii (2007)

Figure 2.6: Examples of TUI for geometry.

However, critics have been expressed about the use of tangibles in learning. In a meta-review Sowell (1989) reported inconsistent and limited effects of manipulatives on learning outcomes. For example, Uttal et al. (1997) explained that manipulatives are note necessarily different from symbols. The children also need to associate a meaning to the objects and their relations, which is not very different from written symbols.

2.1.4 Digital Pens

Digital pens are another kind of interface especially relevant in the domain of geometry education. They enable the precise extractions of the writings done on a piece of paper. Paper has supported the development of mathematics by allowing the representation of abstract objects (with formulas or diagrams) (Laborde, 2002). Laborde (2002) reported the example of a teacher whose epistemological view of geometry fundamentally links understanding with paper. Schorr et al. (2007) described mathematics as symbolism which one operates. In contrast to written natural languages, which are meant to be read, a mathematical formula is meant to be transformed.

Vygotski (1978) generalized: “The invention and use of signs as an auxiliary means solving a given psychological problem (to remember, compare something, report, choose, and so on) is analogous to the invention and use of tools in one psychological respect. The signs act as an instrument of psychological activity in a manner analogous to the role of a tool in labour.” This idea is in line with Piaget’s constructivism, in the sense that what Vygotsky calls the practical tools, i.e. cultural symbols, are to be transformed into psychological tools, i.e. individual symbols (Laborde, 2003). This is called the internalisation: knowledge comes from the ability to have an internal discussion. This is the difference with constructivism: Piaget stresses the importance of individuals (building knowledge from experiences); Vygotsky’s socio-constructivism stresses the importance of cultural artefacts (building knowledge from social interactions).

There are several examples of digital pen and paper systems used for education. They are often based on Anoto technology, which consists of printing a microscopic pattern on paper which is then detected by a camera embedded in a pen, such that the strokes can be acquired. For example, Pietrzak et al. (2010) presented a system augmenting note taking in the classroom. They concluded from interviews with teachers that augmented paper could address the separation of digital content shown in the classroom, oral explanations, and notes. CoScribe (Steimle et al., 2009) integrates pen-and-paper-based interaction techniques that enable users to collaboratively annotate, link, and tag both printed and digital documents.

Digital pens are also less distracting because they are less associated with novelty. Oviatt et al. (2007) brought forward this intent of making the interface as *quiet* as possible. They compared how students worked on geometrical problems using regular pen and paper, a digital pen using the microscopic pattern, a pen tablet, and a graphical tablet. They showed that that the closer to the familiar work practice (i.e. pen and paper), the better the performance. This in line with the vision of Weiser and Brown (1996), who sees the upcoming age of ubiquitous computing as a solution to information overload, by continuously providing located, peripheral information via computing devices embedded in every aspect of our familiar lives.

2.1.5 Paper Interfaces

The tangible and document nature of paper interfaces are not mutually exclusive. For example, Song et al. (2007) used a paper cube and a digital pen to combine the advantages of physical and digital modeling tools in the early stages of architectural training. Zufferey (2010) proposed an interface that combines tangible and paper elements to teach logistics. The mix of tangibles and paper is particularly powerful in education, where manipulation is often key to experimentation, and documents help make reflection become persistent.

In other words, paper interfaces allow for the combination of features from TUI and interfaces based on digital pens. Similarly, the features of DGS can be found in a paper interface, since they can control dynamic figures. Bringing computers and paper together also solves the debate about whether paper is just a context for geometry or not (Laborde, 2002): there is no choice to make. DGS and paper become two sides of the same coin: paper can be a simple interface for DGS, or DGS can augment a regular paper context. Finally, Perlman (1976) showed how a paper interface could even be used to explore the Logo microworld, by allowing cards to be ordered in rows, which represent commands in sequences.

Several approaches based on paper have already addressed geometry. Martin-Gutiérrez et al. (2010) showed the design of an augmented book combined with a screen to develop the spatial abilities of engineering students by showing them variable views of 3D models (See Figure 2.7). The researchers measured a positive impact on the students' spatial abilities, and the users found the system to be easy-to-use, attractive, and useful. Similarly, Cuendet et al. (2011) integrated a tabletop to augment exercises on paper to develop the spatial abilities of carpenter apprentices (See Figure 2.8).

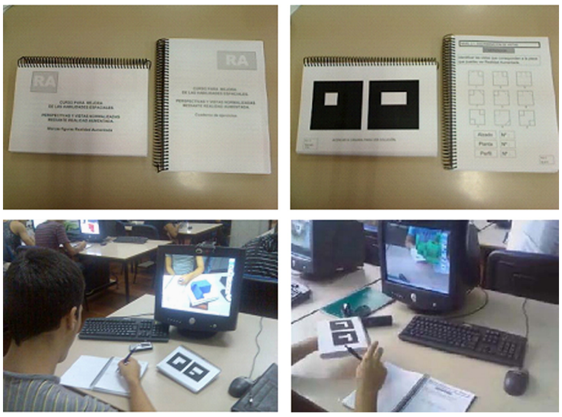

Figure 2.7: An augmented book to develop spatial abilities.

Paper interfaces can further rely on their tangible nature by using artefacts that are more important for their presence and position than for their content, such as gaming cards. For example, Horn et al. (2009a) mixed the use of tangible and virtual images to explore phylogenetic classifications. Do-Lenh et al. (2009) used paper tokens on a tabletop augmented by a camera-projector system in order to build a concept map.

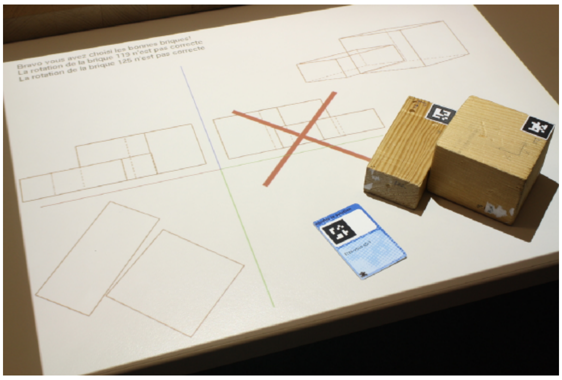

Figure 2.8: Augmented tabletop to develop spatial abilities.

2.2 Addressing the Challenges of Learning Geometry with Paper Interfaces

Laborde (2000) highlights a specific ambiguity of geometry education: it is based on both discourse and spatial-graphical representations. Her paper also mentions other dualities: figure versus drawing, microscopic surfaces versus macroscopic spaces, and empirical versus deductive approaches. We follow this epistemological approach of geometry education, illustrating the possibilities offered by paper interfaces that would address the challenges of this duality.

2.2.1 Language and Diagrams

Laborde (2000) notes that learning geometry involves two semiotic systems: discourse and spatio-graphical representations. Both registers allow for the description of objects and relations as defined by their geometrical meaning, but they are not organized in the same way. Figures have a global meaning, but verbal descriptions are read linearly, from start to the end.

These two semiotic systems can be independent. For example, a proof can be done purely with a discourse, or purely with a figure, as shown in Figure 2.9. More frequently, however, diagrams are seen as help for the discourse: they can be used as heuristic exploration preceding a hypothesis, or as an illustration of the discourse of a proof. However, some researchers insist on the importance of the tight integration of the different semiotic systems in the social debate (Richard et al., 2007b) and in teaching (Mariotti et al., 1995; Duval, 1998). With this integration, experts are able to show to novices how information can be extracted from a diagram relevant to the geometrical discourse.

The tight integration of language and diagrams are obviously a strength of paper interfaces, as they integrate with freely annotated documents. Practically speaking, this means that a regular pen can be used to write a text, or draw a figure, without having to issue complicated commands for special symbols, which are very frequent in mathematics.

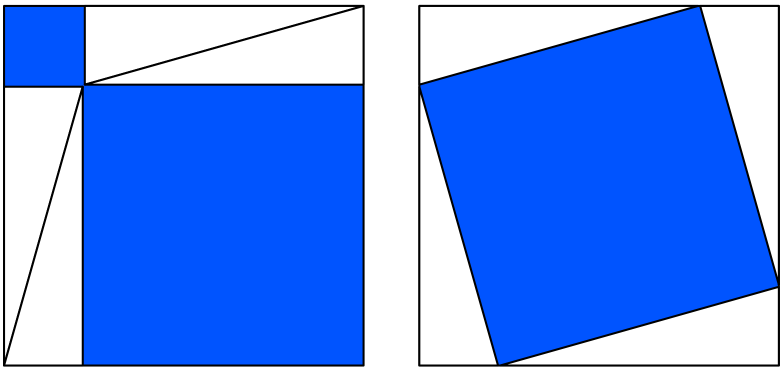

Figure 2.9: The “proof without words” of the Pythagoras theorem. The diagram is simple, but its meaning is not obvious. Explaining this figure would defeat its purpose, and thus, is left as a riddle to the reader.

2.2.2 Figure vs. Drawing

The spatio-graphical register itself bears another duality: the difference between a drawing and a figure. As explained by Pratt and Ainley (1996), a figure is an ideal, an infinite set of cases, while a drawing is fixed as a single case.

A figure requires the pertinent properties to be separated from the drawn artefacts. In other words, some of the spatio-graphical properties are matched to geometrical properties. Many spatio-graphical properties are to be disregarded in a figure, e.g. thickness of the line, imprecisions, etc. Laborde (2000) illustrated this with the fact that 11 out of 15 children fail at recognizing symmetries if the axis is not vertical, and cannot easily distinguish a rhombus from a square if the shapes are not in their prototypical orientation, as illustrated in Figure 2.10).

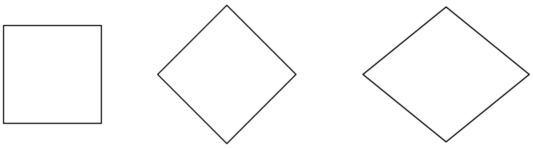

Figure 2.10: A square (left), when printed “on its edge” (middle) can be mistaken for a rhombus (right).

Hölzl (1996) explains that the drag-mode of DGS is a mediator between the concepts of drawing and figure. More precisely, the mediation of the spatio-graphical distinction happens as a result of resistance and feedback (Laborde, 2000). In other words, if a pupil tries to drag a point that belongs to a line, the point will resist movements that do not follow the line, giving the pupil visual feedback of the point staying on the line and not under the cursor. This results in pupils giving meaning to reactions, such as “put some glue in the middle of one piece and then on another and glue them together” (Pratt and Ainley, 1996).

Paper interfaces can be leveraged to highlight the differences between figure and drawing, this time taking advantage of the computer component rather than the paper component. The variability of computers can be leveraged to distinguish properties of a figure from drawing artefacts, similar to DGS. For example, if a line touches a circle, moving one of its control points can help determine whether the line is actually a tangent or if the contact was just an artefact. This change can be controlled by a piece of paper used as a tangible control. Furthermore, the mix between drawn and projected content can be a useful way to help highlight the differences between figures and drawings. For example, two lines drawn approximately parallel can be overlapped with projected parallel lines.

2.2.3 Microsurface and Macrospace

The next duality is actually twofold, and refers to an epistemological paradox of geometry. Firstly, geometry is aimed at grasping spatial situations, but pupils are mostly trained using problems concerning 2D situations. Even 3D problems are modelled using flat diagrams.

Secondly, geometry addresses various scales: Brousseau (1983) described three levels: micro-space, meso-space and macro-space, corresponding respectively to the scale of objects that can be manipulated, human scale and successively wider scales. The macrospace is not restricted to large areas or measures, but also unbounded objects like angles. Berthelot and Salin (1998) showed that pupils fail in microspace: for example, they cannot reproduce a rectangle that is of the size of a bench (in the mesospace) because they fail at reproducing the angles with enough precision, which result in a quadrilateral that is not a parallelogram.

Regarding the first challenge, paper is admittedly limited to representing surfaces. However, the surfaces exist in space and can be combined together: a castle can be made out of cards. This is an interesting way to concretely divide a spatial problem into one or more planes, which are more easy to comprehend. Also, paper is not necessarily a plane: its surface can be one of a volume, effectively adding a third dimension to paper artefacts. For example, it is easy to create a “cylinder” or a “cube” with paper.

Paper interfaces also offer opportunities to address the second challenge. Microspatial elements can be contained in a sheet. Large (or several) sheets together can cover mesospatial elements. As for the macrospace, the mobility of paper can be leveraged: sheets can be placed all over the world if need be.

2.2.4 Empirical and Deduction

Since Euclid, demonstrations have a central place in geometry. Demonstrations present another ambiguity for pupils: they should be inspired by what they see, but cannot base their reasoning on it. This example presents other dualities of geometry: sensible versus intelligible and empirical versus deductible. Regarding the empirical domain, a spatial reality is validated by the experience while a theory is validated by proofs. Chevallard (2001) talked about the former as experimental geometry and Houdement and Kuzniak (1999) referred to it as natural geometry. Regarding the reasoning, Laborde (2000) detailed the difference between a demonstration and an argumentation. Demonstrations rely on substitutions of proposition with an operating value (hypothesis or conclusion). Argumentation, on the other hand, relies on an accumulation of arguments which have more or less force, as opposed to the propositions in a demonstration which is always true.

Researchers showed a wide range of relationships between empiricism and deduction. Richard et al. (2007b) set them at the same level, by describing ways to graphically carry reasonings and make inferences based on figures. Rouche (1999) described an order in this duality: first perceive, then conceive, finally infer. In this case, the empirical is detached from the reasoning, but Douek (1999); Bartolini Bussi et al. (1999) promoted their cognitive unity, integrating exploration and proofs.

Such cognitive unity can be found in paper interfaces because of the persistent aspect of paper. For example, a series of drawings and figures can progressively lead to a proof, by allowing the side by side comparison of as many drawings as wanted. This comparison supports the perception, conception, and inference process. Empiricism and deduction can be discriminated using techniques similar to DGS, thanks to the computer side of paper interfaces; integrating exploration and proofs is facilitated by the integration of discourse and figures on a sheet.

2.3 Elements of Teaching

In this section, we discuss educational topics that are less specific to geometry but act as a framework to develop ecologically valid technologies. Again, we briefly present this framework, and how it fits in with paper interfaces.

2.3.1 Collaborative Scripts: Orienting Toward Productive Interactions

Learning gains are not intrinsic to a technology, in the sense that students do not learn simply from interacting with it (Dillenbourg and Evans, 2011). Similarly, students do not learn from just working in groups; they learn because they engage in explorations, explanations, elaborations, knowledge elicitation or conflict resolution (Dillenbourg, 1999). Collaborative learning aims at constructing shared knowledge through interaction (Dillenbourg, 1999; Roschelle and Teasley, 1995).

Pedagogical scripts offer an approach to foster productive interactions (O'Donnell et al., 1992; Aronson and Patnoe, 1997; Dillenbourg, 2002). They consist of a set of rules as to how the learning group should interact, collaborate, and solve the problem (e.g. "draw a trapezoid, share it with the pupils of your group, and discuss how to sort them by increasing area"). A *macroscript* (Jermann and Dillenbourg, 1999) is a higher level scenario, that aims to structure the activity. For example, it can consist of classifying students according to their opinions, and pairing according to their divergence.

Lakatos et al. (1984) showed the importance of social debates and discursive logic in mathematical discoveries: a class often tests the robustness of the solutions to a math problem proposed by the teacher. Reid (2002); Benbachir and Zaki (2001) showed how frequent it is that a student expresses a disturbing (apparent) contradiction. In this social debate, the teacher is crucial for the evolution of meaning and for the move from description of activity to definition (Falcade et al., 2007).

Debating a geometrical artefact is a very powerful learning mechanism. For example, Laborde (2000) noted how effective it is to be able to explain why a construction is impossible. Furthermore, explanations can be productive even if they don't correspond to the expected answer. Hölzl (1996) gave the example of a learner who was asked to find a given property of a class of triangles. Failing to do so, the learner did come up with a procedure to transform any triangle into a triangle with the expected property.

Paper interfaces are adapted to collaborative scripts in two ways. First, paper is well adapted to structuring an interaction, as it supports content in an ordered way: sheets are read from left to right (in our case), from top to bottom, and one page of a booklet after another. Second, the advantages of tangible interfaces for collaboration also apply to paper interfaces: the interaction happens in a shared space, which is horizontal and thus open to everyone around it, and the roles in the group are identified by the ownership of paper artefacts.

2.3.2 The Three Circles of Interaction with DGS

In general, the instructions of a pedagogical macroscript can have three kinds of recipients: the individual, the group, or the whole class (see Figure 2.11 for an illustration). The scripts correspond to the geometry classroom organization (Falcade et al., 2007): lab activities, where students are usually placed in pairs on computers; individual writing, where students write their own formulation of the solution; and collective discussion, where the teacher formalizes the discoveries.

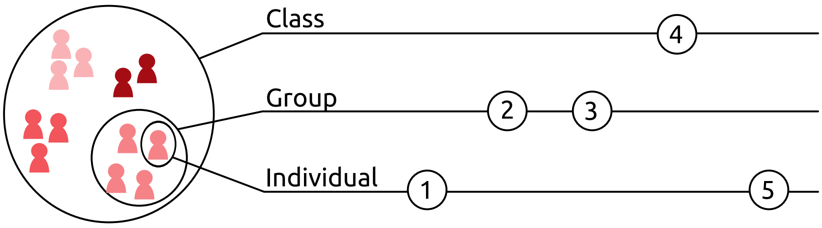

Figure 2.11: The three levels of the ArgueGraph script (Dillenbourg and Hong, 2008): (1) Students take a quiz. (2) Students are paired according to the quiz and discuss. (3) Pairs write a common questionnaire. (4) The teacher reviews and formalizes the answer of the whole class. (5) Each student writes a structured synthesis.

The class-level activities are often overlooked in DGS based education. The teacher introduces notions, validates mathematical objects by formulating them, and links the technical symbols between the pen/paper environment and the DGS environment (Laborde, 2002). DGS acts only at the individual and group levels, giving meaning to the introduced notion, allowing to explore problems and apply solutions.

2.3.3 Orchestration: the Real Conditions

With all the effort invested in these systems, approaches, and theories, one could wonder why teachers remain reluctant to integrate technology (Guin and Trouche, 1998). According to Artigue (1991), new technologies are not integrated because of specific class management issues. Chevallard (1992) explained further: "The introduction of new tools in the didactic system must not be taken for granted... It is necessary to take into account didactic permanences, problems specific to teaching and learning that even new technologies cannot avoid. Quite often, computer-based strategies take rather little account of the logistic problems left to the teacher in his/her class".

According to Rogers and Ellis (1994), the lack of consideration about how tasks are performed in situ has led to the design of many computer based systems that are unable to support the tasks designed for. The field of Human Computer Interaction often takes the individual as a unit of evaluation and collaboration-centred approaches take the group as unit of evaluation. A recent trend in Computer-Supported Collaborative Learning takes the class as an object focus (Kollar et al., 2011; Dillenbourg et al., 2011; Alcoholado et al., 2011; Do Lenh, 2012). Orchestration refers to real time classroom management of multiple activities and multiple constraints conducted by teachers. Luckin (2008) saw the "classroom as a complex technological ecosystem". This ecosystem is defined by a rigid time structure, a physical environment, and a content structure (the curriculum).

Paper interfaces can support the orchestration of a class because they are flexible: for example, if a piece of paper is used to display a hint, the teachers can keep it or leave it to the pupils, depending on what they think is best. Paper interfaces can also improve the awareness of the teachers, because as previously mentioned, the elements of a tangible interface help identify the role and activity of the users.

2.3.4 Integration

We review the adoption of technologies for learning based on two factors facilitating orchestration: integration in the classroom environment, and control of the activity by the teacher. The first criteria refers to how the system fits within the existing curriculum and practices. The second criteria refers to the simplicity and flexibility of the system, which allows the teacher to remain in control of the activity.

Regarding DGS, for example, the main issue about integration is raised by (Laborde, 2002, 2003). Laborde assumes that a tool in general is not transparent, in the sense that using it requires an appropriation of its meaning. Using another tool means using another geometry.

With DGS, new notions appear, others disappear, and others are modified. For example, constructing a parallelogram with a compass is based on congruent lines, but DGS makes it easier to use parallel lines. And yet, for some teachers, it is not possible to do geometry without pen and paper (Laborde, 2002). Furthermore, in addition to the changes in the domain knowledge, DGS modifies the teacher-student interaction, making face-to-face communications become face-to-screen communications.

The teachers may be uneasy, or simply reject DGS if they don't feel in control (Laborde, 2002; Clarou and Laborde, 2000). The most important reason for this rejection concerns the technical mastery and the knowledge of the given DGS. Complexity or unexpected reactions can easily scare off someone who has to manage more than twenty pupils perform a task in a limited time. Even if the system is simple, teachers are lost if the practices are too different from normal practices. In any case, any kind of teaching innovation provokes time inflation (Schneider, 1998).

Beyond the design of the technology, its use has to be taken care of. As noted by Ruthven (2002), the difficulty of designing tasks in computer environments is finding a balance between two extremes: too high of a technical or conceptual demand, or too much control.

The danger of pushing the teacher from "sage on the stage" to a "guide on the side" emerges, however. Since the teacher is the first user of a technology in classrooms, the adoption of a new technology is an obstacle. It is thus important to focus not only on the individual and groups of students, but also on the teacher and the class.

There are other constraints to take into account, even if the teacher is convinced. Most notably, the curriculum is an important criteria to satisfy when designing a pedagogical system. This is especially critical in our context, as computer environments shape students' solution strategies (Noss and Hoyles, 1996). It is thus important not to invent a new geometry but rather augment the existing one (i.e. the Euclidean one, based on straight edge and compass construction).

In terms of integration, paper interfaces are most pertinent. All activities in the classroom are based on paper: exercise sheets, books, etc. The paper used as an interface not only does not replace traditional paper items like these but can actually be used in combination with these types of items. Another advantage is that paper used as an interface can be brought home. Furthermore, several aspects of the geometry curriculum directly involve paper. For example, pupils have to learn how to use a protractor, or draw a given figure.

2.3.5 Indicators

Another aspect to take into account in order to support orchestration is the acquisition of information by the teacher in the ongoing activity, and the class in general. A technological tool can do this in two ways: passively or actively.

Passive acquisition of monitoring information can happen easily with a paper interface due to its tangible aspect: it is easy for the teachers to observe who does what by simply looking at who manipulates which part of the interface.

The active way of supporting the monitoring routine of orchestration is by delivering indicators to the teacher. Interaction Analysis (Dimitrakopoulou et al., 2006) consists of automatically analysing a computer-mediated activity to produce indicators to be reused or to qualify or adapt the activity. These indicators are typically computed from the traces of an interaction between learners and computers via a mouse and keyboard. However, this is also possible with tangible interactions, and these indicators are even richer because interaction happens in the real world.

These indicators are not automatically applicable, because they do not directly qualify the learning activity. For example, an indicator showing hesitation is a good thing if it corresponds to reflection, but too much hesitation is not productive; too little hesitation is acceptable if it is the result of an exercise being too easy for a group but not if it is the result of a lucky guess. Indicators are thus the basis of a retroactive loop on the learning group. The indicators need to be interpreted within their context, and this interpretation can be used to generate feedback about the group.

Many TEL approaches are based on this principle, explained in more detail by Soller et al. (2005). They show three families of systems that support the management of collaborative learning interaction, which differ in the exploitation of the indicators. First, with a mirroring tool, pupils are shown the indicators and left free to decide what to do with them. Second, meta-cognitive tools, similar to mirroring tools, let the pupils decide what to do with indicators but also show a desired value for the indicators, so that the pupils aim at a desired state of interaction. Third, guiding tools actually process the indicators and propose a remedial action for the pupils to reach a desired interaction state.

Paper interfaces can draw from their tangible nature to extract indicators. By tracking the movements of the interface in the physical space, the information that can be extracted is much more complete than only the movement of a mouse for example. Furthermore, paper interfaces make it easy to duplicate the tangible components for each pupil. This way, it is possible to easily identify not only which piece of interface is being used but also by whom. This opens up a myriad of possibilities in terms of indicator extraction.

2.3.6 Map: Teaching Geometry with our Activities

Our second map illustrate the various topics that have been discussed in this section by situating our activities against them:

- Most of our activities are simple collaborative *scripts*. However, SpaceJunk is the most advanced on this topic. It builds on the SWISH principle Dillenbourg and Hong (2008) to foster an interaction about angles: by preventing pupils to share the vision of an angle, they have to describe it verbally, hence leading them to describe angles.

- The *three circles of interaction* are best illustrated by the activities based on StarrySheets (AngleHunt, Sympliage, Triangram and Messangles): pupils work individually on an exercise given by their peers, they may have to discuss with them about who is wrong, and the teacher arbiters and drives the whole activity.
- *Orchestration* was an important component of all the activities that have been tested with a whole class at once. After a first try with SpaceJunk in a controlled environment with a full class, we ran the experiments based on Kaleidoscope and StarrySheets directly in the classroom, to investigate for example how to manage a shared, limited digital resource (the camera-projector system).
- *Integration* is one of the main goals of this dissertation. The most successful activity in this domain is Kaleidoscope: teachers designed their own activity independently, ran them autonomously, and even reused the exercises in a non augmented setting. Angoli also shows how paper interfaces integrate with the curriculum, because it allows to learn the usage of a protractor.
- *Indicators* are the main result of the study involving Quads. Using simple data from the manipulation of the paper interface, we could draw several kinds of information about the learning process.

2.4 Conclusion

We reviewed work related to interfaces in geometry education. Four kind of interfaces (command line, graphical user interfaces, tangibles, and digital pens) were used to expose various notions of geometry and pedagogy. We illustrated how paper interfaces could combine the advantages of these approaches. We then reviewed ambiguities that were challenging for learning geometry, and again related them to paper interfaces. Finally, we reviewed more general pedagogical concepts, and showed how they could be deployed with paper interfaces. Paper interfaces thus seem like a promising approach to geometry learning.

The importance of paper and computers for geometry is stated in the conclusion of Laborde (2003):

> “Paper technology and printing technology certainly played an important role in the development of mathematics by facilitating the representation of mathematical ideas and the expression of relations by spatial configurations on sheets of paper, not only in geometry but also in arithmetic and algebra. Dynamic geometry environments introduce a spatial representation of another key feature of mathematics, the variability. We are only at the beginning stages of taking advantage of the semiotic mediation potential of this new dimension in the teaching

> of mathematics. But even in this era of great possibilities offered by dynamic geometry environments, there is a need for pursuing the reflections on the design of interface.”

On these grounds, using paper as an interface for computer-based geometry is most adapted, because it combines the contributions of both technologies to address the dualities in geometry education. Paper has supported the two semiotic registers used by geometry (language and diagrams) for centuries. The variable diagrams made possible with computers makes it possible to exploit drawings and figures more extensively. Computers are not limited to a “microsurface”: projectors create mesosurfaces and the variability allows the exploration of macrospaces. This exploration is one of the possible benefits for the empirical aspect of geometry, and the computational power of computers are instrumental for assisting deductions.

The organisation of this chapter also reflects three questions to investigate about the use of paper interfaces for geometry. The first section concerned purely interfaces, i.e. the interaction aspects. The second section concerned how pupils learned. The third section concerned the use of technologies from the teacher’s perspective. We explained how paper interfaces related to these topics, but remained hypothetical. The goal of this dissertation is to answer these three questions: how do pupils use paper interfaces? How do pupils learn geometry? How do teachers use a paper interface to teach geometry? We explain how we will address these questions in the next chapter.

3 Research Questions and Method

This dissertation is the result of the exploration of paper interfaces for learning geometry. In this chapter, we expose how we carried out this exploration. We start by defining the goals of the exploration, by stating three research questions. In the rest of the chapter, we describe the means of the exploration. First, we present the design space that provided us a framework to create activities that make use of the possibilities of paper. Second, we describe our approach, which is inspired by a family of research methods that suits our objectives: Design Based Research. Design Based Research aims at grounding theoretical work in real contexts, to improve it, and to reuse it in subsequent experimentations. Such an approach is thus well adapted to exploratory work. Finally, we present the system used to implement and test our designs, as well as the tools developed to analyse the collected data.

This chapter is a good opportunity to complete the maps given in Chapter 1 and Chapter 2. In this chapter, we will situate the activities developed in relation to the participants of the Design Based Research, i.e. how the teachers were involved in the design, and how the pupils tested the various iterations. We also situate the activities based on the methods of analysis that were applied in the studies.

3.1 Research Goals

As we saw in the conclusion of Chapter 2, our exploration can be phrased around three topics: the use of paper interfaces by pupils, their effects on learning, and how teachers are involved in this process. We thus formulate the following three research questions:

1. *How do pupils use a paper interface?*
2. *How do pupils learn with a paper interface?*
3. *How do teachers use a paper interface?*

The first question aims at studying how children, as users, interact with paper interfaces. The second question aims at studying the effects of paper interfaces for learning, e.g. how children as pupils can use them to explore a problem, collaborate, etc. The last question aims at studying how a teacher can use a paper interface, from the design of an activity to its management in a classroom.

These questions remain very open. In the studies that we will report on in this dissertation, we will actually investigate more specific questions. We will mark these specific questions with (Q_u) , (Q_l) , and (Q_t) , depending on whether they relate to the first, second, or third general question, related to usability, learning, and teaching, respectively. In Chapter 9, we will reassemble all the specific answers into one answer for each of the general questions.

3.2 Framework of the Design Space

In order to direct the answers to our research questions, we devised a framework based on two orthogonal axes: the properties of paper and the properties of interaction. This section describes these two axes, and illustrates the space they cover.

3.2.1 Computational Model for Paper Input Devices

Paper can be used as an input device. This has already been illustrated with approaches based on digital pens, for example, which emulate the keyboard to input text in some way. In this case, paper is used as a black and white scanning device: the computer only captures the image of the annotations written on the paper, via the digital pen.

Digital pens are also used as pointers, in the same way that graphical tablets can replace computer mice. For example, a tap with a digital pen on a specially delimited area of an augmented sheet can trigger an action. Similarly, augmented reality approaches often use paper elements as three dimensional cursors. The markers printed on tangible elements (e.g. an augmented textbook) specify an input device with up to six degrees of freedom (a three dimensional position and a three dimensional orientation), which has been extensively studied by Card et al. (1991).

Paper is rarely used beyond pointing and ink capture. However, paper is more than an object in space, or a black and white scanning device. In addition to being moved and written on, it can be cut, folded, flipped, shown, burned, etc. While burning seems like a problematic type of interaction, the other actions can be detected and used as input.

In order to ground our exploration, we can define a model of paper as an input device, and investigate each component. For example, we can model paper as a surface delimited by a set of curves by defining the edges, with a three dimensional position, or with a three dimensional rotation. We can define the curvature of paper by adding a stress function applied to the surface, the folds by a sequence of axes and rotations, and the content by a mapping from the surface to the ink. Then we would need to make the model extensible for more exotic transformations, such as burning, scratching, getting it wet, etc. Finally, we can derive other properties by transforming the components of this model: speed is the derivate of the position, semantic relationships could link the content of different pieces of paper, several pieces of paper could be combined, etc. Such a model would be impractically complex.

Instead, we define seven attributes of pieces of paper:

- The *presence* of a piece of paper refers to its existence in some referential. Mostly, the presence corresponds to the visibility, by a human eye or a computer camera. However, a piece of paper also exists if it is not visible; interactions can also be built on the absence of a piece of paper, i.e. if a user of a system knows that it exists, but does not see it. For example, a binary option can be activated depending on the absence or presence of a piece of paper.
- The *position* of paper is two dimensional in our context: for flat geometry, we only need to know the 2D position of a piece of paper on the interaction surface. However, paper can also be stacked. In this case, a third dimension, the altitude, or the order in which the pieces of paper are stacked, can provide another kind of input.
- The *orientation* of paper is a simple angle in our context: again, in flat geometry, it is enough to know the orientation of the axis relative to the interaction surface.
- The *side* of a piece of paper is technically given by another rotation. However, it is more intuitive to define a separate property, as it is used differently than the orientation previously defined.
- The *folds* of a piece of paper simply correspond to the places where it is folded.
- The *edges* correspond, of course, to the outside borders of the piece of paper, but also to the holes cut inside.
- The *ink* corresponds to an image of the content on the sheet. It can be printed, annotated, or even added ephemerally by a projector.

These characteristics are quite intuitive for a user. For example, users will easily understand what “set this piece of paper to this position/orientation/side/fold” means; users know how to change the edges or the ink on a sheet if they are given the right tools (scissors or pens, for example). Furthermore, these characteristics can be captured by a computer device. They can also be used to compute other kinds of input. For example, the derivate of the position can give the speed and acceleration; the derivate of the presence or of the side tells a computer system whether a piece of paper has appeared or whether it has been flipped; the distance between two positions can tell whether two pieces of paper are close to each other.

Such a list of characteristics aims at revealing interaction techniques. When we design an activity, and need to define an action, we can browse this list and imagine intuitive ways of carrying out the action that may not be obvious.

3.2.2 Temporal Structure of Paper Interfaces

The attributes of paper defined in the previous section are not directly applicable – in our case, not directly translatable into a pedagogical activity. The attributes of the model simply allow the computational device to read the state of a piece of paper, but do not define the interaction. The interaction is defined by the temporal sequence of states. For example, if the edges of a piece of paper change, this means that the piece of paper has been cut.

We break down the temporal structure of an activity into three classes: permanent components, persistent components, and ephemeral components.

- Permanent components define the activity: the printed *content* of the page, the *edges* of the pieces of paper, the axis where paper should be *folded*, etc.

- In contrast, ephemeral components of the activity correspond to the atomic actions done by the user and the temporary feedback given by the system: the *content* that is projected, the distance between the *positions* of two pieces of paper, the *side* on which a piece of paper is placed, etc.

- In between, persistent components of an activity correspond to the unerasable traces left during the activity: the *content* written by the pupils, the *folds* imprinted on the pieces of paper, the *positions* fixed by sticking pieces of paper together, etc.

Figure 3.1 illustrates the life time effect of interactions for each of these three categories. Ephemeral actions can happen only during the period of interaction. Persistent actions happen during the interaction, but leave marks that exist afterwards. Permanent actions have been done before the interaction, and last afterwards.

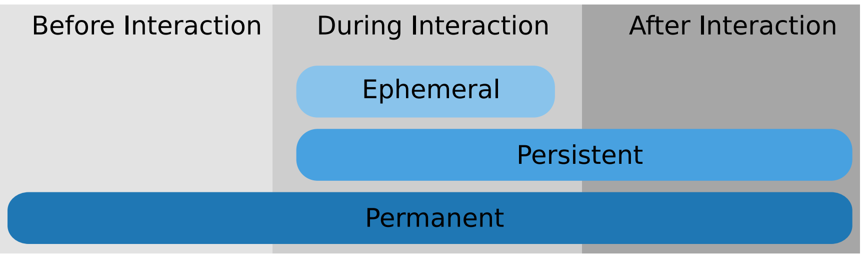

Figure 3.1: The life time of our three categories of paper interaction.

3.2.3 A Map of Paper Interaction

In Table 3.1 we illustrate how these two previously defined dimensions of the design space cover paper interactions: we give an example for each combination.

Table 3.1: Examples of actions covering the space defined by the framework.

Property	Ephemeral	Persistent	Permanent
Presence	A pupil *shows* a card to start an exercise.	A pupil *keeps* an exercise card if it was solved.	A teacher *creates* another card that shows the answer.
Position	A pupil *moves* a card that controls a pointer.	A pupil *tapes* a paper cut-out on a grid.	A teacher *binds* a sequence of exercise sheets in a booklet.
Orientation	A pupil *rotates* a card to increase a parameter of a simulation.	A pupil *staples* a paper clock hand inside a drawn clock.	A teacher *lays out* the expected orientation of the clock hands.
Side	A pupil *flips* a card to trigger a feedback.	A pupil *sticks* a paper coin facing head or tail up.	A teacher *prints in duplex* the controllers of mutually exclusive feedback.
Folds	A pupil *opens* a folded sheet to display the symmetric image of a part of a sheet.	A pupil *folds* a paper shape along a symmetric axis.	A teacher *defines an axis* of symmetry on an exercise sheet.
Edges	A pupil *juxtaposes* complementing cardboard shapes.	A pupil *cuts* a paper trapezoid to reassemble it into a rectangle.	A teacher *defines the shape* that will be cut-out and manipulated.
Ink	A pupil *draws with pencil* construction lines of a figure.	A pupil *writes* the answers on an exercise sheet.	A teacher *prints* the instructions of an exercise sheet.

Each activity will be positioned on this map, but in the following chapters. Doing it now would result in an overload of details on the interfaces that are not pertinent to this chapter.

3.3 Method

In order to answer our research question guided by the framework, we carried out Design Based Research. In this section, we first present this approach. We then report on the recruiting of the teachers who collaborated with us. We also present their classes, i.e. the pupils who used our designs. Finally, we present the material of the experiments, i.e. the geometry curriculum.

3.3.1 Design Based Research

Technology Enhanced Learning (TEL) is a research field that aims at supporting education through technology. It ranges from developing new educational theories to implementing and deploying systems for education. In this context, the detachment of research from practice is problematic: theories without practical applications are worthless, but it is hard to apply theories to produce a system that works in any condition. In other words, there is a tension between fine-tuning a system for a given experimental set-up, and proposing a method to create systems adapted to a wide range of situations. The difficulty is to identify whether a TEL system has been successful because of the theory driving the development, the particular adaptation of the system to its set-up, or a combination of both.

Design Based Research (DBR) addresses this issue. It finds its origins in the seminal papers of Brown (1992); Collins (1992). We refer the reader to Wang and Hannafin (2005) for a comparison of the variations in DBR, and for selected work.

DBR blends theory-driven design of environments and empirical educational research in real contexts in order to understand how the theoretical claims can be transformed into effective learning in educational settings (Collective, 2003). In order to create usable knowledge on how, when, and why educational innovations work in practice, DBR approaches iteratively gather theoretical knowledge about a topic, implement the theory into the design of an artefact, test the system in real settings, and refine the theory according to the findings.

DBR promotes the deployment into real, complex settings, because of the strong link between cognition and learning. Barab and Squire (2004) explained that cognition is not located within an individual thinker but rather corresponds to a process distributed between the knower, the environment in which knowing occurs, and the activity. With this perspective, isolating variables – as intended in laboratory or otherwise controlled experiments – leads to an incomplete understanding of their importance.

In addition to this *contextual* characteristic of DBR, Wang and Hannafin (2005) identified the following:

- pragmatic: in addition to asking whether a theory works, DBR investigates how well it works;
- grounded: DBR identifies and addresses existing problems and issues;
- interactive, iterative, and flexible: DBR stresses the collaboration between researchers and participants: while other research approaches use non-researchers as simple subjects, DBR involves them in the design process and its evaluation;
- integrative: DBR is not a rigid method; rather, it is based on several existing approaches for the design and evaluation.

DBR promises four kind of benefits (Collective, 2003). First, it allows the exploration of novel learning and teaching environments, and favours their adoption thanks to the involvement of the real actors (e.g. teachers), placing them in direct ownership of designs. Second, it allows the development of better theories about the components of the learning context. Third, it increases knowledge on design. Fourth, DBR improves the capacity of educational innovations by providing insights emerging from complex settings.

Of course, there are challenges to keep in mind to fulfil these promises. A first challenge concerns the evaluation: since the researcher is both the advocate and the critique of a theory, it may be difficult to distinguish fundamental failures from a need of adaptations. One possible solution is to triangulate multiple sources of data and results to abstract from the context. This approach also helps capture the context of an activity, and the myriad of independent decisions intervening in its design or its application. It is important to capture as much information from the context as possible, so that the model created by an analysis can be refined in subsequent studies, but still evaluated with the data captured in the previous analyses.

In order to address these challenges, Wang and Hannafin (2005) formulated nine principles to follow in DBR:

1. Support design with research from the outset, i.e. identify relevant resources from the related work
2. Set practical goals for theory development and develop an initial plan
3. Conduct research in representative real-world settings
4. Collaborate closely with participants
5. Implement research methods systematically and purposefully
6. Analyze data immediately, continuously, and retrospectively
7. Refine designs continually
8. Document contextual influences with design principles
9. Validate the generalizability of the design

3.3.2 Participants

Design Based Research depends on the involvement of the participants, and the ability to deploy and test the design iterations in real conditions. Our first step was hence to start a collaboration with teachers, who would provide us with their insights on the designs, and implement these insights within their classes. This section reports on the various channels that we used to deploy our system in real contexts.

In the Swiss Confederation, each *canton* is the main authority for its school system. We invited a delegation of the head of pedagogical matters of the *canton de Vaud* (where our university is located), and the *canton de Neuchâtel*, a neighbouring *canton*. We also invited a delegation of the pedagogical services from the *canton de Genève* and the director of School 1[1]. Each of these delegations came on separate occasions.

The socio-cultural level in Switzerland is generally rather high. The schools we collaborated with were usually in the better half, since teachers and administrations interested in collaboration with a university have already meet their primary needs. Still, we collaborated with a socially heterogeneous set of schools: a private school in a small city, a public school in two small small cities, and a public school in the suburban area of a big city. The small cities have slightly higher socio-economic levels than the big city.

We gave the delegations a demonstration of paper interfaces with a sample activity in order to show the possibilities offered by the system. Their main concern was the accordance to the curriculum, and the coverage of its other domains, not just geometry or mathematics. Even if it would be a longer term goal, they were easily convinced that the approach of paper as an interface could be easily adapted to these other domains. The delegations from Neuchâtel, Vaud, and Geneva recommended us to voluntary teachers, while the director of the private school invited us for an informal pilot study in his school.

We started our collaboration with five teachers from schools in School 2, School 3, and School 4. Each of them were invited for demonstrations in our laboratory. The subsequent collaboration depended on the activity, and is reported in the following chapters. All their classes tested our designs in their schools.

We were also able to evaluate our designs in conditions relatively close to those of the classroom from our collaboration with the *Equal Opportunity Office* of our university, with which we organised two workshops. The first one was part of an open door event, in which we received seven classes from various schools from the French speaking part of Switzerland. The second workshop was part of a *Take Your Child To Work Day*, where we received two groups of employees' children.

Table 3.2 gives an overview of the groups of pupils who tested our designs.

3.3.3 Maps: the Iterations of our Activities

One of the goals of this dissertation was the credible integration of paper interfaces in a valid context. Several iterations allowed us to complete this goal on two levels: the involvement of the participants, and the physical place where the activities were run.

---

[1] We numbered the schools for anonymity reasons.

Table 3.2: Summary of the groups of subjects who tested our designs.

School	Year	Grade	Age Range	Number of pupils
School 1	2011	5	10-11	8
School 1	2011	5	10-11	9
School 1	2011	6	11-12	5
School 2	2011	5	10-11	19
School 2	2011	5	10-11	20
School 2	2012	5	10-11	20
Open Door Workshop	2011	5 & 6	10-12	7 classes of ~20
Open Door Workshop	2012	5 & 6	10-12	7 classes of ~20
School 3	2011	6	11-12	20
School 3	2011	6	11-12	21
School 3	2012	(mixed) 3 & 5	8-9 & 10-11	9+10
School 3	2012	4	9-10	18
Take Your Child To Work Day	2012	-	10-13	20
Take Your Child To Work Day	2012	-	10-13	9
School 4	2012	5	10-11	20

Involvement of the Participants

The participants were gradually involved in the creation of the activities:

1. Quads was a simple proof of concept to illustrate the possibilities of a paper interface. It had been developed before the first contact with teachers. No other participant was involved.

2. During a design session with teachers, we decided to create a problem-based activity involving the protractor. The teachers recommended that we precede this problem with an introductory exercise to get familiar with the tool. The teachers only identified the difficulty of using a protractor, and we were left free to design the rest of Angoli. It was thus a very short punctual involvement.

3. The rest of the design sessions with teachers were spent discussing the problem-based activity, which became SpaceJunk, thus the result of a co-design. It was thus a longer punctual involvement.

4. Kaleidoscope allowed the teachers to develop the activities themselves, independent of us. Furthermore, they ran the activity without us intervening.

5. The StarrySheets activity allowed the pupils to define their own exercises and give them to their peers.

Spatial Configuration

The location of the studies was also increasingly pertinent:

1. Quads, Angoli and SpaceJunk were experimented with in a small room away from the regular classroom.

2. SpaceJunk was further experimented in a setting that is more ecologically valid, but still controlled: a room big enough to accommodate a full class, but still in our laboratory. This was an intermediary step to the final destination.

3. The ultimate goal was the regular classroom. It was achieved with the activities based on Kaleidoscope and StarrySheets.

3.3.4 Material

Our Design Based Research does not aim at revolutionizing geometry didactics but rather at evaluating the benefits of paper interfaces for geometry education as it exists already. Moreover, it was clear from the first visits of the various stakeholders that their main concern was the ability to abide by the curriculum. In this section, we hence give an overview of the geometry curriculum.

On May 21^{st}, 2006, all the French speaking *cantons* of the Swiss Confederation approved a reform of the education, unifying the curricula. The *Plan d'Études Romand* (PER) gives high-level descriptions of the progression and expectations in 23 disciplines from 8 domains (Conférence Intercantonale de l'Instruction Public de la Suisse romande et du Tessin, 2010).

In the discipline of mathematics, 5 categories of skills are defined: space, numbers, operations, measures, and modeling. Appendix A summarizes the PER in the categories relevant to geometry, i.e. space and measures. We decided to target a young age range of pupils (9-12 year olds), because the subjects are not taught at a very abstract level, which means the manipulability of paper can be explored more intensively.

3.4 Metroscope

In this section, we present Metroscope, the system for supporting geometry that we iteratively developed during our research. We start by presenting the Metroscope, the system used to augment paper, and justify its choice by comparing it to other augmentation technologies. We also report on the technical contribution regarding the precision of the mapping between the captured image and the projected image. Finally, we describe the technical infrastructure used to collect and analyse data.

3.4.1 Metroscope: A Camera-Projector System

The Metroscope is shown in Figure 3.2. This system has already been developed in the context of two other doctoral theses in our laboratory, under the name of TinkerLamp (Zufferey, 2010; Do Lenh, 2012). It incorporates a camera and a projector directed to the tabletop surface via a mirror held above the desktop by two metallic arms. The camera and the projector are connected to a computer embedded in the case, so that the interaction with the hardware is minimum for the end user: switch on or off. Its only requirement is an electric outlet.

Figure 3.2: The Metroscope, our camera-projector system, along with various types of objects which can be augmented on a table: sheets, cards, tools and wooden blocks.

The projector has a resolution of 1280 × 768 pixels, which is mapped to an area of about 70 × 40 centimetres on the tabletop holding the lamp. The camera has a resolution of 1280 × 768 pixels and captures a region larger than the projection area by a few centimetres. It can capture up to 15 gray-level images per second using a global shutter, which means that the image is recorded all at once, as opposed to rolling shutters capturing lines one after another and assembling them. This is important, because the projector updates its image line by line, and color channel after color channel, which means that a rolling shutter would capture disturbing artefacts, as those shown in Figure 3.3

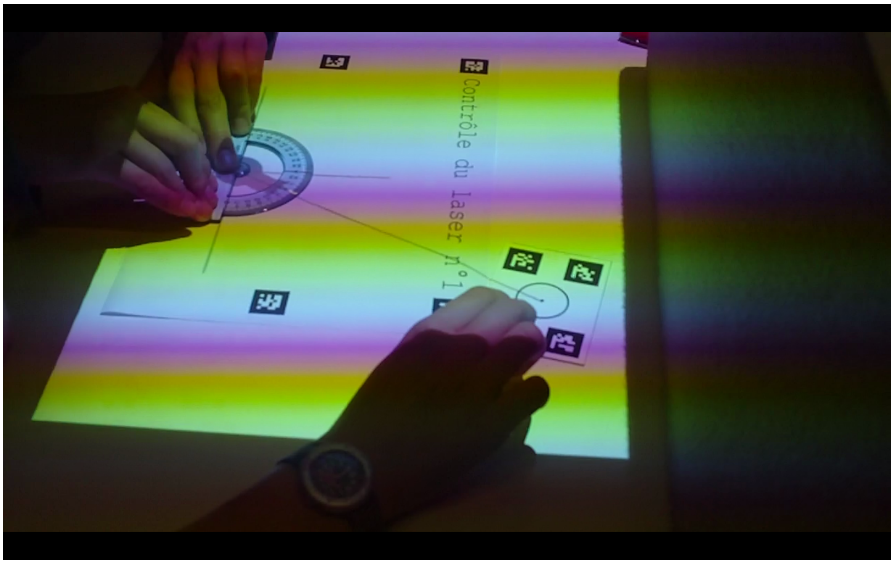

Figure 3.3: Artefacts resulting from capturing the picture of a projection using a rolling shutter camera. The Metroscope uses a global shutter which avoids these artefacts.

The Metroscope has already been used in a mature environment in the context of education (logistics). Its main use was to simulate a warehouse by projecting augmentations on tangible mock-ups of shelves. However, we needed one main technical adaptation to use the Metroscope in the context of education in geometry: the calibration.

Geometry requires very high precision: if a point is projected more than one millimetre away from its expected position, the system is not usable by pupils who are taught to be precise with their tools. We adapted the existing calibration method to achieve a sub-millimetre precision, under the condition of working on one plane only. The technical details are given in Appendix B

3.4.2 Comparison with Other Augmenting Technologies

The Metroscope is a variation of the DigitalDesk (Wellner, 1991) in the sense that it consists of a camera and a projector directed at a tabletop from above. It shares many similarities with the system described by Wilson (2005). In Chapter 1, we saw many technological alternatives to a camera-projector system. We placed these alternatives in the decision tree in Figure 3.4 to justify our choice.

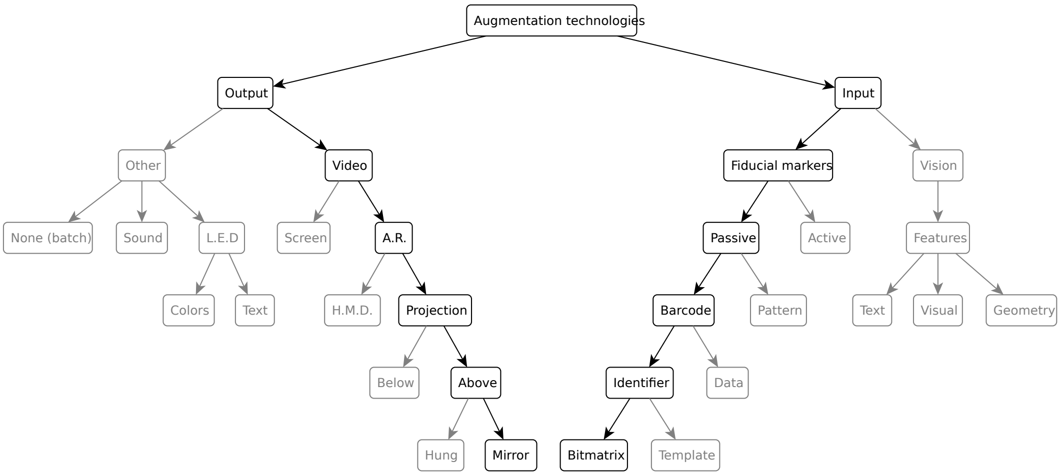

Figure 3.4: A decision tree showing the technological alternatives to augment paper.

Output Technology

The first choice concerns the definition of the output. For example, a lot of digital pens have no interactive output (the annotations are gathered for offline batch processing), or a low definition output, like LEDs, or sound. Since geometry is strongly graphical, projecting images is very beneficial.

Images can be displayed on a screen, or projected directly on the paper or physical elements. The former suffers from the *split attention* effect, i.e. an extraneous cognitive load resulting for a learner having to switch between two sources of information. Head Mounted Devices provide a more integrated experience, but Kaufmann and Dünser (2007) reported that such displays can cause *simulator sickness*, similar to motion sickness; they concluded that one hour is too long to use such a system, which causes exhaustion and eye strain.

At the expense of the mobility of the system, projectors can be used to display information on the interaction surface. Dillenbourg and Evans (2011) compared the advantages of projections from above versus projections from below, on a semi-transparent surface. Projecting from below restricts the projection to the interactive surface, while projection from above allows for the display of information on other elements, such as paper or tangible objects. Conversely, these other elements can provoke unwanted occlusions and shadows. We chose the projection from above system, and used a mirror instead of hanging the projector above the table in order to increase the focal distance and thus the interactive surface.

Input Technology

There are several ways to track paper, or more generally the actions of the user. A first design decision concerns whether the augmented element can be tagged with a fiducial marker, or if the system should use computer vision techniques to recognize unprocessed elements.

Beyond the debate between seamless and seamful augmentations, fiducial markers are, by design, reliable. Reliability is important in a pedagogical context to avoid distracting the pupils or wasting the teacher's time. For example, lighting conditions cannot be controlled, but the system should not break if the day is not as sunny as expected (see Figure 3.5).

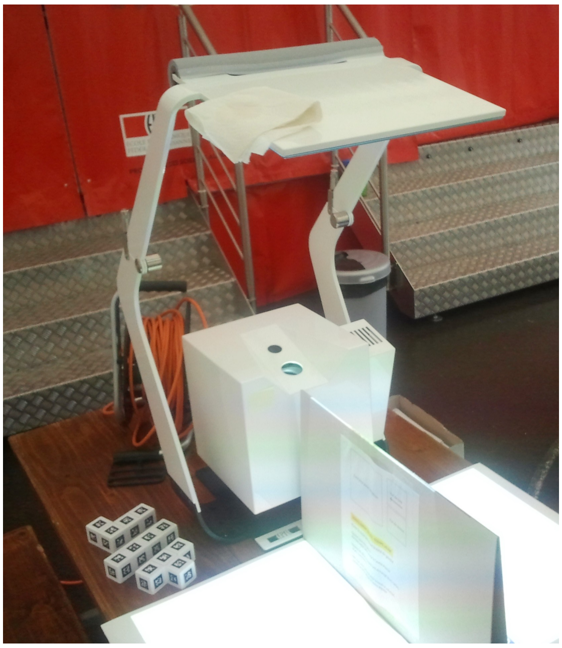

Figure 3.5: The Metroscope in less than ideal (lighting and humidity) conditions. Note the tissue on the top to absorb the water dripping from the roof of the tent.

The markers can be active, e.g. RFID or other electronic beacons, but this immediately raises the production costs. In contrast, it is very easy to produce tagged paper, using regular printers. This does not limit the augmented elements to paper: the barcodes can be printed on paper, and stuck on anything else.

Some fiducial markers (e.g. QR-Codes) encode arbitrary data (e.g. a URL or other alphanumeric content), but they are not designed to be precisely located. In contrast, markers for augmented reality simply code identifiers, in the form of a bitmatrix or template easily recognizable with computer vision techniques. The software then maps the identifiers to their meaning in the application. We use fiducial markers inspired by ARTags[2] for historical reasons: the ARTag system has been the most reliable with the previous versions of the Metroscope, and we redeveloped the system internally when the support ceased.

3.4.3 Data Analysis

Finally, since Design Based Research relies on data collection to capture the context of experimentation as much as possible and analyse it several times, we introduce the research

---

[2] http://www.artag.net/

prototypes that we designed, describe the instrumentation of the system to collect data, and present a tool to explore it.

Prototype Development

There are two sides to the paper interfaces used as prototype designed and evaluated in each iteration of our Design Based Research. The paper elements of the interfaces have to be designed, and the corresponding software developed. We strived to use each side of the paper interface where it is best: static elements are best printed in high resolution, while dynamic elements are controlled by the software via the projection.

The paper side can be done with any computer editor, or even by hand if there are not tags. We favoured the svg format, because it was easy to manipulate from an interface generation script, which would generate both a document to print, and a configuration file to run. Such a configuration files allows to describe to the software the arrangement of markers that will be detected, and links them to the software components that will generate the augmentations.

The software was written in C++. Programming a paper interface has implications beyond the components related to the interaction: the computer can use an additional memory, stored in the paper. For example, the state of a running program can not only be stored in the main memory of the computer, but also in the position of the pieces of paper relatively to each other, or in the orientation of another piece of paper. This is an interesting property to move the activity away from the computer into the paper, where the Metroscope is simply a reactive tool rather than the central component to interact with.

Data Collection

For almost all the experiments, the only input to our system was the positions of the detected fiducial markers. This is an opportunity to simply log these positions along with a timestamp, and be able to rerun the software with the exact same input as the one received during the experiment. This is particularly useful to instrument the software with additional outputs, such as the logs of a particular event, even if this event was not determined before the experiment.

In one experiment, the software also needed the image captured by the camera to reproject a part of it on the table. For this reason, and also to capture the hands of the pupils manipulating the interface elements, we saved a snapshot of the image captured by the camera every second.

We also recorded two kinds of videos. The first one was taken by a panoramic camera placed between the lamp and the projection area. It enabled us to see the pupils around the interaction surface. The other kind of video was taken by the webcam of laptop computers, and captured a more global view, i.e. the room. The panoramic camera also captured the audio around the Metroscope, and the computers were linked to a high quality ambient microphone. Figure 3.6 shows one of the set-ups in a classroom.

(a) A computer facing the class was placed on the teacher's desk; its webcam filmed the class from the front.

(b) Another laptop computer was placed on top of one of the classroom computers; its webcam filmed the class from the back. An ambient microphone was placed near the laptop in the back to capture the sound in the classroom. A panoramic camera was placed under the Metroscope.

Figure 3.6: Pictures of the experimental set-up at the front of a classroom (left) and the back (right).

Analysis Tool

In order to explore all this data, we developed a simple tool in order to synchronize the display of the various multimedia content: the captured videos, the snapshots taken by the camera, and a reproduction of what the lamp was projecting, generated from the log of the positions of the markers. Additionally, this tool was used to generate more data, e.g. by providing a simple GUI to annotate video. Figure 3.7 shows a screenshot of the software.

Map: the Analyses of our Activities

There is no standard way to analyse data collected with a paper interface. We borrowed methods from other domains or invented ad hoc methods. This map gives an overview of the analyses that we did with our activities:

- The learning outcome is a standard measure of learning: a test is given before and after the use of an interface, and both scores are compared. We used this method in Quads and Angoli.
- We wanted to investigate the movements of the cards in SpaceJunk. To do so, we borrowed techniques used by the eye-tracking community to segment eye movements, and applied them to the extraction of trajectories.
- We decomposed the data of the manipulation of the various components of the interface of Quads until we found meaningful indicators about the activity.

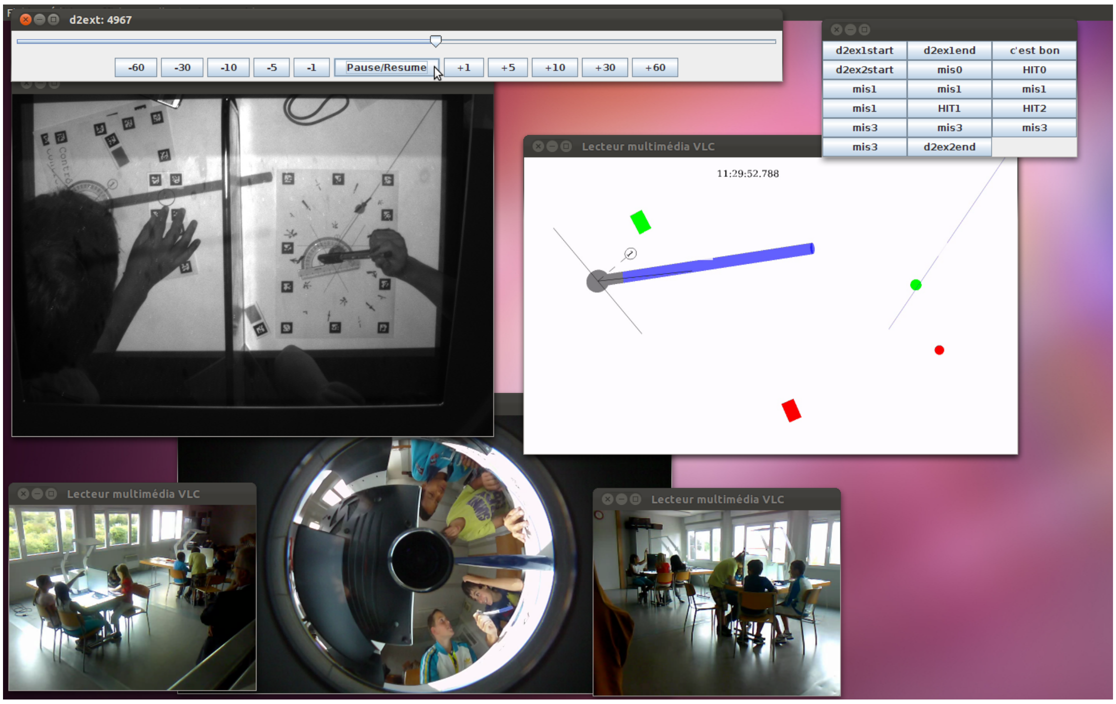

Figure 3.7: Screenshot of the tool, with which a single time line (top left) or user defined bookmarks (top right) allow the user to synchronize the display of several multimedia data. In this screenshot, from top to bottom, and left to right: a series of snapshots taken from the camera of the Metroscope, a screencast of the projection by the Metroscope, a video from the left side of the room, the panoramic video from below the lamp, and a video from the right side of the room.

- We applied various statistical analyses to check whether an observed phenomenon could be modelled and to ensure that it was not random. We did so to check the relevance of the previous points.
- When we found pertinent results, or to guide hypotheses, we created various visualisations. This was especially important for SpaceJunk.
- Qualitative observations were also a rich source of information about the use of paper interfaces. We applied them to all the activities.
- We also completed observations with direct feedback from pupils: we collected simple comments on sticky notes for SpaceJunk and AngleHunt.
- On the teacher side, we interviewed some of them formally, in addition to the casual debriefing after each study.

4 Classifying Quadrilaterals

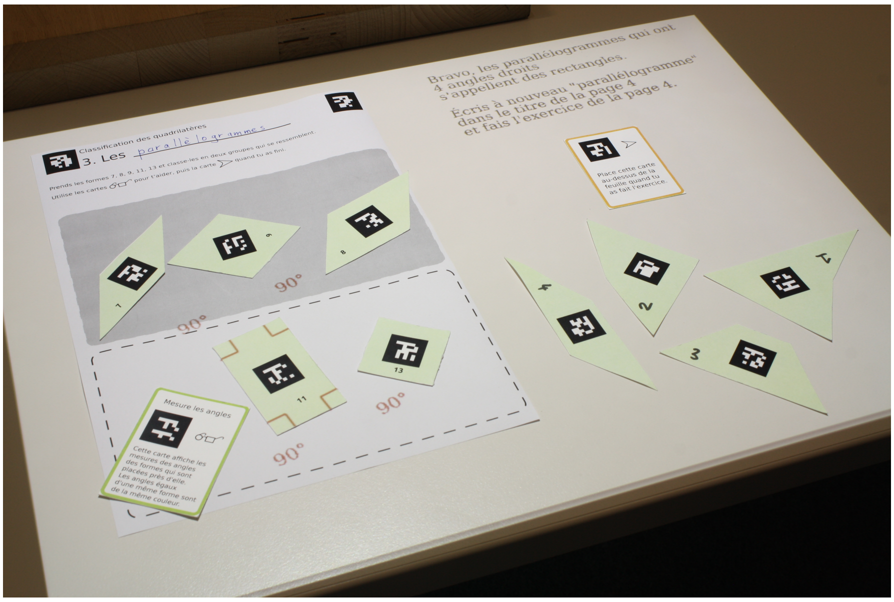

Figure 4.1: The components of Quads. (Left) Five cardboard quadrilaterals are classified into two groups on the instruction sheet, a card displays the measure of the angles of a rectangle. (Right) The feedback card displays the validation text.

This chapter presents the first paper interface tested by pupils, in the form of a pedagogical activity named Quads. It was developed as a demonstration of paper interfaces to the various stakeholders we invited to collaborate with us (i.e. directors of pedagogical matters and teachers). We had the opportunity to test it during two studies in primary schools. We first describe this activity. We then present the studies and their results and conclude by summarizing the findings. This chapter illustrates that one advantages of paper interfaces is that they allow to extract information on the activity of pupils.

4.1 Description of the Activity

Quads is a pedagogical script to introduce the classification of quadrilaterals: squares, rectangles, rhombuses, parallelograms, and trapezoids. Its interface is shown in Figure 4.1. The materials include a leaflet, a set of numbered, quadrilateral cardboard shapes, and four cards. The activity is a sequence of seven exercises - one per page of the leaflet. The main task consists of placing shapes into sets, e.g. rhombus versus non-rhombus. Following, we present each element of the interface separately, describe the scenario for using these elements, explain the design goals behind this activity, and position the activity in the framework.

Figure 4.2 provides a link to a video demonstration of the activity.

Figure 4.2: A QR Code of the url http://short.epfl.ch/quads which points to a video demonstration of Quads.

4.1.1 Interface Elements

Each of the elements of the interface has a fiducial marker. All of these elements are produced with a regular printer.

Cardboard Shapes

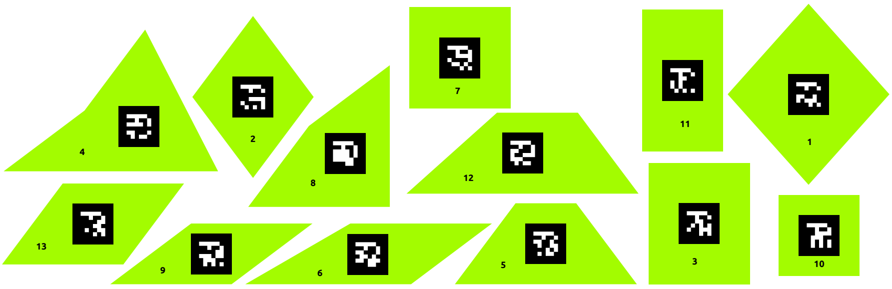

Figure 4.3: The numbered cardboard quadrilaterals used in the study.

Quads allow the pupils to manipulate quadrilaterals with paper artefacts. These cardboard shapes are shown above in Figure 4.3. The thickness of the cardboard makes the manipulation of the shapes easier than if they were made out of paper. However, the shapes can be produced arbitrarily, because cardboard is as easily printed as paper, as opposed to more tangible shapes, e.g. out of wood or plastic.

Cards

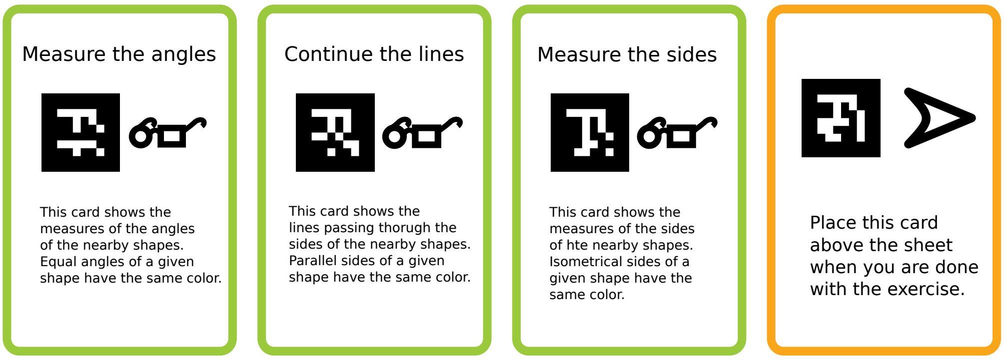

Measure the angles

This card shows the measures of the angles of the nearby shapes. Equal angles of a given shape have the same color.

Continue the lines

This card shows the lines passing thorugh the sides of the nearby shapes. Parallel sides of a given shape have the same color.

Measure the sides

This card shows the measures of the sides of hte nearby shapes. Isometrical sides of a given shape have the same color.

Place this card above the sheet when you are done with the exercise.

Figure 4.4: The English translation of the three tool cards (left, with a green outline) and the feedback card (right, with an orange outline)

The four cards in Quads, shown in Figure 4.4, have a small text describing their function. Three of the cards are tools. When brought close to a cardboard shape under the Metroscope, these cards will each display a given characteristic of the shape: the length of the sides, the measures of the angles of the shape, or which sides are parallel. This information is projected around the related cardboard shape. If measures are equal, they are projected in the same colour. If lines are parallel, they are projected in the same colour. The cards are shown in action in Figure 4.5.

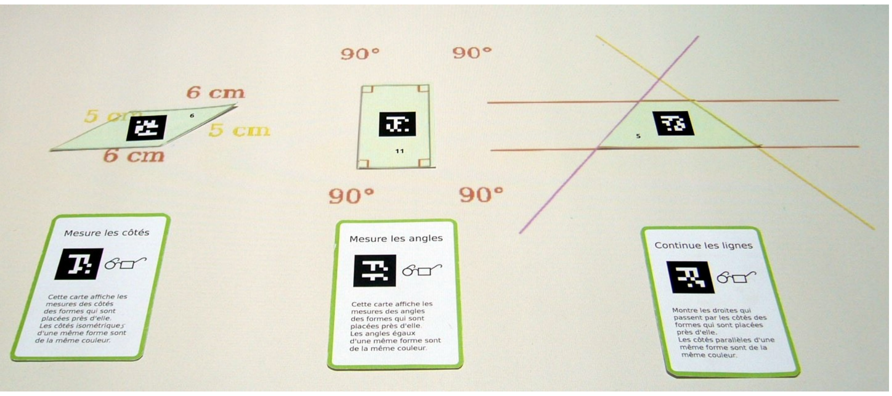

Mesure les côtés

Cette carte affiche les mesures des côtés des formes qui sont placées près d'elle. Les côtés isométriques d'une même forme sont de la même couleur.

Mesure les angles

Cette carte affiche les mesures des angles des formes qui sont placées près d'elle. Les angles égaux d'une même forme sont de la même couleur.

Continue les lignes

Montre les droites qui passent par les côtés des formes qui sont placées près d'elle. Les côtés parallèles d'une même forme sont de la même couleur.

Figure 4.5: An example of the information displayed by each of the cards.

The fourth card is the feedback card, which triggers the display of a text when it is placed next to the worksheets. If all of the expected shapes have not been placed on the sheets, a reminder to do so will be displayed. If the classification of the shape is wrong, the learner will be invited to try another one. If the classification of the areas is correct, a confirmation of the solution will appear, e.g. “Good job! Quadrilaterals with a pair of parallel sides are called trapezoids.” The feedback is more detailed for page six, where pupils have to classify a square: if the square has been classified as a rectangle, the pupil will be hinted to the fact that squares also have four equal sides like the rhombus, etc.

Sheets

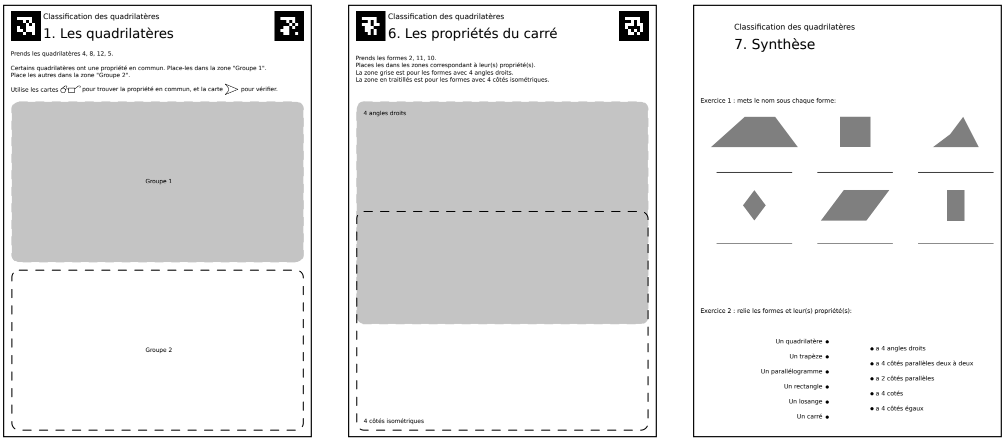

Classification des quadrilatères

1. Les quadrilatères

Prends les quadrilatères 4, 8, 12, 5.

Certains quadrilatères ont une propriété en commun. Place-les dans la zone "Groupe 1".
Place les autres dans la zone "Groupe 2".

Utilise les cartes pour trouver la propriété en commun, et la carte pour vérifier.

Groupe 1

Groupe 2

Classification des quadrilatères

6. Les propriétés du carré

Prends les formes 2, 11, 10.
Places les dans les zones correspondant à leur(s) propriété(s).
La zone grise est pour les formes avec 4 angles droits.
La zone en traitillés est pour les formes avec 4 côtés isométriques.

4 angles droits

4 côtés isométriques

Classification des quadrilatères

7. Synthèse

Exercice 1 : mets le nom sous chaque forme:

Exercice 2 : relie les formes et leur(s) propriété(s):

Un quadrilatère ●
Un trapèze ●
Un parallélogramme ●
Un rectangle ●
Un losange ●
Un carré ●

● a 4 angles droits
● a 4 côtés parallèles deux à deux
● a 2 côtés parallèles
● a 4 cotés
● a 4 côtés égaux

Figure 4.6: The lay-out of pages 1 (left), 6 (middle) and 7 (right) of the leaflet.

Seven pages form the leaflet of Quads. The first five pages have the same lay-out: a short text and two areas labelled "Group 1" and "Group 2" (see Figure 4.6, a). For the first exercise, where two trapezoids are expected to be grouped together, the text reads:

> Take quadrilaterals 4, 8, 12 and 5.
>
> Some quadrilaterals have a common property. Place them in the "Group 1" area. Place the others in the "Group 2" area.
>
> Use the [tool icon] cards to find the common property, and the [feedback icon] card to check your answer.

Page six is similar, but the two group placement areas have an overlap; this illustrates the fact that the square shares characteristics of rectangles and rhombuses (see Figure 4.6, middle).

Page seven is a recapitulation exercise, in two parts (see Figure 4.6, right). In the first part, the pupils have to write the names of the shapes under a figure representing them. In the second part, pupils have to match quadrilaterals and their properties, by linking a list of quadrilateral names (quadrilateral, trapezoid, parallelogram, rectangle, rhombus, and square) and a list of properties (four right angles, two pairs of parallel sides, one pair of parallel sides, four sides, and four equal sides).

4.1.2 Scenario

The pupils follow the instructions given by the leaflet one page after another. For example, they have to use the six cardboard quadrilaterals specified by their numbers for the exercise on page 1. Three of these quadrilaterals are trapezoids. They can use any of the tool cards to help them classify the trapezoids into a group, but for this exercise, the useful one is the one drawing the lines passing through the sides of the quadrilateral. For the three trapezoids, the pair of parallel lines will have the same colour.

After trying the various cards and cardboard shapes on this exercise, the pupils place the shapes among the two areas delimited on the sheet. They place the feedback card next to the sheet. If it is correct, they go on to the next page; otherwise, they try again. This continues until the last page, which does not require the camera-projector system; it is a regular exercise on paper.

4.1.3 Design Goals

This activity was originally developed as a demonstration to the teachers of the possibilities offered by paper interfaces. It aimed at combining the advantages of paper and computers in the form of paper interfaces.

Cards act as virtual tools that do more than regular tools for two reasons. On one hand, they can act as a scaffold for skills that the pupils have not yet mastered (e.g. measuring angles). On the other hand, they can assist the pupils by doing tasks that they have already mastered and are unrelated to the particular activity. For example, it would be very time consuming to have the pupils measure the sides of the shapes themselves, only for them to realize that they are equal.

The feedback card can be seen as a tool for the teachers. It relieves them of the task of providing trivial feedback, such as the reminder to place all of the shapes in a given exercise.

One advantage of augmented paper versus simple paper is that the level of feedback can be adapted. For example, the negative feedback given on pages one to five is voluntarily vague (“You can do something else.”) in order to foster exploration. In contrast, the negative feedback on page 6 provides more guidance towards the expected classification, by naming which specific shapes have been misclassified.

A paper interface is expected to be easier to use than complex computer interfaces. The teacher can always be in control of what help is to be given to pupils, because the cards can be physically given or taken away at any time. Back to the example of the measure card, if the teacher decides that it would be beneficial to practice ruler measurements after all, it is simply done by taking away the measurement card. Similarly, if the teacher wants the learners to skip an exercise, it is as simple as skipping a page. If the teacher wants to follow pupils more closely, the feedback card can be taken away, so that the pupils have to call him to check their work.

Paper interfaces are also more robust in the sense that they offer a natural fall-back: if something goes wrong with the system, the cardboard shapes still exist, the virtual tools can be replaced by manual measurements, and the feedback can be done by the teacher.

The mobility of paper can also be exploited: pupils write their answers on paper, and then they can carry them back to their desk or even take them home. If an exercise is too hard, the work flow of the activity can be changed so that, instead of the lamp, the pupils will receive face to face feedback from the teacher.

The quadrilaterals could have been represented as images on cards, but using an actual cardboard shape instead of the uniform shape of the cards is expected to present the advantage of lowering the abstraction level: it is easier to associate the idea of a trapezoid with a trapezoidal object rather just an image. This can be the starting point of a smooth transition toward abstract concept, where regular cards can replace the real objects.

Furthermore, the cardboard shapes are free of artefacts linked to their representation. For example, squares are easily mistaken for rhombuses as when they printed in an orientation different from the stereotypical one (See Figure 4.7).

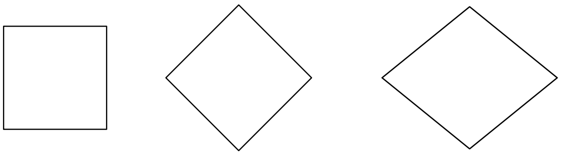

Figure 4.7: A square (left), when printed “on its edge” (middle) can be mistaken for a rhombus (right).

4.1.4 Position in the Framework

Table 4.1 shows the characteristics of Quads in the frameworks. The most important property of paper artefacts used during the activity is their position (the position of the shapes relative to the sheets or the other tool cards) and their presence (for the feedback card). The displayed properties of the shapes also follow the card orientation. The position of paper sheets relative to each other (their sequence) is also very important for the design of the sequence: the leaflet defines the sequence of activities.

The position and presence are ephemeral information used for the interaction between the user and the Metroscope. The answers of the pupil can be written persistently on the sheets. The other components of the activities are permanent definitions by the designer (e.g. which shape are to be used at which page or the features of each shape).

4.2 Pilot Study

In this section, we report on the first deployment of the Metroscope outside of the laboratory. School 1 agreed to host a demonstration of the system to its pupils. The demonstration occurred on a Friday afternoon, during the period of time normally allocated to homework. Three classes participated: two fifth grade classes and one sixth grade class; i.e. the pupils were 10 to 11 years old.

Table 4.1: Position of Quads in the framework

Property	Ephemeral	Persistent	Permanent
Presence	The pupils show the feedback card.		The pedagogical designer defines which shapes will be used in which activity.
Position	The pupils move the tool cards and the cardboard shapes. The tool cards and the cardboard shapes are placed close to each other.		The pedagogical designer defines the sequence of sheet.
Orientation	The pupils rotates the paper shape.		The cardboard shapes are free of stereotypical orientations.
Side			
Folds			
Edges			The pedagogical designer defines the edges of the cardboard shapes.
Ink		The pupils write their answers on the last sheet.	The pedagogical designer writes the instructions on the sheets.

4.2.1 Objectives

The objective of this first study was primarily to explore the reception of paper interfaces by pupils, in order to validate the approach in general. It also acted as a pilot study for the following study, which aimed to refine the Quads activity in particular.

The research questions investigated in this study were thus the following:

- (Q_l) *Does the activity support exploration?* The pedagogical purpose of such a short activity is to support exploration. It aims at letting the pupils get familiar with the topic and generate observations that can be shared with the whole class, and formalized by the teacher.

- (Q_u) *How do the pupils use the paper interface?* In other words, what are the differences between the way the interaction was designed, and the way pupils naturally used the interface? As a first study, this research question is broadly exploratory.

4.2.2 Procedure

10 groups of variable sizes (1 monad, 6 dyads, 2 triads and 1 tetrad) were given five minutes to go through the first page of Quads. They were given the sheet, the five cardboard shapes shown in Figure 4.8, and the three tool card, but not the feedback card. They were simply asked to use the tool cards on the shapes and place the shapes into two groups. After five minutes, we asked them why they grouped the shapes in the way that they did and to explain the purpose of each card.

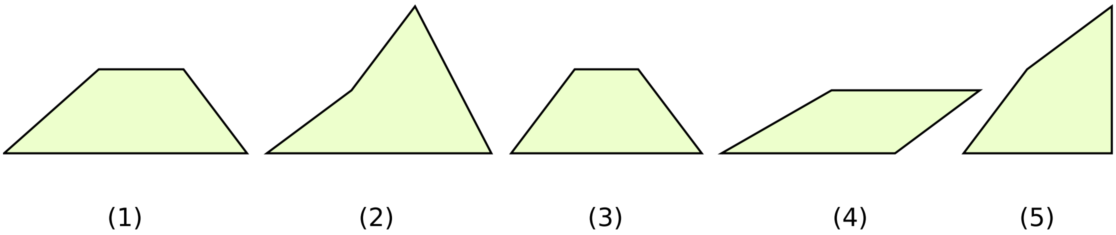

Figure 4.8: The five quadrilaterals used in the pilot study.

4.2.3 Results

(Q_I) **Does the activity support exploration?**

The classification expected by the design was the distinction between trapezoids (Shapes 1, 3, and 4 in Figure 4.8) and non trapezoids (Shapes 2 and 5 in Figure 4.8), using the card that shows the lines passing along the sides of the quadrilaterals where parallel lines are coloured the same. Of course, any classification can be made as long as it comes with a justification. We investigate here the diversity of the proposed solution, and their soundness in a geometrical context.

Only one group found this intended classification and justification. Three other groups found the same classification, with different justifications. Only one of these justifications was mathematically wrong: the group thought that Shapes 2 and 5 were triangles. Another group used the fact that Shapes 2 and 5 can be assembled like a tangram, and that they both had 5 centimetre sides as classification criteria. Another group used an artefact of the projection: the arc symbolising the angles was not projected correctly when it was too close to the fiducial marker. This happened only on Shapes 1, 3, and 4.

The second most popular grouping (found 3 times) separated stereotypical trapezoids (Shapes 1 and 3) from the others (Shapes 2,4, and 5). It was unanimously justified by the appearance of the feedback created by the card showing the lines passing through the sides. Figure 4.5 shows an example on the right: the “middle” line appears to be in the middle of a triangle formed by the other lines. This is a typical example of how children have trouble extracting mathematical properties form figures: the lines are not interpreted as lines, but as segments.

Another group also based its classification on the feedback of the 'lines' card, adding Shape 5 because it looked like a triangle with the projection of the extended lines, which is wrong. The last two groups each proposed a classification with a mathematically correct justification: first, Shape 5 has a right angle (as opposed to 1, 2, 3, and 4); and second, Shapes 2, 3, and 5 are the only shapes with a pair of sides having the same length, as opposed to 1 and 4.

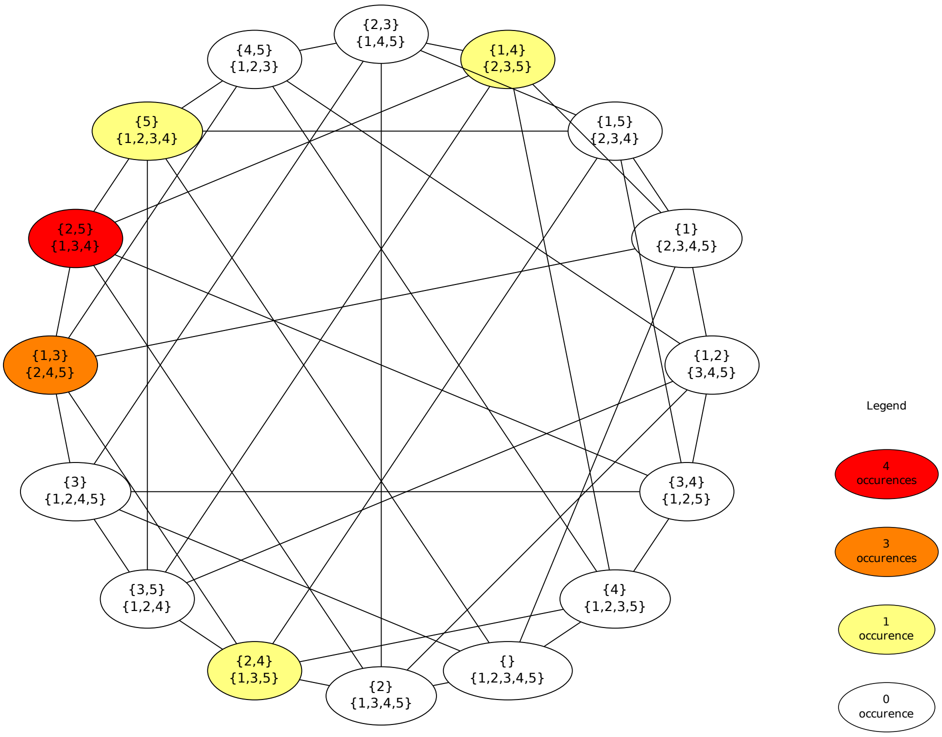

Figure 4.9: The classifications proposed by the pupils out of the 16 possible classifications. Each node represents one possible classification. Their colour indicates how often they were proposed by a group. Each classification is linked to the classifications that are possible after changing one element.

This five minute activity could have been followed by a rich discussion among the pupils, directed by a teacher. The augmentation contributed to the diversity of the answers: it is likely that pupils would have based their classifications on the appearance of the shapes rather than the characteristics revealed by the tool cards (e.g. the length of the sides). The artefacts produced by the augmentations, such as the line continuing the sides which formed a triangle, were also interesting from a pedagogical point of view: the teachers could explain the difference between the pertinent properties of a *figure* (e.g. the sides of a cube remain parallel in an isometric representation) and the artefact linked to the *drawing* (e.g. not all the sides of a cube have the same length in an isometric representation). This distinction is central in geometry education.

(Q_u) How do the pupils use the interface?

The interactions were based on the position of the paper elements: when a card was brought close to a shape, it would display additional information. There were two ways to do that: either move the shape toward the card, or move the card toward the shape. The latter seems more intuitive, as cards are usually kept in the hands. However, four groups out of ten moved the shapes towards the cards.

Within the groups moving the shapes rather than the cards, an interesting interaction emerged. Three groups created a test bench by placing all of the tool cards together, and bringing the cardboard shapes in the common neighbourhood of the cards to show all of the related information at once, as shown in Figure 4.10. This is the first appearance of what we call *creative appropriation*.

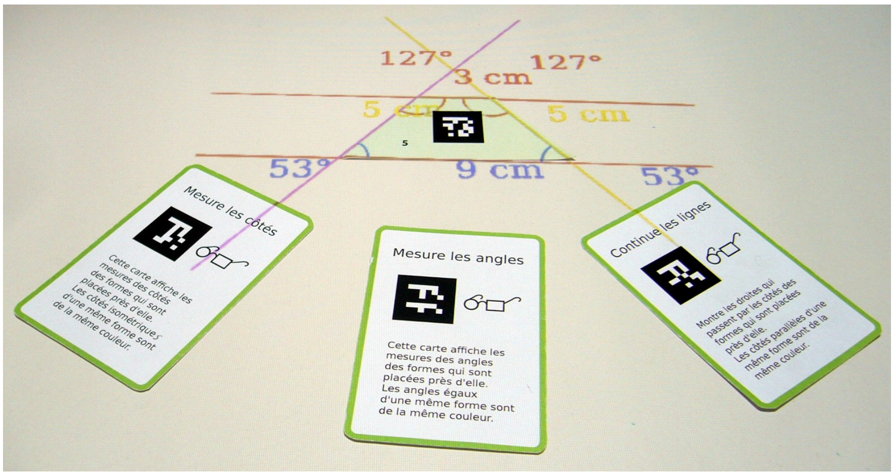

Cette carte affiche les mesures des côtés des formes qui sont placées près d'elle.
Les côtés isométriques d'une même forme sont de la même couleur.

Cette carte affiche les mesures des angles des formes qui sont placées près d'elle.
Les angles égaux d'une même forme sont de la même couleur.

Montre les droites qui passent par les côtés des formes qui sont placées près d'elle.
Les côtés parallèles d'une même forme sont de la même couleur.

Figure 4.10: The single display of the cards shown in Figure 4.5 can be combined into a "test bench" (right), where cardboard shapes are placed to display all of their characteristics at the same time.

We also observed another emerging interaction pattern that addressed technical limitations of the system. Indeed, there was a perceptible lag (of a few hundred milliseconds) between a card being close to a shape, and the display of the related information. Moreover, the single marker on the cards and shapes was often occluded (e.g. by the hands of the pupils), preventing the detection of the elements of the interface, and thus the resulting display. To address this, some groups interacted in two steps: they would first manipulate the cards and shapes, and then leave the interface alone for the system to display the information. We suppose that this technical limitation can have a pedagogical value, and, for this reason, we decided not to 'fix' this limitation for the in-situ study. Indeed, the fact that the pupils stop their manipulation to wait for feedback might actually foster an observation time that could be beneficial for their reflection Cuendet et al. (2012).

Finally, we observed that the size of the group had a strong impact on the course of the activity. On one extreme, one or two pupils manipulating the interface seemed very intimidated by our presence and did not dare interact too much. On the other extreme, the four pupil group was easily distracted and focused less on the task. Three pupils per group seemed like the best compromise.

4.2.4 Lessons for the Following Iterations

This first pilot study provided us with important lessons for the following iteration of the Quads activity. First, pupils need a guide for explorations, which is the reason we introduced the feedback card. This card structures the exploration by ensuring that the expected milestones (the expected classification) have been reached in the expected way (by displaying the expected justification).

From a technical perspective, this pilot study revealed a display problem : arcs of angles were hidden by tags. Because this problem does not have any pedagogical value, we fixed it by drawing a smaller arc. However, the lag between manipulation and feedback can have a pedagogical value. In fact, it could result in more reflection from the pupils during the observation phase.

4.3 In-Situ Study

Figure 4.11: The physical set-up of the experiment: three pupils sitting under the Metroscope, in a spare room of their school.

We deployed the activity in another study, which took place in a spare room of another primary school (see Figure 4.11). The pupils came from two sixth grade classes (i.e. 11-12 year-old pupils). We could only extract pupils from workshop-like activities related to mathematics. As a consequence, our study was limited to two periods (of 45 minutes) per class. We decided to have the groups go through the whole activity, which meant that only one group per period could take part in the study. Two groups from each class tried the system during one period each, after a brief presentation to the whole class.

4.3.1 Objectives

Studying the tangible interaction of a group brings useful insights about the learning activity. This is the field of research called Interaction Analysis (Dimitrakopoulou et al., 2006): it consists of automatically analysing a computer-mediated activity to produce indicators to be reused, to qualify, or to adapt the activity. These indicators are typically computed from the traces of an interaction between learners and computers via a mouse and keyboard. However, this is also possible with tangible interactions, and these indicators are even richer, because interaction happens in the real world. This is what we investigate in this study: (Q_I) *How can we interpret the manipulations of pupils?*

4.3.2 Procedure

The study took place in a spare room of the school, shown in Figure 4.11. The four groups who tried Metroscope during the study were composed of three pupils. For each teacher, the first group was composed by higher performing students than the second group. Each group had 40 minutes to go through the seven pages of the activity. Pages one to six, the classification exercises, were completed as a group under the lamp; page seven, the recapitulation exercise, was completed individually on separate desks. The expected classifications for each of the pages are shown in Figure 4.12. We did not reserve time at the beginning of this study for the pupils to get familiar with the system, because in the pilot study, they immediately understood how to use the interface.

We provided the feedback card to the pupils, with a mechanism to make intensive trial and error strategies difficult: the feedback card had to be placed on the sheet, i.e. on top of the shapes, which prevents from moving them. The experimenter gave one exercise sheet and its related cardboard shapes to the group at a time, after the feedback card validated the classification for that particular exercise.

4.3.3 Results

Performance

We used page 7 of the activity as a pre-test and post-test. The score was computed as the number of quadrilaterals that could be named correctly. For nine out of twelve pupils, the score did not change, mostly because of a ceiling effect: the answers were already correct in

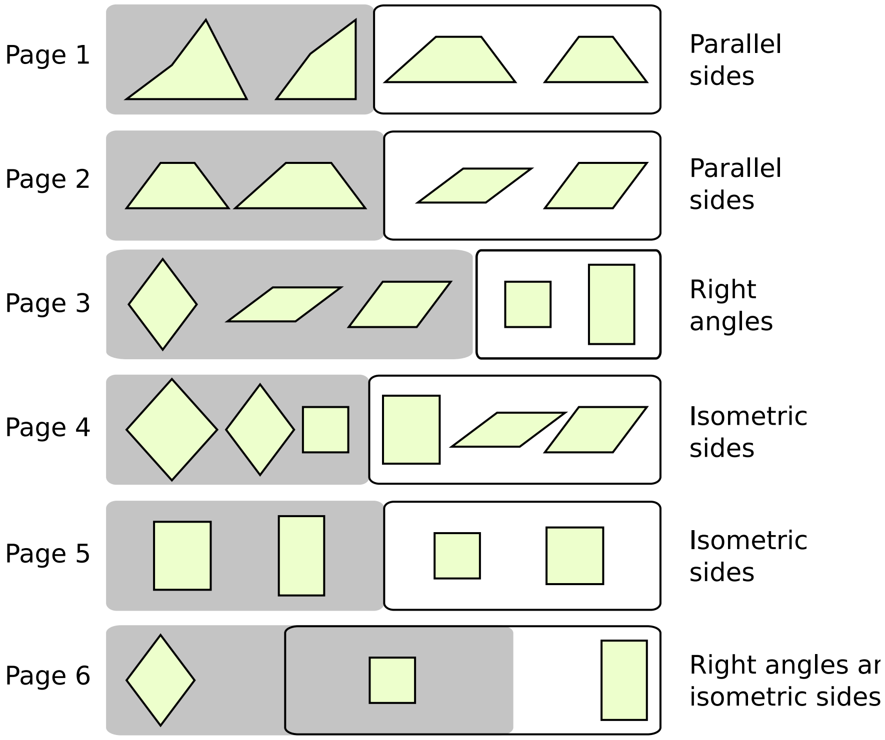

Figure 4.12: Expected classifications for the first six pages of the study. The criteria for each classification is shown on the right. The classification of page 6 is not a partition: a square is in both the rectangle and rhombus classes.

the pre-test. Three pupils improved their score from four to six, from two to four, and from three to four, respectively. These three pupils were in groups 3 and 4; groups 1 and 2 scored perfectly on both tests. The explanation lies more in the group formation than in the activity: groups 1 and 2 came from one class, while groups 3 and 4 came from another.

More importantly, the teacher purposely formed the groups of homogeneously performing pupils: pupils from groups 1 and 3 were higher performers than those from groups 2 and 4. This is reflected in the completion times: the time to solve the exercises (excluding exercise 4) was 515, 631, 819, and 911 seconds for groups 3, 1, 4, and 2, respectively. Exercise 4 was excluded from the calculation, because group 3 did not receive this exercise: the experimenter erroneously skipped the page.

All the groups managed to go through the activity in less than 40 minutes, which is important, because it fits the 45 minute periods between the two breaks. Figure 4.13 shows the time spent on each exercise, from the presentation of the page by the experimenter to the positive feedback given by the validation card.

Figure 4.13 highlights the differences between the exercises: pages 2, 3, and 5 were more obvious, since they could be completed on a perceptual basis (compare with Figure 4.12). This helped the pupils formulate a hypothesis very quickly, but they double-checked their

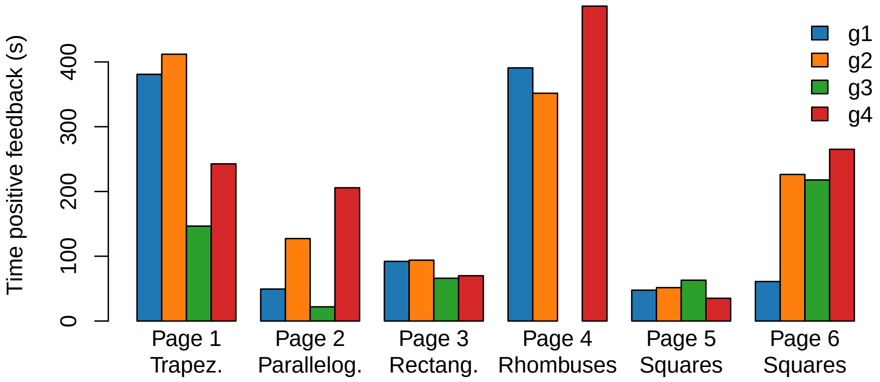

Figure 4.13: Time spent by each group on each activity.

hypothesis with the tool cards in most cases before viewing the feedback. Only one pupil tried to validate a classification immediately without using the tool cards.

Page 4 stands out in completion time. This is due to the fact that the appearance was misleading: the first intuition of the pupils was to group the square and the rectangle together because they have equal angles, but the expected classification was actually based on the length of the side. Pupils tried several hypotheses before correctly classifying the shapes. Page 1 also stands out, mostly because it was the first exercise: pupils were still discovering the activity and the interface of the system to solve it. Finally, page 6 was grouped with the more challenging exercises, because it was slightly different: the two areas of the page in which to classify the shapes were actually overlapping, and one shape (the square) was in the overlap of these two regions. We now present indicators that corroborate these explanations.

Use of the Augmentations as Problem Solving Step Indicator

The most basic information regarding the manipulation of an interface element is its use, i.e. whether it is activated or not. In the case of a tool card, it means that its functionality was being used by being close to a shape. Not all of the cards were useful for completing each exercise. For example, the card showing the lengths of the sides was useful to distinguish rectangles from squares, but not to distinguish squares from rhombuses. A mixed effect ANOVA with the group as a random factor showed that the useful cards were used significantly more than the others ($F[1,64] = 7.24$, $p < 0.01$). In general, it informs us that the pupils were not toying with the system, and understood which cards could help them. More particularly, it is a first indicator on the learning activity: by detecting which tool card is activated, it is possible to determine whether the pupils are trying to *infer* the expected classification, or whether they are still guessing.

The associations between the tool cards and the cardboard shapes can bring more information. Theoretically, this association is many-to-many: a card can display the attributes of several shapes at once, and a shape can have several of its attributes displayed. In practice, it is a one(-card)-to-many(-shapes) relationship: a card often displayed the attributes of several shapes at the same time, but a shape very rarely had several of its attributes displayed at the same time. Throughout the whole experiment, cards were associated to one shape 1.6 times more often than to multiple shapes, while shapes were associated to one card 10.5 times more often than to multiple cards. We assume that when a tool card is activated on one shape, the pupils are still trying to *infer* the rule. If the card is applied to multiple shapes, the pupils are more probably in the *applying* step, because they have grouped shapes together, and can compare the characteristics displayed by a card. Of course, it is easy to know when the pupils are in the *validating* step: the shapes are on top of the exercise sheet.

Average Speed as Hesitation Indicator

We kept a trace of the position of each piece of paper throughout the study. This allows us to recreate the trajectory of the various pieces of paper (see Figure 4.14). This is very rich information. To exploit it, we concatenated the trajectories of paper elements of the same type: shape, validation card, tool card, and sheet, for each exercise of each group. Dividing this cumulated length by the time each group spent on each exercise gives an average speed of a type of interface component. This average speed is a quantity of movement performed by the pupils on the interface, normalized over the length of the exercise; it does not correspond to any actual speed.

Figure 4.15 shows the average speeds of each kind of component. A mixed effect ANOVA with the group as a random factor showed that each kind of component had a significantly different average speed from the others ($F[3,84] = 32.40$, $p < 0.0001$). On average a sheet moved by 7 mm/s. It is the most stable element. The augmented area under the Metroscope, delimited by the projected white background, defines a virtual border for the exercise sheet, out of which it is hard to move. Anyway, moving the sheet away from the centre would mean excluding one of pupils who are sitting on the side. Second comes the validation card, at 18 mm/s, which is simply brought above the sheet when needed. Third come the tool cards, at 32 mm/s, which are manipulated more often. The most mobile element are the shapes, at 52 mm/s; they have to move between the areas of the exercise sheets and the neighbourhood of the tool cards.

We now consider the average speed of the shapes, the most mobile interface element. More time-consuming pages (1, 4 and 6) see a decrease of this average speed. A mixed effect ANOVA with the group as a random factor showed that this decrease is significant ($F[5,14] = 8.34$, $p < 0.001$). The speed can thus be used as an indicator of the hesitation of the pupils: on trivial pages, they will obviously move the cardboard shapes more than on pages where they hesitate. This may be because the exercise is new, as with page 1, or because the exercise is harder, as with pages 4 and 6.

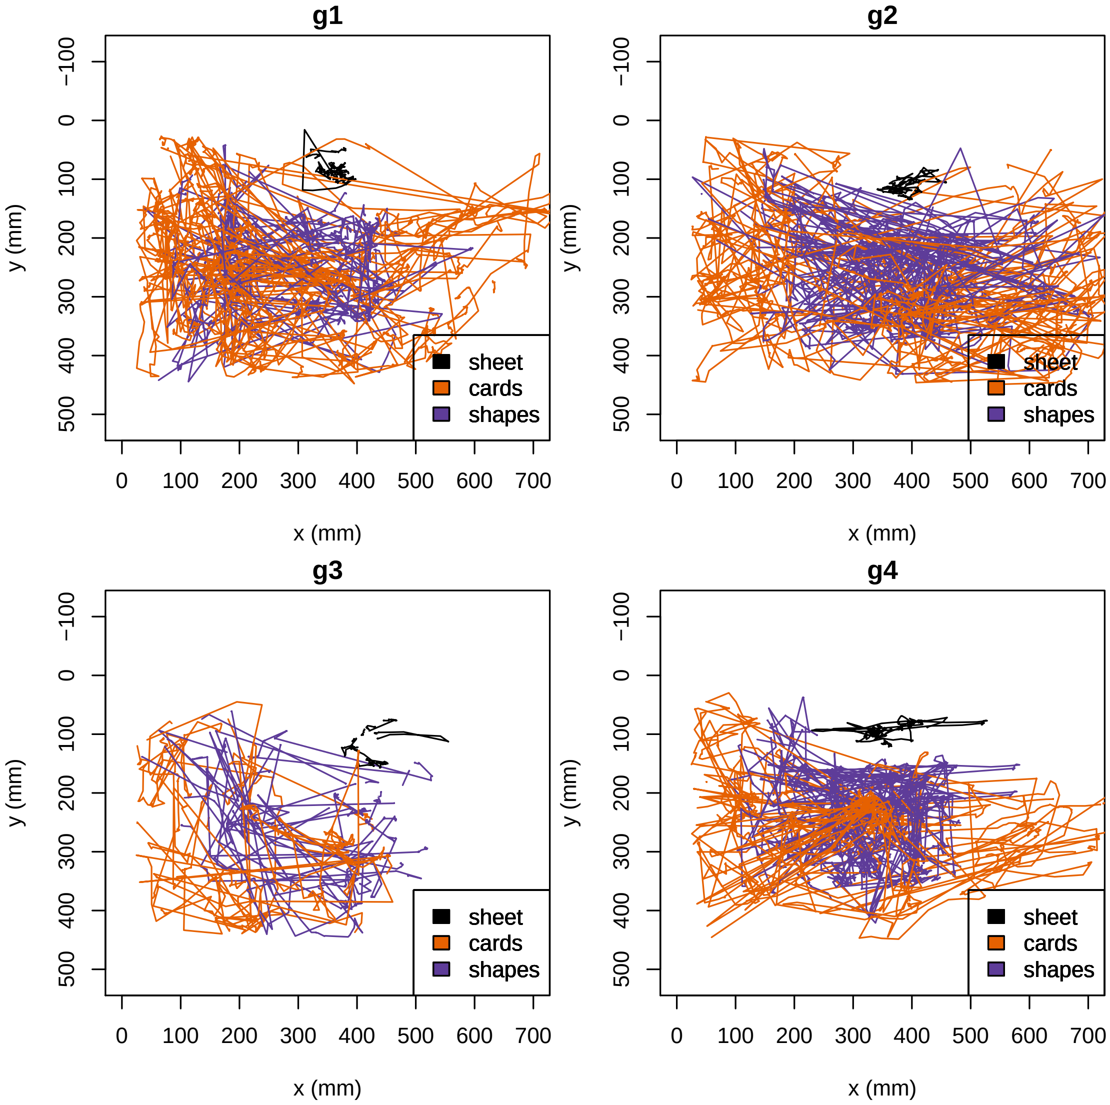

Figure 4.14: The traces of the various pieces of paper.

Verticality as Difficulty Indicator

We can refine the grain of our analysis by only observing the y-coordinate of the positions of the cardboard shapes relative to the y-coordinate of the position of the sheet. Indeed, the goal of the activity is to place shapes in one of two boxes which are on top of each other. Figure 4.16 shows an example of this data. It actually corresponds to the path taken by a group towards a solution. One way to formulate the exercise is the following: shapes are associated to a binary type for each sheet, and the goal is to place the ones of a same type within the boundary of areas that stretch over the whole width of a page. The problem is thus solved when all shapes of a given type are within the vertical range of an area.

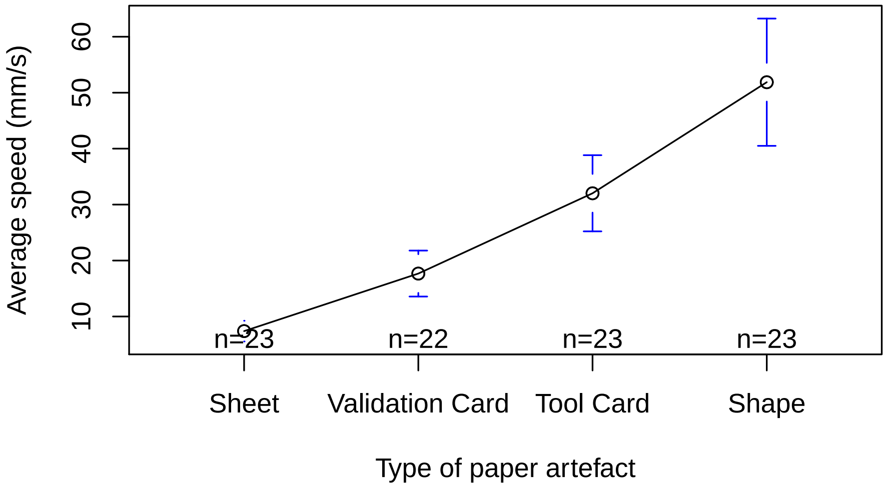

Figure 4.15: Average speeds of the different kinds of paper artefacts.

We can then make the y-coordinate of a cardboard shape discrete by defining one value per classification area on the sheet (i.e. two values for the first five sheets, and three values for the sixth sheet, which has two areas and their overlap). What is interesting then is the number of transitions from one area to another, because it means that a shape was moved to another category. In other words, the more transitions that happen on the discrete vertical coordinate of a cardboard shape, the more changes have been made to a solution. The number of vertical transitions is thus an indicator of the amount of difficulty that the pupils are having.

As seen in Figure 4.17, pages 4 and 6 stand out from the other pages. The finding here is that page 1 does not stand out. This means that the increased time needed for solving page 1, and the increased average speed of the paper elements previously found, are not related to the difficulty of the task. On the first page, pupils are beginning to discover the activity and the interface and do not yet know how to manipulate the interface elements.

The fact that we did not reserve time for the children to familiarize themselves with the interface allowed us to make this distinction: pupils can spend more time on an exercise because they are having difficulties with the exercise (as indicated by the number of vertical transitions) or because they are hesitant about how to use the interface (as indicated by the lower speed of manipulation).

We fitted several models to explain the number of vertical transitions, and the best fit is a mixed effect model (the group being a random variable) with the difficulty of the exercise sheet being a factor ($F[1,87] = 126.51$, $p < 0.0001$), the group performance level being a factor ($F[1,2] = 43.92$, $p < 0.05$), and the interaction of the two of these factors ($F[1,87] = 25.78$, $p < 0.0001$). This model is visualized in Figure 4.18. It categorizes the exercise sheets into two difficulty levels (reflected in Figure 4.17), and the groups into two performance levels (based on the group's composition by the teachers, corroborated by previous observations). We see that the number of vertical transitions significantly increases on difficult pages, and with lower

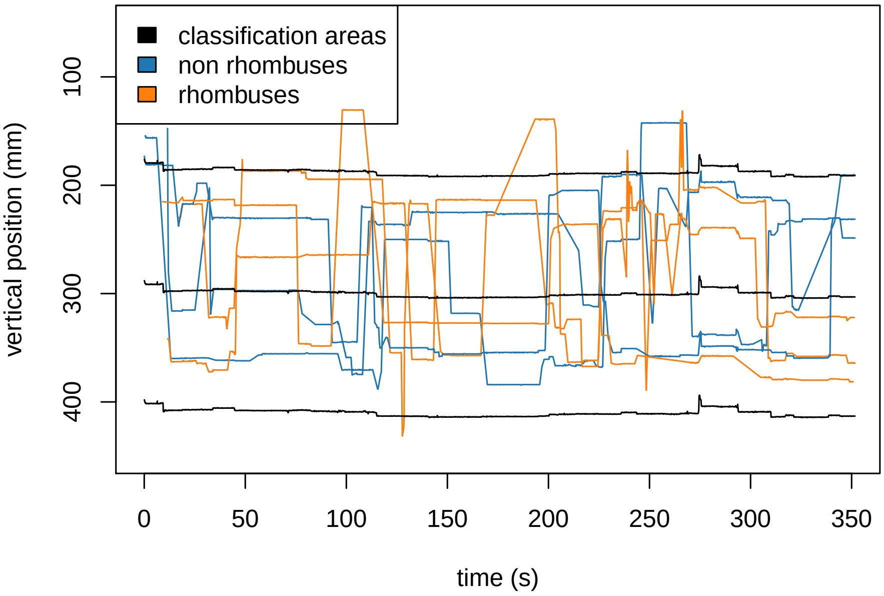

Figure 4.16: The y-coordinate of the position of the various elements of the paper interface as a function of time, for one group and one exercise. This allows us to visualize the position of the cardboard shapes (blue and orange, depending on the expected classification) relative to the areas on the exercise sheet (delimited in black).

performing groups. Furthermore, when these two conditions are met, the increase is even higher. The number of vertical transitions is thus a measure of the difficulties that a group of pupils faces on an exercise.

Horizontality as Division of Labour Indicator

Similar to the verticality of pieces of paper that we just investigated, we define the horizontality as the x-coordinate of the position of a piece of paper (card of shape) relative to the sheet. We show that it allows us to visualize the participation of each pupil. As shown in Figure 4.19, the three pupils sat next to each other. The exercise sheet was always in front of the MIDDLE pupil. The experimenter stood to the LEFT of the group, distributing the cardboard shapes from the LEFT of the group. The position of the pieces of paper was hence a hint at the balance of participation: the pupil on the RIGHT had to request or otherwise fetch the interface elements.

For example, Figure 4.20 shows the positions of the cardboard shapes and cards relative to the first exercise sheet. This allows us to observe the difference in collaboration between groups 2 and 4: the use of the cards was more balanced between the members of group 2. In both cases, the pupil on the RIGHT started far away from the pieces of paper. However, in group 2, the pupil on the RIGHT grabbed a cardboard shape after three minutes, and a card after 4 minutes, in order to do measurements. The pupil on the LEFT took them back, but the pupil on the RIGHT took the validation card, and used it twice, before the pupil on the LEFT took it back. In group 4, the pupil on the RIGHT was more shy and did not participate in the exercise.

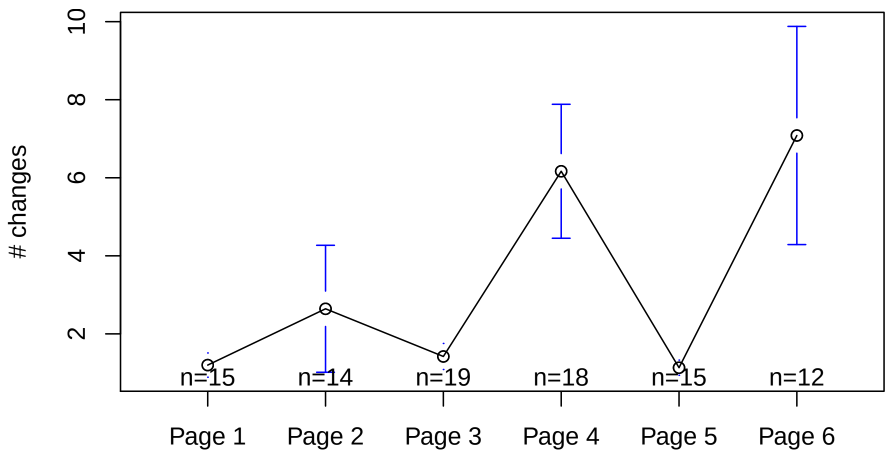

Figure 4.17: Number of vertical transitions per exercise sheet.

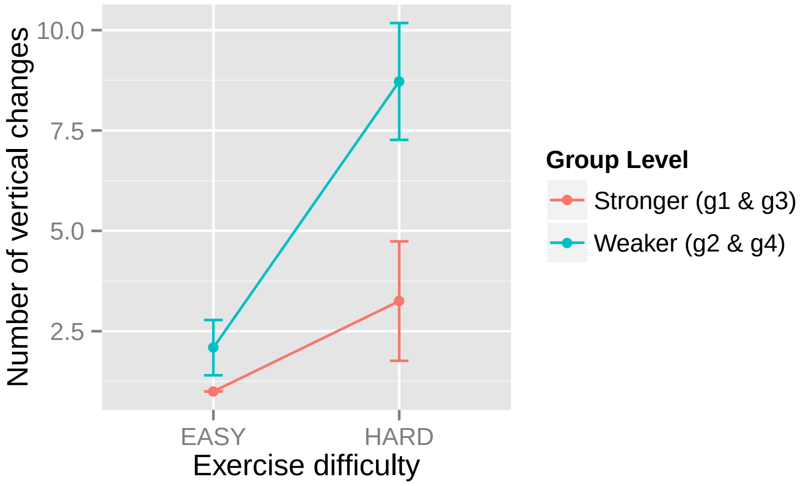

Figure 4.18: Number of vertical transitions per exercise sheet.

First, observing the time spent by the pieces of paper in the three sections relatively to the position of the sheet (LEFT, MIDDLE, and RIGHT), we can see three kinds of collaboration types (see Figure 4.21). The collaboration of groups 1 and 2 happened naturally; it is unbalanced in favour of the LEFT pupil because the shapes and sheets were initially placed by the experimenter on the left side. Group 3 shows another kind of collaboration: the RIGHT pupil did not manipulate anything in her section. Instead, she was monitoring and leading the manipulations of the other members of the group. Finally, Group 4 was artificially balanced: the experimenter imposed a turn-taking rule for the pupils to manipulate, after realizing that the RIGHT pupil was not participating at all.

Similar to the vertical changes, we can define the horizontal transitions as the number of time a piece of paper changes sections (LEFT, MIDDLE, or RIGHT). This corresponds to a

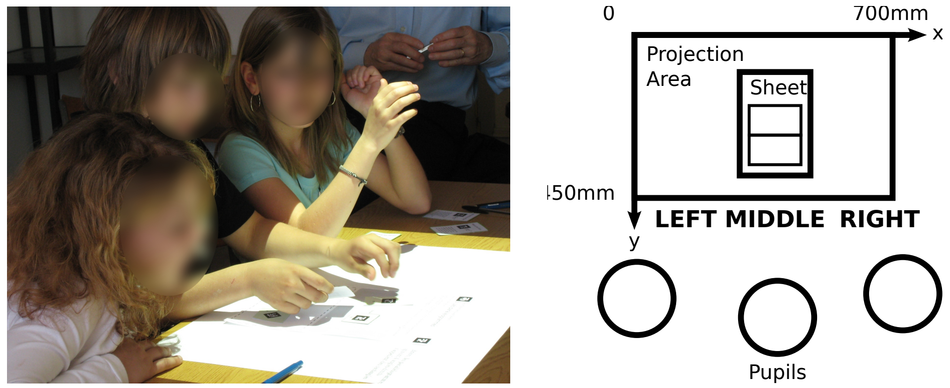

(a) Three pupils sitting in front of the augmented area, with the sheet in the middle.

(b) A schema of the positions of the pupils revealing the three areas defined by the position relative to the exercise sheet.

Figure 4.19: The positions of the pupils.

pupil exchanging (either by taking or giving) pieces of paper. This provides an important complement to the previous data, by indicating how much the pupils cooperate, i.e. interact jointly. This is different from the division of labour: if the pupils evenly dispatch the elements of the interface among them and work independently, there is an even distribution of labour, but no cooperation. Group 2 showed the other extreme: they exchanged the pieces of interface about twice as much as the other groups. This is to contrast with Group 4, who had a more even repartition of the interface elements due to the intervention of the experimenter, but this did not significantly increase the cooperation. Note that a higher cooperation is not by itself a sign of a good collaboration. In the two higher performing groups, two pupils were manipulating, and one was overseeing, while the two other groups were less coordinated and worked more separately. Indicators can only be used after being interpreted, as we will discuss in the next section.

There is a difference, however, in the associations between the cards and the cardboard shapes. Theoretically, this association is many-to-many: a card can display the attributes of several shapes at once, and a shape can have several of its attributes displayed. In practice, it is a one-to-many relationship: a card often displayed the attributes of several shapes at the same time, but a shape very rarely had several of its attributes displayed at the same time. Throughout the whole experiment, cards were associated to one shape 1.62 times more often than to multiple shapes, while shapes were associated to one card 10.5 times more often than to multiple cards.

4.4 Conclusions

The studies on the Quads activity provided insights about how the form of paper influences its usage. We saw that cards function as actions, that are applied to objects, in this case cardboard

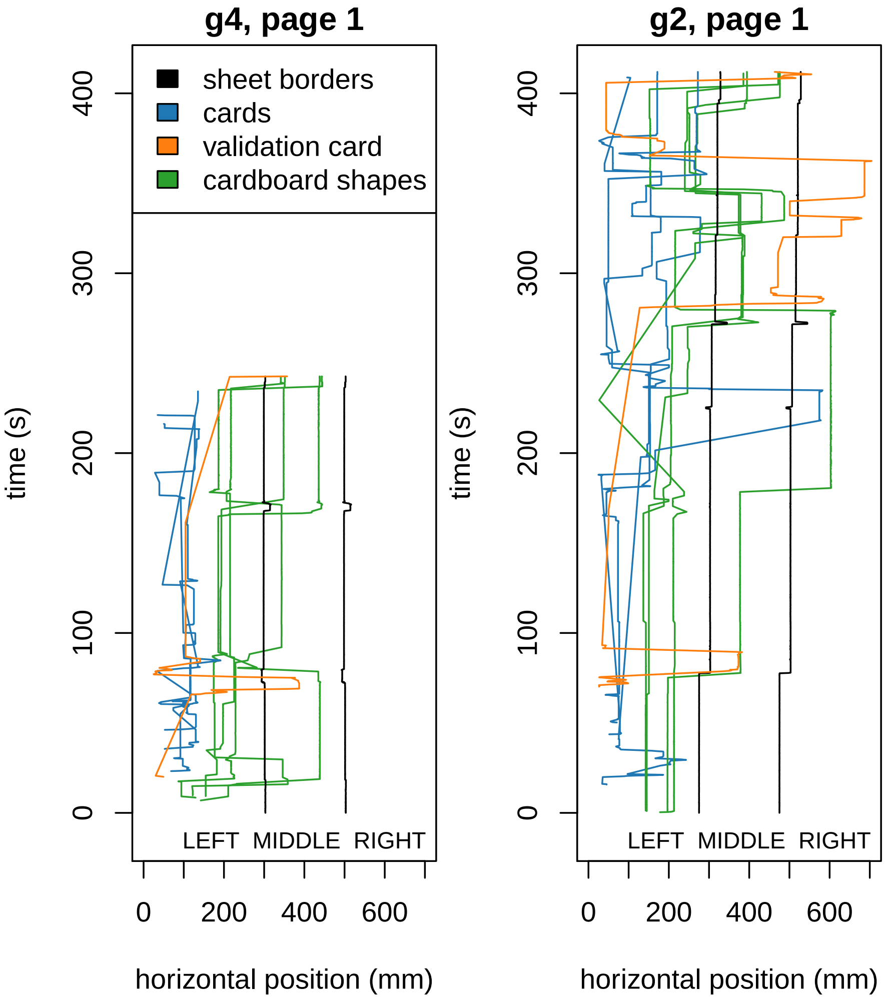

Figure 4.20: The x-coordinate of the position of the various elements of the paper interface as a function of time, rotated horizontally. This shows the position of the pieces of paper relative to the first exercise sheet (see Figure 4.19b)

shapes. The sheets allowed us to structure the interaction in time, by ordering exercises in pages in a leaflet. The sheets provide a spatial structure, as illustrated by the classification areas. The sheets structure the activity at a higher level than the cards, which define the immediate state of the interaction. This is confirmed by the fact that the sheets move less.

In a preliminary study, we illustrated how Quads supported this exploration, and observed how pupils used the paper interface. In a following in-situ study, we investigated how the manipulation of a paper interface can be used to compute indicators on the learning activity. We compared these indicators to various performance data (tests, completion time, and evaluation of the groups from the teachers) to explain how they can be interpreted.

In summary, we found four indicators that could be extracted from the manipulation of the paper interface for Quads. These indicators are derived from the position of the various

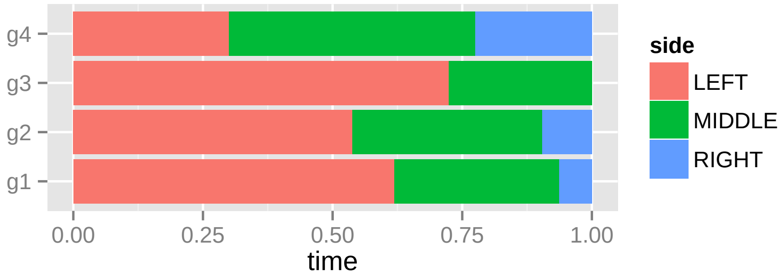

Figure 4.21: Time spent on each section of the collaboration surface by the pieces of paper.

pieces of paper. First, the position of the tool cards relative to the cardboard shapes gives an indication about the problem solving step. Using the right tool card is a hint that the pupil is close to the solution, and applying it to multiple shapes indicates that the pupil has a hypothesis on how the quadrilaterals should be grouped. Second, the speed of the various elements of the interface shows how much the pupils hesitate. Increased hesitation was seen when doing the first page, when the activity was unknown, and when doing the more challenging pages. Third, a projection of the position of the cardboard shapes on the vertical axis of the exercise sheet indicated how much difficulty the pupils had in the construction of the solution, even if they did not request feedback. We showed that this was different from the time spent on or the hesitation when doing an exercise: the first exercise was trivial, i.e. the pupils did not have much difficulty, but they hesitated to use the interface because it was new to them. Fourth, the horizontal projection of the position of the pieces of paper was an indicator of the distribution of the roles among the groups of pupils.

These indicators are not automatically applicable, because they do not directly qualify the learning activity. For example, hesitation is a good thing if it corresponds to a reflection, but too much hesitation is not productive; too little hesitation is acceptable if it is the result of an exercise being too easy for a group, but not if it is the result of a lucky guess. Indicators are thus the basis of a retroactive loop on the learning group. The indicators need to be interpreted within their context, and this interpretation can be used to generate feedback on the group.

Many TEL approaches are based on this principle, explained in more detail by Soller et al. (Soller et al., 2005). They show three families of systems that support the management of collaborative learning interaction, which differ in the exploitation of the indicators. In a mirroring tool, pupils are shown the indicators and left free to decide what to do with them. Meta-cognitive tools are similar, but also show a desired value for the indicators, so that the pupils aim at a desired state of interaction. Finally, guiding tools actually process the indicators and propose a remedial action for the pupils to reach a desired interaction state.

Orchestration can also profit from the extractions of indicators from the manipulation. Orchestration refers to the tasks of teachers to identify and exploit learning opportunities and constraints in real time. This task is very challenging in a classroom: teachers have to integrate tight time constraints, expectations of the curriculum, practical matters (e.g. a forgotten book), the energy of the pupils, etc. A big part of this task consists of monitoring the class to acquire information. In this context, computing indicators from the manipulation of pupils on a tabletop can support the teacher in the orchestration task. It would even give a sense of activities based on paper without augmentation: the added value could be the tracking of the various pieces of paper, in order to provide indicators to the teacher in real time about the learning activity comparable to the one achieved if a teacher was monitoring only one group continuously.

By observing how pupils manipulate a paper interface, we extracted some features that can be used as basic design guidelines. For example, cards are easily used as a function that can be applied to several objects. More importantly, the way Quads was designed allowed for easy extraction of the indicators we described. For example, the exercise sheets are laid out in a way that the y-coordinate of cardboard shapes maps directly to the solution state, and the changes of this value map to the difficulty. The workspace is set up in a way that the x-coordinate maps directly to the ‘owner’ of a piece of the interface, and the changes of this value show the coordination among the group. Finally, the tools are dispatched into independent pieces of paper, which allows us to know which function is being used and how, indicating where the group is on the path to a solution.

We do not claim that our interpretation of the indicators is perfect. Instead, our contribution is to inspire designers of TEL to use paper interfaces as a way to easily design activities that highlight targeted characteristics of a TEL activity. For example, the simple design of the exercise sheets supported the extraction of features related to the difficulty of and the collaboration in an activity. We hope that this approach will inspire other TEL researchers to tailor interfaces to help them investigate their own questions.

We believe that it would be possible to use further indicators from the manipulation of a paper interface. We could exploit the scattered aspect of paper interfaces to model the cognition of the pupils with a higher granularity, revealed by which element of the interface is being used, or how. Furthermore, there are many other features that can be exploited from the manipulation of a paper interface, such as the trajectory of the pieces, how they are stacked, how they can be combined, etc. This opens a wide range of possibilities as to the design of paper interfaces that makes indicators easily collectable.

In these studies, the two roles of the teacher –pedagogical designer and conductor of the classroom– have been fulfilled by the experimenter. However, we could easily imagine studies that would extend the coverage of the framework by Quads, as shown in Table 4.2.

Table 4.2: Possible extensions of Quads in the framework

Property	Ephemeral	Persistent	Permanent
Presence	A teacher could remove a card that is over used by a group. Moreover, skipping a page (like with group 3) is not necessarily a flaw: it can be a feature allowing the adaptation of the activity to current conditions.		
Position	The teacher can reorder the sequence of the exercise sheets by stacking them differently.	To make their answers persistent, the pupils could paste the cardboard shapes to the answer sheet.	
Orientation			
Side			
Folds			
Edges			The group in the pilot study who classified the shapes according to how they could be combined added a permanent component to the script: the edges as classification criteria.
Ink		Another way to make the answers persistent would be to have pupils use a pen to outline the shape in the corresponding area.	

5 Discovering the Protractor

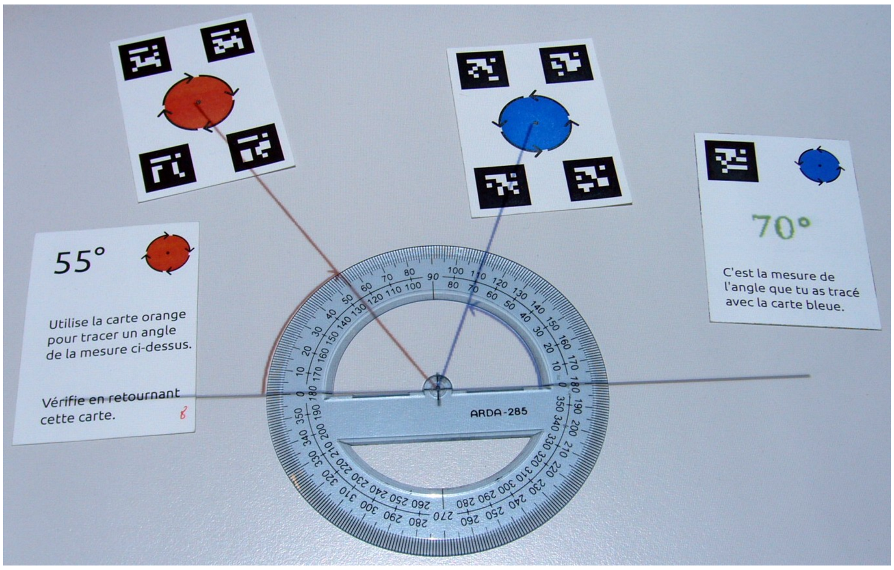

Figure 5.1: The various elements of Angoli: the two control cards and two of the angle measure cards. The one on the left is flipped and shows the measure of the angle constructed with the blue control card (70°).

This chapter presents Angoli, an activity to introduce the use of protractors. It has been created as an introductory activity for the one presented in Chapter 6. This chapter is short, because Angoli and the two related studies are very simple. However, they allowed us to gain some insights in the manipulations of a paper interface implemented in the form of a deck of cards.

5.1 Description of the Activity

5.1.1 Interface Elements

Angoli consists of a deck of cards of two types: ten angle measure cards and two angle control cards (an orange and a blue one), shown in Figure 5.1.

The orange and blue control cards correspond respectively to clockwise and counter-clockwise angles. When a control card is shown to the system, an angle is projected, with its origin in the center of the projection area, an extremity on the center of the control card, and the other extremity fixed horizontally on the left or right side of the origin, depending on whether the card is orange or blue.

Each angle measure card has a different value in degrees, an icon representing one of the control cards, and instructions to use the corresponding control card to construct an angle by measuring the given value. The angle measure card further instructs to flip the card to check the value of the constructed angle. The other side of the card contains a text explaining that the projected value is the measure of the constructed angle.

When the angle measure card is flipped, the measure of the constructed angle is displayed in a color depending on the error (green for correct, yellow for close enough, red otherwise). The angle can still be adjusted by moving the card accordingly; in this case the measure of the angle is updated in real time.

Figure 5.2 provides a link to a video demonstration of the activity.

Figure 5.2: A QR Code of the url http://short.epfl.ch/angoli which points to a video demonstration of Angoli.

5.1.2 Scenario

The intended use of these cards is also very simple: the pupils take one measure card after another, use a protractor to build the corresponding angle with the corresponding control card, and flip the measure card to check. If the pupils are wrong, they can try again with an immediate feedback, or flip back the measure card to hide the feedback.

The side on which the feedback cards is placed thus acts like a switch for the feedback. Since the feedback is updated when the switch is on, Angoli can be used in two ways. In one case, the pupils build an angle with the control card and check their solution afterward. In the other case, pupils can build an angle while at the same time having the angle measure card switched to have continuous feedback. Of course, pupils can use one way, and then the other. For example, they can start with the interactive feedback, then try to switch it off for the next exercise, and come back to continuously updating feedback if they have too many difficulties.

5.1.3 Design Goals

This activity has been designed with two teachers from School 2. After the experiment described in Chapter 4, we invited the teachers to a two hour brainstorming session in our lab.

Our objective for this session was to design an activity that would combine paper interface elements with a physical tool, in order to illustrate the fact that paper interfaces are a good fit with the regular geometry teaching. We decided together that this tool would be the protractor, because it had not been introduced in the classes of the teachers so far; this introduction was planned for the end of the year, or the following year. This way we could assume that the pupils had no or little prior knowledge on the topic of the activity.

The brainstorming session focused mostly on a scenarized activity involving protractors, which we describe in Chapter 6. The teachers recommended that this scenarized activity be preceded by a simple introduction exercise, in order to learn the basic usage of the protractor. Angoli is this prologue.

The teachers told us that beyond the basic usage of protractors (i.e. placing it correctly and reading the measure), the main difficulty for learners was to determine which graduation to use between the clockwise and the counter clockwise ones, i.e. between the one having its zero on the left and the 180° mark on the right, or the contrary. A common mistake is indeed to measure the supplement of an angle because it shares the same mark on the graduation, e.g. 60° clockwise angles are on the same mark as 120° counter-clockwise angles. The two control cards address this issue, by forcing the use of one graduation or the other, since each control card follow a different graduation.

Concerning the implementation, we chose to base the activity entirely on cards. Cards allow to design activities that are modular on two aspects. The first kind of modularity concerns the pedagogical design: these cards have a simple, generic function that could be used in a different scenario than the one previously described. For example, the pupils could be asked to pair the cards corresponding to angles that have a common side, e.g. the line controlled by the blue card drawing a 60° angle would match the line controlled by the orange card drawing a 120° angle (see Figure 5.3). Then the pupils could be asked to infer a rule to determine which angles will have a common side (their sum is 180°), in order to introduce the concept of supplementary angles.

The other aspect of modularity concerns the management of the activity. Since each card is independent, they can be removed or added on the fly. One can easily imagine a larger deck of cards, with various difficulties, which would allow the teacher to decide spontaneously which pupil should do which measurement.

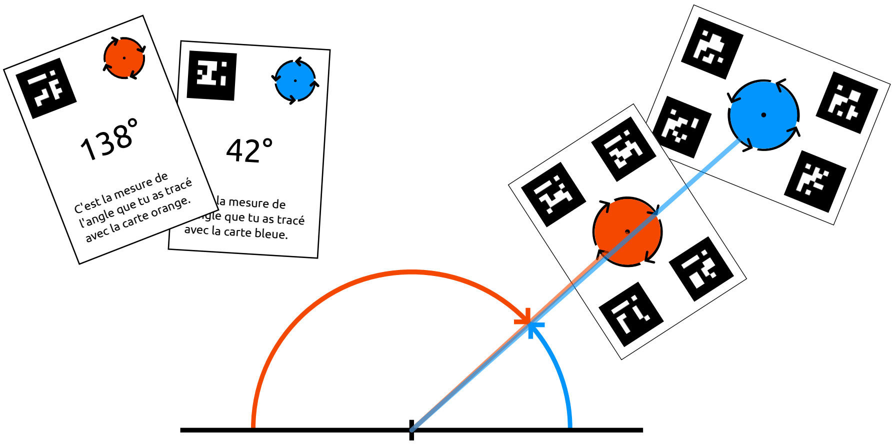

Figure 5.3: The lines building 138° clockwise and 42° counter-clockwise overlap. This relation correspond to the fact that supplementary angles sum to 180°.

5.1.4 Position in the Framework

Table 5.1 shows which parts of the framework can be covered by Angoli. The pupils interact in two ways with Angoli. They either manipulate the control cards, such that its position defines the angles, or flip the measure cards in order to switch the feedback on or off. The teacher can define which measure cards exist (their permanent presence), and which ones are available during the activity (the ephemeral presence), in order to keep control.

Table 5.1: Position of Angoli in the framework

Property	Ephemeral	Persistent	Permanent
Presence	The teacher can remove measure cards.		The pedagogical designer defines the angle of the measures card.
Position	The pupils move the control card relatively to the protractor.		
Orientation			
Side	The pupils flip the measure card to switch the feedback.		
Folds			
Edges			
Ink			The angles to build are described on the cards.

5.2 Controlled Study

We first ran a controlled study to evaluate Angoli alone. In this first study, we controlled the conditions in order to isolate one factor: whether the feedback was continuous or not.

5.2.1 Objectives

This first study aimed at answering the following questions:

- (Q_l) *Which kind of feedback is most helpful to learn how to measure an angle?* As explained previously, the pupils can use Angoli to display continuously the measure of the angle being built, or display it only when they think they are correct. We want to know which way is the most productive in terms of learning.
- (Q_u) *How do pupils use feedback switches?* Angoli makes it very easy to switch feedback on and off, which gives two possibilities: either constantly showing the live feedback, or only punctually checking an answer.

Both questions deal with the feedback. The first one focuses on the learning gain, the second one on the self-regulation of the feedback.

5.2.2 Procedure

This study involved two fifth grader classes (10-11 years old) from School 2. The teachers of these classes were thus not the same as those who designed the activity with us. One class consisted of 19 pupils and was split into five triads and two dyads. The other class had 21 pupils, split into seven triads. The teachers ran a workshop-like activity, and sent each group one after another for ten minutes. The lamp was installed in the corridor next to the classes.

Before and after the activity on the lamp, we asked the pupils to fill a pre-test and a post-test to evaluate the learning gain. Each test consisted of two exercises as the one shown in Figure 5.4. Each exercise consisted of two measures to read on the reproduction of the graduations of a protractor, one in each direction. Both post-test sheet had three acute and one obtuse angle to measure.

During the activity, we gave the measurement cards one by one, to a different pupil each time, in the order shown by Table 5.2. The pupils used a circular protractor shown in Figure 5.1. All but two groups went through the 15 measurements: one group stopped at the tenth measurement, and the other group at the fourteenth, because of the limited time.

We demonstrated how to build the first angle, and to display the feedback. We then repositioned the control card, and let the pupils proceed. Due to a technical issue, the precision of the feedback on the clockwise angles lacked precision (they were shifted by three degrees).

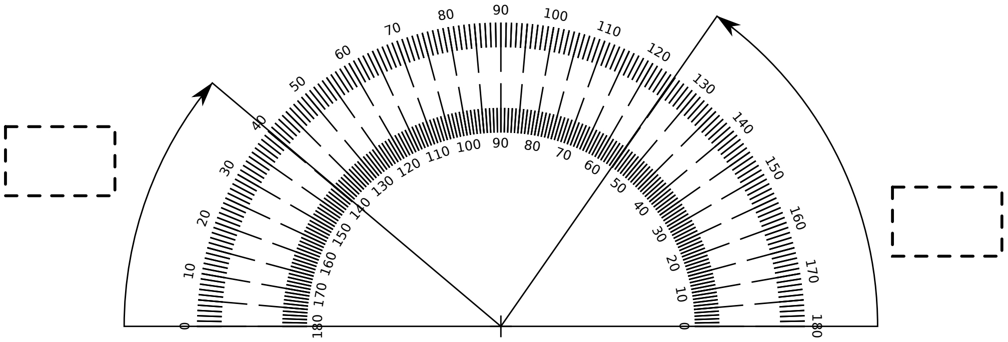

Figure 5.4: One exercise representing a protractor used on the pre- and post-test.

Table 5.2: Order of the angles to be built. (CCW and CW respectively stand for counter-clockwise and clockwise).

Order	Measure	Orientation
1	90°	CCW
2	60°	CCW
3	135°	CCW
4	42°	CCW
5	150°	CCW
6	177°	CCW
7	360°	CCW
8	120°	CW
9	90°	CW
10	45°	CW
11	138°	CW
12	30°	CW
13	3°	CW
14	180°	CW
15	270°	CCW

The experimenters thus intervened to validate a measurement for these angles, when the feedback was wrongly negative.

We split the groups into two conditions: *continuous* and *restricted.* In the first, we allowed the pupils to display the feedback continuously, even when they were changing the angle. In the second condition, we prevented them to move the control card when the measure was displayed. In the latter condition, we gave them only two tries. It happened only once that the pupils were not able to build the angle (360° CCW) within two tries; all other groups succeeded on all the angles within two tries. This was the only time where the experimenter needed to intervene to enforce the use of the feedback in the restricted condition. In the continuous condition, the experimenters reminded the pupils during the activity that they could use the feedback as a guide to build the angle.

5.2.3 Results

(Q_l) **Which Kind of Feedback is Most Helpful to Learn how to Measure an Angle?**

Treatment Check The restricted condition was more limited than the other: the pupils were not allowed to move the control card while the feedback was switched on. Conversely, the pupils in the continuous condition could emulate the restricted condition by not flipping the feedback card when they were moving. We could not prevent this; we could not force the pupils to look at the feedback, so in the worst case, they could always ignore the continuous feedback. We will report later on the details of the usage of the feedback, but in general, not all pupils in the continuous condition used the feedback continuously.

Significance On average, the restricted condition had a beneficial effect on the learning gain[1], but a mixed effect ANOVA with the group as a random factor showed that this was marginally significant ($F[1,32] = 3.80$, $p = 0.06$). Actually, in the continuous condition, the learning gain was negative ($M = -0.3$, $SD = 0.15$); the learning gain in the continuous condition is 0.1, with a standard deviation of 0.2. Figure 5.5 situates these scores on the range of possible learning gain (from -4 to 4).

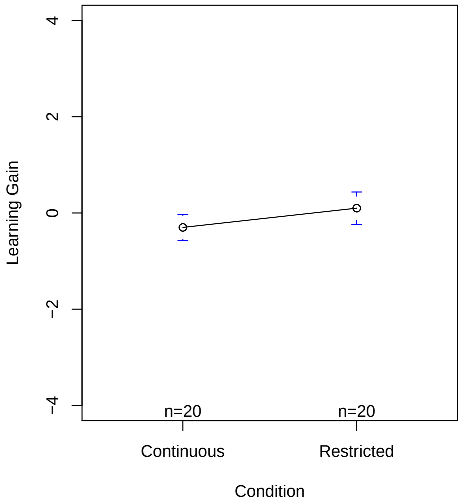

Figure 5.5: Effect of the condition on the learning gain.

Discussion These results are not strictly conclusive on whether the restriction to two feedback requests had an impact on the learning gain: the learning gains are close to zero, the difference between the two conditions is marginal, and not statistically significant. However, we find these results noteworthy.

---

[1] learning gain = $\text{score}_{\text{post-test}} - \text{score}_{\text{pre-test}}$

First, we did not expect such a short activity to fundamentally change the knowledge of pupils. Thus, learning gains close to zero are not surprising. Second, the difference between the two conditions is not huge. They have the same content, and as explained in the treatment check, some subjects of the continuous condition behaved as in the restricted condition. Third, the ANOVA does not allow us to be conclusive, but given the small sample, it lets us think that the effect is not random.

In any case, it is not our objective to prove that controlling the way pupils can check their answer is beneficial for learning. It is not hard to imagine that telling the pupils that they have a limited number of trial will push them to take a greater care about what they do. Here, we illustrated the fact that an interaction based on paper allows subtleties that can be exploited in effective pedagogical designs.

(Q_u) How do Pupils use Feedback Switches?

In the continuous condition, the pupils did not always use the feedback in a continuous way, i.e. displaying the measure of the angle being constructed, and adapting the movements of the controller to the value being projected. We thus analysed how the feedback was used in the continuous condition. We could do it post-hoc, based on the traces of the interaction. Figure 5.6 shows an example of visualisation that we used in this purpose. It shows which measure was displayed to one group (in the continuous condition) along with their target measure, and the most likely error (the supplement of the target angle)

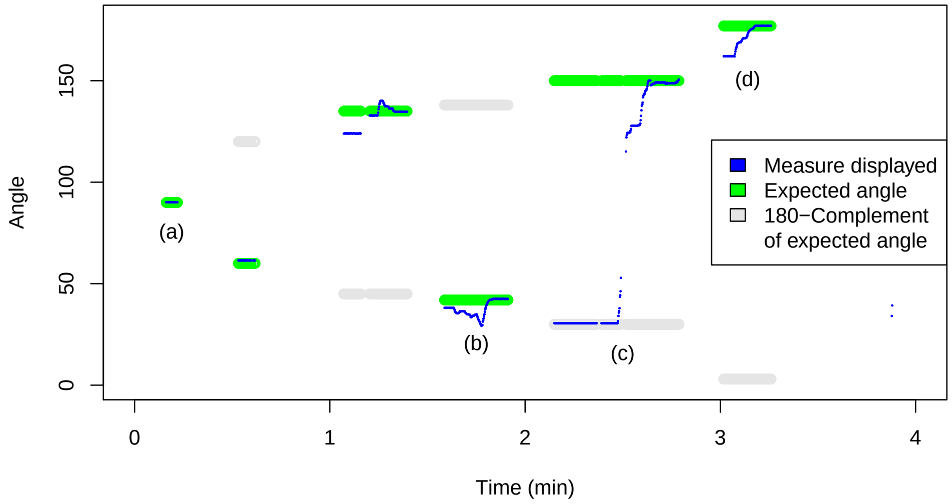

Figure 5.6: Visualization of the feedback given to one of the groups in the continuous condition. The value of the angles built, together with the angle expected and their supplements, are plotted against time. The various values are plotted only when a feedback was shown.

We distinguished four ways of using feedback switches. Figure 5.6 also shows an example of each:

1. Punctual feedback: pupils stop moving the control card, flip the measure card, read the feedback, and flip back the measure card before moving the control card again (see Figure 5.6, a). This way corresponds to the interaction enforced in the restricted condition, but also the most popular in the other condition: it concerned 61 of the 100 total requests of feedback requested by all the groups on any angles.
2. Fine-tuning. In some cases, the measurement displayed was not exactly the one expected; there was a tolerance of three degrees to validate an angle. In 24 cases, the pupils used the continuous feedback to interactively tune the angle to reach a perfect measurement (see Figure 5.6, b).
3. Correction of errors: the pupils would hide the feedback while preparing their answer, and when they saw that it was wrong, they corrected it while keeping the feedback displayed (see Figure 5.6, c). This happened in 9 cases.
4. Continuous feedback: the most intensive use of the feedback, i.e. continually displaying the measure of the angle built (see Figure 5.6, d), happened only 6 times.

This was surprising because we were expecting the pupils to “play” with the interaction. Actually, the pupils considered it more playful to challenge themselves by trying to fulfil the exercise without any help. One of the pupils even hid the feedback saying: “I want to see whether [the manipulating pupil] succeeds”.

5.3 Complementary Study

The previous, controlled study highlighted the fact that the pupils were not using Angoli as expected. It thus seemed like an interesting complement to run another study where pupils were free to use Angoli as they wanted. We report on this study hereafter.

5.3.1 Objectives

The main difference between the first and second study involving Angoli is that the activity is not regulated by the experimenters. This was done in order to investigate the following question: (Q_u) *What organization emerges from the unconstrained interaction of a group of pupils with Angoli?*

Angoli is composed by many small elements (cards and a protractor) which are easily distributed in space and among the pupils, in order to dispatch roles, responsibilities, tasks, etc. Without the intervention and presence of the experimenters, pupils are able to distribute these elements of the pedagogical script as they see fit.

5.3.2 Procedure

This study involved two fifth grader classes (10-11 years old) from School 3. The 33 pupils who did not participate in the Quads study were split into seven tetrad and one pentad. The teachers were running a workshop-like activity, and sent two groups each time for ten minutes. Two lamps were installed in a spare room of the school.

This study used a different set of cards interleaving clockwise and counter-clockwise angle, as shown in Table 5.3. No pre-test and post-test was done. The pupils were shown the activity with an example 90° angle. Then, they were given a deck of cards sorted in the order given in Table 5.3. Moreover, the cards were numbered.

Table 5.3: Order of the angles to be built. (CCW and CW respectively stand for counter-clockwise and clockwise).

Order	Measure	Orientation
1	150°	CCW
2	120°	CW
3	70°	CCW
4	30°	CW
5	60°	CCW
6	110°	CW
7	45°	CW
8	55°	CW
9	15°	CCW
10	125°	CCW
11	165°	CW
12	135°	CCW
13	42°	CCW
14	138°	CW
15	180°	CW

5.3.3 Results

(Q_u) What Organization Emerges from the Unconstrained Interaction of a Group of Pupils with Angoli?

Order of the Exercises The most striking observation relates to the order of the angles built. Sheets are usually ordered: their content is read from top to bottom, and from left to right (in our context); several sheets are bound to be ordered one after another. In contrast, the order of the cards is not considered to be fixed.

This study illustrated this fact. We gave the cards in the same (intended) order in the stack. We also numbered them in the bottom right corner, with a red pen. Even then, pupils often chose the angle measure they would prefer. Out of the eight groups, only two followed the designed sequence of cards. Two groups obviously skipped angles of a given orientation (one group built only clockwise angles, the other group only counter clockwise angles), the others skipped one card. Another group even skipped all the cards corresponding to clockwise angles, redoing the same angles at the end.

Cards are thus a good format of paper interfaces when it is expected that the pupils decide on the order of the exercises. They usually carry one atom of content, i.e. they represent an object, or a function, but not several at once. Pupils can pick them when needed, like they would do in a card game.

Feedback Usage Following up on the previous experiment, we also studied how the pupils used the feedback. We found a distribution similar to the continuous condition of the previous study: 44 punctual verifications, 13 tuning, 11 corrections, and 2 instances of continuous feedback. In other words, pupils also preferred the thrill of the challenge rather than playing with the live feedback.

Time Management A possible reason to explain the distribution of the kind of feedback is that time management was strongly enforced. Each pupil wanted his or her share of the exercise, and made sure that the other pupils did not have more. The most clear examples happened on the tuning feedback: the pupils preferred having a perfect result, i.e. the exact measure rather than an acceptable approximation (displayed in yellow instead of green). As a consequence, they were tempted to fine tune their result. Typically, some pupils tried to tune the angle they just built to perfectly obtain the expected measure. If it took too long, the rest of the group pressured them to give up. Afterwards, the pupils also tried to achieve a perfect measure, because “she did it too”.

Distributed Interaction The most important element of the interface was the control card: it is the one that defined who was doing the exercise. All other elements were mostly auxiliary. The pupil manipulating the control card would rarely fetch the measure card; another pupil would do it. Sometimes, a third member of the group would flip the measure card to display the feedback. More rarely, a pupil would adjust the position of the protractor even though he or she was not controlling the angle.

However, everyone was focused on the activity, even the pupils who were not manipulating. Most of the group argued over whether the manipulating pupil was right or not. This led the manipulating pupil to verbalize, at least to justify his or her actions.

Example Figure 5.7 illustrates the previous points. It is the transcript of the translation of a discussion between a tetrad, at the beginning of the activity. One pupil (C) had already understood which graduation of the protractor was to use given the colour of the angle (clockwise for orange, counter-clockwise for blue). He is to build an angle of 120° clockwise. The other members of the group had not understood the correspondence between color and graduation. Since the previous angle used the counter-clockwise graduation, they thought that C was wrong – he was actually correct.

(C is controlling the angle; he has to build a 120° angle.)
B: It's 120.
A: The other way...
D: The other way...
A: The other way...
D: The other way...
D: It's gonna be 60.
D: Go ahead, try!
C: Do you want to see?
D: Yeah I want to see!
B: Let's bet!
C: *(With a very satisfied tone)* 118.
D: Yep.
A: Well it's not 120...
A: In fact there is a 120 on both sides.
B: Wait let me see.
D: Because it's yellow.
B: Move this card away.
B: Wait...
C: 120.
D: There you go!
C: 60!
C: 60, I put it.
D: 59
B: Ah OK, I got it!
C: Good job, you're improving...
B: Hey!
C: Whose turn is it now?
D: *(Pointing at one of the zeros of the protractor)* Ah yeah, it's because it comes from here.
C: Whom is this for?
B: For me.
D: 70.
B: Because actually it's here.
C: No, because it's blue, so it's here.
B: *(Opens her mouth wide, but does not say "Ah!")*
D: *(Pointing at the two graduations)* When it comes from here, it is here, and when it comes from here, it is here.

Figure 5.7: Translated transcript of an episode of conflict between pupils A,B,C and D in Angoli.

C is sure of his answer. He uses the punctual feedback to build some tension (line 9) and has an enjoyable victory (line 12). Had he used a continuous feedback, the pupils would not have built this conflict, and they probably would not have reflected as much on the difference between the graduations of the protractor in the exercise. C was able to drive the discussion, because he was recognized as the one manipulating, and the rest of the group was just observing and assisting (e.g. B reading aloud the measure in line 2).

As a counter measure to C's lack of humility, A remarks that the measure is not exact (line 14), but she won't use the feedback to fine-tune the answer, because it is not her turn to manipulate. She also won't insist on the tuning, because she wants her turn to come faster.

The mix of clockwise and counter-clockwise angles, along with laconic explanations about the difference between the two control cards, provided the potential for interactions. The tangible interface supported this potential in two ways. First, the pupils could point at the objects they are talking about to illustrate their points. Second, the feedback allowed to authoritatively rule who was right or wrong. It is thus a good way to foster learning from conflicts between peers.

5.4 Conclusion

In this chapter, we presented Angoli, an activity to master the usage of a protractor and measure angles. Angoli makes use of a real protractor, which is part of the curriculum and not replaceable by a virtual one. An anecdote concerning the pre- and post-test of the first study could show the importance of using a real protractor. One pupil counted every graduation of the representation on paper instead of directly reading the measure. During the activity, she used the protractor correctly. At the post-test however, she counted again every mark of the graduation. It is as if for her, there was no link between a virtual protractor and a real one.

Angoli consists of two control cards, and a series of measure cards, which are used to control the display of feedback. A first, controlled study hinted toward the importance of controlling the feedback: restricting it seemed to have a positive effect on learning gains. This study also taught us that the pupils preferred punctual feedback rather than continuous feedback. A second study confirmed it.

This is surprising compared to what Do Lenh (2012) called the manipulation temptation, i.e. learning being hindered by the temptation to use the augmented reality interface for the sake of using it. We would have expected the pupils to prefer seeing the measure being modified accordingly to the movements of the control card. On the contrary, they preferred limiting the usage of the augmentation. However, had not the pupils shown this responsible behaviour, Angoli would have made it easy to regulate it, by taking away the feedback cards from the groups who abuse them.

We also illustrated how a scattered interface like Angoli supported collaborative learning. The roles within the groups are clearly defined and materialized by the various elements of the interface. They can be used to support the discussion, and arbiter conflicts.

We have not investigated voluntarily the teacher side of this activity. However, we can add an anecdote on this topic. For the second study, a deck of cards was missing. We were able to replace it by photocopying the one we had a few minutes before the experiment started. This misadventure is actually an illustration of a powerful feature of the activity: it is robust to perturbations as critical as half of the interface missing.

Angoli was also a good opportunity to gather general usability observations on the manipulation of cards as a (flat) tangible interface. These observations are not specific to one of the two studies, and are reported hereafter.

The control cards were used as a cursor. Their orientation did not matter; only their position did. The precision they allowed was sufficient for the high requirement of the activity: an angle is a mathematical object that is very sensitive to small imprecisions. Moving a card on the desktop is well adapted. It moves easily, but remains in place without sliding away when no movement is applied anymore.

The measure cards were used as a switch. Their position or orientation did not matter, only their side did. Detecting the side of a piece of paper is very reliable; the manipulation is however not as easy. It is sometimes problematic to pick up a card from a flat surface without damaging it. One fail safe technique consists in sliding the card on the border of the desk top in order to grab it. This difficulty could also explain why the pupils were assisting the one manipulating by holding the measure card, and flipping it only on his or her request. The distributed interface allows this kind of adaptation.

Two questions remain open regarding the manipulation of cards as interface. First, in which condition should they be used as pointers? It is doubtful that emulating a computer mouse is a good use of paper interfaces. Cards do seem adapted as tangible controllers when their trajectory corresponds to the logic of the application; in our case, we saw the pupils move the cards in circular patterns. We will investigate this point further in the next chapter.

Second, we saw that cards are well adapted to embody atomic contents, i.e. to map one logical item to one physical object. It is sufficient for simple activities like Angoli, which have only two kinds of content: measures and control. It is not clear however how it would scale to more complicated interfaces: too many cards may become cumbersome to manage. Moreover, not all the activity can rely exclusively on cards. The next chapter will also help investigate this point, in combining cards with other paper artefacts.

6 Describing Angles

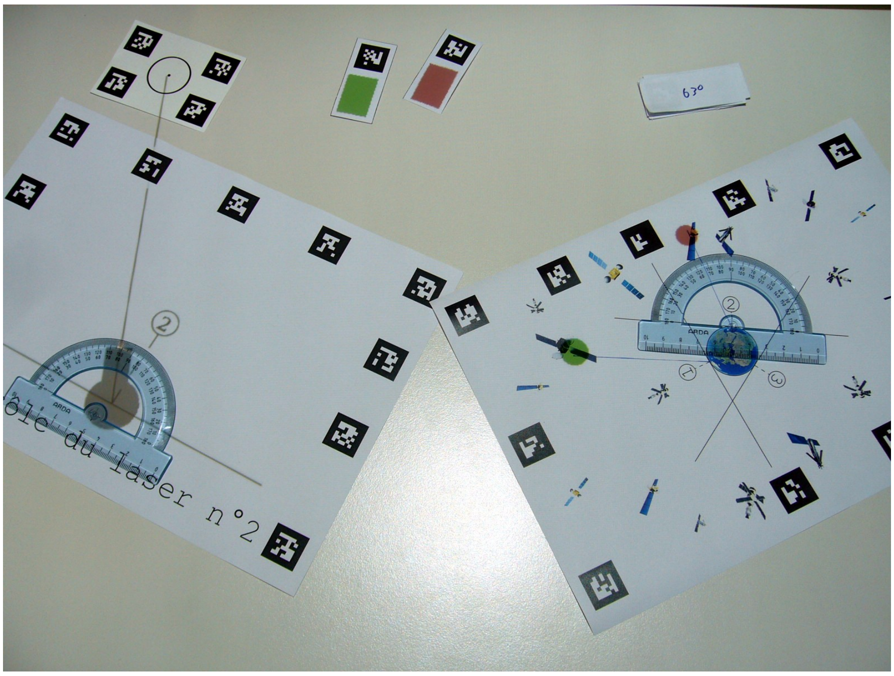

Figure 6.1: The components of the interface of SpaceJunk.

In this chapter, we present SpaceJunk, a problem based activity. As discussed with the teachers, SpaceJunk gives the pupils the opportunity to apply the use of the protractor, after Angoli introduced the basics (see Chapter 5). The problem tackled in SpaceJunk is the description of an angle to someone who can not observe the same representation.

We first present the activity, and then four studies to asses the strength and weaknesses of SpaceJunk in pedagogical context. We used each of these studies to improve iteratively the next one. The first one helped in making the activity viable, and the last one used the observations drawn form the two main studies to refine the interface.

6.1 Description of the Activity

“Space junk” refers to the non-functional artificial objects orbiting around Earth, such as retired satellites. The goal in SpaceJunk is to clean space of this waste with lasers. There are two roles in the completion of this task: *controllers* orient laser cannons from the ground towards the space waste, and *observers* state the inclination given to these lasers. The observers and the controllers are physically separated by a wall, which acts as a blinder (see Figure 6.2). We further present the elements of the interface, the scenario, the design objectives, and the position in the framework.

Figure 6.2: The observer team (left) measures the orientation to give to the laser, and communicates it to the controller team (right).

Figure 6.3 provides a link to a video demonstration of the activity.

Figure 6.3: A QR Code of the url http://short.epfl.ch/spacejunk which points to a video demonstration of SpaceJunk.

6.1.1 Interface Elements

There are three laser cannons to control in SpaceJunk. For each of them, the controllers are given a sheet, like the one shown in Figure 6.4a. Only the identifying number (1, 2, or 3) of the laser is printed on each sheet, and fiducial markers allow the projection of a cannon, and a line representing the ground (see Figure 6.4b). The number of the cannon is also projected to label the projection. Projecting the origin of the cannon rather than simply printing it is a technical concession to allow a more precise correspondence between the projected and measured angle.

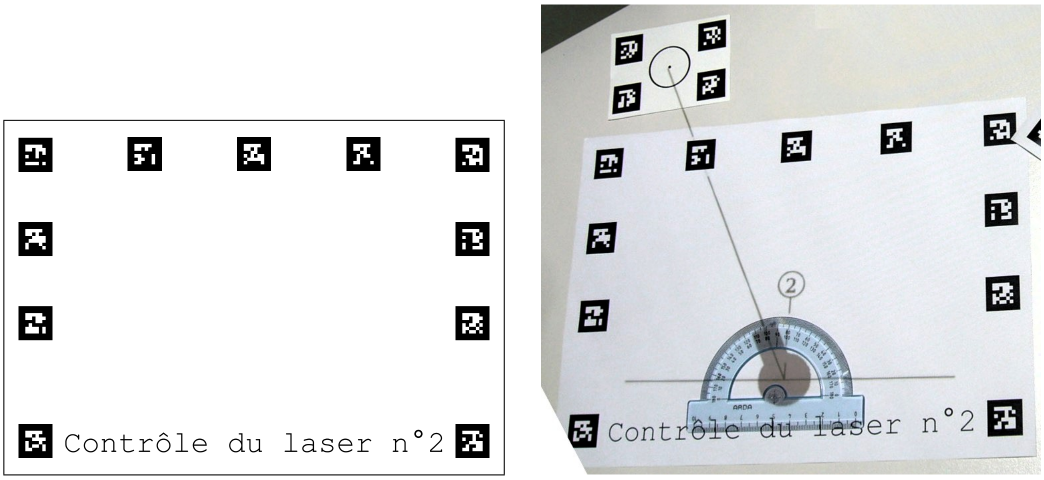

(a) The printed content of the sheet for laser number 2.

(b) A laser sheet and the augmentation linking it to the control card.

Figure 6.4: Laser Sheets.

The inclination of the projected cannon follows a line originating at its base and ending on a card, as shown in Figure 6.4b. This card is similar to the one used in Chapter 5 but is not associated with a determined orientation (clockwise or counter-clockwise). The controllers are provided with a protractor to measure the inclination of the cannon, i.e. the angle between the ground line on the sheet and the line between the cannon and the card.

The inclination of the cannon is determined by the observers. To do so, they are given a protractor. The observers are also given a sheet representing Earth and orbiting satellites, as shown in Figure 6.5a. The positions of the laser cannons on the surface of Earth are marked by a label of the corresponding number. The cannons are equally distributed around the surface of Earth, i.e. every 120°. The tangent to Earth is printed for the position of each cannon, to depict the horizontal baseline of each cannon. The fiducial markers on this sheet allow the projection of a green circle for the current target, and red circles for satellites already hit by a laser (see Figure 6.5b).

The last element of the interface is the *ammunition*, tiny sheets of paper with a single fiducial marker printed on them (see Figure 6.6). The observers use these ammunition papers to write the angle of inclination that will be given to the controllers for their cannons. The ammunition is also used to trigger the shot. When the fiducial marker of the ammunition is detected, the projection of a yellow rectangle grows for 3 seconds, which offers the opportunity of a last minute shot cancellation by flipping the ammunition paper back over. Each ammunition can only be used once, except for a special kind of unlimited ammunition, which has the form of a single fiducial marker on a longer piece of paper. When used, the ammunition turns green or red, depending on whether the target was hit or missed.

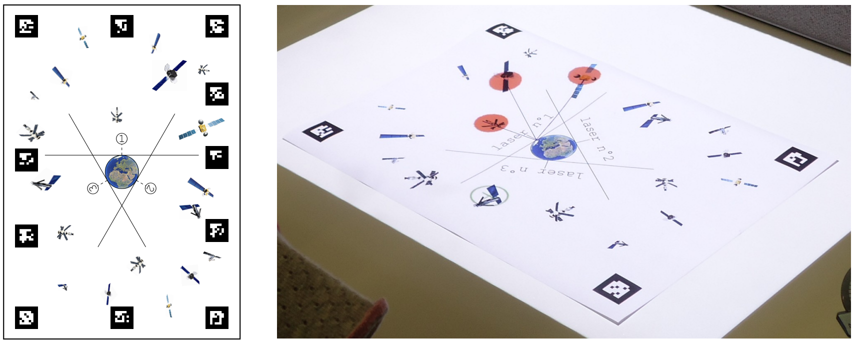

(a) The printed content of the satellite view.

(b) The view of the satellite, after three satellites have been hit and marked in red. A fourth satellite, the current target, is circled in green (bottom left from Earth in the picture).

Figure 6.5: Satellite View Sheets.

Figure 6.6: Two used ammunitions on the left, and a stack of remaining ammunitions (the first of which has already been annotated on its back). The green augmentations mean that the shot was a hit, and the red augmentation means that the shot was a miss.

6.1.2 Scenario

The observers are instructed to draw a line originating from the position of one of the 3 laser cannons to the target satellite (using a ruler on the view of Earth). They use the protractor to measure this angle with respect to the horizontal axis for this laser. Finally, the observers have to describe this orientation to the controllers. Optimally, the observers give the value of the angle, the direction of measurement, and which laser to use. They write this information on the back of an ammunition paper and give it to the controllers.

The controllers can change the inclination of the appropriate laser using the control card. They reproduce the angle received from the observers using a protractor. Finally, the lasers can be activated by flipping the ammunition to uncover the tag.

The trajectory of the laser is shown for three seconds on both the sheets of the controllers (laser sheet) and the observers (Earth view), with a fading blue line. If the satellite is hit, the ammunition turns green. Otherwise, it turns red, indicating a missed shot.

The lasers shoot through Earth, i.e. below the horizons. However, in some cases, two lasers have a same target in sight; any laser can be used. The order of the satellites to destroy is given by the system, and only the target highlighted in green can be hit.

6.1.3 Design Objectives

The main objective of SpaceJunk is to illustrate the fact that protractors can be used with the Metroscope in a problem situation. Moreover, the observers can use the protractor as expected in traditional geometry lectures: with a pen and a paper sheet.

An angle is useful for two types of situations: applying theorems (e.g. in trigonometry) or for describing and communicating orientations. Since the former will come later in the education of the pupils, we aim at leveraging the intrinsic motivation of the pupils by creating a situation where the measure of an angle results from a need of communication. There are three intrinsic difficulties that the observer has to face in this communication: the precision of the measure of the angle, the direction of measurement (clockwise or counter-clockwise) as well as the origin of the angle (which cannon).

To orient the pupils towards establishing the convention, we used the SWISH principle described by Dillenbourg and Hong (2008) to design pedagogical scripts. We split the interaction on the visualization of the inclinations of the cannon: the observers see the target inclination and the controllers see the actual inclination, but not vice versa. By doing so, we force each part of the group to interact about the description of the angle.

6.1.4 Position in the Framework

Table 6.1 shows which parts of the framework are covered by SpaceJunk. The main piece of paper for the interaction is the control card, which is used as a cursor. Only its position is used, not its orientation. Finally the laser sheet, which selects one of the three lasers, and the ammunition, which triggers the shot.

6.2 Preliminary Study

Because SpaceJunk is more complex than the previous activities, we carried a preliminary, small-scale study to test the viability of the activity:

- (Q_I) Is SpaceJunk adapted to pupils' learning to communicate angles?

This study involved a fifth grade class (10-11 years old) from School 2. The teacher was thus not involved in the design of the activity. The class had 23 pupils, which were split into four dyads, one triad, and three tetrads. The tetrads tried the activity for 20 minutes, while the smaller groups had 10 minutes. The teacher was running a workshop-like activity, and sent each group one after another to the corridor next to the classroom where the lamp was installed. Regardless of group size, the experimenter intervened when the students were stuck on a given target for too long.

Table 6.1: Position of SpaceJunk in the framework.

Property	Ephemeral	Persistent	Permanent
Presence	The controllers show the ammunition to trigger the shot. They show the laser sheet to select the corresponding laser.		
Position	The pupils move the control card relative to the laser sheet.		
Orientation	The pupils can rotate the satellite view sheet for easier measurement; the target indication follows.		
Side			
Folds			
Edges			
Ink	The observers measure the angle they draw on the satellite view sheet.	The observers draw the expected laser shot on the satellite view sheet. They also write the expected inclination on the ammunitions.	The angles to build are defined by the targets printed on the satellite view.

The main observation drawn from this study concerned the influence of the group size on the performance. The dyads and triads barely managed to hit two targets in ten minutes despite intensive help from the experimenter. On the contrary, the tetrads usually hit their first target before the five minute mark with little or no intervention from the experimenters.

We hypothesized two explanations for this observation. First, the presence of the two experimenters was more intimidating, because there were as many of them as the subjects. The pupils did not dare experimenting, and asked for confirmation or waited for a feedback at each step.

Second, and more importantly, the dynamics of the group is clearly different between the conditions, as illustrated in Figure 6.7. In dyads, the observer sits idle while the controller works, and vice versa. In tetrads, one sub-team can be active while waiting for the other sub-team to be ready. They discussed the next measurement, or reviewed the previous one. The pupil manipulating the protractor was also receiving comments from his or her teammate. Finally, a small competition between the two sub-teams emerged, and fostered discussion. As a result, the pupils tried more strategies and succeeded faster.

(a) A pupil day-dreaming through the window while his team mate does his part.

(b) Two pupils discussing the task while their team mates do their part. Within a sub-team, one pupil observes and gives feedback to the one acting.

Figure 6.7: Typical group dynamics depending on the size of the group.

In the end, with tetrads, the activity seemed not too easy yet do-able for pupils learning to communicate angles. The study further revealed other technical adaptations for future studies, such as a bug in the validation computation, superfluous and disturbing graphical elements (both projected and printed), and an insufficient number of fiducial markers resulting in occlusions that the system could not compensate for any more.

6.3 Informal Study

We organized a workshop during an open day event targeted at primary schools. Seven classes visited our workshop, which gave us the occasion to study how *a whole class* uses SpaceJunk.

6.3.1 Objectives

Such a study is very far from a traditional, controlled experiment. The number of subjects is high, but time is a hard constraint because we had to respect the schedule of the rest of the event. In these conditions we could investigate the following topics:

- (Q_u) What do pupils think about this activity in more realistic conditions? The difficulty in asking for feedback from the pupils about the system comes from the novelty effect. Pupils are enthusiastic about the novelty (especially if it replaces something from the school) and their feedback is artificially positive. With a high number of pupils, we can expect to extract some useful information. However, in the context of an open day event, we cannot ask the pupils to fill out a full usability questionnaire; this is not what they came for. Instead, we used sticky notes for them to quickly express the most salient comments.

- (Q_l) How does SpaceJunk support the emergence of ad-hoc conventions? The main point of SpaceJunk is to have the groups elaborate a convention to describe an angle without a shared vision. The pedagogical goal is to let them explore the difficulties of such an elaboration, so that the pupils understand the benefits of a standard. With so many subjects, we can expect a wide variety of conventions, captured on the ammunition sheets, used as communication between the observers and the controllers.

- (Q_t) Is SpaceJunk manageable with a full class? This study is also an opportunity to deploy the Metroscope in a configuration where the whole class uses it at the same time, as opposed to one or two groups at a time.

6.3.2 Procedure

The study took place in the meeting room of our laboratory, shown in Figure 6.8. We ran one session for each of the six classes, every hour. Each class stayed for 45 minutes. During 10 minutes, we introduced the laboratory, Metroscope, and the activity. For the first four classes we gave verbal instructions, and helped the groups for their first shot. For the last three classes, we gave a direct demonstration. The activity itself lasted 30 minutes, and the last 5 minutes were reserved for the collection of feedback. We told the pupils to exchange roles (i.e. observers become controllers) after a few shots.

The pupils were 11-12 years old, and had already been introduced to the protractor. We split the first class into six tetrads and triads, and the last seven into five tetrads or triads, for a total of 36 groups. Each group used a different Metroscope. The first 3 classes had a special ammunition, which was indefinitely reusable. It consisted of a tag on a piece of paper, like the regular ammunition, except that the piece of paper was bigger, to allow more annotations. The last four classes had the regular, one-use ammunitions.

To collect the feedback of the pupils, we distributed sticky notes at the end of the activity. The first three classes were asked to write down what they liked or disliked about the activity. Then, the pupils had to stick their note on a white board split in two halves, one with a smiling face, and the other with a sad face (see Figure 6.9). We cleared the board after each class so as not to influence the following one. Since we found that there was too much positive feedback (on the smiling face side), we asked the last four classes to write one sticky note for each side of the board, i.e. one thing they liked, and one thing they did not like, or that we could improve.

6.3.3 Results

(Q_u) What do Pupils Think About this Activity in Real Conditions?

In general, the pupils were very engaged and enthusiastic about SpaceJunk. The vast majority of sticky notes showed a sign of appreciation; the sticky notes reserved for bad feedback also often stated “nothing” as an answer to “what did you not like/what would you change?”.

Figure 6.8: The first class of the study listening to the explanation by one experimenter. Photo Alain Herzog

This is probably linked to the playful aspect of SpaceJunk: 15 sticky notes qualified SpaceJunk as a game, versus one as an “electronic protractor”. Several sticky notes proposed game-related improvement: more colors, make the satellites explode, more details (such as characters), shooting something else, or “make the game touch-based”. It is encouraging to see that the engagement in SpaceJunk is comparable to that of a game; however, it is not its purpose, and more recreational elements could reduce the focus on the topic to be learned: angles and their description. We thus find the level of engagement satisfying, and do not think that we should increase the playfulness.

On a related note, the pupils clearly preferred the role of controller to the role of observer (ten sticky notes on the former, versus one for the latter). Similarly, a pupil noted that he or she did not like to wait. While it would be possible to give concurrent tasks for the observer and the controllers, it would probably hurt the communication between them, which is the main point of the activity.

The sticky notes did reveal several points as to how SpaceJunk could be improved. First, as a result of the observations made in the preliminary study, more tags had been added on the sheets to increase the stability of the projection. It seemed sufficient during the development, but eight pupils disagreed, writing on their sticky note that the projections were not stable.

Another possible improvement concerned the explanation of the task, which three pupils found unclear. This is understandable, given the difficulty that the experimenters had of explaining the task without giving a precise procedure that would spoil the exploration of the pupils in terms of communicating angles.

Figure 6.9: Pupils sticking their feedback note on the whiteboard. Photo Alain Herzog

Three sticky notes also revealed a usage detail that escaped the notice of the developer: the physical positions of the pupils across the projection surface makes one of them very close to the hot exhaust of the Metroscope, which is not comfortable. Another pupil complained about the vibrations of the lamp.

Last but not least, some sticky notes commented on some of the design decisions of SpaceJunk. Three sticky notes complained about the form of the ammunitions: they did not like the fact that they could be used only once, and that it took time to write on them. However, two pupils wrote that they liked using the protractor, as permitted by SpaceJunk.

(Q_l) How does SpaceJunk Support the Emergence of Ad-Hoc Conventions?

Using free annotations on the ammunitions to communicate the description of the inclination to give the laser cannon resulted in a surprisingly wide variety of conventions, as illustrated in Figure 6.10. The convention to describe a shot needed three elements: the value of the angle, its orientation, and its origin, i.e. which of the three lasers should shoot. As simple as it seems, the descriptions created by the pupils revealed seven characteristics, which we describe hereafter:

1. how the angle was coded
2. how the laser was coded
3. how the orientation was coded
4. the order in which these elements came

5. the verbosity of this information
6. the representation of this information
7. the evolution of this information

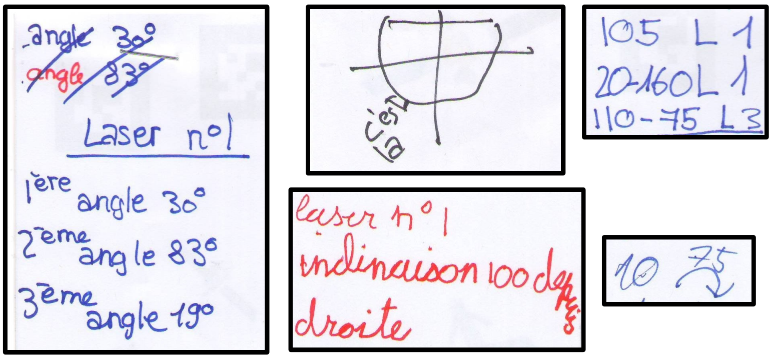

Figure 6.10: A sample of the various formalisms invented to describe an angle.

At first, the value of the angle may not seem like a place of creativity: one would expect to simply read the value off the protractor. Most pupils did so, with varying precision (sometimes sub-degree, sometimes rounded to the next multiple of five or ten). The system allowed a tolerance of ±10° for the first satellites, to ±4°, to ±2°. One group also mixed up units and decimals, writing for example “10,2” instead of “12”, which actually makes sense if read as “the second graduation after the “10” mark”. Another group took their measures starting from 90° rather than 0°, presumably by rotating the protractor by 90°.

The lasers, when identified, were invariably identified by the number labelling them. This is the only constant we found in all the descriptions. Only 6 out of 36 groups did not mention the laser in their written communications.

In contrast, the orientation of the inclination spawned many different conventions. A first observation is that no group used the term “clockwise”[1]. Instead, the most popular way to disambiguate the orientation was right or left (on the protractor). This convention came after a short negotiation on whether it was the observers’ or the controllers’ left, most often to realize that it did not matter. Two groups used “from right to left” and “from the left” which is actually closer to the clockwise concept than left and right. The second most popular convention for the orientation was to indicate the two marks on the protractor, i.e. “120/60” versus “60/120”, where the external graduation of the protractor came first. Within this convention, there were further variations, e.g. the second angle could be rounded off differently from the first, or the internal graduation could be used first. In addition, graphical representation could express the orientation, either by an arrow, or by a schema of the protractor, annotated or not, with a mark on the approximate position. Finally, one group just indicated that "you have to do it incorrectly" if the orientation was not correct at first.

---

[1] Note however that the French idiom for clockwise, “dans le sens des aiguilles d’une montre”, is long to write.

The order in which this information was presented was also a source of variation. In most cases, the angle measure, the orientation and the number of the laser came in this order. Many groups used other orders, in particular placing the laser first. Some groups in the unlimited shots condition, who thus had more ammunition and more paper space, organized the shots per laser rather than writing the angles one after another.

Not all ammunitions gathered the three pieces of information in a written or graphical form. The missing information was communicated verbally. Often, the annotations of the ammunition were simply redundant: the controllers had already started setting the inclination of the laser cannon before receiving the ammunition, because they heard the discussion between the observers.

In their written form, the communication between the observers and the controllers covered various levels of verbosity: The unit of measurement of the angles (degrees) is a good illustration: it could be written as “°”, “degree”, both previous ones, “d”, “"d"” or simply omitted. The ammunition could also carry more messages, such as “hurry up”.

Finally, the communications recorded on paper allowed us to see the evolution of the convention between observers and controllers. Evolution of conventions, however, was rare: only four groups changed their convention, and one other did not use a consistent convention. The evolution did not necessarily converge toward a complete description of an angle. For example, one group started by writing only the angle and its supplement, then added the laser but reduced the precision of the angles.

This big diversity that resulted from a seemingly simple task illustrates another potential of paper interfaces in education. After such an activity, a teacher could easily exploit the experience of the class to explain the need for a convention, and justify the standard by looking at the flaw of alternatives.

(Q_t) **Is SpaceJunk Manageable with a Full Class?**

Figure 6.11 shows the distribution of the number of satellites hit by each group. There was a maximum of twenty targets, but no group achieved it. This graph shows that 32 (out of 36) groups were able to hit at least four targets, which means that each pupil had the opportunity to hit a satellite. 35 groups were able to hit at least three targets, which means that each pupil tried both roles (observer and controller). Only one group did not have such opportunities. In general, this is a good result, because the activity was rather ambitious in terms of class management: the task was complex, engaging, collaborative, voluntarily ill-defined, and based on the newly acquired knowledge of using a protractor (in other words: noisy and messy).

The difficulty in managing a full class using SpaceJunk at the same time mainly resides in getting the group started. Once they understood the task, the pupils were autonomous and only required interventions for technical problems or discipline issues. This means that most interventions were needed mostly at the same time, i.e. at the beginning. The annotations on

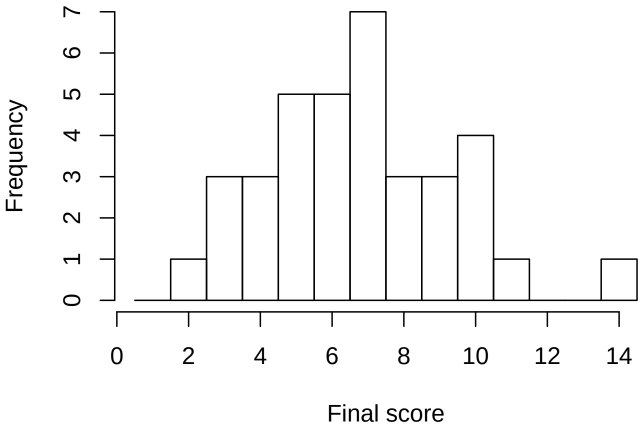

Figure 6.11: Histogram of the final scores of the 36 groups.

the various pieces of paper (the line drawn on the satellite view and the angles described on the ammunitions) helped the experimenters recover the context of the exercise instantly: what the possible problem was, which sub-team was responsible for the problem, etc. The tangible aspect of the interface also helped with the explanations, because the experimenters could show the pupils how to manipulate the elements. Finally, it was easy for the experimenter to keep an up-to-date status of the whole class by wandering around and having a quick look at the interactive surface and the interface elements.

6.4 Formal Study

6.4.1 Objectives

In contrast with the previous study, we conducted another one with only one or two groups at a time. This allowed for a more precise analysis of the interaction between pupils and the Metroscope. The questions we wanted to answer were:

- (Q_l) What are the difficulties encountered by pupils in SpaceJunk? SpaceJunk is an ill-defined activity that aims at letting the pupils explore the problem of describing angles. As such, it is interesting to see whether SpaceJunk allows for different kinds of productive errors.

- (Q_l) How do pupils collaborate in SpaceJunk? SpaceJunk is focused on the communication of an angle. The groups are composed of two sub teams of two. This allows for various group dynamics, which we investigate with this question.

- (Q_u) How do pupils interact with each element of the paper interface? The paper interface of SpaceJunk has different types of components: sheets, a card, and a tiny piece of paper used to trigger shots. We investigate the characteristics of the components of the paper interface.
- (Q_u) How do pupils manipulate pieces of paper in a polar referential? The pupils have to move a control card to orient the laser. By definition, the most efficient move is a straight line. However, what counts from the control cards is the angle it forms in a referential set by the laser sheet. The value controlled is thus the angle of a polar coordinate. In this study, we investigate the influence of these two considerations on the actions of the pupils.

6.4.2 Procedure

This study involved two fifth grade classes (10-11 year olds) from School 3. In this study, the pupils started with Angoli, and continued with SpaceJunk. The first part was reported in Chapter 5, and we report here the second part. The 33 pupils were split into seven tetrads (d1, d2, d4, g1, g2, g3, and g4) and one pentad (d3). The letter of the group (g or d) correspond to which lamp they used; the number of the group (1, 2, 3, or 4) corresponds to the experiment session. The pupils only experience with a protractor was within the context of Angoli. The teachers were running a workshop-like activity, and sent two groups at a time for twenty-five minutes. Two lamps were installed in a spare room of the school, shown in Figure 6.12.

Figure 6.12: The sub-teams on the left are switching their roles, and the group on the right has just hit another satellite.

We used the experience from the previous study to tune the explanations given to the pupils. After explaining the context of SpaceJunk and the overall goal, we gave a short demonstration of what was expected. However, it was too rich for the pupils to remember all the details, which let them explore the problem by themselves.

We told the pupils to switch their roles after two satellites are hit, and gave them 20 ammunitions, without mentioning whether they were limited or not (and they did not ask). We did give the unlimited ammunition, similar to the one from the previous study, to the one group (d3) who used up all its ammunitions. Otherwise, we decided not to intervene, except for technical questions. When the pupils asked what to do, we asked the other members of the group to answer.

6.4.3 Results

(Q_l) **What were the Difficulties Encountered by Pupils in SpaceJunk?**

The number of satellites hit gives a score for each group. Its evolution is shown in Figure 6.13. It gives an overview of the performance, but there is much more to say by describing the different kind of errors that led to missed shots.

For example, d1 and g1 seem to have comparable performances, but it is interesting to note that d1 never missed a shot, while g1 missed five. The last four missed shots are actually due to the confusion between the two graduations of the protractors. This highlights the fact that the duration of the activity should not be too short. Indeed, had the activity stopped earlier, the pupils would have had a perfect score, while they simply had not reached the core of the problem yet. The four missed shots showed them the ambiguity in the description of the angles, even if had not been problematic from the start.

Figure 6.14 helps visualize how pupils missed their shots. It shows that confusing the two graduations of the protractor is not the only explanation for missed shots. In many cases, the inclination given to the shot is neither the expected one, nor its supplement. Only one case (the last missed shot of g2) corresponded to a lack of precision in the reading of the protractor. Most other cases were linked to a wrong usage of the protractor, as shown in Figure 6.15: it happened on the observer side for 28 missed shots, and on the controller side for 16 missed shots. One reason to misplace the protractor is if the line drawn by the observers on the satellite view is missing (as in 4 missed shots) or incorrectly drawn (as in 15 cases). Figure 6.16 shows examples of incorrectly drawn lines. The error produced by incorrectly aligning the protractor was not within the tolerance allowed by the system.

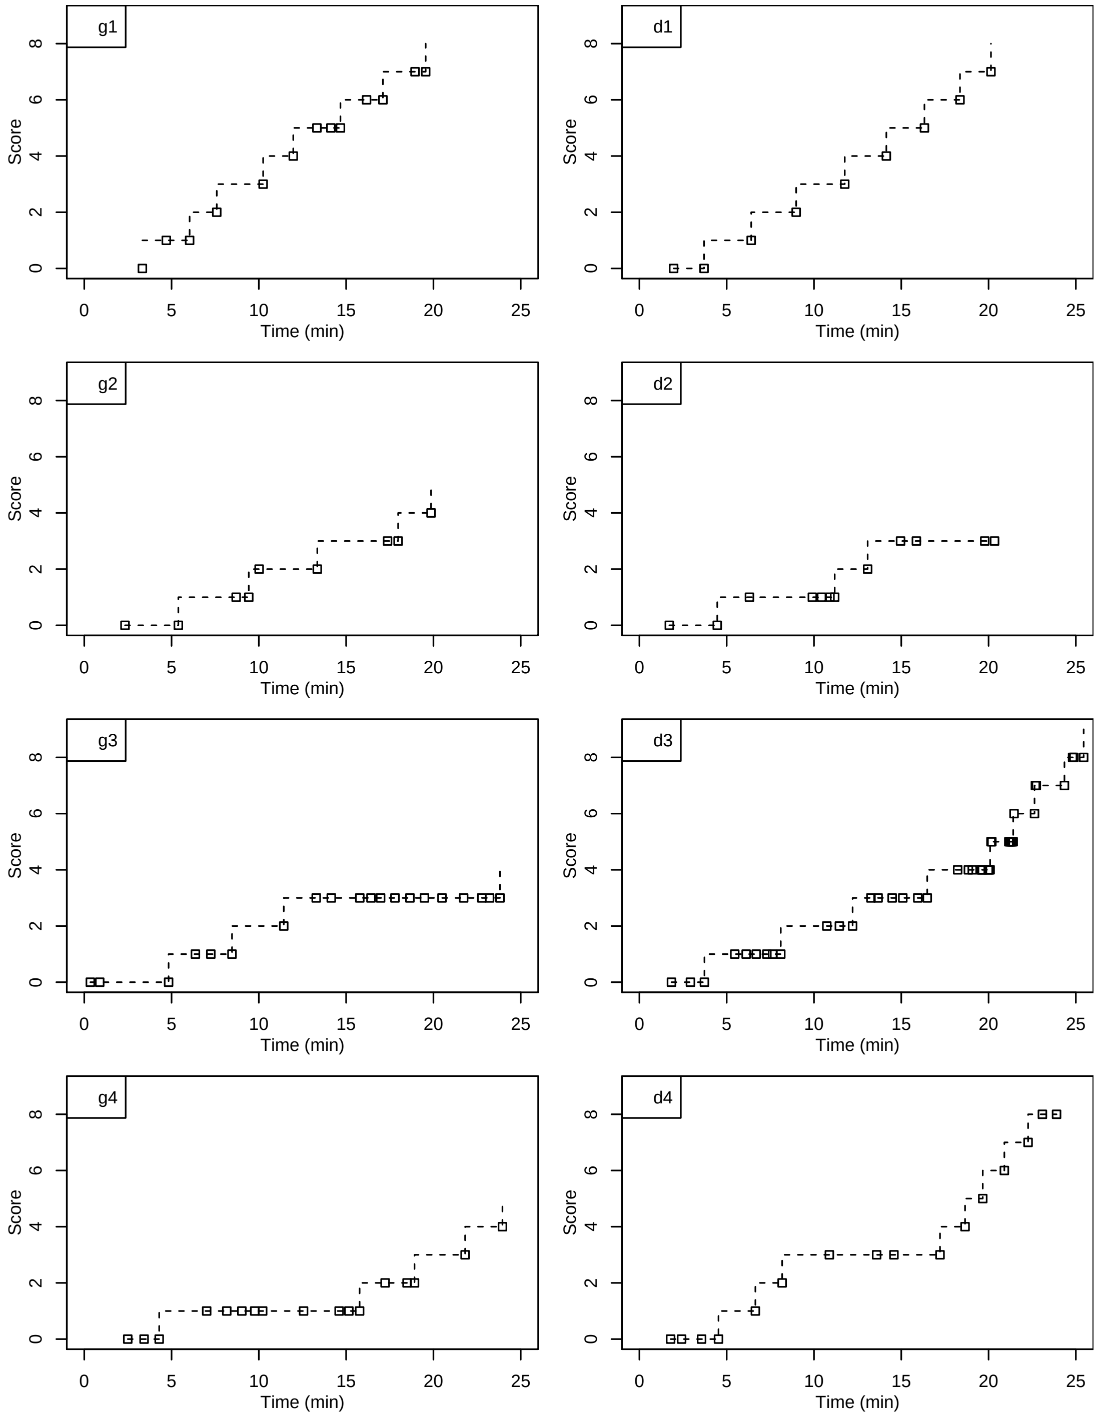

Figure 6.13: Evolution of the scores. Each marker is a shot. The line shows the score (i.e. the number of hit targets) as a function of time (in minutes). Markers in the middle of a segment correspond to missed shots; markers on a corner correspond to hits.

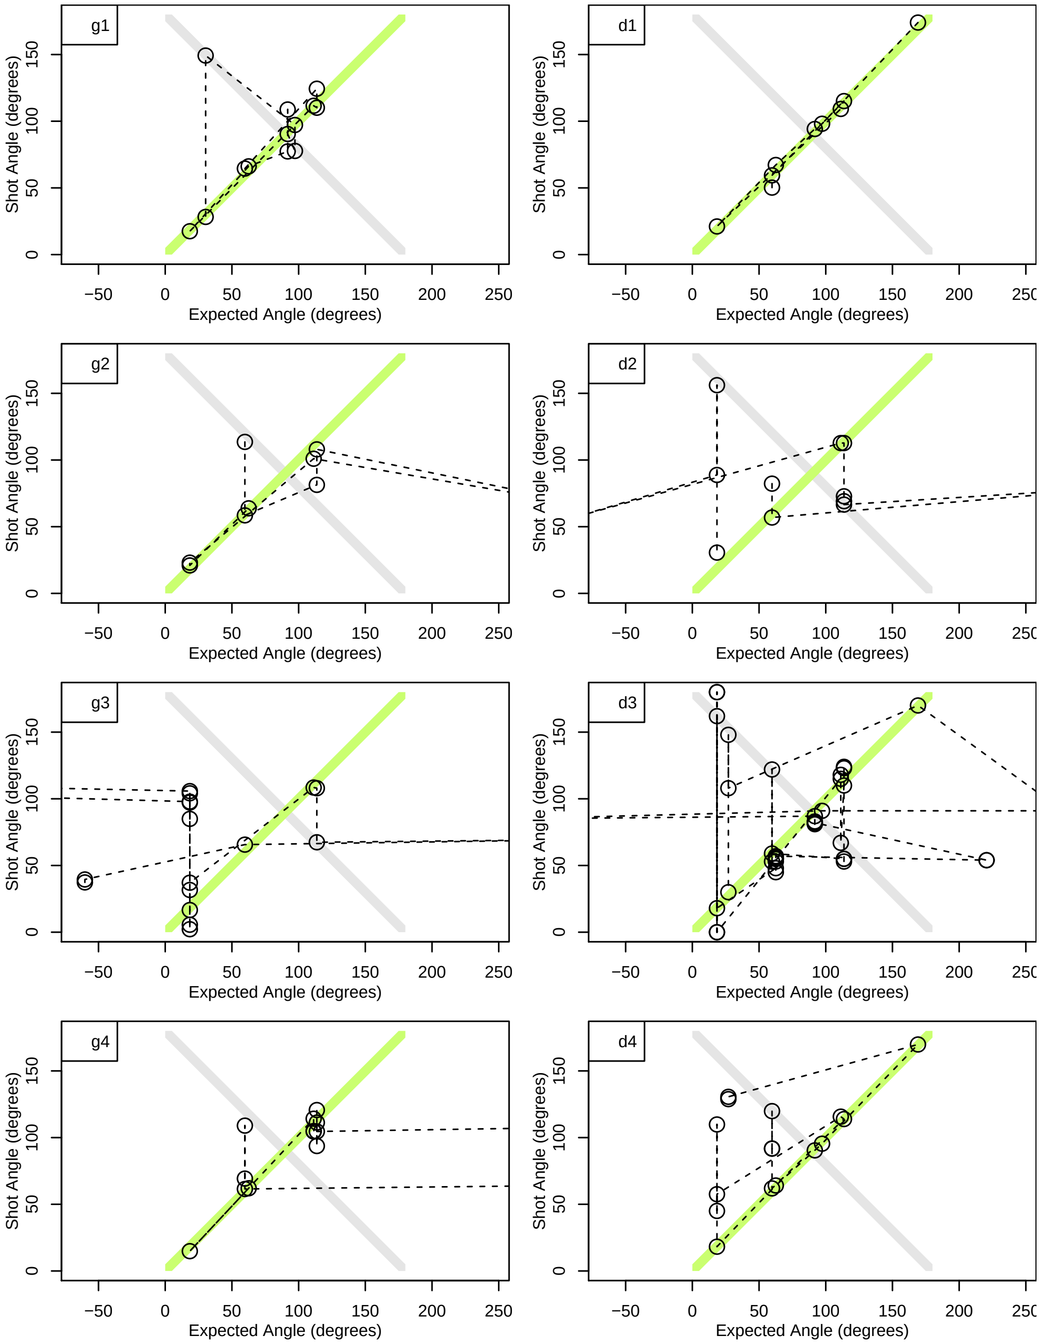

Figure 6.14: The angles shot by each group as a function of the expected target angle. Correctness can be estimated by the closeness of the line $y = x$ (light green, from bottom left to top right). If a shot is close to the line of equation $y = 180 - x$ (light grey, from top left to bottom right), it means that the angle is the supplement of the expected angle. An expected angle outside the $[0, 180]$ range means that the satellite was not attainable with the laser used for the shot. A vertical line with a lot of points signifies a lot of errors on a specific target. The dashed line shows in which order the shots were made (two consecutive shots are linked by a dashed line).

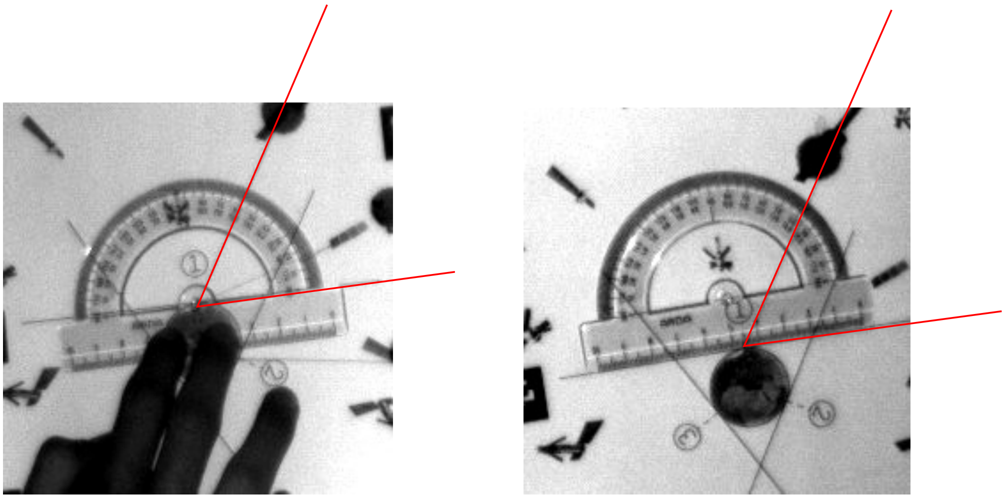

Figure 6.15: Examples of a correctly (left) and incorrectly (right) aligned protractor. The angle is outlined in red.

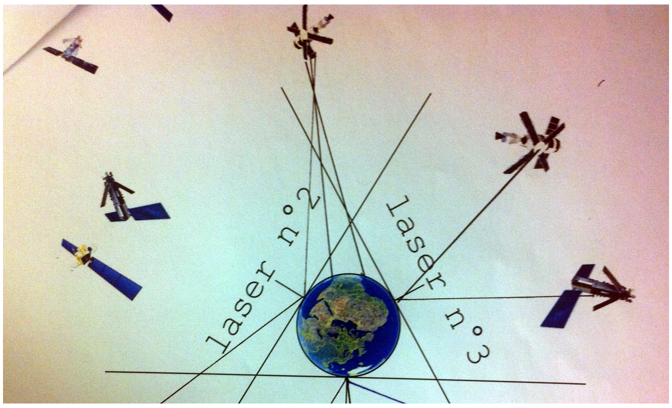

Figure 6.16: Close-up of a satellite view where multiple lines towards the satellite on top show that the origin of the angle was not immediately understood.

Finally, we observed two more causes of missed shots, which have less pedagogical value compared to the previous ones, because they are linked to the interface. First, the shot was performed by the wrong laser, i.e. the wrong laser sheet was selected, which happened 17 times out of 116 shots. Second, the shots could be triggered unintentionally, when the ammunition was detected but the controllers where not ready to shoot. This happened 13 out of 116 times (plus another time, but the shot was actually a hit).

Each group explored the problem space more or less differently:

- d1 wrongly positioned the protractor once, but corrected itself and steadily hit targets afterwards.
- g1 hit targets at a similar pace, but erroneously shot at orientations supplementary to the expected one for the last three targets.
- g2 regularly hit targets too, but at a slower pace. However, their mistakes are more constructive. Each sub-team misused the protractor on their first shot, but corrected themselves. The two other missed shots do not have so much pedagogical importance: one is due to an unintentional shot, and the other to an imprecision in the reading of the protractor.
- d4 followed the same pattern as g2, i.e. each sub-team misused the protractor. However, one sub-team took three shots to correct itself. The other gave up after three missed shots, and switched roles again.
- g3 covered all possible mistakes before hitting the first target, but learned to correct them. However, when the sub-team switched roles, the observers randomly gave a measure that happened to be correct. This was catastrophic for the rest of the activity, because they were driven on a misconception, and missed the next targets 12 times. The last hit was done because a member of the controller sub-team came and explained how to place the protractor at the very end of the activity.
- d2 never drew the correct line, but eventually managed to hit targets, using a trial and error adjustment for the second target.
- d3 lacked precision in angle measurements, having more trials than the other groups. They were the only group to use up their 20 shots after hitting 5 targets in about 20 minutes. They were then given the unlimited ammunition, which further intensified the trial and error.
- g4 switched their sub-teams 4 times. The first controller sub-team was not positioning the protractor correctly but managed to hit targets with trial and error adjustments. The second controller sub-team forgot to select the right satellite, and failed many times, until they switched roles again. Once again, the first sub-team overcame their misuse of the protractor with adjustments. On the third switch, the two sub-teams exchanged their practice, and hit one target each without mistakes.

This large variety of errors reflects the complexity of using a real protractor, compared, for example, to a virtual protractor in which a simple click selects its graduation. In a sense, this is another example of *creative appropriation*, because the designer did not think of all of the complexity of the activity. From a usability point of view, it sounds like a negative point to create issues, but in a learning context, failures constitute opportunities for learning.

(Q_l) How do Pupils Collaborate in SpaceJunk?

SpaceJunk is very rich in terms of collaboration strategies. Theoretically, the whole group should cooperate to fulfill a common goal. In practice, we did not observe full collaboration within the group: the observer and controller sub-teams competed on various levels. This competition can be explained by the physical separation, and the fact that the individuals who made up each sub-team were never separated. Moreover, we observed a stronger separation between sub-teams of opposite genders (e.g. two boys working with two girls).

This competition could be more or less intense, and more or less productive. d1 showed an intense, productive competition. One of the pupils zealously enforced the rules of the script: she shouted at another pupil from the other team for looking on the other side of the separation via the mirror of the Metroscope. She also refused to tell more than what she wrote on the ammunition. However, she congratulated the other sub-team when the shot succeeded, explained what to do when they switched roles, and reminded the other sub-team to help each other if need be.

g4 demonstrated an intensely unproductive competition. All groups were very engaged and often insulted each other on small mistakes, but g4 gratuitously insulted each other. The climax of this unfruitful cooperation was when the observer sub-team purposely gave wrong information to have the controllers miss the shot.

g3 was an example of non-intensive collaboration. After the first switch, one of the observers gave a random measure, and it worked by chance. This led the observers to a sequence of failures that they were not able to understand. At some point, the controllers took the lead, and told the observers what to write, because they thought that they needed to write the measurement on the ammunition for the system to know where to shoot. By their passivity, the observers gave the lead to the controllers, who did not have the information to initiate the laser cannon process. This resulted in the lowest score of the study.

Within a given sub-team, the roles were clear. Since each member had to do one shot, one pupil in the observer sub-team manipulated the protractor and the pen and one pupil in the controller sub-team manipulated the protractor and the control card, while the other member of each sub-team assisted, e.g. passed an object, or supervised what the other was doing. Only one pair of pupils manipulated together, e.g. both moving the control card. In two other pairs, one pupil did all the work for the sub-team, while the other passively waited.

(Q_u) How do Pupils Interact with each Element of the Paper Interface?

One of the designed features of SpaceJunk was the ability to keep track of the progress of the group throughout the task. The line the observers drew on the satellite view sheet would give a hint as to whether they positioned the protractor correctly, and the ammunition was supposed to keep track of the communication between the observers and the controllers.

With the satellite sheet it was easy for the experimenters to spot misunderstandings about the origin of the shot (see Figure 6.16). Moreover, the satellite sheet leaves a persistent trace as to how far in the activity the group went: during the activity, red circles indicated which satellites had been hit so far. After the activity, the lines between the lasers and the targeted satellite made this information persist.

In contrast, the ammunition did not consistently help keep traces throughout this study. All the pupils wrote the angle on the ammunition, because they thought, at least at the beginning, that the Metroscope would not recognize the angle if it was not written on the ammunition. However, only one group wrote which of the three lasers to use on the ammunition; the other groups used verbal communication instead.

We mentioned that the pupils thought that the ammunitions would not fire if they were not annotated, or even crossed-out. This was not the only stressful characteristic of the ammunitions. Since the pupils could see their stock of ammunitions, they could see this resource being used up. Even when we did not tell them that there were only 20 shots available, all but one group tried to use the ammunition sparingly, often blaming the other sub-team for the loss of one.

The unintentional shots were another stressful characteristic of the ammunitions. Since the pieces of paper were small, it was easy to forget them and involuntarily show their tag to the Metroscope, which would interpret this as a fire command. One pupil even played with this tension to tease a pupil of the other sub-team, by hiding the ammunition at the last moment.

We also compared the characteristics of the sheets to the ones found in Chapter 4. In Quads, we saw that the sheet was a referential stable in space. Here, we also observed that the pupils did not rotate the satellite view, even though it would have facilitated the drawing and measuring. Only one group rotated it by more than 30°. This can be explained by the reaction of one pupil after his sub-teammate moved the satellite view: he yelled at him, stating that the measure had to be redone. However his sub-teammate was able to show him that the red spot followed the satellite targeted on the sheet, and that the measure was still valid.

(Q_u) How do Pupils Manipulate Objects in a Polar Referential?

From the logs of the marker positions, it is easy to compute the position of the control card relative to the position of the laser sheet. Figure 6.17 shows the result of such transformation. One can already guess the circular form of the movements. Some groups kept their control card off of the laser sheet, others kept the control card on the laser sheet, and some groups did not have a clear preference.

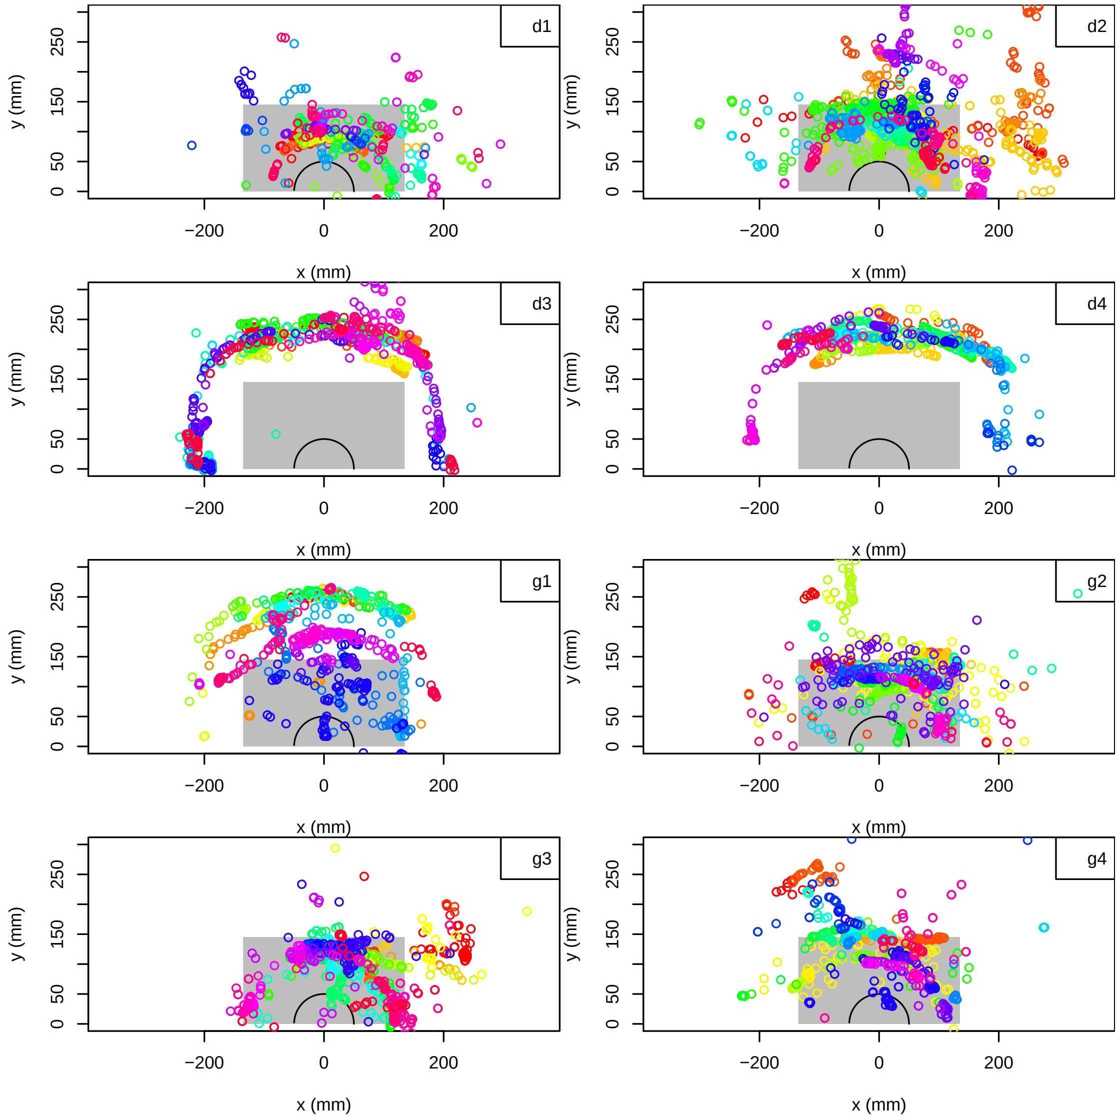

Figure 6.17: Positions of the laser control relative to the laser sheet (represented as a grey rectangle) for each group. The arc gives the scale of the protractor. The colours of the rainbow are used as an indication of time (from red to purple).

We used algorithms borrowed from eye tracking research to extract movements from these positions (Salvucci and Goldberg, 2000). The process is as follows. First, we extract *fixations*, i.e. sequence of coordinates where the control card stays in the same place for a sensibly long enough time. In our case, we kept fixations of at least three seconds within a radius of one millimetre.

Saccades are those movements which happen between these fixations. In eye tracking research, the saccades are very fast (less than 20 milliseconds) and are considered linear. The saccades we consider in our case are much longer, which allows us to record the positions of the consecutive steps on their trajectory. We kept the saccades with at least 10 positions recorded, and whose start- and end-points were at least 5 centimeters apart. An example is shown in Figure 6.18a.

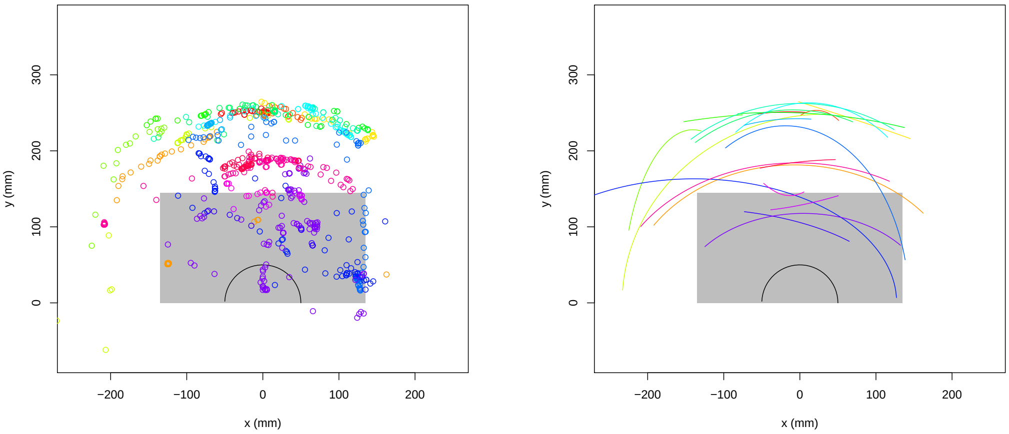

(a) The steps of several saccades extracted from the logs of g1.

(b) Example of splines interpolating the saccades of Figure 6.18a

Figure 6.18: Exploitation of the traces of manipulation.

We then interpolate these saccades by finding their best fit to splines using the parabolic form:

$$y^2 = a.x^2 + bx + c$$

This form is convenient, because it can easily be transformed to the equation of an ellipse of radius (r_x, r_y):

$$\frac{y^2}{r_y^2} + \frac{x^2}{r_x^2} = 1$$

with $r_x = \frac{\sqrt{c}}{\sqrt{-a}}$ and $r_y = \sqrt{c}$. In this transformation, we assume the ellipses to be centred on the laser sheet, and that the focal axes are perpendicular to the axes of the sheet. This means that $a < 0$, $c >= 0$, and $b \approx 0$. This is the case for most splines, and thus a good way to filter out noisy data. In practice, this means that pupils move the control card on an elliptic course around the sheet, which is not surprising.

Figure 6.19 shows the distribution of the radius of the ellipses approximating the movements of the pupils with the control cards relative to the laser sheet. It is clear that while the width of the movement (if it was a half ellipse) is around 40 centimeters (2 × 200 mm) and the height is either around 10 centimeters or around 25 centimeters. This is easily explainable by the fact that the movement was either on the sheet or off the sheet, but not across its borders.

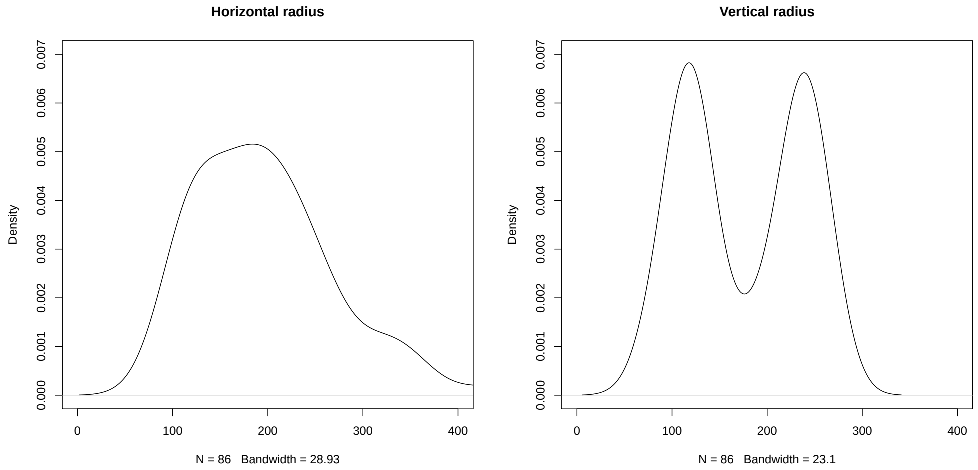

Figure 6.19: Distribution of the radius of the ellipses approximating the movements of the pupils with the control cards relative to the laser sheet.

We further analyse the circularity of the movements, i.e. whether they are close to a perfect circle, or an ellipse, i.e. a “compressed” circle, and in which direction. We define the circularity measure of an ellipse of radius (r_x, r_y) as

$$circ(r_x, r_y) = \frac{r_y}{r_x}$$

Thus, $circ = 1$ for a perfect circle, $circ < 1$ for a vertically compressed ellipse, and $circ > 1$ for a horizontally compressed ellipse.

We found that the circularity varies as a function of the lateral radius r_x ($circ = -0.0017248 \times r_x + 1.28$, $p < 0.001$), which means that when the card is 16 centimetres away from the center of the angle, its trajectory is a perfect circle, and a trajectory with a horizontal radius of 30 centimeter would have a vertical radius of 23 centimetres. Figure 6.20 gives a visualization which helps interpret this result: the border of the sheet influences the trajectories of the control card when it is on the sheet. When the control card is off the sheet, ergonomic reasons can explain the vertical compression: the arm of the pupil has a limited length and the perception of the trajectory is altered by the perspective.

We investigate the influence of the sheet further, by defining three zones, drawn in Figure 6.21: on the sheet (trajectories with a radius smaller than 15 centimeters), on the border (trajectories with a radius between 15 and 20 centimeters), and off the sheet (trajectories with a radius greater than 20 centimeters).

Figure 6.2 shows that this model indeed explains the trajectories better.

The most important result is simply that pupils move the card controlling an angle in the form of a circle or an ellipsoid. It shows that the pupils enact the angle, i.e. they control a dynamic

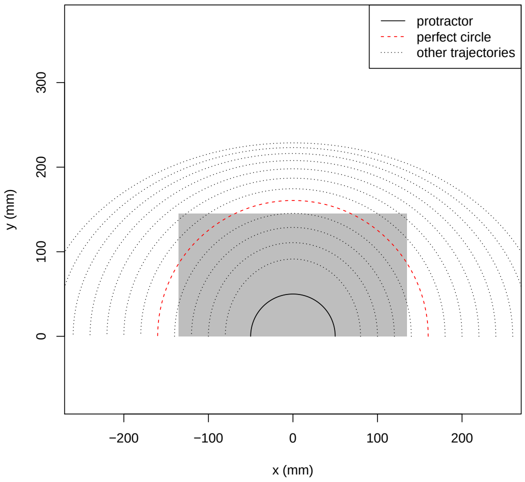

Figure 6.20: A visualization of the model of the trajectories followed by the control card. The protractor and the sheet (in gray) are drawn for reference.

Table 6.2: Statistical models of the influence of the sheet on the trajectories of the control card. A model defining three areas (in, on the border of, or outside of the sheet) explains best the fact that the elliptic trajectories of the controller card are less and less circular as the movements are wide.

Model	Df1	Df2	F	p	
Circularity is random.		85			
Circularity depends on the radius.	1	84	48.4	<0.001	***
Circularity depends on whether the radius is smaller or larger than the sheet.	1	83	171.0	<0.001	***
Circularity depends on whether the radius is smaller than, on the border of, or larger than the sheet.	1	82	8.5	<0.005	**

representation of the angle, which is a constructivist way of learning angles. Indeed, pupils can construct the concept of angle from a concept they already knows: moving in circle.

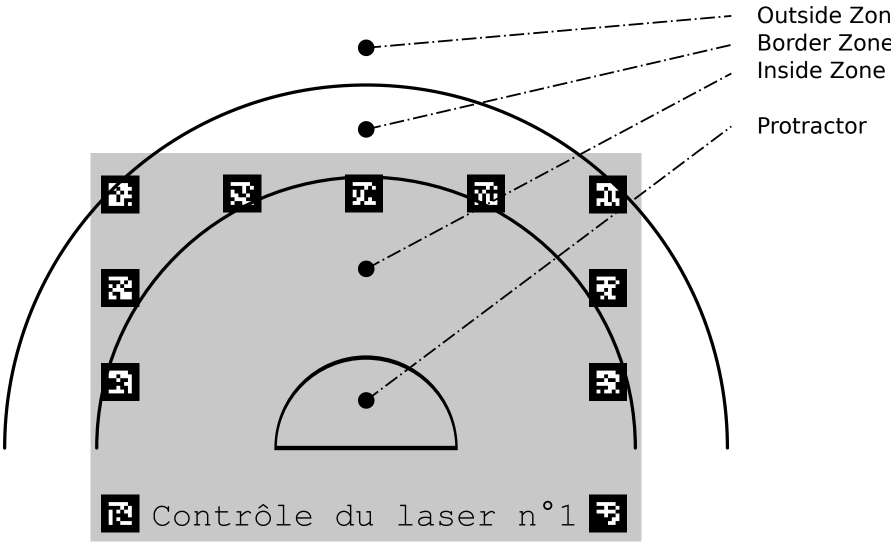

Figure 6.21: Three zones defined to refine the model of the trajectories made by the control card.

6.5 Complementary Study

We decided to improve SpaceJunk by implementing some of the feedback from the informal study, and solving some of the non-constructive issues encountered by the pupils in the formal study. Figure 6.22 shows the improved interface.

6.5.1 Description of the changes

Table 6.3 summarizes the differences between the two versions of the interface.

Table 6.3: Differences between the two versions of SpaceJunk.

First version	Second version
Each shot is described on a different, small piece of paper.	All the shots are written on a single A4 sheet of paper.
Shots are triggered by showing one side of the same small piece of paper.	Shots are triggered by flipping a dedicated, reusable card.
The control of angles are relative to a sheet of paper.	The origin of the angle is fixed on the interaction surface.
The laser is selected by showing the corresponding sheet.	The laser is selected by orienting a card
Instructions are given verbally.	Instructions are written on sheets.

The pupils in the first study complained about the small size of the ammunition, which made the ammunitions cumbersome to annotate and manipulate. In addition to the frustration by the pupils as users, the small size can lead them to stop annotating the ammunitions anymore,

Figure 6.22: An improvement of the interface for the activity shown in Figure 6.1. The laser (1) is at a fixed place on the desktop. It is triggered by flipping a card (2), indicating its readiness (3). The selection of the laser happens by rotating another card (4). The two sub-teams communicate with a full size sheet (5).

which means that there is no trace left of the communications between the observers and the controllers. We thus replaced the ammunitions by an A4 sheet whose sole purpose is to keep track of the shots.

The other function of ammunitions was, of course, to trigger the shot by simply showing them. This was a source of frustration for the pupils of the informal study who did not like the fact that the ammunitions could be used only once. It was also the source of non-productive errors in the formal study, because several shots were unintentional, only triggered because the pupils did not notice that the ammunition was flipped in time to stop it.

Thus, we replaced the tiny pieces of paper with a card, shown in Figure 6.23a, which triggers the shot when it is flipped. This is a better alternative than just showing the tiny piece of paper representing the ammunition, because the process of showing is a duration, in contrast to flipping, which is an intentional instant, just like triggering a shot. To prevent intensive trial and error strategies, the card displays a “loading bar” which takes three seconds to fill up. When the bar is not full, flipping the card has no effect.

We also observed in the studies that when two laser sheets were displayed, two cannons appeared, which caused confusion as to which cannon would be fired. The only purpose of the laser sheets was to select the cannon. Thus, we replaced it by a single card, shown in Figure 6.23b. Orienting this card (e.g. North, South-East or South-West) corresponds to activating one of the 3 lasers. The improvement here is that cards are less fragile and more easily manipulated than sheets. The position of the cannon is fixed in the projection area, which makes for a more stable environment to place the protractor.

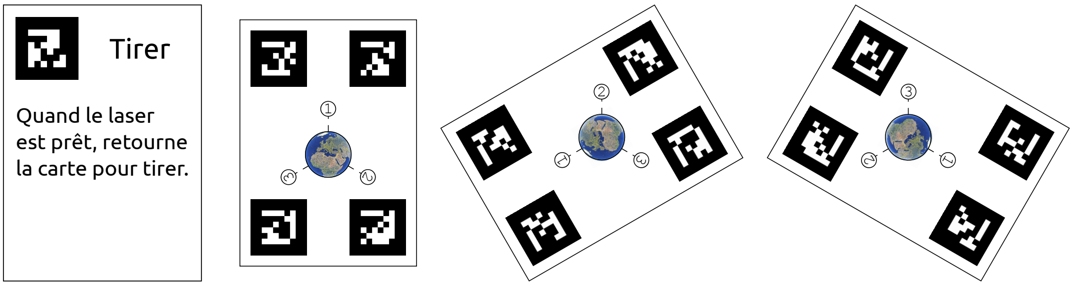

(a) One side of the trigger card. The other side is similar, with a different tag.

(b) The three orientations of the selection card, each corresponding to a different laser.

Figure 6.23: New elements of the interface.

Finally, we tried to replace verbal instructions by step-by-step instructions for each sub-team. These instructions are written on two sheets, one for the observers, and one for the controllers.

Table 6.4 places these changes in the framework.

Table 6.4: Position of the new version of SpaceJunk in the framework

Property	Ephemeral	Persistent	Permanent
Presence			
Position	The pupils move the control card relative to an absolute position in the augmented area to control the inclination of the laser.		
Orientation	The pupils orient the selection card to specify which laser to use.		
Side	The controllers flip the trigger card to shoot.		
Folds			
Edges			
Ink		The pupils write the expected inclination on the ammunition sheet.	

6.5.2 Procedure

We ran another workshop as part of a "Take Your Child To Work Day" at our university. The children were ten to thirteen years old and did not know each other. The workshop lasted for two sessions, the first with 22 children, the second with 12 children. The activity lasted for 30 minutes.

Again, we asked the children to write down their feedback on sticky notes: one for what they liked, and one for what they did not like. Due to a technical problem, we could only save the logs for three tetrads in the first session session, and for one tetrad in the second session.

6.5.3 Results

This section aims at answering a general question: (Q_u) *What are the consequences of the changes made to SpaceJunk?*

We can infer about the performance of each group even without the logs of the activity because we kept the annotated satellite views. We assume that each line between a laser and a satellite means that that specific satellite was targeted. In the first sessions, the different groups managed to hit 8, 5, 1, 2, and 8 satellites. In the second session, they hit 15, 8, and 15 satellites. This wide variance can be explained by the fact that the groups were very heterogeneous: some pupils had never used a protractor, others had already mastered it.

Figure 6.24 shows the evolution of the scores for the group whose data could be saved. In this figure, we can see one of the consequences of replacing the ammunitions by a trigger card: this creates the possibility for trial and error, i.e. by shooting repeatedly and successively adjusting the orientation. There were several sequences of trial and error, e.g. when s1g3 was aiming at the fourth target. This example seems to be the only one in the groups shown in Figure 6.24 where the trial and error was successful.

This leads us to mention the importance of the trade-off between pedagogical effects and usability while designing an interface for learning. Often, effective ways of learning something are cumbersome. In our case, if the pupils could shoot continuously, the system would be easier to use (the laser would just have to swipe an approximate position), but it would not be interesting at all from a pedagogical point of view. However, we could argue that the trial and error approach allowed by the trigger card is still constructive: Figure 6.24 showed that the pupils stopped using it. The pupils probably realized that it was not the most effective method, because of the three second cool-off period. This may explain, at least in part, why a lot pupils complained about the trigger card on the sticky notes.

Figure 6.24 also shows that the time before the first shot seems longer than in the previous version of SpaceJunk (around five minutes in the previous version, around ten minutes in this version). This is probably due to the replacement of the demonstration on how to use SpaceJunk by the instruction sheet, which, due to all of the text to read, takes up the beginning of the designated activity time. The instruction sheet was not a popular change: the lack of explanation was one of the two main critiques from the sticky notes. Moreover, the experimenter had to intervene more often to help the groups than in the other sessions. However, from a class management point of view, this solution seems viable: the interventions were not

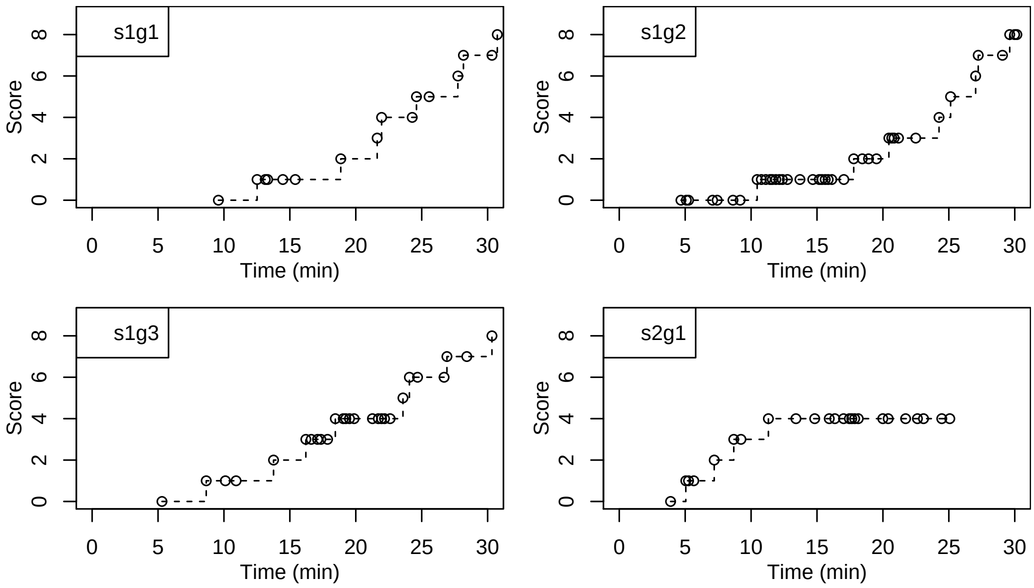

Figure 6.24: Evolution of the scores. Each marker is a shot. The line shows the score (i.e. the number of hit targets) as a function of time (in minutes). Markers in the middle of a segment correspond to missed shots; markers on a corner correspond to hits.

needed at the same time for each group, or not needed at all for some of the groups, which made it possible to distribute the teacher's interventions among the class.

We did not observe a major impact on the way of communicating angles between the two subteams, i.e. the change from several small ammunition sheets to one sheet with all shots. Out of the eight groups, one wrote precise, verbose text descriptions of the angle, one alternated between writing the angle measure plus right or left, four groups wrote only the angle measures, and two groups wrote nothing. There is, however, a change in the persistence of the written information: since the angles were described one after another, it is possible to recover the sequence, as opposed to independent tiny ammunitions of paper that are shuffled. One can recognize the writing, and guess how the pupils alternated their roles. For example, the verbose angle description was always written by the same person, which means that she did not try the role of controller.

Finally, we were able to confirm the hypothesis that the laser sheet was responsible for the two values around which the vertical radius of the elliptic movements of the control card oscillated. In this version, there is no sheet disturbing the movements of the control card, and Figure 6.25 shows that the distribution of vertical radius of ellipses done by the saccades is more uniform. We also observed that the circularity of the movements decreases as a function of the horizontal radius ($circ = -0.0063147 \times r_x + 2.3567383$, $p < 0.001$): when the card is 21 centimetres away from the center of the angle, its trajectory is a perfect circle. If the card is closer, its trajectory is horizontally compressed; if it is further, its trajectory is vertically compressed. This makes us think that the laser sheet from the previous version of SpaceJunk was simply an annoyance. Indeed, pupils did not use the fact that it could be rotated or moved to control the referential in which the angle was defined, but it disturbed their movements.

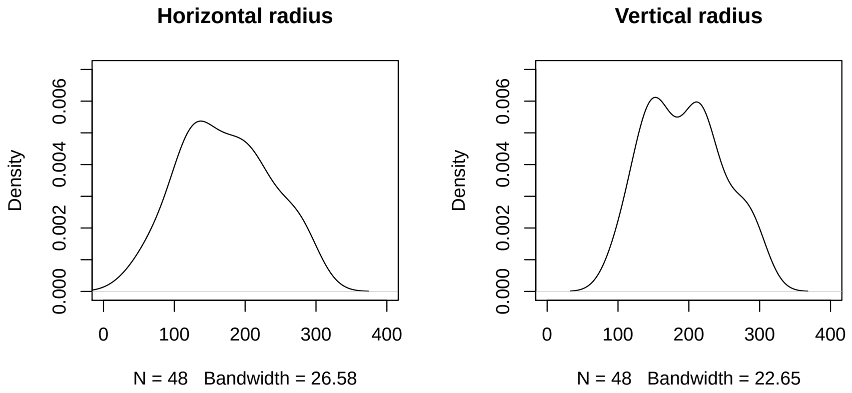

Figure 6.25: Distribution of the radius of the ellipses approximating the movements of the pupils with the control cards relative to the laser sheet.

6.6 Conclusion

SpaceJunk went beyond the proof of concept of Quads and the small exercise represented by Angoli. SpaceJunk is a more ambitious activity allowing many degrees of freedom in its use. Describing angles, the core of SpaceJunk, happened to be the source of a lot of creativity. We showed how the simple task of communicating information could make use of paper to unravel a myriad of possibilities. Paper is a good compromise between communication modes that are too informal, like speech, ones that are too rigid, like those allowed by a mouse and keyboard.

Paper also supported collaboration in SpaceJunk in different ways. Each element of the paper interface defined the roles of the participants. Since the elements are on a horizontal surface, they are easily shared and manipulated. The paper can retain traces of the collaboration, which encourages thinking on common ground.

The creativity fostered by paper also concerned errors. In a learning context, it can be productive to guide pupils towards errors when they can be isolated and addressed by the teacher, so that they are internalized for later. SpaceJunk is rich in this domain.

The errors allowed by SpaceJunk does not mean that it was a frustrating, traumatizing experience for children. They were all enthusiastic about the game aspect, which was made possible by incorporating the capabilities of computers into paper.

Yet again, the playful aspect of SpaceJunk did not hurt the pedagogical goals. Pupils were still able to use a protractor as expected, and even use it in a complex scenario rather than over-simplified exercises.

In summary, SpaceJunk is a good example of taking the best of different worlds: game and learning, paper and computers.

The related studies also allowed us to pinpoint some features of paper interfaces. First, the elements have physical interactions that influence their use. For example, the laser sheet influenced the trajectories of the control card. Since this influence did not bring anything to the interaction (the borders of the sheets are just an artefact, without any relation to the activity), it was just an annoyance. It could be used as a guide, however, to make sure that the control card does not go between 180 and 360 degrees, for example.

Second, paper can be used in many ways, and some are better than others. For example, large sheets such as the laser sheets are not adapted to select a simple value; a lasting action such as showing is not adapted for instant-based interactions such as triggering a shot. However, the advantage of a paper interface is that it is relatively easy to modify: it is as simple as printing a new version of a document.

7 Exploring Symmetries

Our next step in exploring paper interfaces for education consists of increasing the participation of teachers in the design of an activity rather than simply having them be involved in the brainstorming. Indeed, teachers are obviously experienced in pedagogical design in particular, and especially in tailoring activities for their class.

The best way to involve teachers is to separate the technical design from the pedagogical design. This is what we did with Kaleidoscope, which we present in this chapter. Further, we report two activities designed by teachers based on Kaleidoscope, and how they were used in classrooms. Additionally, we used Kaleidoscope to design a controlled study, which aimed at investigating the differences between a tangible way of making answers persistent on paper, and the more traditional way of annotating. Finally, we shortly present a more creative activity, aimed at exploring how a paper interface can help in the sharing of a computing device among the class.

7.1 Description of the Demonstration Activity

Kaleidoscope is not exactly an activity. It is actually a generic component that can be instantiated into a variety of exercises. This component is a single sheet which reflects its content according to one or more symmetry axes. We first present this sheet and a sample activity that we developed to demonstrate Kaleidoscope to the teachers. Figure 7.1 provides a link to a video demonstration of the activity.

7.1.1 Generic Sheet

Kaleidoscope sheets contain a rectangular area that serves as input, e.g. $20cm \times 20cm$ in the centre of the page. One or more axes are defined within this input area. An axis splits the input area into a smaller input area and an area showing the reflection of the remaining input area, e.g. with a vertical axis, the left side of the original area will be shown reflected on the right. More axes further split the input area and show more reflections. For example,

Figure 7.1: A QR Code of the url http://short.epfl.ch/kaleidoscope which points to a video demonstration of Kaleidoscope.

a vertical and a horizontal axis will cause the reflection of the top-left corner onto the three other corners, where the bottom right corner will be the result of two reflections. See Figure 7.2 for an illustration. Many markers on top and bottom of the sheet allow for a very precise localisation by the system. This is especially important to check whether an input and its reflection overlap.

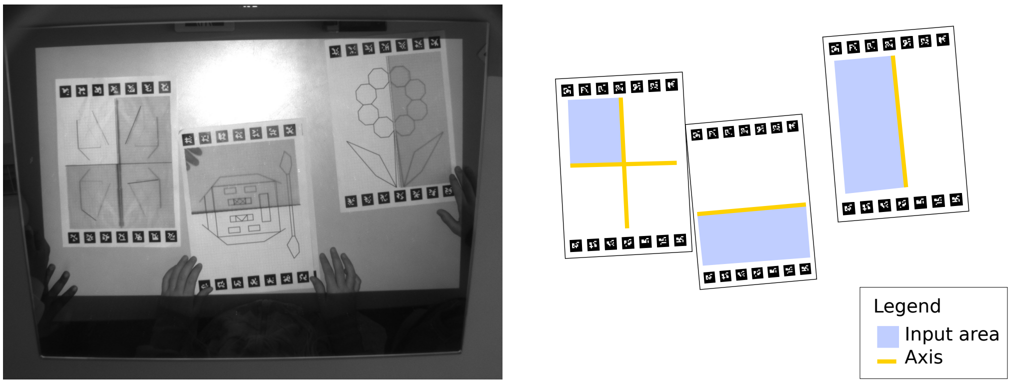

Figure 7.2: A view from the top of the Kaleidoscope sheets, with a schema indicating the axes reflecting the input area. The greyer sides of the sheets are actually the projection of the other side of the axes. Note that anything in the input area is reprojected, such as fingers, not only the ink on the paper.

Technically, any number of axes can be defined, but there are practical limitations. More than four axes make the input area too small to be used. Also, two axes are hard to conceive if they are not perpendicular; the number of axes are usually a power of two; and four axes usually intersect at a single point. Figure 7.3 shows practical instances of axes.

7.1.2 Example Sheets

In order to show the teachers how these sheets could be used, we prepared a sample activity (see Figure 7.4) with six exercises:

1. The first exercise consisted of drawing directly under the Metroscope to illustrate the symmetric movements of the pen. Half of a castle is pre-printed on one side of a vertical

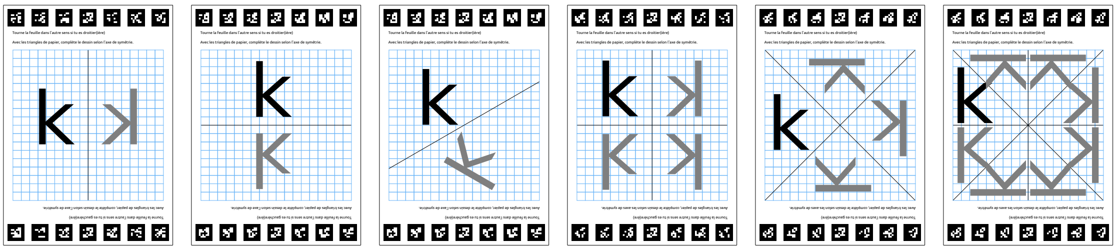

Figure 7.3: From left to right: a vertical axis, a horizontal axis, an oblique axis, plus-axes, cross-axes, and star-axes. The grey k's illustrate how the Metroscope would augment the black k's on each sheet.

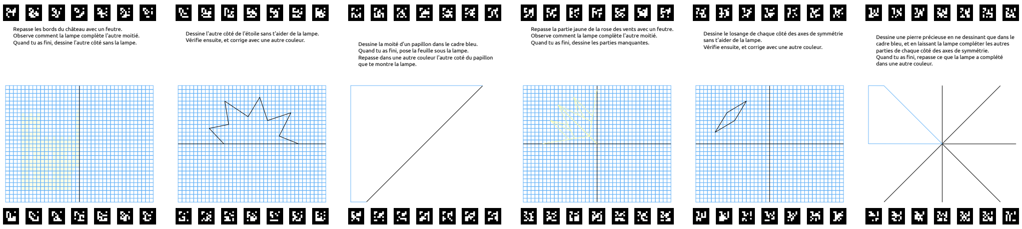

Figure 7.4: The sheets of the sample activity.

symmetry axis. When the pupil draws over this pre-printed pattern, the other side of the castle appears, projected in real time by the Metroscope. The castle is pre-printed in a light yellow color, which does not contrast enough for the Metroscope to notice this annotation and reproject it. However, when the pupil draws over this pre-print with a blue or black pen, the ink is detected. This way, the castle is reflected only when it is drawn with the ink of the pupil.

2. The second exercise consisted of drawing the symmetric image of a pre-printed figure without help from the machine and bringing the result to the Metroscope in order to check the answer. If the drawing was correct, the drawing of the pupil would overlap with the projection of the symmetric image of the pre-printed pattern.

3. The third exercise consisted of drawing one side of a free symmetric figure, to be completed by the Metroscope. Thus, the pupil was expected to think of a symmetric shape, and draw one half relative to the given axis.

4. The fourth exercise was similar to the first, but with two perpendicular axes.

5. The fifth exercise was similar to the second, but with two perpendicular axes.

6. The sixth exercise was similar to the third, but with two pairs of perpendicular axes.

7.1.3 Design Objectives

Kaleidoscope illustrates how a paper interface can leave the entire pedagogical design of an Augmented Reality activity to teachers. For each kind of axis, a teacher can be given a blank page containing only fiducial markers, the axes, and the limits of the input area, but without instruction text.

A teacher can then customize this blank Kaleidoscope sheet without any help from the developer. It is enough to add free text (instructions), a drawing, or a grid to the paper. The teacher does not even need a computer: everything can be hand-drawn and created with regular tools (ruler, transparent papers, etc.) and placed in a copy machine.

7.1.4 Position in the Framework

Table 7.1 shows which parts of the framework are covered by Kaleidoscope. Kaleidoscope covers the ink properties of paper. It uses the ephemeral information produced by the apparition of the ink on paper to progressively display the symmetric image. The ink can obviously persist on paper. The permanent aspect of the ink corresponds to the drawings and instructions made by the teachers on the sheet, when they design the exercise.

Table 7.1: Position of Kaleidoscope in the Framework

Property	Ephemeral	Persistent	Permanent
Presence			
Position	Anything placed in the input area is reflected according to the symmetry axis of the sheet.		
Orientation	The pupils can rotate the sheet to change the side of the input area.		
Side			
Folds			
Edges			
Ink	The pupils see the evolution of the symmetric image of their drawing.	The pupils leave a trace on the sheet, which serves as an answer.	The teacher can specify instructions and pre-printed patterns in the input area, such as a grid or a figure.

7.2 First Teacher Design

A teacher from School 4 used the generic sheet to create a customized activity, which she ran in her class. This section reports on the design process, the resulting activity, and its use.

7.2.1 Objectives

Kaleidoscope sheets are the basis components of activities allowing pupils to explore symmetry. The main point of the sheets is to be easily customizable, so that a teacher can create such activity. Thus, we study the following points:

- (Q_t) How can a developer transfer the pedagogical design of a paper interface to the teachers? In other words, does Kaleidoscope succeed in shifting the design of the content of the activity to the teacher, leaving only the technical part to the developer?
- (Q_t) How do teachers use their own design? If the shift toward the teacher as the designer of a pedagogical activity is successful, it is interesting to observe the behaviour that emerges from the teachers.
- (Q_u) How do pupils use teacher designed activities? Consequently, the way teachers use their own activity may change the way they interact with the pupils.

7.2.2 Procedure

We invited the teacher to our laboratory to demonstrate the sample activity. The original idea was to illustrate the possibilities of Kaleidoscope, and co-design an activity to be used in her class. To this end, we wanted to produce a draft of the paper interface, beginning with a blank Kaleidoscope sheet, and later drawing figures and writing the instructions developed by the teacher. The description of the actual design session follows in the next section. However, as we will report, the co-design evolved into a design fully controlled by the teachers.

We ran the designed activity in the classroom of the teacher. It lasted two periods of 45 minutes, and involved 20 fifth graders (10-11 year olds). The Metroscope was placed on a table in the back of the classroom. The teacher managed the class and the usage of the Metroscope by herself; the experimenter only observed, eventually answering questions unrelated to the Metroscope.

7.2.3 Results

(Q_t) How to Transfer the Pedagogical Design of a Paper Interface to the Teachers?

The teacher did not consider the Metroscope an interactive tool but rather more of a verification tool: “For me, [the Metroscope] allows them to control what they have done... But I can’t see what they could do directly under [the Metroscope].” The teacher insisted that the pupils do not work under the Metroscope, and only use it to check their answers. We repeatedly asked whether any exercises should be done interactively, but the teacher preferred that the pupils stay at their desks to work.

Her goal was to use the Metroscope as a way to make errors productive. If the pupils were wrong, Kaleidoscope made the errors stand out. The pupils then realize what is wrong, why, and how they can fix the errors. They can try again at their desk, and check their answer again. If they are still wrong, then, and only then, can they work directly under the Metroscope.

The teacher proposed different ways of using Kaleidoscope that were not considered in the sample activity. For example, instead of drawing only one side of a symmetrical figure, the pupils could draw the whole figure, and check that the outline overlapped with the symmetric image of the outline. Moreover, multiple axes of symmetry could be checked by rotating a Kaleidoscope sheet containing only one axis. However, the teacher noted that the advantage of Kaleidoscope sheets compared to the mirrors normally used is that two axes can be superposed at the same time.

Her most interesting addition to the sample activity concerned the use of shapes printed or drawn on pieces of paper independent from the Kaleidoscope sheet (see Figure 7.5). This allowed the pupils to check the presence of axes of symmetry of geometrical shapes (e.g. a square has four symmetry axes) by aligning the symmetry axes on the Kaleidoscope sheet and on the shape. If the outline of the shape and its symmetric projection matched, the axis was correct. Moreover, the use of sheets maintains the goals of the curriculum more closely: geometrical shapes correspond more to a mathematics course than free drawing does.

The exercise of finding the axes of symmetry of shapes led her to define various levels of difficulty within the activity. For instance, a slower group could be given shapes that have only one axis of symmetry, while the rest of the class works on shapes with multiple axes.

Adapting the difficulty is a major aspect for integrating Kaleidoscope in her class. Indeed, she admitted to having heterogeneous pupils. In addition to the difficultly adaption previously mentioned involving shapes with different numbers of axes of symmetry, she also thought of an extra, more challenging exercise for the highest performing pupils. This exercise involved the reconstruction of a shape by reflecting other parts of a shape over different symmetric areas.

The teacher also planned further integration of Kaleidoscope in her class. First, she decided to introduce symmetries a few days before the introduction of Kaleidoscope so that the exercise would be a complement. This influenced the design of the activity, because it did not need to contain deep explanations or examples. Second, since the pupils were used to working in workshop activities, she planned to integrate the Metroscope inside the classroom with each pupil working independently and using the Metroscope as needed. She intended to support this workshop style by allowing only two pupils at the Metroscope at the same time.

The session steadily evolved towards an appropriation of the design by the teacher. Our original goal was to draft activities, i.e. writing down instructions and drawing figures on “blank sheets”, i.e. sheets with only markers and an axis for symmetric augmentation. This way, we could obtain a live prototype to be used as a common ground between the teacher and the developer. However, the teacher preferred keeping sheets without instruction, as it "allowed her to explain things". Then she asked whether it was essential to use these particular sheets, or if the developer could just send her the canvas for her to prepare exercises herself. Eventually, she agreed to do all the work herself, so the developer simply sent her the digital version of the "blank sheets". Since she is a very computer enthusiastic (and knowledgeable) teacher, she produced the digital versions of the augmented sheets herself.

She sent us her activity in one week after a few emails were exchanged. She used some of the patterns given in the official curriculum. The only change we made to her design was making the colour of the grid light yellow instead of blue. We thought that a lighter colour would prevent the Metroscope from detecting pre-prints and only show the symmetric images that the pupils drew.

The teacher not only designed the pieces of paper (shown in Figure 7.6) but also the course of the activity, in three steps:

1. The pupils would cut out outlines of figures and classify them according to the number of axes of symmetry (no axis, one axis, or multiple axes). They would use the Metroscope to check their classification by manipulating the outlines on a blank sheet; if the axis was correct and matched the one on the sheet, the outline would cover itself (see Figure 7.5).
2. The pupils would be given one side of a symmetric figure (a flower or a house reflected in water, see Figure 7.6, left.) and be asked to draw the other side. They could check their answer by placing the sheet under the Metroscope to see if the symmetric projection matched their drawing.
3. The pupils would complete a figure whose strokes were scattered on each side of the axes (see Figure 7.6, right.). This was the challenging activity aimed at higher performing pupils.

This design process illustrated several strengths of paper interfaces:

- Paper can be carried: this allows the activity to be designed to be performed at the pupil's desk, while involving the control of a single, shared device in the classroom.
- The persistence of paper allows pupils to keep a trace of their previous errors, making the correction process easier.
- The augmentation comes from a computer, which is not limited by physical constraints. This can be contrasted with a mirror, for example, which can not easily emulate multiple axes of symmetry.
- Various levels of difficulty can be defined and implemented by simply distributing certain sheets to less or more advanced pupils.

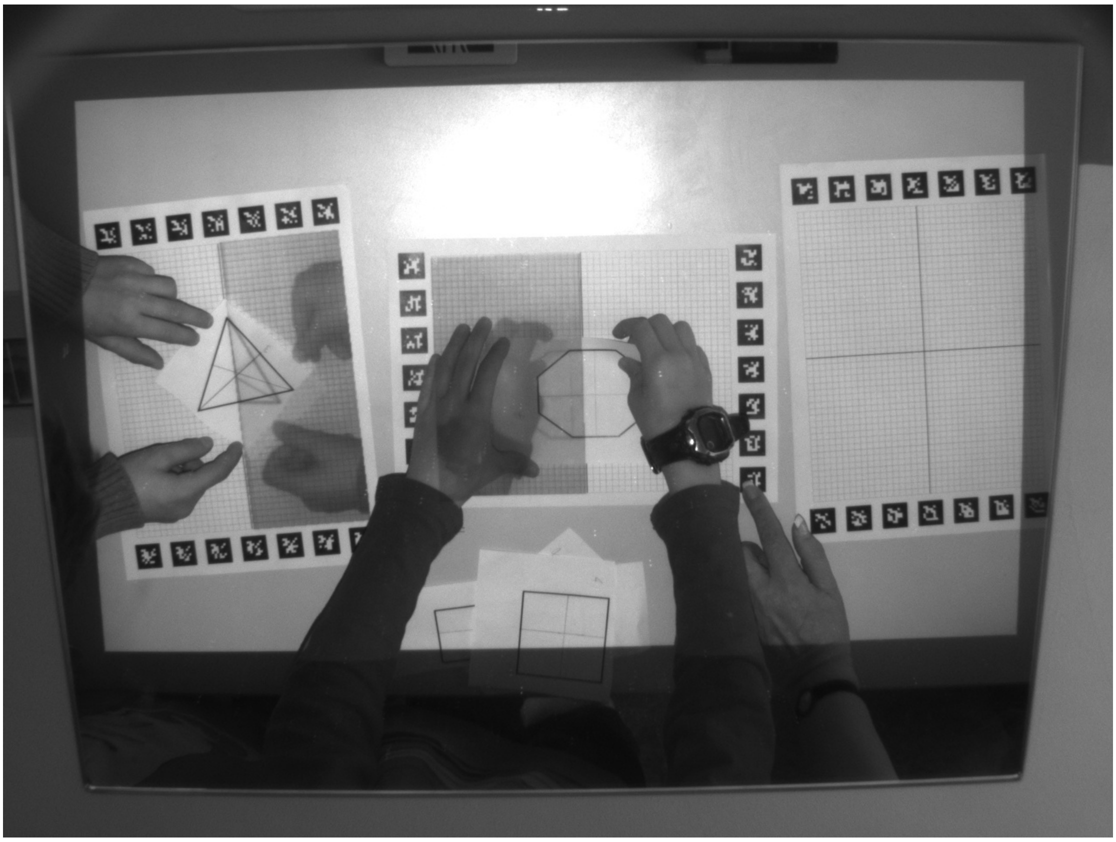

Figure 7.5: An overview of pupils manipulating outlines placed on paper cut-outs. The pupils align the symmetry axis they find on the outlines with the symmetry axis of the symmetry sheets. If the projection of the outline matches their outline, their axis is correct.

- The interface can easily move from a physical form to a digital form. The paper form of the interface allows for grounding of a design session and prototyping. The digital version allows for easy modification and transmission (in our case, by email).

(Q_t) How do Teachers Use Their Own Design?

The teacher started the activity by explaining the three exercises at once. She did so in front of the class, while the Metroscope was in the back of the room, i.e. without demonstration. However, she showed the sheets she was talking about (see Figure 7.7).

She gave the instructions twice, emphasizing the fact that the Metroscope should only be used if it is really needed. The teacher indicated which exercises where harder, and told the pupils to do only one for each of the three steps. The others could be done if time remained upon completion of the first exercise. Finally, the activity was to be done with a pen and a ruler, and the teacher emphasized the need for precision, which the grid could help provide. The teacher did not enforce a particular order in the three steps.

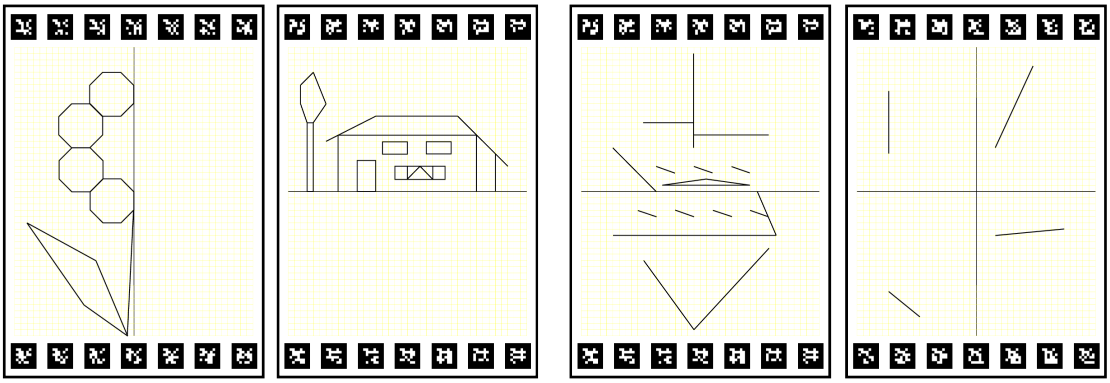

Figure 7.6: Figures to complete by reflecting the outlines across the axes. The Metroscope projects the symmetric projections of the pre-printed figure and the pupils' answer.

Figure 7.7: The teacher explaining the exercise to the class.

The teacher also made two rules to ensure a successful use of the Metroscope. First, not more than four pupils could be using it or queued at a time. In the meantime, if the pupils could not do the other exercises, they could complete the drill exercises about addition and subtraction. Second, the easy exercises were only allowed to be checked at the Metroscope once, because the machine showed the answer directly.

It took nine minutes to give the instructions and five minutes to distribute the various sheets of the activity. The teacher could check that everyone had each sheet by showing the sheets one after another, and giving copies to the pupils who did not have one. The sheets were easily identifiable from afar given their content.

When the first pupil came to check her answer, the teacher gave her a demonstration of how to use the Metroscope. This also happened to be the first time the teacher was using the Metroscope. She also gave a demonstration to the second pupil, and validated the work of these first two pupils. Indeed, the Kaleidoscope simply shows the reflection, and does not give feedback about the correctness. The teacher gave this feedback in the form of sentences like “perfect”. Eight minutes into the activity, she left the Metroscope to take care of the rest of the class, coming back regularly to check that the pupils did not need assistance.

The teacher used the Metroscope to illustrate her explanations with live examples. On multiple occasions, she went past the queue, bringing a pupil from his/her desk to show what was wrong or unclear. For instance, she could dynamically show what, on a figure, goes up/down or left/right depending on the axis of symmetry. She could also give a concrete explanation of why precision is important in geometry: a lot of pupils drew the symmetric axes of the paper shapes approximately, and they could see that aligning them with the symmetry axes gives improper results because of the lack of precision.

The teacher appropriated the use of the Kaleidoscope sheets herself. There was a sheet with a vertical axis, another one with an horizontal axis, and a last one with a pair of perpendicular axes. We thought that the former ones would be used to check the shapes classified as having one axis, and the latter one would be used to check the shapes with two axes of symmetry. Instead, the use of the sheets was more flexible: a single axis could check multiple axes by simply rotating the shape, and a pair of axes could be used as a single axis if the shape was not placed at the intersection of the axes (see Figure 7.8).

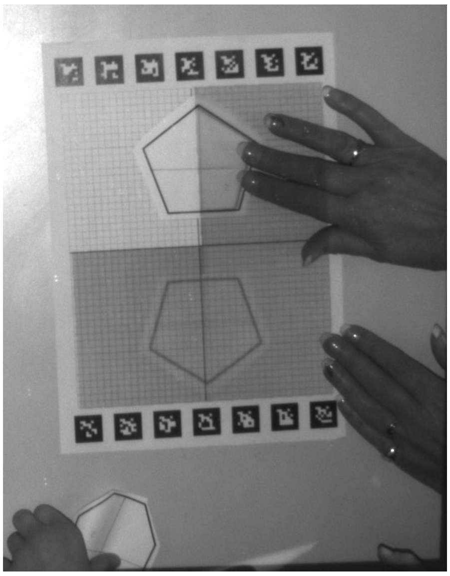

Figure 7.8: The teacher using the +axis to show a vertical symmetry axis on a pentagon.

At the end of the activity, it took five minutes to tidy up the classroom. The sheets were placed in a cardboard folder where the pupils place their work of the week. The teacher took five additional minutes to collect feedback from the pupils: they found the Metroscope to be both enjoyable and helpful, and agreed with the teachers about the difficulty levels of the various sheets.

The overall comment from the teacher about her experience was that it was rather exceptional to have the pupils focus on the same topic for two periods.

(Q_u) **How do Pupils Use Teacher Designed Activities?**

The first pupil to try Kaleidoscope did so four minutes into the activity. She repeated the action showed by the teacher, i.e. rotating the shape on the Kaleidoscope sheet to check whether the reflection could be aligned with the outline of the shape. She took more care adjusting the shape, however. The second pupil not only imitated the actions of the teachers, but also her comments, saying the same words that the teacher used to validate an answer. After the teacher left, she was called by a pupil at the Metroscope only to confirm that he made a mistake.

The pupils did follow the rules delegated by the teacher in order to fairly share the Metroscope: they only checked one stack of shapes at once, and only needed to be reminded once that not more than four pupils could be using or waiting for the Metroscope. The queue was, however, constantly full.

The pupils successfully shared the augmentation area. They did not use it for tasks that could be done at their desk, i.e. correcting the mistakes revealed by the Metroscope. They used the augmentation area in parallel when possible, i.e. when they were using different sheets. The unused sheets were simply placed to the side of the augmentation area.

The pupils waiting in the queue to use the Metroscope were paying attention and giving feedback as to what happened in the augmented area. This resulted in natural peer corrections, where the teacher was only the arbiter for disagreements in the verdicts. The paper interface favoured this in two ways. First, the horizontality of the display allowed the other children to watch the interaction. Second, it was easy to point at and manipulate the various interface elements in order to illustrate explanations.

This visibility of paper had a downside for the learning aspect of the activity, however: the teacher left the solutions on her desk unattended, and the nearby pupils peaked at them several times. These quick glances simply gave them a hint, however, and they still had to precisely reproduce the solution, which was the point of the exercise. Another case of "cheating" occurred when a pupil used the projection of the Metroscope to subtly mark the corners of the figures to draw, but she was scolded by other pupils who saw her do it.

The fact that the sheets were not bound, and that the teacher indicated the estimation of difficulty, allowed pupils to organize their work. For example, one pupil directly started on the hardest exercise, because he likes challenges, as the teacher explained.

During the break, we saw a nice example of engagement turned into exploration by creative appropriation: the pupils brought various objects, such as their snacks, to the Metroscope, as shown in Figure 7.9.

The major usability default was a consequence of the only change we made to the teacher's design: we changed the grid from blue to light yellow, so that it would not interfere with the reflection. This turned out to be a bad decision: in addition to the Metroscope, the yellow

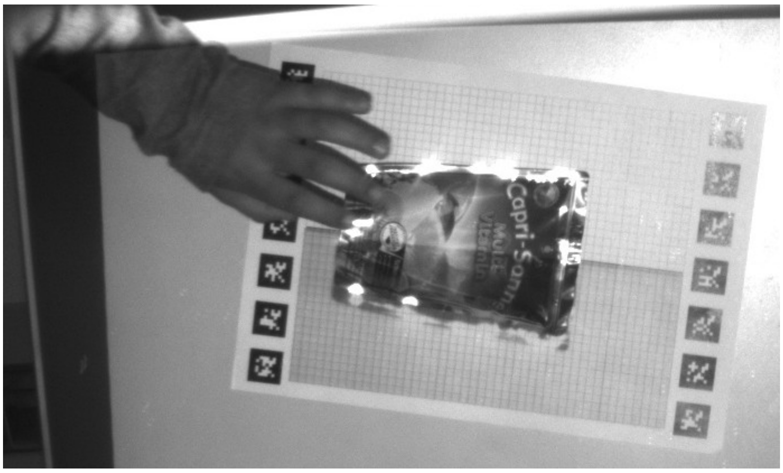

Figure 7.9: A pupil checking the symmetry axis of his lunch during the break.

colour was too light for the pupils to see, and they had trouble identifying the pre-printed figures. Otherwise, the pupils easily overcame several instances of technical glitches that occurred when the same sheets were shown at the same time, causing the reflection to be random. They verbally passed on the solution of this problem, which consisted of removing one of the sheets from the projection area.

7.3 Second Teacher Design

We repeated the study described in the previous section with two teachers from School 3. We asked them to design a pedagogical activity, and run it, with minimal intervention from us. The results are very different, as can be expected with customizable components. The differences are striking both in the resulting design and in its usage, as we will report in this section.

7.3.1 Objectives

The objectives of this study are the same as the previous ones:

- (Q_t) How to transfer the pedagogical design of a paper interface to the teachers? In other words, does Kaleidoscope succeed in shifting the design of the content of the activity to the teacher, leaving only the technical part to the developer?
- (Q_t) How do teachers use their own design? If the shift toward the teacher as designer of a pedagogical activity is successful, it is interesting to observe the behaviour that emerges from the teachers.
- (Q_u) How do pupils use teacher designed activities? Consequently, the way teachers use their own activity may change the way they interact with the pupils.

7.3.2 Procedure

For practical reasons, we met the teachers for a design session in their school rather than our laboratory. The consequence of this is that we showed them a video demonstrating the sample activity rather than a live demonstration. However, the objective was the same: to illustrate the possibilities of Kaleidoscope, and have the teacher design an activity to be used in his/her class.

We ran the designed activity in the classrooms of the teachers. The first teacher had a fourth graders class. He used Kaleidoscope during two periods of 45 minutes, with half of the class in each (the other half was in a knitting class). Each half consisted of 9 pupils.

The second teacher had a mixed class of third and fifth graders. While one group was using the Metroscope, the other one was autonomously performing a drawing activity. The group of third graders consisted of 9 pupils, and the group of fifth graders consisted of 10. Each group used the Metroscope for 45 minutes.

In both cases, the experimenter introduced the Metroscope at the beginning of each session, but only observed afterwards, leaving the teachers to manage the class and the Metroscope alone. The Metroscope was placed on a table in the back of the classroom (see Figure 7.10).

Figure 7.10: The lamp in the back of the classroom.

7.3.3 Results

(Q_t) **How to Transfer the Pedagogical Design of a Paper Interface to the Teachers?**

The design session had three periods. For 10 minutes, the teachers discovered Kaleidoscope by watching the video of the demonstration activity, and then by asking questions about how it works. For the next 30 minutes, the teachers brainstormed about the different possibilities in terms of using Kaleidoscope. In the last 50 minutes, they specified the exact activities for the third, fourth, and fifth grade classes, i.e. planning the exercises, dictating instructions and deciding which figures to pre-print on which sheet. We strived not to intervene in the design process, asking only for clarifications about the activity from the two teachers, e.g. whether the axis they were talking about was horizontal or vertical.

The video could replace the live demonstration to illustrate the features of Kaleidoscope: the teachers noticed that the pen was also reflected along with the pen ink. They further inquired about this, asking what would happen if someone drew outside of the input area (it is simply ignored). They also asked whether the grids had to be preprinted (they did not), because its use is more adapted to the older pupils. Another important point for them was whether pencils could be used, as opposed to the permanent markers used in the demonstration. Indeed, pencils allow for greater precision, and cleaner erasure. They also wanted to know whether the Metroscope could indicate what was wrong and what was correct about the drawings, but Kaleidoscope is designed to show only the expected solution along with the given solution, leaving the user to infer and decide about the correctness himself/herself. We justified this by the fact that automatic correction would be harder to implement, and consequently less reliable.

The teachers were very enthusiastic about Kaleidoscope ("With this... Can you imagine? They'll understand in a sec!"). They repeatedly expressed the advantages they saw, which we report in the next three paragraphs.

The general advantage of Kaleidoscope is that it allows users to instantly see the effect of symmetry on their drawings. For example, with a vertical axis, the pupils can see that a pen moving up and down on one side does the same on the other side of the axis, but that drawing closer to the axis makes the reflection come closer too. More particularly, when the pen touches the symmetry axis, the reflection also touches the axis. The teacher expected the pupils to be able to construct the notion of symmetry more easily with such a tool.

The teachers also compared Kaleidoscope with two existing alternatives. First, it replaces demonstrations done at the blackboard, with imprecise tools and difficult manipulation. Second, Kaleidoscope replaces using a mirror placed on the symmetry line to discover symmetries. With the mirror, it is difficult to draw, while Kaleidoscope allows for precise drawing. This lack of precision also hurts self corrections: pupils often quickly check that the reflection is what they expect, and move on, even if it is incorrect.

Comparing Kaleidoscope to other tools, the teachers noted self correction as another advantage of the exercise. Transparent paper is often used to provide corrections similar to those which Kaleidoscope proposes. The problem is that the interaction is modal: if it is not correct, the pupils have to remove the transparent sheet to correct their answer, and the lack of precision can be similar to the one related to the use of a mirror. With Kaleidoscope, the expected answer is projected directly on the answer given by the pupil, which can be corrected right away. Additionally, a time limit on the feedback could be implemented with Kaleidoscope, in order to let the pupil think rather than simply copying the solution.

The teacher noticed these advantages throughout the design session, however, Kaleidoscope inspired them from the start. During the video demonstration, they had already mentioned the possibility of creating a mandala[1], and letting a pupil draw to observe what happens. These are the two main exercises in the activity resulting from the whole session, which we present hereafter.

The teacher of the fourth graders led the first session. The pedagogical scenario he had in mind for the pupils was the following (see Figure 7.11):

1. Draw something freely on a symmetric sheet under the Metroscope and observe what happens.
2. Return to your desk and write down your observations. Share them with the class.
3. Draw a geometric shape in one quadrant of a symmetric sheet with two perpendicular axes. Add details to build a mandala that will be printed, and that you will later color.
4. Draw the other side of the butterfly at your desk, and check your answer under the Metroscope.
5. Draw the reflection of the castle in a lake at your desk and check your answer under the Metroscope.

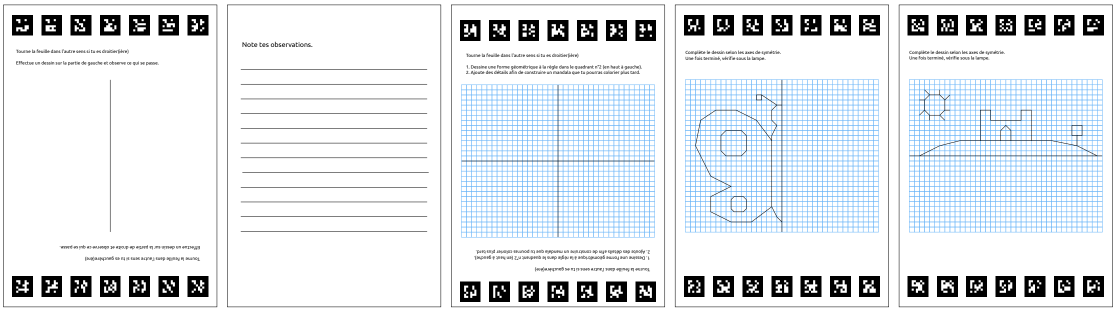

Tourne la feuille dans l'autre sens si tu es droitier(ière)

Effectue un dessin sur la partie de gauche et observe ce qui se passe.

Tourne la feuille dans l'autre sens si tu es gaucher(ère)

Effectue un dessin sur la partie de droite et observe ce qui se passe.

Note tes observations.

Tourne la feuille dans l'autre sens si tu es droitier(ière)

1. Dessine une forme géométrique à la règle dans le quadrant n°2 (en haut à gauche).
2. Ajoute des détails afin de construire un mandala que tu pourras colorier plus tard.

Tourne la feuille dans l'autre sens si tu es gaucher(ère)

1. Dessine une forme géométrique à la règle dans le quadrant n°2 (en haut à gauche).
2. Ajoute des détails afin de construire un mandala que tu pourras colorier plus tard.

Complète le dessin selon les axes de symétrie.
Une fois terminé, vérifie sous la lampe.

Complète le dessin selon les axes de symétrie.
Une fois terminé, vérifie sous la lampe.

Figure 7.11: The sheets of the booklet for the fourth graders.

---

[1] A mandala is a circular pattern often consisting of the repetition of one or more simpler patterns.

The activity designed for the third graders is derived from the first three steps of the activity for the fourth graders (see Figure 7.12):

1. Draw something freely on a symmetric sheet under the Metroscope and observe what happens.
2. Return to your desk and write down your observations. Share them with the class.
3. Draw one half of a house. Observe what happens. Return to your desk and complete the house.


Tourne la feuille dans l'autre sens si tu es droitier(ière)

Effectue un dessin sur la partie de gauche et observe ce qui se passe.

Tourne la feuille dans l'autre sens si tu es gauchère(ère)

Effectue un dessin sur la partie de droite et observe ce qui se passe.

Note tes observations.

Tourne la feuille dans l'autre sens si tu es droitier(ière)

1. Dessine sur la partie de gauche la moitié d'une maison avec deux fenêtres et une porte.
2. Observe bien ce qui se passe.
3. Quitte la lampe et reproduis la partie droite de la maison en symétrie.

Tourne la feuille dans l'autre sens si tu es gauchère(ère)

1. Dessine sur la partie de gauche la moitié d'une maison avec deux fenêtres et une porte.
2. Observe bien ce qui se passe.
3. Quitte la lampe et reproduis la partie droite de la maison en symétrie.

Figure 7.12: The sheets of the booklet for the third graders.

Finally, the activity designed for the fifth graders is similar to the one for the third graders, but the figures are geometrical shapes, to be drawn with a ruler, according to two axes of symmetry:

1. Draw a geometrical shape with a ruler in the first quadrant. Reproduce the shapes indicated by the Metroscope.
2. Return to your desk and write down your observations. Share them with the class.
3. Complete the geometrical shape. Check it with the Metroscope.

The teachers found the ability to have different levels of difficulty, which covered a progression of pupils, interesting: the third graders discover symmetry, the fourth grader exercise it, and the fifth graders consolidate their knowledge.

Two additional features were requested. The first concerned the mandala: the drawn quadrant should be scanned from the paper, reflected along the axes, and printed (see Figure 7.13). The second one consisted of an animation showing the construction of the symmetric image of a quadrilateral: red lines perpendicular to the symmetry axis grow until they link an edge of the quadrilateral and its symmetric counterpart. This was supposed to be used as an introductory demonstration for the third and fifth graders.

These activities (K2) are rather different from the one created in the previous study (K1). Table 7.2 summarises their differences. First, K2 is designed for half classes, and it targets

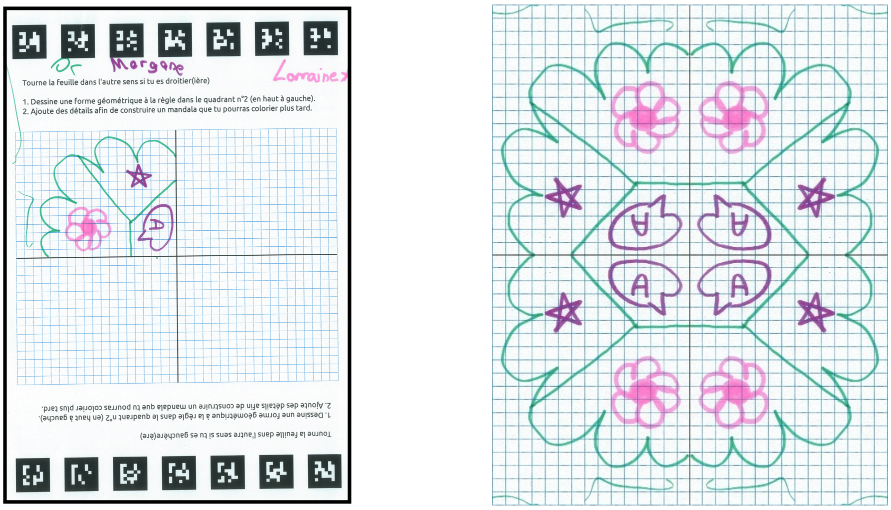

Figure 7.13: Left: the input for the generation of a mandala, i.e. a drawing in a quadrant of a Kaleidoscope sheet. Right: the output of the drawing after extraction with a scanner, which could be sent back to the teacher to distribute to the class.

small groups (triads) with whom there is close interaction. In contrast, K1 targets a whole class of autonomous pupils, assisted by the teacher on demand. Given this, it is surprising that K1 relies on oral instructions, and K2 contains oral instructions, even if the teacher is always with them. It is also interesting to note that K2 relies on live feedback, even if the teacher is present to explain and comment on the work of the pupils. Naturally, since K1 is designed for a whole class, the use of the Metroscope had to be kept short, resulting in simple feedback. Another difference between K1 and K2 concerns the origin of the figures: K1 contains official figures from the curriculum (e.g. the flower pattern), and K2 contains original patterns designed by the teacher. Finally, each teacher invented a different functionality that was not demonstrated with the sample activity: cutting out and placing shapes on blank sheets in K1 and printing mandalas and showing the dynamic construction of symmetric images in K2.

Table 7.2: Differences between resulting designs.

Characteristics	First design (K1)	Second Design (K2)
Group size	whole class together	small groups one after another
Medium of instructions	oral	written text
Role of the Metroscope	correction	live feedback
Source of content	reused official content from the curriculum	original content
Innovations	cut-out shapes	mandala printing and animation of the construction

(Q_t) How do Teachers Use Their Own Design?

Fourth Grade The teacher of the fourth graders organized the half classes into three triads. While one triad was using the Metroscope, the two other groups were given an unrelated activity that they could perform autonomously.

Each session began with an introduction of the Metroscope by the experimenter, and an introduction to the activity. As in the first study, the teacher showed the sheets that were going to be used to illustrate his explanations.

The rest of the session was organized as follows: 10 minutes were reserved for the first sheet, 10 minutes for the sharing of the observations with the class, and the remaining time (10 to 15 minutes) was used for the mandala sheet. During the session the teacher decided to skip the last two sheets because of time constraints.

The teacher remained at the Metroscope the whole time. He had the instructions read aloud by the pupils, and managed the activity closely. Most importantly, he reminded the pupils to pay attention to what was happening so that they could write down their observations.

The teacher also precisely specified whose turn it was, and hinted at what the pupils could draw when they were taking too long. He also controlled the start and the end of the interventions of each group by distributing and removing the symmetry sheets.

The teacher also prevented misuses of Kaleidoscope. For example, he did not allow pupils to reflect the images of their hands, and ensured that the head of the drawing child did not block and hide the augmentations for the other pupils.

The Metroscope caused several interruptions of the pupils' attention on the task. For example, pupils noticed that the Metroscope was generating a lot of heat, or that the projection moved when too many tags were hidden. The teacher tried to ignore these distractions, but was forced to give a short explanation to move the focus back to the task. For example, he explained that the small lag between the movement of the pen and the projection (of a few hundreds of milliseconds) was just a technical detail, and not as important as what was to be observed.

The teacher spent most of his time taking care of the interactions at the Metroscope, but obviously did not neglect the rest of the class. For example, he checked the status of the autonomous activity while the triads were switching, or called out the remaining time before the sharing of the observations to the class.

The sharing of the observations occurred in two steps. First, the teacher asked the pupils what they observed, and tried to elicit the expected observations, such as "A double image? Can you explain? Exactly the same thing? Like in a mirror? I'll write down 'mirror' ". The second step consisted of a little exercise on the blackboard, where the teacher asked pupils what would happen under the Metroscope if the pupils drew a given shape. This shows how Kaleidoscope can easily integrate into the regular class workflow.

After the observation sharing, the teacher did not give further details on the following activity. He continued managing the activity closely, however, e.g. specifying the level of detail to put into the drawing, or stating that not drawing straight lines was acceptable.

Finally, the teacher asked the first half of the class to keep the details of the activity a secret from the other half of the class. This was not respected, because the pupils obviously knew more about the Metroscope during the second session than during the first.

Differences Between the Half Classes This study is particularly interesting because the same activity was repeated. According to the teacher, this allowed for a more polished session. We review the differences corresponding to this polishing in the three following paragraphs.

First, the teacher improved the efficiency of the interaction by performing trivial tasks: he skipped the first instruction, taking on the task of rotating the sheet depending on whether the pupil was left- or right-handed himself. In this more polished version, the teacher asked the pupil whether he/she was right- or left-handed and then rotated the sheet accordingly. He also held the sheet in its place to make the writing more stable. The teacher also directly showed an example stroke on the mandala, taking the pen at the beginning when the pupils were unsure about what they should draw.

Second, the teacher liked the idea that emerged in the last group of the first session. One pupil had written her name, which was interesting for the teacher to comment on when sharing the observations, because not only was the order of the letters reversed, but the letters were also mirrored. Thus, in the second session, the teacher had every pupil write their name in the first exercise.

Third, when the teacher gathered the observations of the pupils after the first part, the questions to the class were more focused, and the guidance of the teacher was more focused.

The teacher also proposed a problem, which was left on the whiteboard during the second phase with the Metroscope: the teacher drew two perpendicular axes with an arrow in one of the quadrants. The question was how the symmetric image of the arrow would look if it was reflected in the fourth quadrant. Several pupils proposed three different solutions. The teacher took a few minutes from the break to discuss the solution with the class.

Third Grade The teacher of the third graders organized the activity a little differently. He also organized the half classes in triads, but each pupil performed the activity alone while the others watched.

The first part of the activity lasted for 10 minutes. Once all the groups had performed it, the teacher called all of them around the Metroscope to have them share their observations. Indeed, since the rest of the class was working on an unrelated topic, the teacher could not use the blackboard.

For the third graders, this activity was just a first encounter with symmetry. The teacher did not expect too much from the pupils, and spent relatively less time sharing observations (about two minutes). Another reason for such a quick common phase was that the teacher intensely strived to complete the whole activity. This involved strict time management. The teacher did manage to go through the whole activity, but did not use the animation for the construction of symmetries that he requested, and used three minutes of the break time to do the corrections.

The teacher encouraged the pupils to draw towards the axes of symmetry, so that he could ask them what happened (the pen stroke and its symmetric image meet). He then asked the pupils to go write their answer. Like the first teacher, the second teacher constantly reminded the pupils to observe what was happening.

The second phase lasted 15 minutes, and the pupils went to the Metroscope to check their drawings in an unstructured order, i.e. whoever finished first rather than organized in triads. Finally, the corrections took up the last 3 minutes of the class.

The teacher kept control of the activity by distributing the sheets. There were three types of sheets, kept in three separate stacks. Two of them where on a table behind him, inaccessible by the pupils. The last one was located directly on the knees of the teacher. This is an interesting example of how paper has an established use: this situation was not the most convenient for the teacher, but it is still acceptable.

This session was more challenging in terms of class management. After a while, the teacher left the third graders to draw by themselves, so that he could assist some fifth graders that needed help. He also invited a fifth grader who finished the exercises early to come to the Metroscope so that he could watch.

Fifth Grade The second session of the second teacher cannot be compared to his first session in terms of content, because the activity was different. However, we noticed a few changes in the organization. The teacher tried to manage time more efficiently. He streamlined the activity so that one pupil started an activity when the previous one was using the Metroscope. He sometimes corrected the drawings himself so that two pupils could get feedback simultaneously. He skipped the exchange of observations. Even then, he had to turn down some requests for help from the other half of the class.

Follow-up Both teachers seemed frustrated with the short time spent with the Metroscope. The experimenter proposed leaving the Metroscope at the school and the teachers accepted with enthusiasm. They asked to keep the unused sheets for the fourth graders, and requested blank symmetric sheets for other activities. Theses sheets were actually used as an exercise without the Metroscope, as a consolidation of the exploration made with the Metroscope. One of the teachers tried to use the Metroscope, but had to give up because of a technical problem. After the device was fixed, the teachers tried an interactive activity.

The activity for the fourth graders used the butterfly sheet shown in Figure 7.11. The pupils used the lamp to check their answer. The teacher ran the activity autonomously, without the researchers. Judging from the logs captured automatically, they switched on the Metroscope half an hour before the pupils started using it. Judging from the snapshots of the interaction area, the teacher remained next to the Metroscope, pointing at the mistakes made by the pupils. After five minutes, the teacher started correcting the edges of the butterfly pattern with a pen. Ten minutes later, the lamp was switched off again.

The second teacher turned on the Metroscope a few minutes before using it. He gave the butterfly pattern to the third graders, and the castle pattern to the fifth graders (see Figure 7.11). The pupils were autonomous in their use of the lamp. The activity lasted 25 minutes.

After this independent activity without the researchers, the teachers sent us a simple email telling us that “the pupils loved this kind of autonomy and control”.

(Q_u) How do Pupils Use Teacher Designed Activities?

The pupils enjoyed the activity. They were not very free to explore by themselves, because both of teachers closely controlled the procedure. Regarding control, one episode is interesting to note. The teacher of the fourth graders had paused the activity by removing the sheet from underneath the Metroscope to explain something to the whole class. However, a pupil had such a strong desire to draw something that he took the sheet right out of the teacher’s hand, which interrupted his explanation. This is an illustration of the fact that paper objects can act as a central token in the group.

We saw another kind of ownership materialized by paper, but on the content rather than on the tangibility. When drawing the shared mandala, one pupil asked another pupil in his group whether he could draw a line on the mandala. The third pupil asked him why he would ask that, since it was a ‘shared’ mandala after all. The first pupil answered that he wanted to draw the line in a place that would overlap something already drawn by the second pupil.

One usability issue was related to the projection from above. Most pupils drew with their face very close to the sheet, in order to see what they were doing more closely. However, this would block the projection of the symmetric image, which made the teacher intervene, telling the pupils not to occlude the projection with their heads. It is also interesting to note that the very first observation gathered concerned the time lag between what the pupils drew and what was projected. This means that they were, indeed, observing the progressive construction of the symmetric image.

The pupils seemed frustrated by the guidance of the teachers. They often asked whether they could do some variation of the intended activity, such as colouring, or drawing a different pattern than the expected one.

The Metroscope was also very helpful for the pupils. A clear example is given by the exclamations declaring the drawings to be correct, as soon as they were placed under the Metroscope. Also, even though it was not expected from the third graders, they spontaneously started to correct each other. This was very beneficial, as highlighted by the case when one of the pupils saw a mistake that the teacher missed.

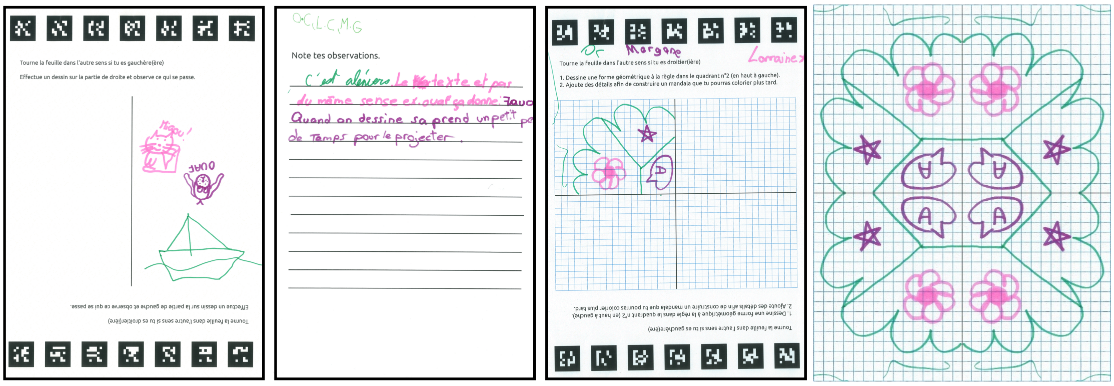

Tourne la feuille dans l'autre sens si tu es gauchère(ère)

Effectue un dessin sur la partie de droite et observe ce qui se passe.

Effectue un dessin sur la partie de gauche et observe ce qui se passe.

Tourne la feuille dans l'autre sens si tu es droitier(ière)

O.C.L.C.M.G

Note tes observations.

C'est aleniers Le ~~te~~texte et pas du même sense ex.ouaf ça donne: Jauo
Quand on dessine sa prend un petit pe
de Temps pour le projecter.

Or Morgane Lorraine

Tourne la feuille dans l'autre sens si tu es droitier(ière)

1. Dessine une forme géométrique à la règle dans le quadrant n°2 (en haut à gauche).
2. Ajoute des détails afin de construire un mandala que tu pourras colorier plus tard.

1. Dessine une forme géométrique à la règle dans le quadrant n°2 (en haut à gauche).
2. Ajoute des détails afin de construire un mandala que tu pourras colorier plus tard.

Tourne la feuille dans l'autre sens si tu es gauchère(ère)

Figure 7.14: From left to right: the output of one group of fourth graders for steps 1, 2, and 3, and the resulting mandala rendered from a scan of the third sheet.

7.4 Controlled Study

The activity designed by the first teacher had two novel features. First, the activity used the Metroscope in a way that we did not think about. We had proposed the use of either enough devices for a whole class working simultaneously or a single device for a group to work with it while the others were working on something else. The teacher from School 4 proposed an activity where the whole class made use of one device.

Second, the first activity made use of the fact that Kaleidoscope sheets reflected anything placed on them, not only ink. The pupils could check that shapes had a given axis of symmetry by aligning them on the axis of a Kaleidoscope sheet.

We thus designed another activity based on Kaleidoscope to investigate these two points. We created a series of exercise sheets where pupils had to place the symmetric image of pre-printed right triangles. They could do this in two ways: by drawing the symmetric triangle, or by taping a pre-cut triangle shape on the page. The pupils could use a Metroscope shared among the class in order to verify their answer, by checking that the symmetric projection of the pre-printed triangles overlapped the placed triangle.

7.4.1 Objectives

This activity allows us to investigate the two points discussed previously:

- (Q_l) Which persistence is better for learning? In this study, the pupils gave their answer by using two persistent interactions: either drawing, or taping. The end results seem to be the same, but the process can have implications on the learning outcomes for several reasons: whether the triangle is manipulated or not, if one mode is slower, etc.
- (Q_u) How do pupils use the persistent characteristic of paper? Since the interaction (drawing or taping) is persistent, the pupils can carry their own paper interface from their desk to the Metroscope and back. Of course, this takes time, but this can have positive effects, because it encourages the pupils to prepare a good answer before moving, or even to think about their answers during the trip to the Metroscope.

7.4.2 Procedure

We ran the study with a fifth grade class of School 2, i.e. 10-11 year olds (n=22). In the (45 minute) period before the study, the teacher introduced symmetry. The study spanned over two sessions, each using a 45 minute period in which 10 minutes were spent for instructions, and 30 minutes for the exercise itself. The first session occurred during the last period of the morning, and the second session during the first period of the afternoon.

In each session, the pupils were distributed a booklet of four exercise sheets, one for each type of axis, shown in Figure 7.15. The first page had a vertical axis; the second page a horizontal axis; the third page a vertical and a horizontal axis (hereafter referred to as +axes); and the fourth page perpendicular, diagonal axes (hereafter referred to as ×axes). Each sheet had 5 pre-printed triangles, which means that the pupils had to place 5 triangles for the first two pages, and 15 triangles for the last two pages (because of the increased number of quadrants in which the triangles were reflected). In addition to the markers required to locate the Kaleidoscope sheets, each booklet was attributed an identification number, which allowed us to retrace the history of the Metroscope use by the pupil associated with a given identifier.

There were two types of these four-page booklets: one for the taping mode, and one for the drawing mode. The former requires the pupils to tape pre-cut, red, triangle paper-pieces (Figure 7.16) as a reflection of pre-printed triangles. The latter requires the pupils to draw the reflection of the pre-printed triangles with a pencil and a ruler. Pupils in the taping mode were given a band of reinforcement rings[2], so that the time spent taping was minimal: the reinforcement rings are pre-cut, and can be peeled off easily. They were easy to remove and replace if pupils made mistakes, and maintained the cut-out triangles in their intended location. The pre-printed triangles were filled in the taping mode, and only outlined in the drawing mode, so that they would match the appearance of the answers to give.

---

[2] self-adhesive reinforcing pads usually stuck around folder holes to prevent sheets from being damaged

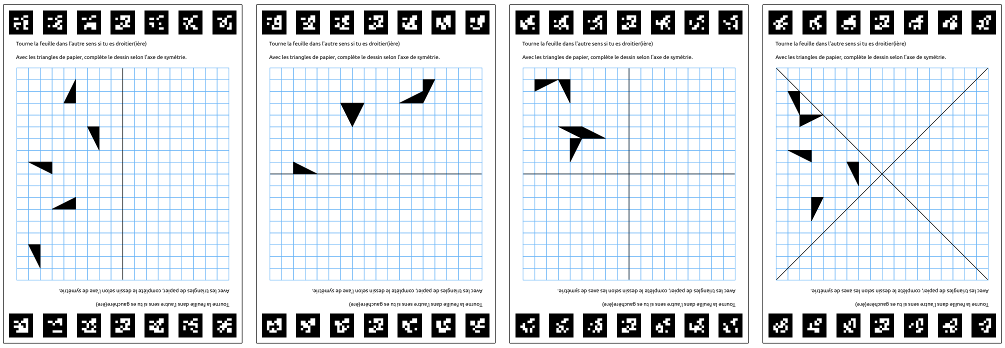

Figure 7.15: Set of symmetry exercises, each with a different axis type. From left to right: vertical axis, horizontal axis, +axes, and ×axes

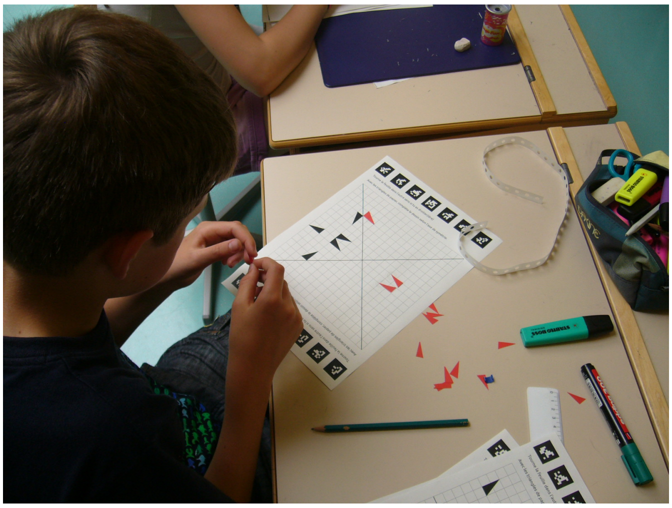

Figure 7.16: A pupil placing triangles by *taping* pre-cut, red, triangle paper-pieces. Note the band of reinforcement rings between the sheet and the pencil case.

The pupils were told to use the Metroscope to check whether they had placed the triangles of a sheet in the correct positions and orientations. Although the Metroscope did not provide explicit information as to whether a worksheet was correct or not, it did provide all of the correct triangle projections on top of the worksheets. In this manner, the pupils could check their answers. The teacher still played her normal role in the classroom, roaming the room to answer questions and provide extra help to the pupils.

Pupils given a *taping* booklet in the morning were given a *drawing* booklet in the afternoon, and vice-versa. This counter-balanced the order effect, and determined the condition: *tape first* versus *draw first*. Pupils were equally dispatched in one of the two conditions, but one pupil was not present in the morning session, and another was not present in the afternoon session. After the afternoon session, the pupils filled out a short questionnaire about the activity. The questionnaire asked the pupils which mode they preferred, and which mode they found easier.

Table 7.3 situates the two persistence modes in the framework. Drawing implies fixing three edges one after another. In contrast, taping requires fixing the position, orientation, and side (which matters for symmetries) of a triangle globally.

Table 7.3: Position of Kaleidoscope in the framework

Property	Ephemeral	Persistent	Permanent
Presence			
Position		The pupils *tape* the triangle in a given position.	
Orientation		The pupils *tape* the triangle in a given orientation.	
Side		The pupils *tape* the triangle on a given side.	
Folds			
Edges			
Ink		The pupils *draw* the three edges and the outline of the triangle.	

7.4.3 Results

Performance

Completion The number of placed triangles per type of axis (see Figure 7.17) shows how far the pupils went at the end of each session. All the pupils completed the vertical axis sheet, and most of them addressed the two following sheets. However, a majority of pupils barely addressed the fourth sheet.

Difficulty of the Sheet Figure 7.18 shows the error ratio per sheet, i.e. the number of incorrectly placed triangles over the number of triangles placed by the pupil. As expected, the difficulty is increasing. Clearly, the vertical axis was trivial, while the horizontal and +axes were slightly harder for the pupils, and the ×axes stood out as being the most difficul. These three classifications are confirmed by pairwise t-tests, yielding p-values below 0.05.

Note however the ceiling effect on the placement errors: when a triangle was placed, it was rarely placed incorrectly, as shown by the fact that the average ratio of errors is about 25%, even for the most difficult sheet.

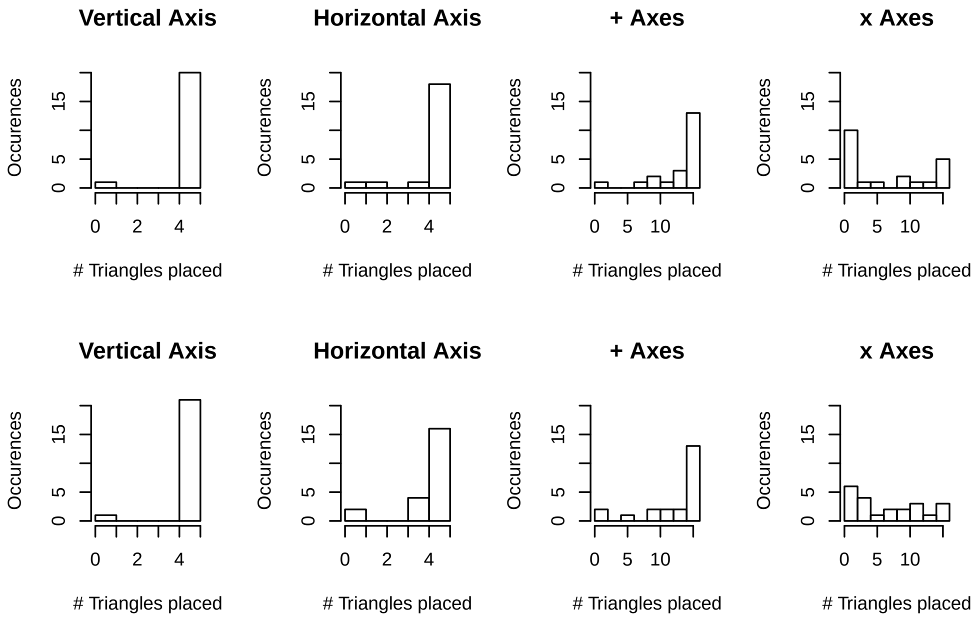

Figure 7.17: Number of placed triangles based on the worksheet axis, for each session. The first row corresponds to the morning session, and the second row to the afternoon. For each type of exercise (vertical, horizontal, perpendicular, and diagonal axes), the histogram shows how many of the triangles were placed. The maximum number of triangles is 5 for the single axis, and 15 for the double axes.

(Q_l) Which Mode of Persistence is Better for Learning?

Performance According to the Mode of Persistence Taping is slower than drawing, but this difference is not statistically significant, mainly because of the ceiling effect: most pupils placed between 25 and 40 triangles in each session. However, 9 pupils drew more triangles than they taped, 4 taped more triangles than they drew, and 4 pupils taped as many triangles as they drew. However, taping and drawing makes no difference at all in terms of correctness.

Types of Errors The errors done by the pupils were different depending on the mode they were using. For example, there were twice as many (six versus three) placement errors in the drawing mode than in the taping mode. A placement error corresponds to a triangle not being the same distance from the axis as its symmetric image. This is probably due to the fact that it is easier to adjust the position of a paper triangle than to entirely redraw a triangle.

The ×axes was unsurprisingly the richest exercise in terms of variety of errors. Two kinds of mistakes happened: either the pupils treated the ×axes like +axes, or they added triangles (see Figure 7.19). The former happened twice as much (six versus three) in the drawing mode as in the taping mode, and the latter happened three times in the drawing mode versus once in the taping mode.

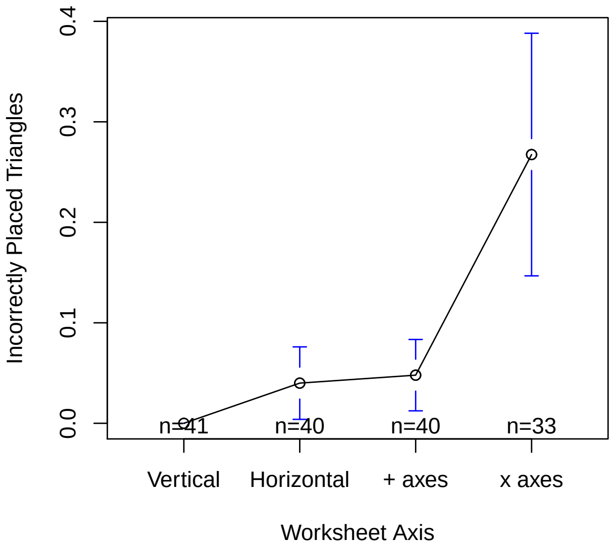

Figure 7.18: Percentage of incorrectly placed triangles based on the worksheet axis.

The last type of mistake, i.e. the triangle having a wrong orientation, e.g. flipped, did not clearly discriminate the two modes.

Order Effect Another way to compare the effect of taping versus drawing consists of observing whether their order has an effect. Figure 7.20 shows the evolution of the global score of each pupil, depending on whether they started with taping or drawing.

The ratio of triangles correctly placed during a session by the pupil improved on average if they started with taping, and then switched to drawing, from 0.60 to 0.69. However, these averages are still lower than the average score of pupils that started with drawing and then switched to taping, which remained almost constant from 0.81 to 0.80. Moreover, the relative improvement, i.e. the ratio of the score of the afternoon over the score of the morning, was more erratic for pupils who started with the taping mode: the standard deviation was 1.25 (with an average of 1.54) versus 0.18 (with an average of 1.00) among the pupils starting with the drawing mode.

The order of the mode has thus an important effect. We explain this by the fact that drawing is a more familiar way to do geometry exercises. In fact, during the introduction to symmetry (carried out by the teacher before the experiment, without the Metroscope), the pupils had just used a pen and paper. We suppose that starting with the drawing mode allowed the pupils to complete the exercises in a way they were familiar to. Thus, the pupils who started with the drawing mode were familiar enough with the notion of symmetry to transpose it to another

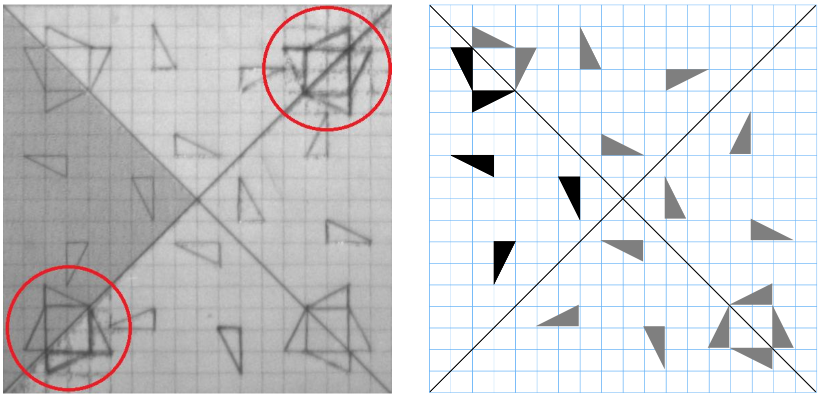

(a) The superfluous triangles are circled in red.

(b) The corresponding exercise, where pre-printed triangles are black, and triangles to place are grey (the triangles are filled for readability).

Figure 7.19: A case of superfluous triangles added by the pupil.

mode afterwards (taping). In contrast, the pupils who started with the taping mode may have been destabilized by the unusual way of working, which distracted them from the content of the exercise. This can explain why the results are so heterogeneous, and generally lower, in the 'tape first' condition.

Perception by the Pupils Figure 7.21a shows the answer to the question: "Which mode did you prefer?" in the questionnaire distributed at the end of the second session, according to their condition. In general, the pupils preferred the drawing condition (13/19). Clearly, the pupils who started with drawing preferred drawing, while the results are mitigated for the pupils who started with taping.

Figure 7.21b shows the answer to the question: "Which mode did you find the most difficult?". This looks similar to Figure 7.21a: the pupils who started off drawing found taping more difficult, and the pupils in the tape first condition are also mitigated. We also measured to ensure that preference and difficulty are not correlated.

Flexibility An anecdote revealed another difference between the writing and taping modes. After arriving at the primary school, the experimenters realized that they had only printed half of the worksheets. Because this augmented paper is just like any other paper, but with fiducial markers. The experimenters simply used the school's photo-copy machine to print the other half of the worksheets. This was only possible because the count of pre-cut triangles was correct. Indeed, cutting the pieces of paper is much more time consuming, and could have hardly happened within the lunch break. In real conditions, this implies that teachers would need to invest a lot of time to prepare an activity based on this kind of pre-cut triangles.

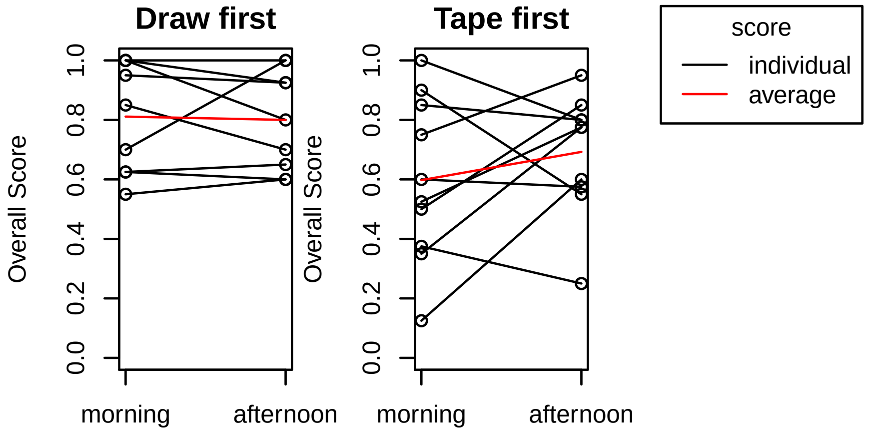

Figure 7.20: Order effects based on 'draw first' or 'tape first' conditions. The left side of each graph represents the morning session and the right side represents the afternoon session, after switching modes.

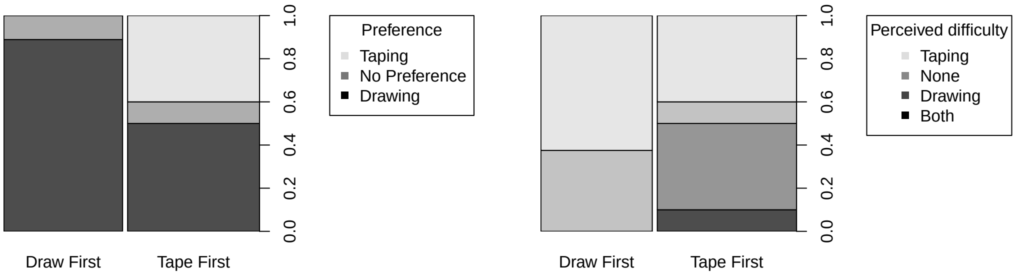

(a) Preference.

(b) Perceived difficulty.

Figure 7.21: Results of the questionnaire.

Conclusion To summarize, this study proves to be rather negative towards the taping mode. The performance, i.e. the number of triangles correctly placed at the end of a session is greater in the drawing mode. At the same time, the mistakes are more varied than in the taping mode, which may be more beneficial for learning, because they correspond to an exploration of the problem space, which the teacher can use as illustrations in her summary. Regarding this aspect, we should also mention the fact that drawing left more traces of previous trials than taping. Pencil cannot be perfectly erased, and the marks it leaves can be exploited to diagnose problems encountered by the pupils.

Taping has been shown to lower the performance of pupils that start with this mode, while it has no influence, on average, on pupils that switched to taping after drawing. This would be interesting if the lower performance was counter-balanced by an increased performance when switching to drawing, but it is not the case. Furthermore, the effect on individuals is a lot less consistent when they are taping than when they are drawing.

We highlighted two other factors that show advantages of drawing over taping. These factors are not directly linked to learning, but definitely have an influence. First, we hinted at the practical difference between drawing and taping as a way to convey answers from the pupil to the teacher. The drawing activity can be produced by simply printing enough copies. The taping activity, in contrast, requires a large investment of time by the teacher to pre-cut the triangles, or unproductive time spent by the pupils to cut out their own set. The other factor to take into account is the preference of the pupil: manipulating paper triangles is not more engaging that drawing triangles with a ruler.

This conclusion is not a definitive judgement on the superiority of annotations over tangibility in general. Rather, it can be the starting point to determine its superiority in this particular exercise involving the reflection over the origin. A first explanation is that drawing is the usual way of working on symmetry, and in geometry in general. Taping tangible shapes thus have an initial disadvantage because of the lack of familiarity that pupils have using tangible items to work on problems. This renders the manipulation awkward, and the resulting distraction, even if weak, can explain, at least in part, the lower performances.

More importantly, the method to solve the problem, and thus the cognitive process is completely different. For drawing, there is no alternative to the “right way” of drawing a symmetric image: each edge of the triangle has to be mirrored one by one, and the outline can be drawn. For taping, the process consists of approximating the placement of the triangle as a whole, and adjusting it. If the pupil fails to see the discrepancies between the expected solution and his, the manipulation does not bring anything.

(Q_u) How do Pupils Use the Persistent Characteristic of Paper?

The Metroscope was shared in the middle of the classroom (see Figure 7.22), and the pupils did most of their work with the paper interface at their desk. The advantage of this situation is that only one device is needed per class, but the disadvantage is, of course, that it takes time to go from a pupil’s desk to the Metroscope. Here, we investigate the implications of this compromise, which exists only because it is allowed by the persistence of paper, i.e. the fact that the paper interface still exists outside of the augmentation area.

Effort to Use the Metroscope The first concern about requiring pupils to go from their desk to the Metroscope is whether the physical effort needed will discourage them from using the augmentation. In the end, the activity could be fulfilled without the Metroscope. We found an opportunistic observation related to this question: a pupil had a broken leg. Although she was temporarily handicapped, she still managed to go to and from her desk and the Metroscope to check her answers. Her crutches did not inhibit her from participating in the activity. She

Figure 7.22: Set-up of the Metroscope in the classroom.

actually checked all her exercise sheets, and except for the last two of the second session, she checked multiple times. This lets think that the effort to move to the Metroscope does not discourage pupils from using it, because even a temporary handicap does not hinder the effort to access it.

Misuse of the Metroscope Conversely, another risk of requiring pupils to go from their desk to the Metroscope is the ability to abuse the playfulness of the device, and waste too much time. We did observe this behaviour, but only once.

A pupil was experiencing a novelty effect visibly more strongly with the Metroscope than the others. Throughout the study, he kept expressing his enthusiasm. After finishing a new exercise, he repeatedly brought his previous, already completed and checked worksheets along with this new one to check all of the answers with the Metroscope again.

While most pupils worked independently with the Metroscope, this pupil needed a lot of assistance from the researchers. The pupil struggled to recognize correct and incorrect triangle placements under the Metroscope (even though the correct triangle placements were projected on the paper). He had difficulty recognizing both when he had drawn an extra triangle and when he had forgotten to draw a triangle altogether. One specific example occurred while he was doing the ×axes worksheet for the 'drawing' mode. One single triangle was supposed to be reflected in the upper middle part of each quadrant. Instead of drawing the triangle placement to the right, however, the pupil drew it to the left. He thought his triangle placements were all correct, even though the Metroscope was clearly projecting a triangle (the correct answer) where he had not drawn one.

One of the researchers intervened, telling him that he had to erase this triangle and draw another one to the right. He responded by saying that he had already done that a few times, but each time the Metroscope displayed the projection of the triangle on the other side. Still, he went back to his desk to attempt to fix the position of the triangles. This time, he returned with the correct triangle placement and the incorrect one (an extra triangle). When he placed the worksheet under the Metroscope to view the projections, he again thought that it was perfect, without errors.

Alternative to the Metroscope We report a third anecdote illustrating the compromise between the time required to use the Metroscope and the value associated with it. One pupil was working on the ×axes worksheet. Unable to visualize the symmetry over the multiple axes, she folded her paper to see where the triangles should be reflected (see Figure 7.23). This is an interesting use of the flexibility of a paper interface between its augmented use and its traditional use.

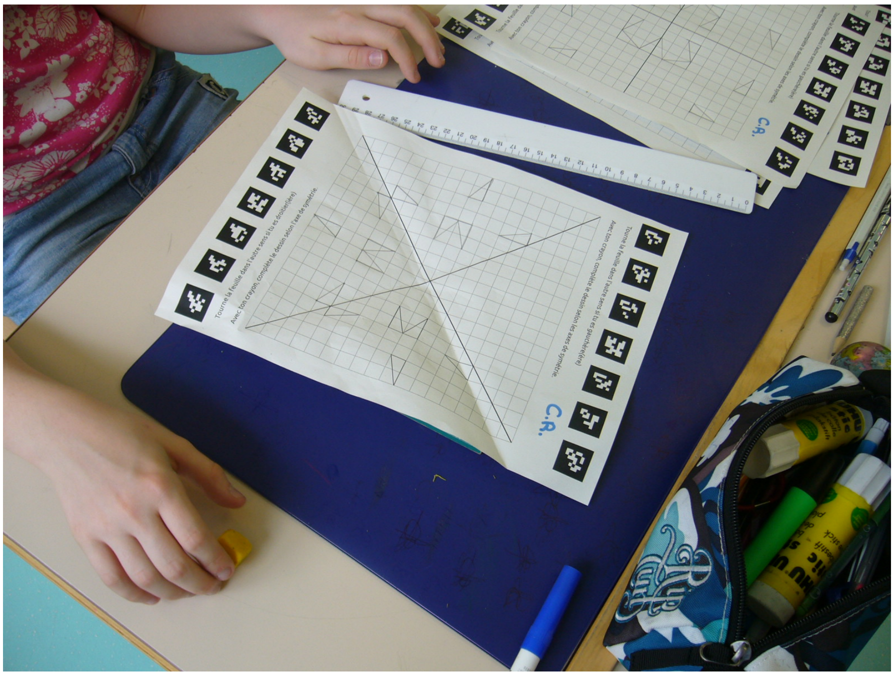

Figure 7.23: The ×axes exercise sheet of a pupil who approximated the Metroscope by folding.

Abuse of the Metroscope Another pupil showed yet another use of the Metroscope, which contrasts with the previous case, in that the pupil used the Metroscope for more than checking. After going back and forth from her desk to the Metroscope to check and then re-correct the position of her triangles a few times, she began marking the position of the correct answers projected by the Metroscope with a slash mark while at the machine. She would then go back to her desk to fix all of these errors.

When she reached more difficult worksheets with multiple axes of symmetry (+ and ×), instead of just marking the correct answers, she would physically move the cut-out triangles to their correct positions directly under the Metroscope. Similarly, during the afternoon session in the drawing mode, we observed this pupil erasing incorrect triangles and drawing correct ones under the Metroscope.

Usage Pattern Unsurprisingly, the pupils used the Metroscope significantly longer for the toughest exercises than for the others (18.7s versus 9.7s, $p < 0.1$). What is more surprising is that the amount of time spent under the Metroscope increases with the number of visits to the Metroscope. Figure 7.24 shows that the average duration of each visit at the Metroscope increases as a function of the number of visits ($p < 0.01$). This can be explained by the fact that pupils who need help also need more time to look at the corrections to understand them.

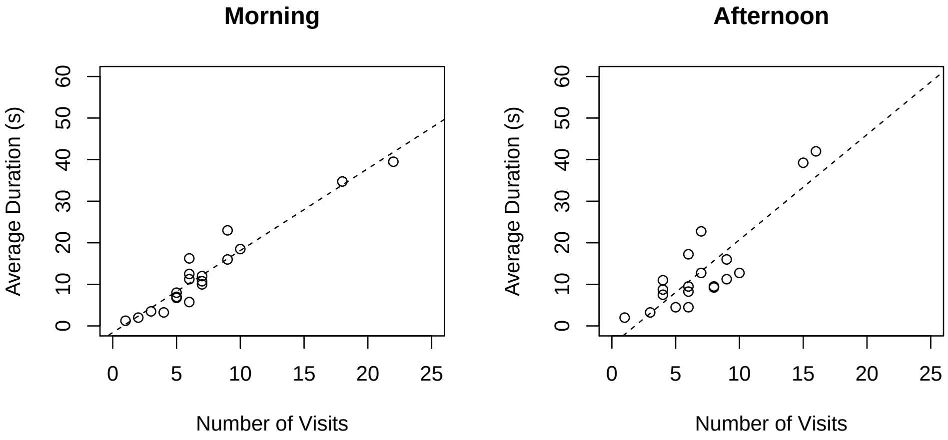

Figure 7.24: Average duration of time spent at the Metroscope per number of visits *per pupil*

.

Performance We found that that the ratio of triangles placed on a sheet where at least one triangle had been placed was significantly higher if the sheet had been brought to the Metroscope. In other words, more exercise sheets had been at least started when they were brought to the Metroscope. 95% of the triangles had been placed on sheets brought to the Metroscope versus 79% if they had not been brought to the Metroscope (t-test, $p < 0.05$). However, this does not mean that the Metroscope helped with the placing of triangles: in fact, the causality relationship is reversed: the exercise sheets brought to the Metroscope were mostly those that were already completed.

It is more interesting to compare the number of triangles placed correctly on a sheet, depending on how often the sheet was brought to the Metroscope. We also observe an improvement in this number (from 82% to 94%, $p < 0.1$) for sheets that have been brought to the Metroscope. In this case, it is reasonable to assume that the consequence is the score improvement, and the cause is the augmentation, because the point of the Metroscope is to correct a sheet.

Figure 7.25 illustrates these results. It also reveals that the variance of the performance is very small when the Metroscope is used, confirming its use as a correction method.

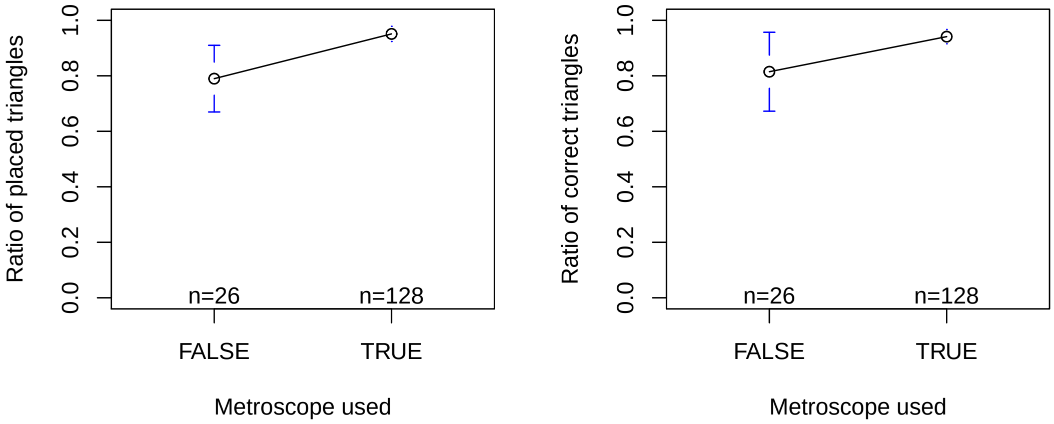

Figure 7.25: (Left) Triangle placements depending on whether the Metroscope was used. (Right) *Correct* triangle placements depending on whether the Metroscope was used.

Summary The fact that an interaction with paper can persist outside the augmented area allows the creation of a spatially distributed system. The spatial distribution of the interaction is not necessarily an obstacle, even with crutches.

Distributing the interaction allows users to work freely in parallel. Actually, it not only allowed the pupils to use the system but also to misuse it (e.g. by using it for the sake of using it, without paying attention to the feedback), to underuse it (e.g. by replacing the system by folding the paper), or to abuse it (e.g. by copying the feedback rather than just comparing the answers).

This distribution of the interface allows the device to be shared efficiently. One Metroscope was enough for a whole class, and we showed evidence that its use was parsimonious: for example, it was used more for difficult exercises, or by pupils who needed more help and therefore more time with the machine. The use of the Metroscope was also effective, in the sense that it fulfilled the goal of helping pupils correct their exercises.

7.5 Creative Activity

Finally, we used Kaleidoscope to explore the creativity of the pupils. After the previous controlled study, we ran a small, informal activity during two periods of 45 minutes. We gave the pupils a sheet with a vertical axis, a sheet with +axes, and a sheet with a combination of + and

× axes. They used these sheets to prototype a partial outline with a pencil in the input area of the Kaleidoscope sheet (see Figure 7.26a). Then, they used the Metroscope to see a preview of the full outline, composed by the partial outline and its symmetric projections. They could adjust the outline, and when they were satisfied with it, the pupils would cut along the outline (see Figure 7.26b).

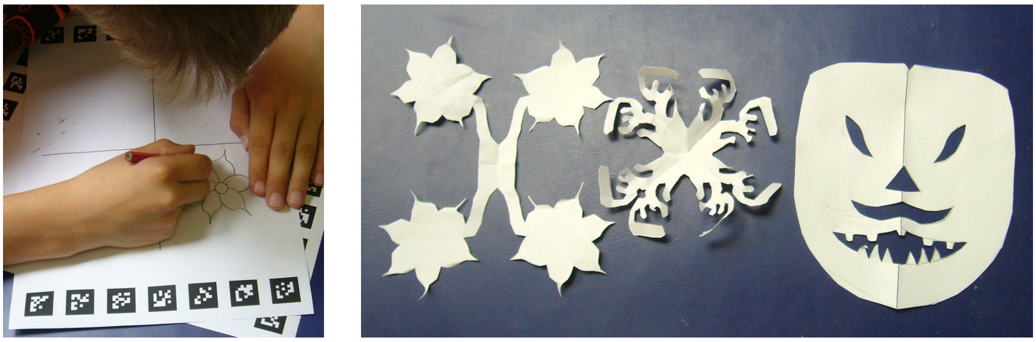

(a) Prototyping the edges. (b) From left to right: Resulting cut out on two, four and one axes.

Figure 7.26: The two phases of augmented creation of paper cut-outs.

We were impressed by the variety of the pupils' creations. Some of them were simple drawings not meant to be cut but simply automatically reflected. The vast majority, however, took advantage of the exercise to actually design cut-outs. A lot of the designs included inner holes, allowing for more complex shapes. The most impressive example of appropriation is shown in Figure 7.27: a pupil designed her cut-out as an object in space rather than a flat object.

(a) Top view. (b) Side view.

Figure 7.27: A cut-out in space.

We expected this activity to create congestion around the Metroscope because of the engagement that could happen. For this reason, we limited the time that each sheet could be augmented: a bar was projected on the sheet, and decreased linearly (see Figure 7.28). When the bar was empty, the Metroscope would stop projecting the symmetric image of the input area, rendering the Kaleidoscope sheet a regular sheet of paper without augmentation capabilities.

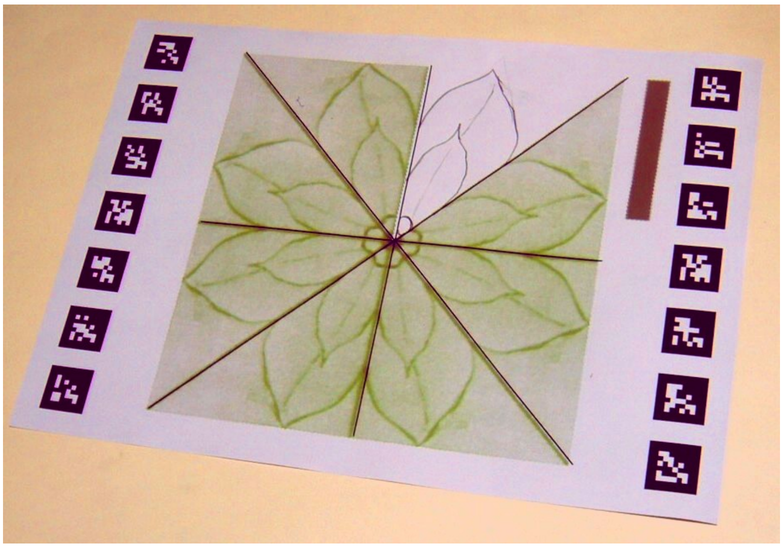

Figure 7.28: A kaleidoscope sheet limited in time. The augmentation disappears when the bar (on the right) is empty.

We computed the expiration time of each sheet as a fair repartition of the whole period, assuming that two pupils would be constantly using the Metroscope at the same time (this worked as long as the axes on the two sheets were different). As it turned out, only one pupil reached the expiration of a sheet: she was drawing a very detailed horse head. Otherwise, this is another example of how the persistence of a paper interface allows for an efficient repartition of a system.

7.6 Conclusion

Appropriation by the Teachers Kaleidoscope showed how a generic component could be used to transfer the pedagogical design of an augmented reality activity to teachers. The classroom is an environment too complex to be modelled by a researcher. Instead, it seems only reasonable to ask teachers, who are most experienced in this domain, to address this issue. Researchers should then focus on implementing modular and customizable components, instead of replacing teachers in their designs or performance.

The resulting activities were rather different. We do not intend to over-analyse these examples, concluding which alternative is better. In fact, we believe that each activity was optimal for the environment defined by the class and the teacher. Also, the teachers owned the activities, and the fact that they designed them themselves probably explains this optimization.

From the start, the teachers had a clear picture of what the activity would look like. For example, the first teacher already knew that her class would need some discipline measures to ensure that not everyone was waiting for the Metroscope at the same time, and thus set a limit to the number of pupils queueing. But the most salient example concerns the difficulty of the exercises. The teachers knew exactly how to adapt the exercises to their classes, even if the class was not homogeneous. On the contrary, the activity we developed for the controlled study had a clear imbalance within the exercises in terms of difficulty levels for the pupils.

Orchestration We also observed several examples of how a paper interface helps a teacher in the task of orchestration. Orchestration refers to the continuous task of identifying and managing the learning opportunities and constraints happening in a classroom in real time.

The strongest constraint in a classroom is time. Even the simplest tasks have a time cost, such as distributing exercise sheets, which takes five minutes. We saw an example of these types of time decisions that teachers have to make constantly in the second experiment, where the first teacher decided to skip the second half of the activity, allowing him to do the first half with more serenity. In contrast, the second teacher strived to finish the activity and managed time intensely. In this context, it helped that the interface could be distributed, because the teacher could give and take the various exercise sheets when he saw fit.

Managing the Metroscope is an example of another type of constraint to handle: the scarcity of some physical resources. It is not always possible to have a given resource per pupil: typically, computers are not always as numerous as the pupils. In this context, it helps that the Metroscope can be integrated in the classroom. For example, the Metroscope allowed the teachers to assign time-filling activities to the pupils who were not using the lamp. Moreover, since it is easy for the teacher to monitor who is using the Metroscope, the teacher could efficiently manage its access.

Orchestration is not only about dealing with constraints. It is also about identifying learning opportunities and using them. We saw several examples of Kaleidoscope supporting this. For example, the explanations could be embodied by showing pieces of the interface. In the first design and the controlled experiment, the teachers naturally brought pupils to the Metroscope when they needed to illustrate their explanations.

Usability The controlled study spoke against the manipulations that are possible with paper. While this is not a general result for all the activities based on paper as tangible, it can highlight another strength of paper. Paper is deeply established in its practices in education. This may explain why the pupils were more effective in the drawing condition than in the taping condition. Even though the taping practice was familiar, drawing figures is a much more familiar way of working in geometry. Thus, beyond the expected gain of a possibility offered by a technology, it is also important to maintain established uses, adding possibilities rather than taking them away and changing everything, as advocated by reality based interaction Jacob et al. (2008). Such a philosophy is clearly in line with the use of paper.

However, the creative activity showed a way to leverage the other possibilities offered by paper as an interface. Since the pupils were free to set their own objectives, rather than one defined in an exercise, the (at least in the classroom) uncommon practice of cut-out folded sheets was not an obstacle. It even stimulated their imagination: it is not so common that pupils do not try to imitate their peers. In the end, paper allows for a creative exploration of its possibilities as much as it allows for the use of already established practices when needed.

Finally, the controlled study revealed another important feature, which was confirmed in the creative activity. Paper interfaces are a very efficient way of distributing a computing resource among a class. No additional time management was necessary specifically regarding the use of a single Metroscope for a whole class. This comes from the fact that the persistence of paper allows the interface to be mobile because it still "exists" outside the interaction area. We will exploit this feature further in the next chapter.

8 Cutting and Folding

There remains some part of the framework that have been little or not explored so far, such as interactions based on the edges or the folds of pieces of paper. For example, changing the edges means cutting some pieces, and reassembling them in a new configuration.

Inspired by the previous chapter, we developed a generic component to address this exploration. Using a generic component aims at addressing a wider range of the possibilities offered by paper as interface.

In this chapter, we first present this generic component, StarrySheets, and illustrate its flexibility with four different activities. These proofs of concept were deployed in classrooms, and we report on these studies. Then, we discuss how the four activities compare to each other.

8.1 StarrySheets

Fiducial markers are a reliable way of detecting the presence, position, orientation, and side of a sheet. However, they lack some flexibility regarding the folds and the edges. Indeed, a marker becomes invalid if it is only partially visible, be it because it is cut, folded across, or folded upon. This reduces a lot the interactions based on cutting and folding, because it can only happen in places where it won't damage the tags too much. In other words, the cutting and folding axes are limited to a predetermined set.

To address this, we developed StarrySheets, based on local descriptors. In this section, we present the principle of the local descriptors we use, and their use as a paper interface.

8.1.1 Principle

Local descriptors are visual features that can identify a part of an object. There is a very large variety of such descriptors (see e.g. (Li and Allinson, 2008) for an overview). We chose a method based on the LLAH (Locally Likely Arrangement Hashing) algorithm developed by Nakai et al. (2006) for several reasons. The main reason is that an implementation was already available in our laboratory. Another good reason is that the descriptor is based on the position of words, as opposed to many other algorithms which based on textures and perform sub-optimally on text. Since paper often carry text, it is an important point.

The basic principle of LLAH can be compared to recognizing constellations in the sky: the arrangement of the neighbourhood of a point allows to recognize it (e.g. the star on the “shoulder” of the constellation that looks like Orion is Betelgeuse). The LLAH algorithm recognizes a constellation by computing its descriptor. This descriptor is a set of local invariants, i.e. values that do not change under a set of transformation, in our case: luminosity, rotation, scale, and perspective. In short, these local invariants are computed from the ratios of the triangles formed by a point and its neighbours.

The original LLAH algorithm uses words as input points. It pre-processes the image of a document to merge close letters into blobs, whose centre is used. In our case, we skipped this step and directly printed random patterns of blobs on pieces of paper (a simplification of the work of Uchiyama and Marchand (2011)).

8.1.2 Implementation

We built StarrySheets around three elements: the StarrySheet itself contains the pattern recognized by the LLAH algorithm, a recording marker to register the constellation, and a checking marker to compare a detected constellation with the registered one. The whole process is illustrated in Figure 8.1.

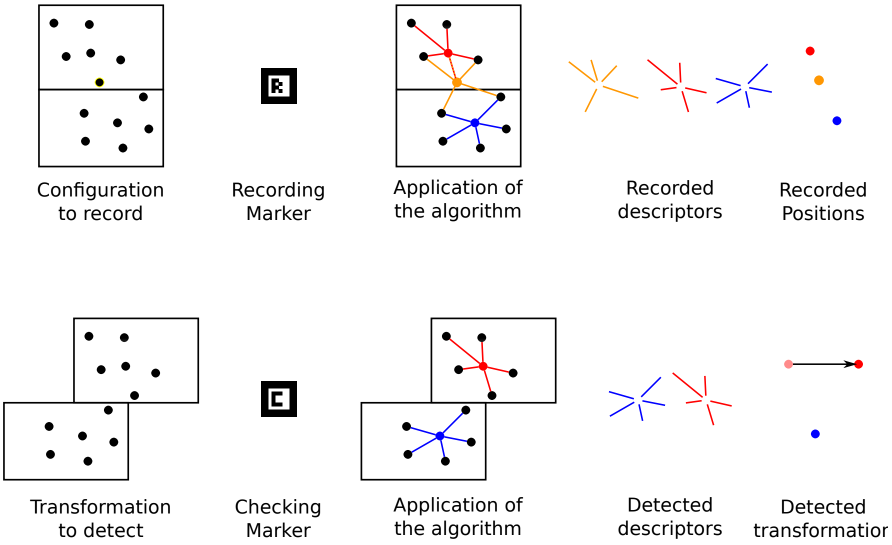

Figure 8.1: Detecting the deformation of a set of blobs using StarrySheets.

Pattern Sheet The pattern sheet is the sheet whose deformation (cutting, folding...) should be detected. It is a regular sheet with hundreds of blobs printed randomly. A blob can be identified using the LLAH algorithm, because its neighbours are most likely surrounding it in a different way than the neighbours of another blob.

Recording Marker When the pattern sheet is in a configuration to be saved, we show a recording marker to the Metroscope, whose presence will trigger the registration process of the pattern. A recording marker is a regular marker, associated with a set of blobs. Practically, the LLAH algorithm will be applied to compute a descriptor for each blob, based on the arrangement of its neighbours. The found descriptor are saved by the Metroscope, as well as the position of the associated blob. All this information is then associated to the recording marker. Another recording marker can save another set of blobs.

Checking Marker Each recording marker is paired with a checking marker. When a checking marker is present under the Metroscope, it will once again use the LLAH algorithm to compute the descriptor of each blob detect. These descriptors are compared to the descriptors registered by the paired recording marker. The checking marker then uses the positions of the blobs with matching descriptors to compute a visualisation of the transformation between the recorded set of blobs, and the checked one.

8.1.3 Feedback

The goal of StarrySheets is to give a feedback on the difference between two sequences of deformations of pieces of paper, such as folds or cuts. To do so, the checking marker triggers a visualisation of the translations that would be necessary to place the detected blobs in the same arrangement as the recorded one. If a blob is where it is supposed to be, a green dot is displayed; otherwise, a red line is displayed, starting from the detected position of the blob, and finishing on its expected position. See Figure 8.2 for a graphical illustration.

The feedback on the difference between the checked blobs and the recorded blobs is rather abstract; it does not explicitly describe the difference between the recorded pattern and the detected one. This abstraction has two advantages. First, it is more generic: the feedback is usable if the fold of a sheet is not the expected one, if the relative position of two pieces of paper is not the expected one, or a piece of paper has not been cut at the expected place. Second, it leaves the interpretation to the user. The computer does not judge whether recorded and checked sets of blobs are similar; instead, the user can appreciate whether the feedback is as expected or not. This also contributes to the genericity of StarrySheets, in the sense that for some applications, some kinds of errors are acceptable, but not in others. Practically, in some cases, it is acceptable, or even unavoidable, to have a small difference between the expectation and the actuality. In this case, the feedback will consist of small red lines.

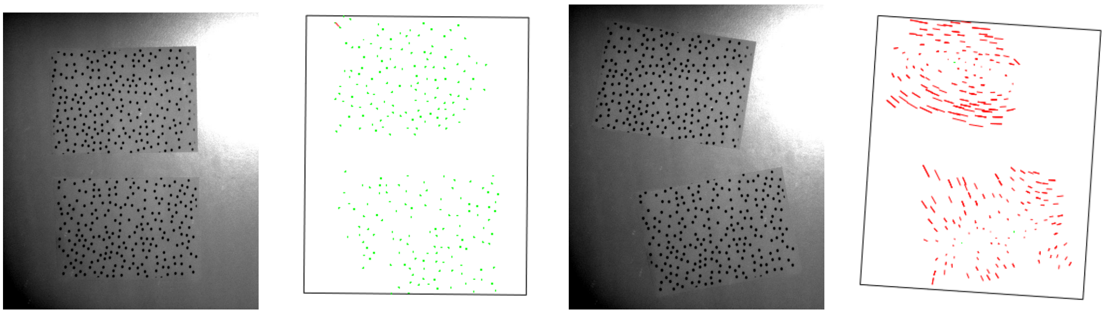

(a) Recorded paper configuration

(b) Feedback

(c) Paper configuration resulting from a different relative position of the two pieces of paper.

(d) Feedback showing the difference between the two configurations

Figure 8.2: Feedback on the relative position of two StarrySheets. Red lines correspond to errors. In this case, the line can be interpreted. The upper cluster looks like a whirl, corresponding to the rotation between the expected setting and the detected one. The lower cluster seems like attracted by the upper cluster, corresponding to the displacement between the expected relative position and the detected one.

8.1.4 Application

It may not be clear how StarrySheets can be applied to detect via the edges and the folds on paper. We illustrate hereafter several examples. Figure 8.2 showed a feedback on whether two pieces of paper are positioned and oriented correctly with relation to each other. In the correct case, there is a green dot for each blob, because each blob is detected in the same position, relative to the others, as when they were recorded. If a piece of paper carrying blobs is not in the same position with relation to the other, all the blobs are misplaced, and the feedback consists of red lines showing the displacement of each blob with relation to to its recorded counterpart.

If the two pieces of paper come from the same sheet, such a feedback can be used to check that the cut was correct (see Figure 8.3). Conversely, if the constellations are recorded in a way that two pieces of paper have a common border, the feedback checks that the assembly of two pieces of paper is correct, i.e. the union of their edges is the expected one (see Figure 8.4). Of course, the pieces of paper can also overlap each other (see Figure 8.5). As for folding, it is enough to print the blob pattern on the two sides of a sheet. Each fold will create a new, unique pattern, that can be checked (see Figure 8.6).

8.1.5 Online Recording

It is important to note that recording a pattern can be done at any time: while designing the activity, or while performing it. This can be used to design pedagogical scripts where pupils prepare a sequence of transformations with geometrical instructions, e.g. “fold along a line

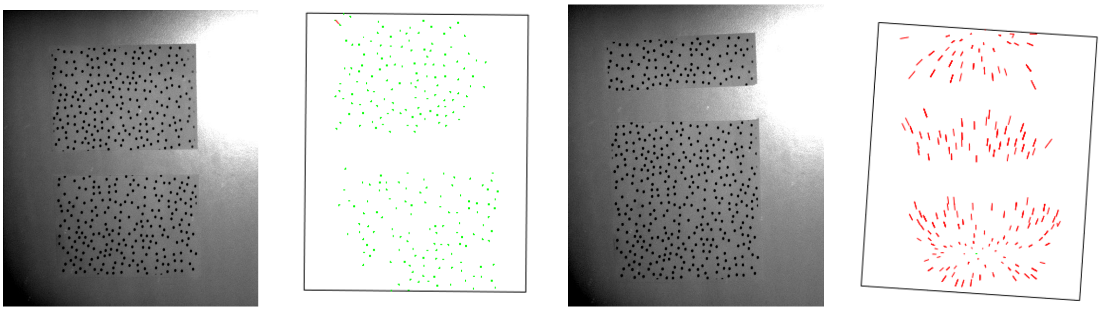

(a) Recorded paper configuration

(b) Feedback

(c) Paper configuration resulting from a different cut

(d) Feedback showing the difference between the two configurations

Figure 8.3: Checking the cut of a StarrySheet.

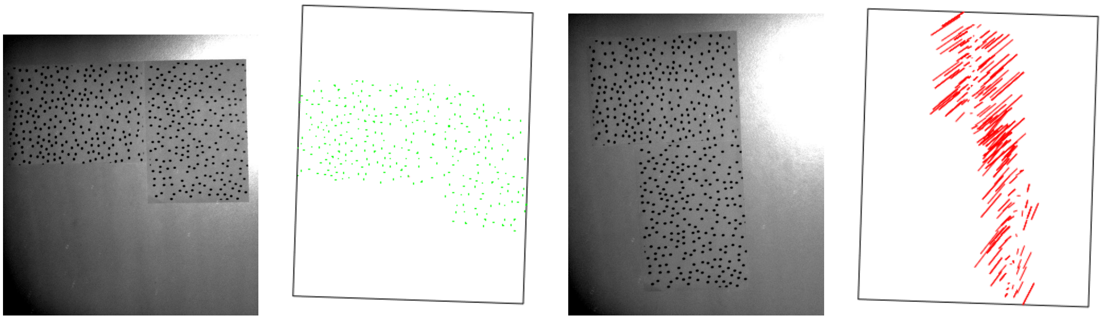

(a) Recorded paper configuration

(b) Feedback

(c) Paper configuration resulting from a different merge of edges

(d) Feedback showing the difference between the two configurations

Figure 8.4: Checking the merge of the edges of a StarrySheets.

that makes a 45° angle with some line". These instructions are pedagogically interesting to produce, and to follow, because pupils have to understand the geometrical meaning behind them. It is even more interesting to exchange the instructions among pupils, because they can then argue over them, and realise why instructions produced are ambiguous, for example.

8.1.6 Design Goals and Context

The rest of this chapter presents four activities based on StarrySheets. Each of them is adapted to a specific topic (e.g. angles or symmetries), and use different properties of paper (e.g. relative position, folds, edges).

The approach we took with these activities is different from the other activities presented in the other chapters. The focus was put on the interaction aspects rather than on the learning aspects. Indeed, it takes several iterations on an activity to make it pedagogically sound. However, in order to investigate the large range of possibilities offered by StarrySheets, it was

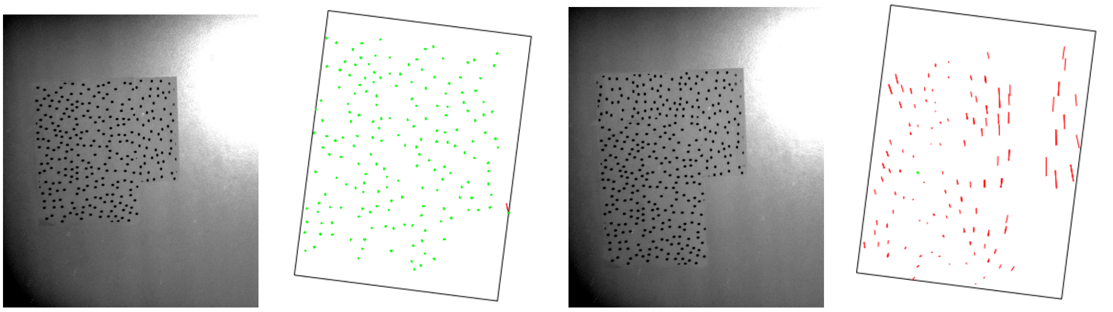

(a) Recorded paper configuration

(b) Feedback

(c) Paper configuration resulting from a different overlap

(d) Feedback showing the difference between the two configurations

Figure 8.5: Checking the overlap of StarrySheets.

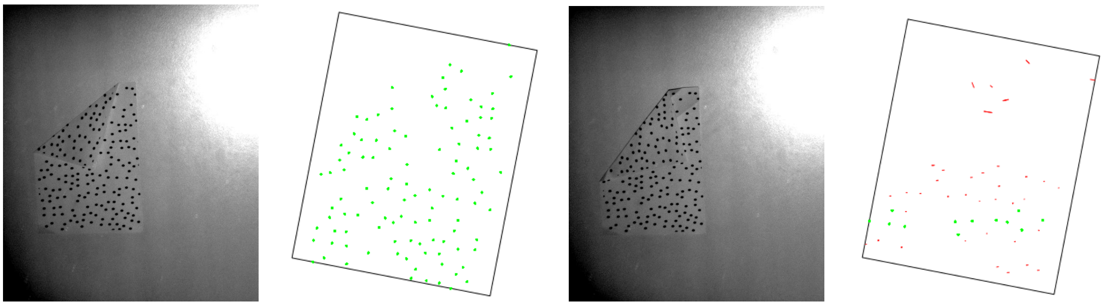

(a) Recorded paper configuration

(b) Feedback

(c) Paper configuration resulting from a different fold

(d) Feedback showing the difference between the two configurations

Figure 8.6: Checking the fold of StarrySheets.

more effective to deploy four different studies, rather than improving an activity over four studies. As a result, the activities were less polished.

We also developed these activities independently from the teachers. We only asked them what the topic should be. Our goal was to investigate whether a generic component like StarrySheets would allow the fast creation of activities. In this perspective, we strived to keep the development time of each activity to one day.

We first report observations made during the study related to each activity, and will discuss them together afterwards.

8.2 AngleHunt

8.2.1 Description of the Activity

AngleHunt is the first activity developed using StarrySheets. It is actually a treasure hunt, where pupils have to recompose a map (see Figure 8.7a). To do so, the class is split into five groups, each receiving one part of the map, and the instructions on how to recompose their part of the map. Each partial map received by the groups was an A4 StarrySheet split into 2 × 4 equal pieces, the actual map was on the other side of the StarrySheet (see Figure 8.7b).

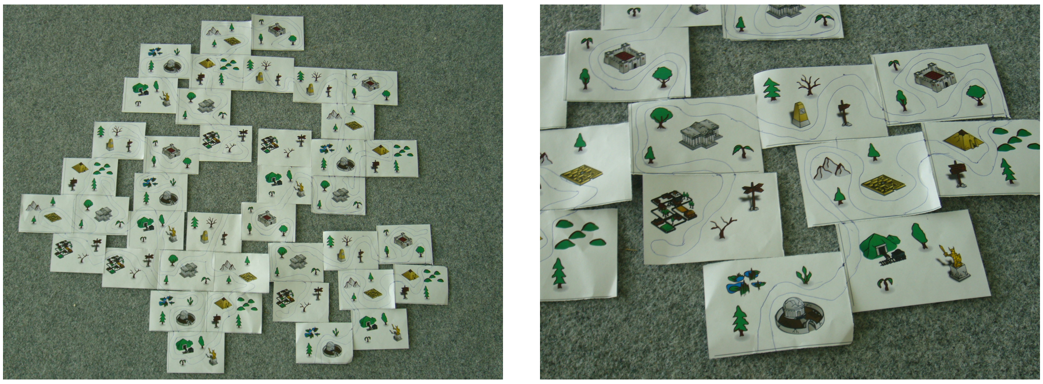

(a) The full map

(b) The (bottom-right) part of the map to be recomposed by one group.

Figure 8.7: The map to be recomposed by pupils.

On each piece of map, in addition to the constellations, three points were marked and named. The instruction sheet to recompose the map was simply the measure of an angle linking two points from one piece of map, and a third point on another piece of map (see Figure 8.8). To avoid ambiguities, the instruction sheet also had a schema of the angle (in order to specify the orientation of the angle), and the pieces of maps were constrained to be horizontal, contiguous, and the names of the points indicated which way was up.

AngleHunt aims at introducing the use of the protractor. The pupils have to use it to position the pieces of map with each other, according to the specified angles. When the position is satisfactory, they could tape the two pieces together, and incrementally add the others. The same principle applied to assemble the partial maps together: a final instruction sheet specified the angles between points from different partial maps. The map was placed between 5 boxes. The map represented a path leading to one of them, which contained candies.

In this activity, pupils interact with the edges of pieces of paper, which is incrementally combined: the class starts with 40 rectangular pieces of paper, and ends up with one monolithic, irregular piece of paper. The Metroscope was used to provide feedback on the correctness of the assembly. The instruction sheets carried a checking tag, which was associated to the expected assembly of the pieces of StarrySheets. Thus, when two pieces of map and the

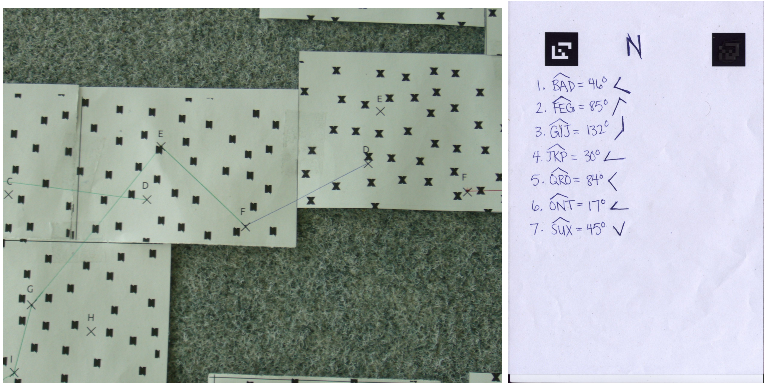

N

1. $\widehat{BAD} = 46^\circ$
2. $\widehat{FEG} = 85^\circ$
3. $\widehat{GIJ} = 132^\circ$
4. $\widehat{JKP} = 30^\circ$
5. $\widehat{QRO} = 84^\circ$
6. $\widehat{ONT} = 17^\circ$
7. $\widehat{SUX} = 45^\circ$

Figure 8.8: Points linking pieces of maps (left) according to the instruction sheet (right). The piece of map in the middle contains the points E and F. The piece of map on the bottom left contains the point G. The instruction sheets defines the angle FEG as 85° wide and gives a schema of it.

instruction sheet were placed under the Metroscope, the feedback was displayed, allowing to check whether the pieces of paper were placed correctly relatively from each other.

8.2.2 Study

We organized a workshop during an open day event targeted at primary schools, as it was the case in Section 6.3. Seven classes visited our workshop, which gave us the occasion to have a first study involving StarrySheets, by deploying AngleHunt.

The study took place in the meeting room of our laboratory, shown in Figure 8.9. We ran one session for each of the seven classes, every hour. Each class stayed for 45 minutes. During 10 minutes, we introduced the laboratory, Metroscope, and the activity. For the following 30 minutes, we ran the activity, and we kept the last 5 minutes to collect feedback.

The classes were fifth grade or sixth grade, i.e. the pupils were 10-12 year old. In general, pupils are introduced to protractors in sixth grade. We dispatched each class on five stations. Each stations had one Metroscope, and a part of the map. Twelve protractors were dispatched among the class, i.e. each station had two or three of them. Each class consisted of around 20 pupils. They were accompanied by one teacher, and sometimes a teaching assistant. They provided assistance to the pupils on the task, along with two experimenters.

To collect the feedback of the pupils, we distributed sticky notes at the end of the activity. The pupils were asked to write down what they liked or not liked in the activity. Then, the pupils had to stick their note on a white board split in two halves, one having a smiling face, and the other having a sad face. We cleared the board after each class not to influence the following one, and the pupils did not know about the candies before filling the sticky notes. We collected 105 sticky notes.

Figure 8.9: The set-up for AngleHunt.

(Q_u) How do Pupils Use a Paper Interface?

General Observations There are two points to note regarding this study before analysing the usability of AngleHunt. First, the novelty effect was strong, because the classes were discovering the Metroscope for the first time. As such, 37 sticky notes mentioned the *machine* or the *technology*. Second, AngleHunt suffered from some technical limitations, being the first activity making use of StarrySheets. Most strikingly, the recognition of the patterns was not reliable: 17 sticky notes mentioned the “buggy” feedback.

Reception of the Feedback Surprisingly however, only 3 sticky notes qualified the feedback as complicated. This was an important result for the next activities, as it validates that the abstract StarrySheet feedback can help pupils. In fact, 11 sticky notes mentioned positively the augmentations.

Addressing Precise Manipulation Regarding the paper interface, the manipulation was not trivial. It required precise moves of the protractor and the pieces of map to link together. Then, pupils had to draw a line across two pieces of paper without them moving. Regarding the precise manipulation of multiple elements, the paper interface allowed for a manipulation by multiple pupils at the same time (see Figure 8.10). We often observed such a collaborative interaction pattern, which seemed satisfying for the pupils.

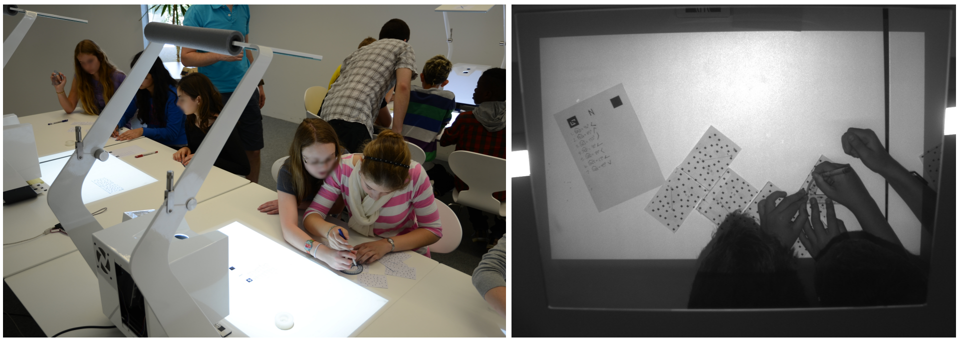

Figure 8.10: Examples of multiple hands manipulations. Photo Himanshu Verma

Improving the Drawing Process Regarding the line to be drawn across two sheets, the teacher of the first group showed us the method naturally developed by a pupil. He would only mark the measure of the angle one piece of paper, and draw the line in a second step, by aligning the mark and the two points. This technique was indeed more reliable than what had been demonstrated by the experimenter, where the mark was done on another sheet, which could move before the line was drawn. For the six other group, we demonstrated the improved method during the explanation of the task. This was a nice example of the appropriation of the interface by the user, who outperformed the designer.

(Q_I) **How do Pupils Learn with a Paper Interface?**

Reception of the Activity A lot (33) of sticky notes declared that the activity was enjoyable. This is a bit surprising, because the task was difficult, especially for the classes who had not been introduced to the protractor. 12 sticky notes did qualify the task as difficult. Interestingly, 7 sticky notes declared that the pupils enjoyed the activity because they learned something (mostly, the use of the protractor). This is almost as much as 6 the sticky notes declaring that they did not like the activity because it was based on mathematics. It means that the activity was more perceived as educational than playful.

Augmentations to Prevent Short Cuts AngleHunt was an interesting example of what augmentations can bring at the pedagogical level, beyond fast calculations. It could have been possible to complete the task without the augmentation, and without really measuring the angles. Indeed, since the instruction sheet gave a schema of the angle to produce, the pupils could roughly position to pieces of map, and flip them to precisely align them using the map pattern. However, the pattern did not allow this strategy: the path draw was not precise enough to help as a guide. That was confirmed by the feedback given by the augmentation. In other words, the augmentations allow to separate the immediate appearance from the appearance of the solution. It allows to let the pupils manipulate the problem directly, while hiding a short cut to the solution.

Paper Interfaces to Help Supervision It was easy to spot “short cuts”, because the tape was on the map side of the sheets rather than on the StarrySheet pattern side. The group dynamics was also easily observable. Figure 8.11 shows a group uniformly engaged in the task compared to a group with an idle pupil trying to hide under his hand. There were several ways to engage everyone in the group: either form two pairs, by giving two protractors, or enforce a turn taking, so that the pupils are forced to follow what happens, in order not to be lost. It is important for teachers to acquire such information on the general progress of the class. It helps optimizing the interventions, and the time dedicated to each group, because it is a scarce resource.

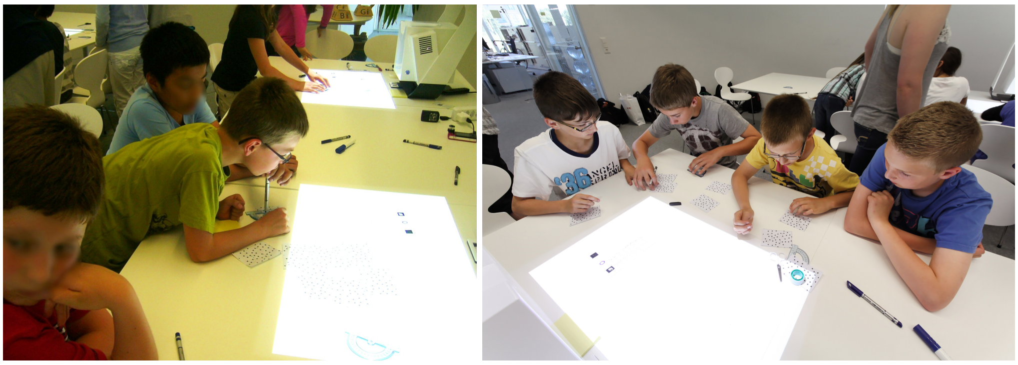

(a) An idle pupil along with a collaborating pair. (b) Two collaborating pairs. Photo Alain Herzog

Figure 8.11: Different group dynamics.

Ownership of the Paper Interface Finally, we were surprised by the discrepancy between the collaboration within groups and the collaboration among groups, i.e. within the class. The goal was to collaborate to assemble a common map. This goal was exactly the same at the group and class level. However, despite our encouragements to do so, no pupils from a faster group helped a slower group. This may be explained by the several sticky notes mentioning the group as a reason to why they liked the activity or note: some had a good laugh, some found the atmosphere within the group unpleasant. In any case, the paper interface being a physical entity, it is associated to a feeling of ownership: pupils take care of the task under their responsibility, not more, not less. In this case, the task is materialized by the pieces of map.

(Q_t) How do Teachers Use a Paper Interface?

Involvement of the Teachers The teachers seemed comfortable in managing such an activity that they had never seen before. They started helping the pupils right after the start of the activity. They did so without being asked (and were of course welcome). They were obviously familiar with the interface element (i.e. pen, paper and protractor) and could use the same cues as we did to monitor and intervene when necessary (see Figure 8.12). This is a clear advantage of paper interfaces over a mouse/keyboard/screen interface for example.

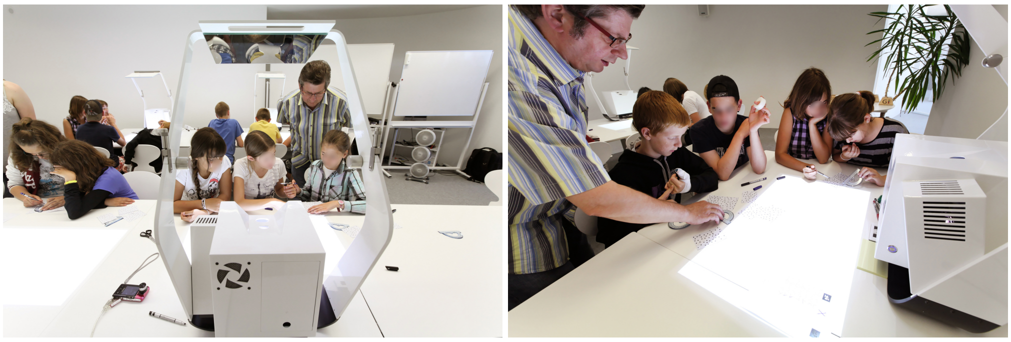

Figure 8.12: A teacher watching and intervening when necessary. Photo Alain Herzog

8.3 Sympliage

8.3.1 Description of the Activity

Sympliage is an activity where pupils find the position of a point by applying symmetry consecutively. Pupils are given a double sided StarrySheet. One of the side has a ($1cm \times 1cm$) grid, and four axes: a vertical, horizontal, and two diagonal (see Figure 8.13).

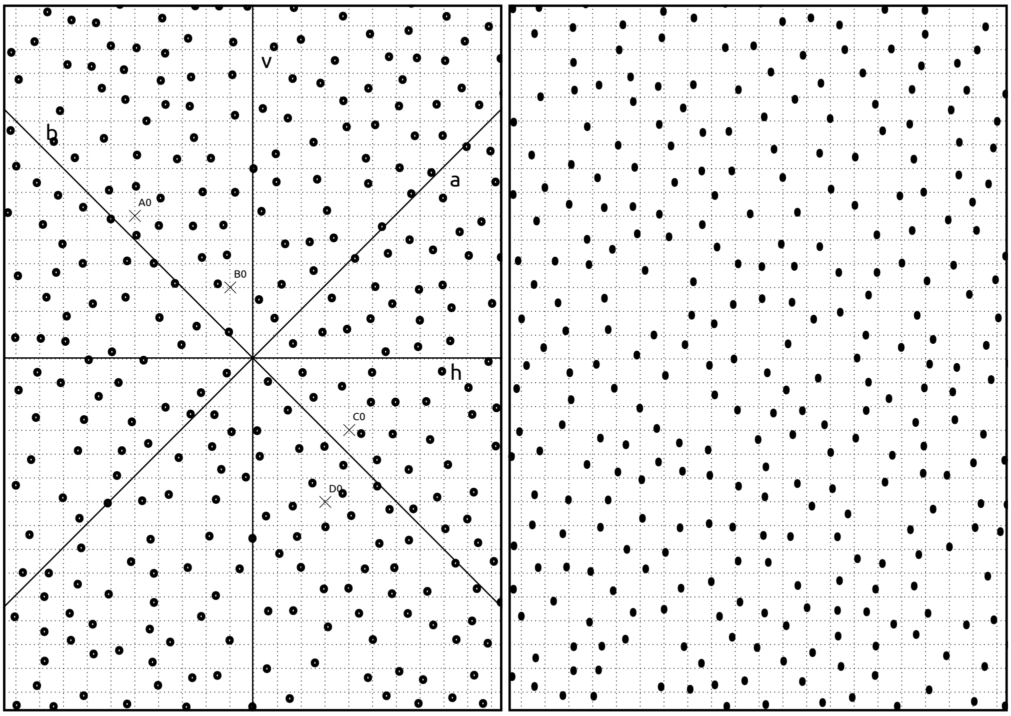

Figure 8.13: Front and back of the sheets used in Sympliage.

Once more, pupils follow geometrical instructions. Starting from a given point (e.g. $A0$), they have to reflect it three times according to the specified sequence of axes (e.g. $A0$ reflected over the h axis results in $A1$, $A1$ reflected over the v axis results in $A2$, etc.). It gives a target point (e.g. $A3$), onto which a corner has to be folded (see Figure 8.14). The instruction sheet also has a checking marker which allows to validate with the Metroscope whether the sheet is folded correctly, i.e. whether the target point is correct. The pupils repeat this with three other instruction sheets, corresponding to one corner each.

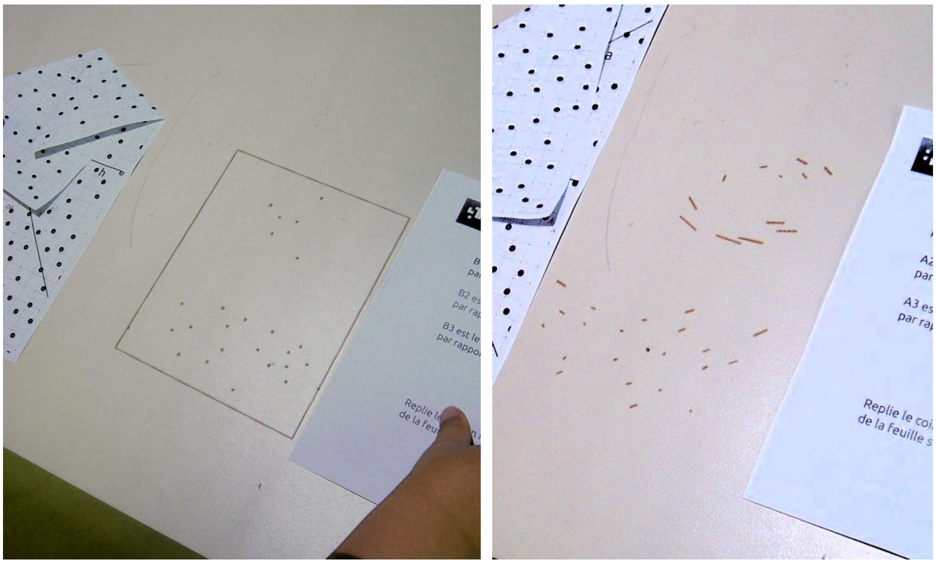
(a) A positive feedback.
(b) A negative feedback.

Figure 8.14: Pupil checking their folds.

Sympliage has also a recording mode. In this case, the axes on the instruction sheet are blank, but the sheet has a recording marker. The pupils have to define their own sequence of reflections, by specifying the axes over which the sequence of points have to be reflected. Then, the pupils record each fold by showing the folded sheet and one of the four instruction sheets. Once the folded StarrySheet is recorded, the pupils can tear up the recording tag from the instructions sheet. This way, the recorded paper configuration is permanently associated to this given instruction sheet, because it it not possible to record another pattern associated to its checking marker. The pupil reproduces the starting points on the original, double-sided StarrySheet with axis and grid and gives this sheet and the read-only instruction sheet to another pupil. The other pupil does the same, and they can follow each other's instructions, as explained before.

8.3.2 Study

We tried this activity in School 4, during two 45-minutes periods. This study took place several weeks after the Kaleidoscope study in Section 7.2. The pupils had thus already been introduced to symmetry, but had not exercised it in between. We used one Metroscope for the whole class, on the same table at the back of the classroom see(see Figure 8.15). The teacher and two experimenters roamed the room to help pupils who needed.

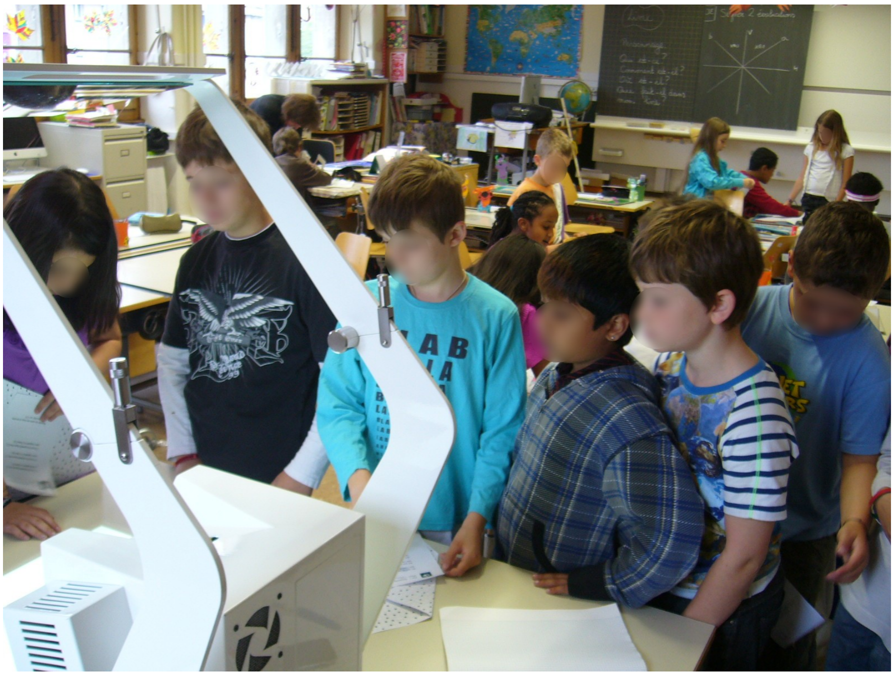

Figure 8.15: The Metroscope, at the back of the classroom.

The pupils started with pre-recorded instructions on how to fold the sheet. The four starting points had been placed on the side of the StarrySheet having the grid and the axes. The pupils had four A5 pages listed the sequence of axes to use to iteratively mirror the points. Once the fourth point once found, a corner could be folded onto it.

In the second part of the activity, the pupils had to prepare such a sequence of axes on a blank instruction sheet. They could record the resulting configuration of the folded StarrySheet by placing it, along with the instruction sheet, under the Metroscope. The blank instruction sheet had a recording marker, which could be torn off to make the sheet read only. Next, the pupils could reproducing the starting point on a blank StarrySheet, so that they could swap it along with the instructions with another pupil. The pair of pupils could then carry on the exercise designed by their peer.

(Q_l) How do Pupils Learn with a Paper Interface?

General Observations The main observation about this study is that the difficulty was not adapted to the class for several reasons. First, the exercise was too ambitious, given that the pupils had not practised symmetry in the previous weeks. Second, the formulation of the exercise was more formal than what fifth graders are used to: point names were indexed letters, axes were identified by a letter rather than by their name (e.g. “v” instead of “vertical”). Third, the pupils had mostly worked with one axis, sometimes two, not four. Lastly, only vertical and horizontal axes had been addressed, not diagonal ones.

The most common mistake was the confusion between a cell of the pre-printed grid, and an intersection (see Figure 8.16). The points to reflect were marked by a cross on the intersections of the grid, but their names were printed inside a cell. The pupils considered the cell, not the intersection. This resulted often in a shift of one grid-unit of the point to fold the corner upon.

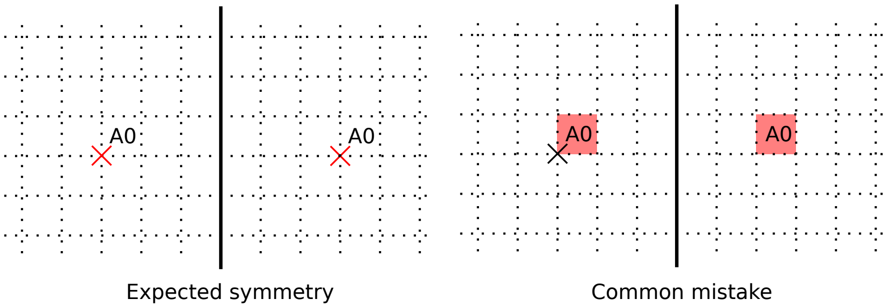

Figure 8.16: A common misunderstanding. A0 was intended to name a corner marked by a cross (see left), but pupils often thought that A0 was a cell of the grid (see right).

Precision of the Folding Detection The Metroscope was actually accurate enough to highlight such a mistake: the feedback was showing short red lines instead of green dot, even if the fold of the corner was off one grid unit.

Triggering Interaction The Metroscope had a positive role in the collaboration between pupils. The video show that as soon as the experimenter went away, the pupils took over his role of interpretation of the feedback. In other words, pupils around the Metroscope were collectively telling to the pupil checking an answer whether it was correct or not. This was due to the fact that many pupils were queueing next to the table where the Metroscope was installed. This table allowed the pupils in the queue to see what happens under the Metroscope, and comment on it. In this case, the Metroscope acted as a spot light on the pupil checking an answer.

(Q_u) How do Pupils Use a Paper Interface?

Slowness of the Feedback Two factors concurred to make the queue so long: the difficulty of the exercise, along with the slowness of the feedback, due to a low technical performance. The feedback could not be updated more often than every second, and it took several update to have a stable feedback. It is however interesting to note that the pupils easily understood this limitation, and could make the difference between a stable feedback, and false negative caused by the slow update of the feedback. This queue had a dramatic effect on pupils who forgot their instruction sheet at their desk to check the answer: they had to queue again, because the other pupils would not let skip the queue.

Intuitive Use Sympliage demonstrated how paper can be appropriated by its users. For example, for the second part of the activity, the teacher explained that the starting point should be reproduced for the other pupil to redo the same fold. To check that the two starting point matched, a pupil naturally used the transparency of paper (see Figure 8.17a). Another example of appropriation of paper is how we used a transparent sheet (for overhead projectors) to hold the folds consistent. Indeed, after several folds, a paper is not really flat any more (see Figure 8.17b).

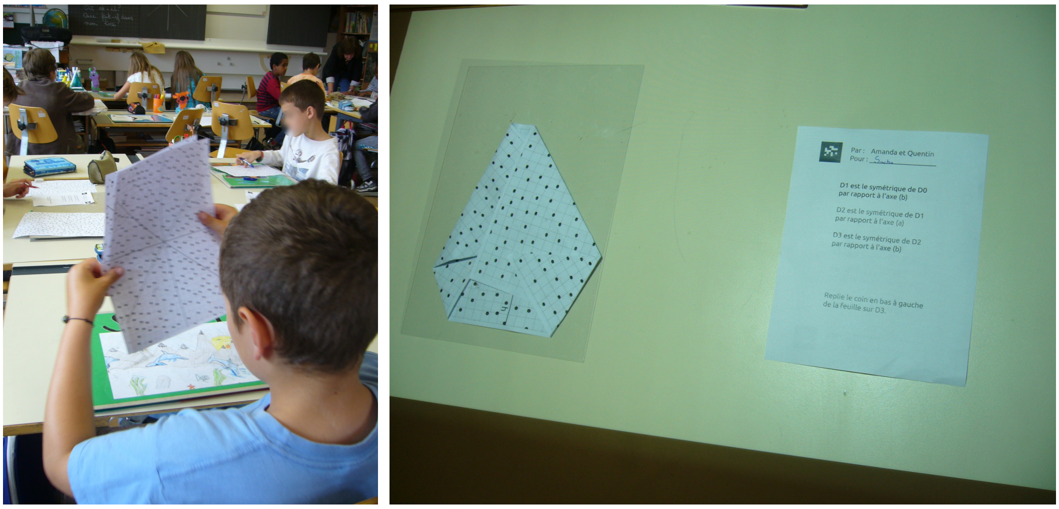

Par : Amanda et Quentin
Pour : Sacha

D1 est le symétrique de D0
par rapport à l'axe (b)

D2 est le symétrique de D1
par rapport à l'axe (a)

D3 est le symétrique de D2
par rapport à l'axe (b)

Replie le coin en bas à gauche
de la feuille sur D3.

(a) A pupil using transparency to check the alignment of points.

(b) The transparent sheet keeping folds flat.

Figure 8.17: Examples of appropriation.

Counter-Intuitive Use In contrast, pupils are not used to tear off pieces of paper in the classroom context. They seemed puzzled by this requirement, when they needed to fix the record of the fold in the second part of the exercise.

The order of the folds was another awkward part of the activity. The pupils had to first fold the top left corner, then the top right one, etc. If they did it in a different order, the edges would be the same, and the corners would be at the same place. However, the pattern would not be recognized any more, because the sheet would overlap itself differently. It is thus a conflict between real paper and this paper interface. However, the folds carried an interesting property: each of them is the persistent mark of a try. A teacher can easily see all the tries made by a pupil for a given corner. This is an indication of the progress done by the pupil in the exercise.

It is interesting to note that matching the edges of a folded paper with a template would have been another way to validate a solution (see Figure 8.18). However, this validation would fail at detecting issues such as the cell/intersection confusion. Similarly, mirrors could be used to visualise symmetric points, but the same limitation applies.

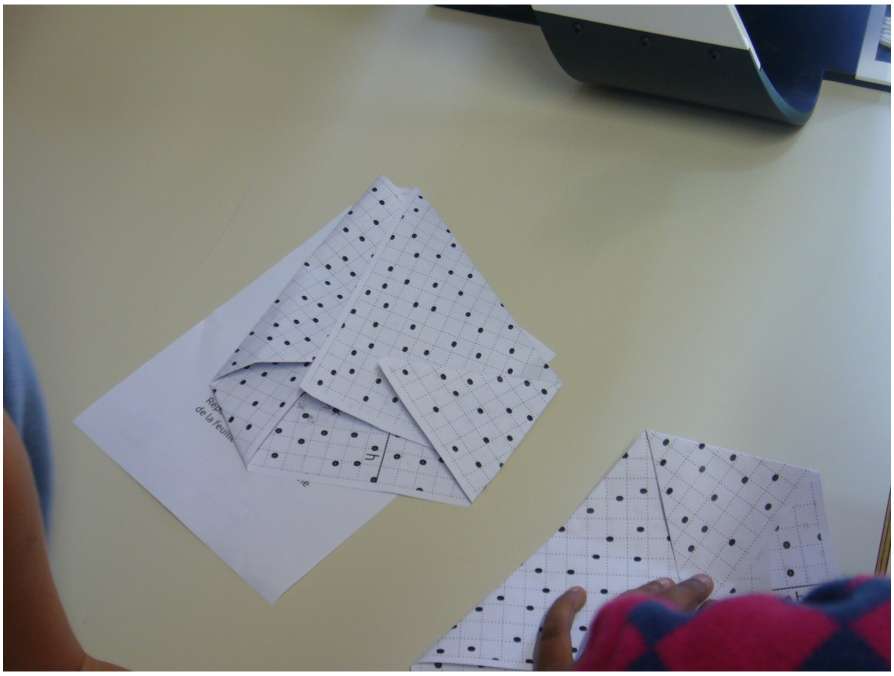

Figure 8.18: The expected solution (left) could be compared visually with a pupil's solution (partial, right).

(Q_t) How do Teachers Use a Paper Interface?

Alternative Validation The teacher actually asked the pupils to try to do the activity without using the mirrors used in previous symmetry exercises. In this perspective, It was a good thing that the Metroscope was used only for simple corrections showing whether it was right or wrong. We did resort to the mirror for pupils who had troubles visualizing the expected solution however, since the feedback from the lamp was not a detailed answer.

Queue Management This time, the teacher did not manage the queue. At some point in the exercise, more than half of the class was waiting . The reason for not managing the queue was probably that there was no alternative activity anyway. Indeed, in the last study with this class, the pupils were given drill exercises to remain busy at any time. This time, the teacher did not know exactly what would happen, and she did not prepare such exercises.

The queue is actually informative for the management of the class. If the queue becomes too long, it means that the difficulty is too high, as it is our case, and the exercise could be simplified. This observations is more related to orchestration in general, but it is a result of the fact that paper interfaces can be scattered among the pupils to let them all use the same device.

Lack of Adoption The teacher did not use the Metroscope at all this time. This may be due to the fact that too much help was needed in the classroom. Indeed, the teacher also qualified this activity as being a lot tougher. Another explanation for the teacher not use the Metroscope could be that she did not own the activity as much as she did in the first study.

8.4 Triangram

8.4.1 Description of the Activity

Triangram is an activity similar to tangrams, and focuses on symmetries. Pupils are given six right triangles, cut out of a cardboard StarrySheet. The pupils arrange the triangles on a grid printed on a cardboard sheet (see Figure 8.19a). Next, they draw the symmetric outline of the triangles on the grid of a regular paper sheet (see Figure 8.19b). The pupils can then record the arrangement of triangles by bringing it under the Metroscope, along with the instruction sheet which contains a recording marker.

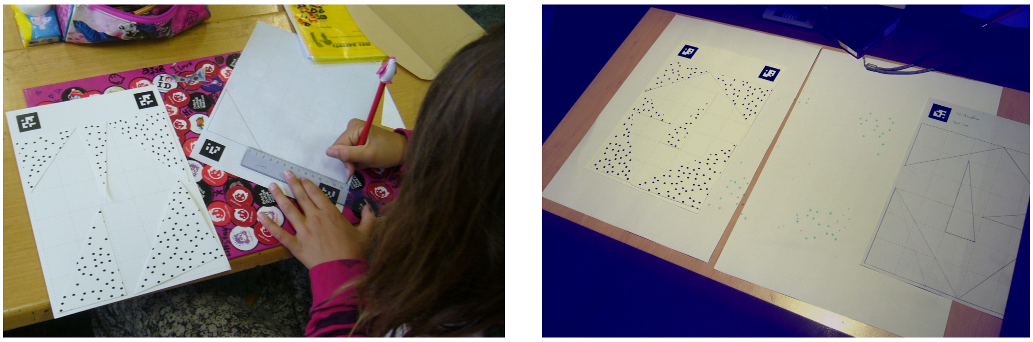

(a) A pupil drawing the symmetric image of the shapes placed on the cardboard grid.

(b) The Metroscope can check that the shapes are in the same arrangement as the one recorded.

Figure 8.19: Two steps of Triangram.

When the arrangement of triangles is recorded, the pupil can tear off the recording marker (to make the instructions read-only), and exchange the instruction sheet with another pupil. The pupils then have to reproduce (with the cardboard triangles on the cardboard grid) the symmetric image of the arrangement drawn on the instruction sheet. When they are ready, they can check their solution by bringing all the pieces of paper under the Metroscope. The instruction sheet triggers a feedback on how the proposed arrangement correspond to the recorded one.

8.4.2 Study

We tried this activity with two classes of School 3, during two 45-minutes periods each. The first class was a fourth grade of 18 pupils, and the other class a mixed third and fifth grade, of respectively 9 and 10 pupils. The fourth grade was split in two halves, one taking part in the study while the other was doing another activity outside the classroom. Similarly, the third graders were performing a painting activity while the fifth graders where using the Metroscope, and vice versa. We used one Metroscope freely accessible for pupils to check their assembly of triangles. The teachers roamed the room to help pupils who needed it.

The fourth and fifth graders used Triangram as described previously: they place shapes on a cardboard sheet, draw the symmetric image of the arrangement, and give it along with the shapes to another pupil who reproduces the original arrangement. However, the third graders were given an introductory exercise, where the symmetric arrangement was already drawn. The goal was to let them familiarize with the exercise before asking them to generate one of their own.

The fourth graders did a first exercise with two shapes, and another one with six. The third and fifth graders started directly with six shapes for the two activities activity.

(Q_u) How do Pupils Use a Paper Interface?

Transportation Issue. The main observation regarding usability concerned the transportation of the paper interface from the desk to the Metroscope, in order to record the arrangement of the shapes on the grid. The shapes were to be placed on a cardboard sheet, but it was not enough to stabilize the shapes during the transportation. The pupils moved very cautiously, but it was not always enough to avoid accidents (see Figure 8.20).

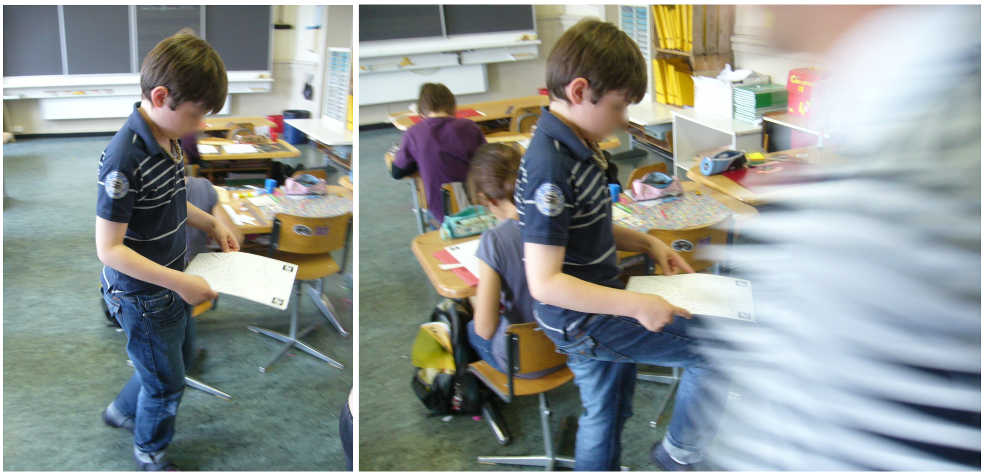

Figure 8.20: A pupil having difficulties bringing his assembly to the Metroscope.

One of the fifth graders had the idea of placing the pieces of paper on a thicker board (see Figure 8.21). She was quickly imitated by the other pupils of her class, because it was a lot more convenient.

The board did not solve the problem completely. There was a congestion around the Metroscope: up to four pupils, i.e. close to the half of the group, where queueing. This is due

Figure 8.21: A board helped carry the pieces of paper.

to the fact tat pupils had to rearrange the shapes that moved before recording or checking the arrangement. In the worst case, for example if the shapes had fallen down, the pupil completely redid the arrangement of shapes under the Metroscope.

Feedback The feedback given by the Metroscope on the arrangement of the shapes had a greater granularity in Triangram than in the other activities. Indeed, the relative position of the shapes could be set with more or less precision. The consequence is that the feedback ranged from green dots for correct solutions, to red dots for small imprecisions, to red lines for larger imprecisions (see Figure 8.22). While we accepted red dots as correct, it was disappointing for the pupils to not get green dots.

On a related topic, the feedback suffered from a technical issue that generated false negatives. If the arrangement of shapes was moved during the recording, the checking would display large red lines (see Figure 8.22). Pupils hardly identified them as false negative, despite the clear distinction with other, non aberrant pieces of feedback.

(Q_l) How do Pupils Learn with a Paper Interface?

Covering the Problem Space In general, the difficulty of the activity was adapted: the third graders made many mistake but completed the task, and the fifth graders easily performed it. Placing flat objects like the cardboard shapes used in Triangram allows to materialize various geometrical transformations. The expected one was of course the symmetry, but there were

Figure 8.22: On the right: the range of feedback. From bottom to top: green dots (correct), red points (small imprecision), a long red line (detection artefact) and red lines (bigger imprecision).

other examples, some wrong, and some not disturbing. For example, all the third graders started by implementing a translation, which is the main confusion related to symmetries. In contrast, some pupils recorded their arrangement upside down; it did not matter as long as the symmetric drawing was coherent with this rotation.

Another transformation allowed by tangible shapes spawned an interesting conflict. A pupil flipped some of the shapes. As it turned out, her partner made a typical mistake in the symmetric description of the arrangement he actually did a translation. The pupil was thus trying to undo a symmetry on a translation, which caused a problem with non isosceles triangles. The pupil knew that it was the wrong side of the shapes, because they did not have the StarrySheet pattern, but still had to flip them to complete the symmetry. Eventually she realized that her partner was wrong.

Conflict was not the only form of interaction between the pupils. Actually, the main one was the collaboration happening around the Metroscope. The pupils collectively corrected the pupil checking a solution, by making explicit the feedback given to the pupil checking an arrangement under the Metroscope.

(Q_t) **How do Teachers Use a Paper Interface?**

The teacher found that this activity was a good complement to the one that they designed with Kaleidoscope (see Section 7.3). The Kaleidoscope activity introduced symmetry, and Triangram deepened the notion with a problem. Unsurprisingly, like in Kaleidoscope, we saw that the teacher of the fourth grade, who had the opportunity to perform the activity twice, was more efficient in the explanation for the second time.

The teachers spent most of their time helping pupils rather than managing the lamp. They did intervene more at the lamp at the beginning, probably to get familiar with it. Afterwards, they focused more on making sure that the pupils preparing an exercise prepared it correctly, so that the partner could work it out afterwards.

On-site Adaptation The teacher of the fourth grade, who was the first of the day, also advised two adaptations of the activity for the following class. First, we gave up starting with only two shapes, because it took too much preparation to rebuild the sets of triangles: pupils inevitably mixed sets with their partners.

More importantly, the teacher advised us to let the third graders start with a pre-drawn arrangement, so as to let them understand the exercise before letting them do their own. It was of course a good advice, since the third graders were less familiar with symmetries. This adaptation was easily done, and illustrated once again the natural integration of a paper interface in schools. During the lunch break, he simply drew a sample exercise, and copied it with the photocopier of the school.

8.5 Messangles

8.5.1 Description of the Activity

Messangles is the last activity we developed with StarrySheet. It focuses on the use of the protractor. The goal is to reassemble a “shredded” sheet, in order to recompose a secret message. The pupils write a sentence on the back of a StarrySheet. Then, they split the sheet into triangles such that one side of the triangle is a part of a side of the sheet (see Figure 8.23).

The measures of these triangles should be different, such that they identify the triangles. Moreover, each side of the StarrySheet is either blue or yellow. This way, the triangles can be coded by a colour and an angle measure. Pupils use this code to write down the order of the triangles (see Figure 8.24). The cut-out triangles are exchanged, along with the sequence on the instruction sheet, with the ones of another pupil. The pattern does not need to be recorded, because everyone starts with the same StarrySheet.

When the triangles are reassembled, the pupils can check the order with the Metroscope. The instruction sheet has a checking marker, pre-recorded with the original StarrySheet. The

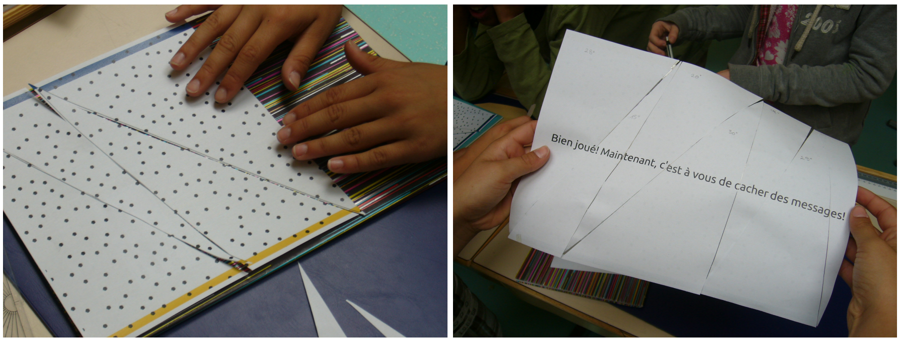

Figure 8.23: Assembling pieces of StarrySheets in Messangles.

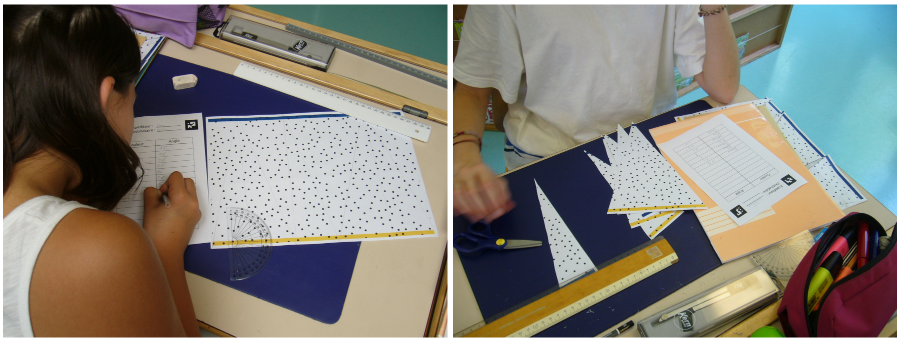

Figure 8.24: Creating a Messangles exercise.

instruction sheet thus triggers the feedback on whether the triangles are in the correct order. If it is correct, they can tape the triangles together, and flip the reassembled sheet to discover the sentence written on the back. The Metroscope allows to check the assembly without having to flip it, which is cumbersome when the pieces of paper are not taped together.

8.5.2 Study

We deployed this activity in School 2, during two 45-minutes periods. In the preceding period, the teacher had introduced the use of a protractor. 21 pupils were present that day. We used one Metroscope freely accessible for pupils to check their assembly of triangles. In the meantime, the teacher normally roamed the room to help pupils who needed.

The pupils started with the second part of the activity, i.e. finding the message on a sheet that had been cut into triangles. These triangles, and the order of their measures in which they had to be reassembled, had been prepared by the researchers. In the second period, the pupils went through the full exercise, i.e. writing a message on the back of a StarrySheet, cutting the sheet into triangles, measuring their angles and coding the sequence on an instruction sheet. The pupils then exchanged the triangles and the sequence of angles with a peer, and recomposed the sheet.

At the end of the activity, we asked the pupils to answer six questions on a blank sheet that we collected:

- Did you prefer this activity, or the activity on symmetries? (see Section 7.4)
- Did you prefer assembling a message, or creating your own?
- Did you use the Metroscope for first part?
- Did you use the Metroscope for second second part?
- What did you like?
- What did you not like?

(Q_u) How do Pupils Use a Paper Interface?

The interaction with the paper interface was in two modes, corresponding to whether it happened on a pupil's desk or on the Metroscope. We first report the "offline" use of the paper interface, i.e. paper alone, and then its use in conjunction with the Metroscope.

Strategies to Use the Paper We explained orally and demonstrated quickly the process to fulfil the task only once. Consequently, the pupils did not remember the exact process, or simply preferred to customize it. We thus saw various strategies to complete the task. For example, some pupils preferred measuring all the angles and then ordering them, while others alternated measuring and placing. We recommended to draw the construction lines of the triangles on the back of the StarrySheet, i.e. on the side of the message rather than on the side of the pattern. However, a significant proportion of pupils drew on the pattern side. When the triangles where correctly assembled, pupils not necessarily used our suggestion to tape triangles one by one with each other, but rather adopted the faster method of placing a long strip of tape on the sheet (see Figure 8.25).

These observations are another confirmation of the fact that paper is a well known medium, that can be easily appropriated by the user. This flexibility is an advantage compared to a more rigid script, be it a traditional exercise on paper, or a computerized exercise on mouse/screen/keyboard, where the learner has to follow the process defined by the designer of the activity.

Feedback from the Metroscope Judging from the comments made by the pupils around the Metroscope, the feedback given by the StarrySheet was helpful. At first, the pupils were asking how to interpret it, but if we answered them by asking what they thought, it was clear enough that green dots are a sign of correctness, and red lines were a negative feedback. To summarise, the pupils were able to autonomously correct themselves.

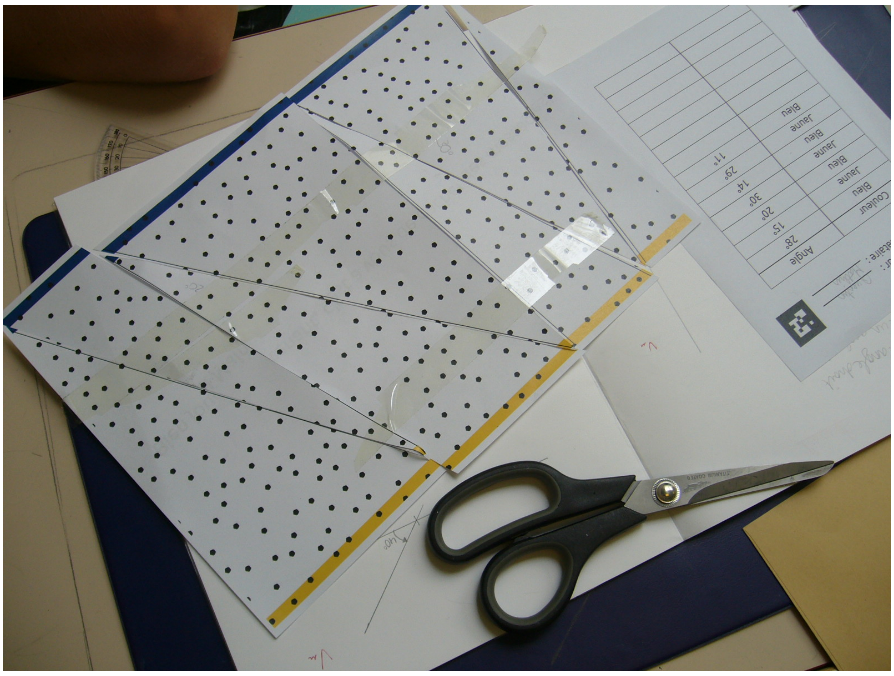

Figure 8.25: Two strips of tape fixing the assembly of triangles at once, which is faster than individually taping each pair of neighbouring triangles.

The major problem with the feedback was actually the fact that it is not compatible to multiple instances of the same StarrySheet. Since all the sheets where actually carrying the same pattern, the recognition would not work if several pupils where trying to check theirs answer at the same time. This limitation was not obvious to the pupils, and we needed to tell them to check again if they had a falsely negative feedback. In the second period however, the pupils were able to police themselves to avoid the presence of multiple StarrySheets under the Metroscope.

Emulating of the Metroscope with Paper The feedback given by the Metroscope was not the only way to check that the triangles were correctly assembled. First, the overall shape of the assembly gave a hint: if the sequence of triangles did not look like a rectangle, there was something wrong. This hint was not enough to do complete the exercise however. For example, the first pupil to use the Metroscope actually “cheated” by guessing the sequence based on the look of the assembly. As it turned out, her guess was wrong, because some triangles were similar enough that they could be switched in the sequence without making the assembly look odd. Of course, the feedback given by Metroscope was clearly negative.

All but one pupil wrote in the questionnaire that they had used the Metroscope to check their answer to the first exercise. This one pupil solved the exercise by reconstructing the message, without using the protractor. This behaviour is however easy to spot because the two sides of the sheets are very different, even from afar. Judging from the tone of the comment in the

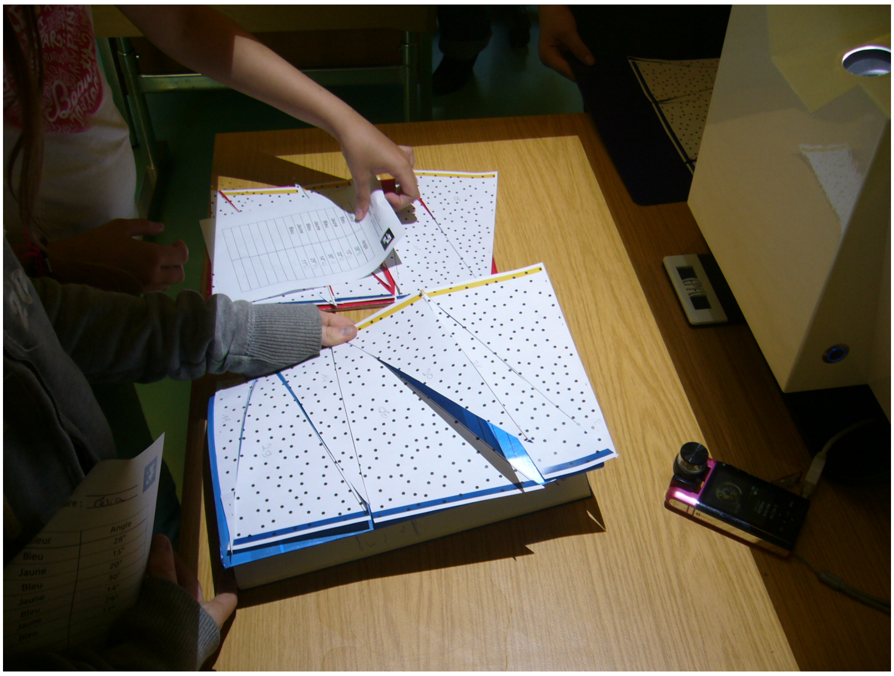

Figure 8.26: Two pupils showing their StarrySheet at the same time, resulting in a falsely negative feedback.

questionnaire of this pupil, it seems that his purpose was more a proof of concept that the protractor was not needed, rather than a cheating attempt.

(Q_l) How do Pupils Learn with a Paper Interface?

Dynamics in the Classroom Running Messangles revealed an interesting point, belonging to both the usability and learning aspects, illustrated in Figure 8.27. When measuring, the pupils are very focused, working at their desk. Cutting seems to be more stimulating, as several pupils stood up to do so. When the manipulative phase was at its highest (i.e. when the whole class was cutting or assembling cut-outs), the class was a lot more dynamic. However, it became calm again when the following activity required measuring.

Peer Correction This dynamics had at least one positive effect: it encouraged interaction between pupils. For example, when one of the pupils could not follow the instructions given by his partner, he could discuss directly with her to find and resolve the problem.

Complexity Messangles was designed as a problem using the protractor, not a simple drill exercise. In addition to measuring angles, pupils had to handle various parameters. First, the angles to measure were not a stereotypical representation on a figure (i.e. two intersecting lines). Instead, the pupils had to measure the angle embodied in a physical object, a paper

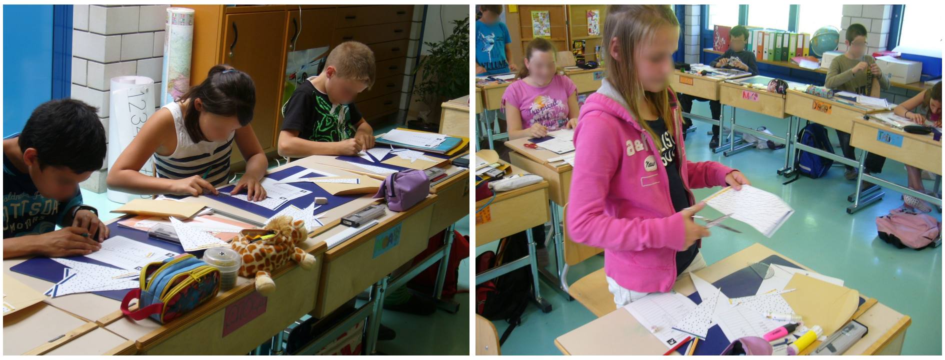

(a) First, everyone is focused.

(b) Pupils stand up for cutting.

(c) The classroom becomes very mobile.

(d) Back to the measuring part, everyone is sitting again.

Figure 8.27: Dynamics of the classroom.

triangle (see Figure 8.28). This was a good opportunity for pupils to deepen their understanding of the protractor. Most notably, not all of them realized where to align the centre of the protractor. It was thus helpful to have another representations of an angle for them to discriminate the invariants of an angle (i.e. an origin and two branches) versus artefacts (e.g. intersections, lines, or named points on a figure).

Second, the measures were not as precise as in the typical exercises. Indeed, the measure were not necessarily multiples of ten or five. Moreover, the precision when cutting a triangle was not perfect either; pupils had to account for this precision too, and were thus force to reduce their own measurement imprecision in order to reduce the overall imprecision.

Difficulty This complexity of the problem obviously has an impact on the difficulty of the activity. All the pupils completed the first part, i.e. reordering the angles according to the sequence of measures given on the instruction sheet. However, not all pupils managed to successfully rebuild the message coded by their partner, and some did not even finish coding their part. Indeed, the second part added another difficulty: drawing triangles with some constraints (the triangles should have different measures; they could not be quadrangles).

Figure 8.28: A pupil measuring the angle of a paper triangle.

Reception The difficulty may explains the reception of the activity by the pupils. The questionnaire also hinted toward a preference for the activities on symmetries (10 pupils preferred it versus 8 preferring Messangles, 3 had no preference). The link with the difficulty is clearer at the question regarding their preferred part: 15 preferred assembling an existing message, versus 5 preferring to create a new one (i.e. drawing the triangles and specifying the sequence of their measures). This is confirmed by the fact that a majority marked the coding of the triangles as one point they did not like in the activity, while they liked reading or writing the message.

(Q_t) How do Teachers Use a Paper Interface?

The teacher corroborated the preference of the pupils. She found the second part too difficult and confusing. She managed the activity herself, and appreciated the fact that it was both "playful and formative". She noted that the strong points of Messangles were the important focus resulting on the pupils and the possibility to use a protractor for both reading and creating angles. These strengths come from the *paper* aspect of the activity.

From the augmentation aspect of the activity, the teacher especially liked the ability for pupils to check their answers by themselves. She declared that she would not use a computer for such a task, showing that augmented paper is a good way to integrate in a classroom the processing power necessary to give feedback to pupils.

8.6 Discussion

The four activities presented previously use different properties of paper to interact with the computer. In AngleHunt, the edges of pieces of paper was modified persistently: small pieces of paper were assembled into a larger one. In Sympliage, the folds of paper was used to specify points. In Triangram, the position and orientation of paper relative to each other was used to code an arrangement. Messangles, like AngleHunt, was another example of edges being modified persistently, but in two ways: a large sheet of paper was actually cut into smaller pieces, and then reassembled. Table 8.1 places these activities in the framework.

Table 8.1: Position of activities based on StarrySheets in the framework

Property	Ephemeral	Persistent	Permanent
Presence			
Position	In Triangram, the relative position of pieces of StarrySheet defines the construction.		
Orientation	In Triangram, the orientation of pieces of StarrySheet defines the construction.		
Side			
Folds		In Sympliage, the point to find is verified by folding a corner onto it.	
Edges		In AngleHunt, pupils have to recompose an larger map from pieces of paper. In Messangles, pupils have to recompose a sheet from triangular pieces of paper.	
Ink		Pupils annotate the various pieces of paper in every activity.	

Besides the property of paper on which the interaction was based, the activities were relatively similar: pupils had to follow instructions expressed in geometrical terms in order to build a paper artefact. This artefact could have been designed by another pupil. The similarity between the activities allows us to compare the consequences of the property of paper used in an interaction.

8.6.1 Validation from the Metroscope

We start by the lowest level of interaction, which is the direct feedback on what the computer detected from the paper interface. The property of paper has indeed an influence on it. For example, in Triangram, the feedback was drawing red lines which length was proportional to the displacement of the shapes: a small imprecision would only result in small red dots. In Sympliage, a small deviation from the expected StarrySheet was amplified by the folding. The feedback was most extreme in Messangles: swapping a triangle with another one in the sequence would completely prevent the pattern to be recognized, and long red lines would appear everywhere.

These differences mattered to the pupils. For example, they were disappointed to have small red lines, even if it meant that they were close to green dots. However, the pupils did not have troubles interpreting the various feedback in the nominal cases: in the four activities, the videos showed that when the adults were gone, the pupils decoded the feedback correctly.

The reaction from the pupils to the technical limitations were more varied. While they easily understood that there was a significant time lag (more than one second), they had more difficulties understanding that only one pattern could be shown at the same time. Moreover, they often failed to realise that one or two long red lines among green dots were simply detection errors from the Metroscope.

In the end, what matters is whether the modifications of the properties of paper are discrete or continuous. In AngleHunt, the pieces of paper could be assembled on the continuum of their edges. The feedback was proportional to the error. In contrast, in Messangles, the position of the pieces of paper are discrete positions in a sequence. The feedback was almost binary: either all green, or all red. In between, Sympliage allowed foldings on a small, discrete grid. The feedback magnified small errors.

8.6.2 Validation without the Metroscope

The main value of the Metroscope is that it magnifies the mistakes of the pupils. Without the Metroscope, pupils were able to check approximatively their answers by comparing the obtained paper artefact to the expected one. For example in Sympliage, when the four corners had been folded, the sheet had a characteristic outline that could be compared to a template. In Messangles, it was easy to see whether the triangles were assembled into a rectangular sheet, as expected.

Such a validation without the Metroscope is acceptable for properties that are changed in a continuum, such as the relative position. In this case, the Metroscope is only used to show imprecisions. It is a bigger problem with discrete changes, as it is the case when ordering the triangles of Messangles. For example, the pupils who assembled a sheet from Messangles only by guessing were wrong on triangles that had similar measures. The resulting sheet looked correct, or maybe slightly irregular, which could be due to imprecision at the assembly as much as an error in the order of the triangles. However, the feedback given by the Metroscope was strongly negative: it consisted only of red lines. This is an example of how the Metroscope can amplify mistake that look small, but are actually important.

Being able to validate an paper artefact having a non figurative shape is another advantage of augmenting paper. For example, one could easily imagine to design a Sympliage that results in a heart-shaped sheet. However, this would give too much hints to the pupils, who may solve the problem with the guidance of the end results rather than the intended process.

However, the paper artefacts did keep a recognisable aspect in their terminal form. This is an advantage for the teachers to have a quick overview of the progress of a pupil, and for other pupils to review the work of their peers.

In summary, a paper interface allows to decide how mans hints will be intrinsically given during the manipulation, from an obvious construction to an abstract artefact.

8.6.3 Mobility

The Metroscope was a hot spot for impromptu peer reviews. Indeed, when a pupil brings a paper artefact to it, it is most likely because the artefact is ready to be checked, meaning that the pupil judges it correct. This implicit judgement can be challenged by surrounding peers.

In this perspective, the mobility of the paper interface is a plus, because it allows to go to the Metroscope and share reviews. Sympliage was the easiest interface to carry, because it was simply a folded sheet and an instruction sheet. The taped pieces of paper of Messangles were also easy to carry, but took a long time to prepare, because taping is not as fast as folding. The least mobile activity was Triangram, where pupils still had to readjust their solution when they were using a board to carry the arrangement of the cardboard shapes.

8.6.4 Appropriations

The use of the board in Triangram was not the only case where the pupils developed by themselves a technique to improve the usability of a paper interface.

In AngleHunt, a pupil improved the technique to draw a line across two pieces of paper. Most groups also manipulated the pen, paper and protractor with more than two hands.

In Messangles, the pupils made small variations from the process explained by the experimenter: they drew the construction line on the other side of the sheet, or made all the measures at once. More importantly, the pupils naturally came up with the idea of taping the whole sheet at once, rather than piece after piece.

The only example of appropriation that we observed in Sympliage was the pupil who used the transparency of paper to check that the starting point was matching on the two sheets.

8.6.5 Intuitiveness

Before appropriation comes intuitiveness. This is the last criteria where we observed differences between the activities. Messangles was almost too intuitive: pupils were tempted to assemble the sheet only on the appearance of the result. In contrast, Sympliage was overly unnatural: folding is not an established way to define a point. Moreover, the order of the fold was just an artefact for the exercise (what mattered was only the point), but it could create falsely negative validation. Triangram seemed intuitive enough, because placing pieces of thick paper is close to placing a card on a game table for example. AngleHunt was a little less intuitive, but did not disturb the pupils either.

8.6.6 Summary

Table 8.2 Summarizes this section. The relatively small sample of four activities shows a wide diversity. The conclusion is that the property of paper used to interact with a computer can only be used to taxonomize a paper interface. There is no inherent characteristics associated with these properties, because they can be used in many ways. It is a glimpse toward all the possibilities that paper offers as an interface.

Table 8.2: Overview of the comparison of the various AngleHunt, Sympliage, Triangram, and Messangles

Activity	AngleHunt	Sympliage	Triangram	Messangles
Property	edges	folds	position	edges
Granularity	continuous	fine	continuous	coarse
Offline Feedback	flip the map	compare edges	none	check edges
Mobility	?	++	-	+
Appropriations	multiple hands, drawing technique	transparency	board	taping, process
Intuitiveness	+/-	–	+	++

8.7 Conclusion

It is important to note that the four activities presented in this chapter have been developed in a day. Once the software component responsible for detecting the StarrySheets had been implemented, it was easy to use it to create activities. Most of the efforts in the creation of these activities were on the content: we tried to satisfy pupils with an engaging activity, teachers with a pedagogical activity, and ourselves with an activity that makes a sensible use of both paper and augmentations.

StarrySheets is thus a generic component that makes it easy to explore paper interfaces. It has been used to develop activities on two different topics (angles and symmetries). Actually, even the activities are generic, as they could be easily transformed to address a different topic. For example, the principle of Messangles can be reused to learn areas, if pieces of paper have to be ordered by area rather than angle.

While interesting as a proof of concept, we should not forget the original goal of our paper interfaces: geometry education in primary schools. The teachers do not necessarily benefit from the high level of genericity of StarrySheets. They have other worries than exploiting the possibilities of a system. And yet, the studies illustrated again the important of their contribution. The activities who did not involve teachers were poorly calibrated in terms of difficulty. For the study on Triangram, the first teacher made precious recommendations to improve the activity for the next part of the study.

StarrySheets did show how a paper interface makes it easy for a teacher to appropriate a computer augmented activity. In addition to the recommendations that we just mentioned, we saw during the study on AngleHunt how teachers easily helped pupils in their task. Indeed, helping a pupil with a paper interface is the same as helping a pupil on a regular, pen and paper exercise, as they are the same elements from the physical world.

Working on the physical world has another advantage: the activities were more *real*, i.e. not with ideal conditions. The problems are closer to real life problems. For example, Messangles showed to the pupils that a protractor is not necessarily used on a stereotypical angle of a geometrical figure, but can also measure a physical object such as a piece of paper. Also, information may be noisy in the real world; in Messangles, the pupils had to take into account that the measurements done by their peers may lack some precision. In Triangram, either pupil of a pair could have erroneously flipped a triangle, and they had to evaluate who was wrong.

We feel that the possibility offered by StarrySheets to pupils to create their own content have been under-exploited. It was possible for each activity, but really implemented only in two of them (Messangles and Triangram). This is due to the time constraints relevant to the experimentation in real condition. In a 45 minutes period, it is only possible to discover and perform an activity, or create and perform an activity, but not to discover, perform, create, perform again and discuss an activity with a peer.

Finally, StarrySheet also revealed the complementarity between various ways for a pupil to validate their work. The Metroscope gives a feedback to be interpreted. This interpretation can be assisted by peers waiting to use the Metroscope. In parallel, the teacher can provide a more adapted and reliable help, but only one pupil at a time. Lastly, there are methods to evaluate a solution by visually comparing it to the expected solution. This is thus another example of paper interfaces augmenting education, i.e. integrating the environment of the classroom while providing additional possibilities.

9 Findings

This dissertation explores paper interfaces for learning geometry at the primary school level. The previous five chapters reported on the exploration that was carried out with five different series of activities; each activity was used for several studies, in various conditions, and with more or less variations.

This chapter aims at summarizing and discussing our findings around the three original questions that initiated our exploration: How do pupils use a paper interface? How do they learn with it? How do teachers use a paper interface? For each of these questions, we aggregate the findings from the studies presented in the previous chapters and summarize them by presenting the lessons we learned.

9.1 How do Pupils Use a Paper Interface?

We first group and comment our findings regarding the *use* of a paper interface by pupils. To do so, we study how each basic property of pieces of paper are used by pupils: their presence, side, position, rotation, folds, edges, and ink. We then discuss how these properties relate to ephemeral, persistent, and permanent changes, and how they influence the practicalities of paper interface implementation.

9.1.1 Presence

Selecting Options by Presence

The presence of a piece of paper is the most basic property that can be used to build an interface. In Angoli, the presence of pieces of paper is used to select options. For example, the presence of the orange and blue control card select whether the angle built was clockwise or counter-clockwise. These options are not mutually exclusive: if both cards are present, both orientations can be built. However, the pupils used only one of them at a time, probably because it was too confusing to build two angles with the same origin at the same time.

Similarly, the presence of the measure cards was a way to select which angle had to be built. The pupils had ten cards corresponding to various orientations and measures in degrees. They placed in the interaction area the card corresponding to their goal when they were fulfilling it. It was thus an intuitive way to tell Metroscope which angle was being built.

In contrast, showing a paper sheet to select the corresponding laser in SpaceJunk was not as successful. In the first version of the activity, the pupils selected whether the shot originated from one of three positions on Earth by showing one of three paper sheets. There were two differences with the option selection from the two previous examples in Angoli. First, paper sheets were used instead of cards. In general, the selection of the laser was more cumbersome with fragile sheets than with cards. Second, the options are mutually exclusive: the shot can have only one origin. Since the sheets were also the referential for the firing angle, it was perceived as being part of the background, and many pupils forgot to actually set the option.

Paper interfaces can make use of the fact that paper is visible to users even if it is not detectable by the camera. For instance, in Angoli, the control card that was not being used (e.g. the clockwise card when building a counter-clockwise angle) was kept out of the interaction area to decrease the clutter.

Mapping Presence to Feedback

In most activities, the presence of a given piece of paper triggered the feedback (hereafter: the *trigger*), but each time with a variation. In order to summarize the set of designs, we propose three kinds of mapping between the presence and the display of the feedback: *continuous*, *segmented*, and *discrete*.

- A *continuous* mapping shows feedback as soon as the trigger is visible by the camera.
- A *segmented* mapping shows feedback when the trigger is visible by the camera and another condition is met. In other words, the trigger can be detected, but nothing changes on the display.
- A *discrete* mapping shows feedback as soon as the trigger is visible by the camera, but only for a limited time.

Kaleidoscope uses a *continuous* mapping between the presence and the display of the feedback. Pupils bring a sheet under the Metroscope, which implicitly triggers the display of the feedback, i.e. the symmetric image of their drawing. This feedback is then updated in real time. The checking marker of StarrySheets also implicitly trigger the display of the feedback, but with a delay, due to technical limitations. These lags (about one second) create a segmented mapping between the presence of the piece of paper and the feedback it triggers. In this case, the condition of the mapping is that a time is elapsed.

Quads is an example of a *segmented* mapping between presence and feedback. In Quads, the feedback is not implicitly triggered in the sense that it is not displayed all the time. Instead, a dedicated piece of paper, the validation card, must be shown. The feedback stays the same until the validation card is removed, the answer is changed, and the validation card is shown again.

SpaceJunk is an example of *discrete* mapping. The ammunition sheet triggers feedback in the form of a shot. It is thus another explicit trigger, however, it is linked to the moment when the piece of paper appears (i.e. when its state changes from absent to present) rather than the whole duration of the presence.

In our studies, the usability of feedback was clearly better in activities using a continuous mapping than in activities using a discrete mapping. Because the presence is a durable state, it is not adapted to triggering a discrete event. In SpaceJunk, using a continuous mapping to trigger a discrete event caused a lot of unwanted feedback, because the pupils did not realize the presence of the ammunition. However, after a while, they did notice that many ammunitions had been used up. The presence of these depleted resources was actually a factor of stress, because the pupils could see how many attempted shots had been made, and knew that they were limited.

Still, from a usability point of view, we noted that continuous mappings are poor when there is a lag between the detection time and the time when the represented information is displayed. In StarrySheet, the computer of the Metroscope was not powerful enough to update the feedback quickly enough. This was disturbing, more so than if we had used a discrete feedback.

However, in a pedagogical context, giving feedback is more than just a question of usability. Assisting the pupils too much will prevent them from thinking on their own and therefore learning. Often, it makes sense to delay the feedback from the action that triggers it in order to foster a reflection from the pupils that can result in learning. Some feedback can be immediate, though. For example, sensori-motor (e.g. showing where a pupil is pointing) or associative (e.g. giving the answer to 6 × 9) feedback does not need to be delayed. For higher level thinking, researchers debate when and how feedback should be delayed (see (Cuendet et al., 2012)).

In Quads, the feedback was delayed by forcing the pupils to place the card on top of the grouped shapes. However, the pupils tended to use the validation card as a continuous feedback, for example when they held it above the sheet, while moving the shapes underneath. This way, they bypassed the precaution that aimed at preventing intensive trial and error, as a continuous mapping between presence and feedback allows.

9.1.2 Side

In the second version of SpaceJunk, we used another property of pieces of paper to trigger the feedback: their side, controlled by flipping. Flipping is a discrete event in time. It was thus discrete actions such as this that were adapted to triggering a shot in SpaceJunk. Moreover, since flipping can be repeated, the shot could be triggered without limit. This relieved the pupils of the stress of limited remaining ammunitions. Flipping the card is also a better solution than the previous attempt at unlimited ammunition, which repeatedly triggered unwanted shots if the card was left unattended in the interaction area.

In Angoli, the feedback, i.e. the value of the built angle, is displayed when one side of a measure card is shown. The measure card of Angoli is always present while building the angle, but carries another piece of information: whether the feedback should be displayed or not. There is a subtle difference between this and the validation card of Quads. In Quads, the pupils tried to make the feedback more continuous, but in Angoli, even when the pupils were told to use the continuous feedback, they preferred switching it on and off.

There are thus two ways to use the sides of a card: either to alternate between two values that are linked semantically by a relation of complementarity, or as a way to specify an instant, corresponding to the moment when the side changes.

This shows how the side of a piece of paper is a property that combines characteristics of continuous and discrete mappings to feedback. A piece of paper can be flipped, which corresponds to the discrete event of switching feedback on or off. However, the possibility to control the feedback continuously remains present, because the piece of paper can be present in the interaction area in a state that does not influence the interaction.

9.1.3 Position

The position of the various pieces of paper was the most commonly used property to design our activities. We distinguish three kinds of positions: the absolute position, the position of a piece of paper relative to another one, and the stacking situation of pieces of paper, i.e. how they are stacked on each other.

Absolute Position

The absolute position is actually the position of a piece of paper relative to the interaction area. It is mostly interesting for its derivative, i.e. the movement. Our studies showed that the semantics of the interface influenced how the pupils moved the pieces of paper. For example, in SpaceJunk and Angoli, the movements were circular, reflecting how pupils think of angles. In Quads, the type of the paper artefacts determined how much they would move: cardboard shapes were the most mobile, followed by cards.

Pupils rarely move sheets intentionally. For example, in SpaceJunk, they were afraid to have to redo the measurement if they moved the sheet that showed the satellites and determined the orientation to give to the laser. They also complained about the instability of the satellite sheet, which did not contain enough markers to be detected reliably. In later iterations, we placed as many markers as we could.

Relative Position

Since sheets are stable, they can be used as a referential for the position of other pieces of paper. However, sheets do move at least a little, and if the relative position of another piece of paper needs a precise adjustment, they are too volatile. This can be frustrating for the pupils. This is why the second version of SpaceJunk had the laser projected at a fixed position on the interaction surface rather than on a sheet. Again, in Triangram, placing cardboard elements on a sheet proved to be cumbersome, especially when the pupils tried to bring the arrangement from their desk to the Metroscope.

A better use of the relative position is for determining whether two elements are close to each other. Typically, in Quads, when a tool card and a shape were close to each other, features determined by the card were displayed on the shape. Once again, the semantics of the interface component determined how the pupils brought elements close to each other. In Quads, the tool cards were functions applied to objects (the shapes). This was reflected by the fact that a card was in the neighbourhood of several shapes more often than a shape was in the neighbourhood of several cards.

We did not further investigate the topic of using the proximity to trigger events. There are many usability questions that can be asked: What is the right distance to consider two objects close? Does it depend on their size? Is the closeness uniform or biased toward a direction? Does it depend on the shape of the objects? Is it better to display a visualisation of the area considered close? When should it be displayed? How can overlaps be disambiguated?

We simply chose to trigger the feedback linked to the proximity of two objects whenever the proximity was detected. Thus, the pupils “felt” the neighbourhood area because the feedback appeared only in this area. We also allowed multiple kinds of feedback to be combined, leaving the pupils responsible for choosing which one they wanted, or leaving the possibility to combine them, as it happened with the tool cards in Quads.

Stacking Situation

The placement of pieces of paper on top of each other is another use of the position. In Triangram, the pieces of cardboard are placed on a sheet that only serves as a support, but in Quads, the position on the sheet actually defines an answer by the pupils.

A sheet of paper acts as a logical container for smaller pieces of paper. However, it is hard to exploit more complex vertical position for at least two reasons. First, it is technically not straightforward for a system like Metroscope to detect pieces of paper placed under others, as they become occluded. This would imply a system that could track various pieces of paper, similar to the one presented by Kim et al. (2004).

Second, not every vertical sequence is meaningful. For example, in Angoli, the stack of the measurement cards was aimed at determining the order of the angles to be built, assuming that the pupils popped up the cards one after another. However, the sequence of a deck of cards is not fixed, and many pupils reordered the cards as they saw fit. This is different from sheets for which the vertical sequence is more often bound by a leaflet. This distinction of the intrinsic order (or lack thereof) between sheets and cards can be used to cover different cases. In some cases, it is necessary to enforce the order of the exercises, for example, if the result of an exercise is needed for the next. In other cases, it can be useful to allow the pupils to make their own order, for example to cover a wider range of difficulty: higher performing students can address harder exercises, and others can focus on easier exercises.

9.1.4 Rotation

Pupils rotated pieces of paper for two reasons: either as part of the interaction, or as a natural manipulation of the paper. In the first case, the orientation is linked to a meaning, as in a rheostat. In the second case, the change of orientation is an artefact linked to how the protractor looks.

Rotation as Command

Using rotation as a command happened only in the second version of SpaceJunk, where pupils rotated the laser selection card to indicate the origin of the shot by orienting the card upwards, downwards to the right, or downwards to the left to select laser 1, 2 or 3, respectfully.

The interaction was comparable to the one based on flipping the shooting card: the selection card can always be present in the interaction zone, and used only when needed. However, orienting is different from flipping, because the orientation is a *continuous* map between the orientation and a value. It can also be transformed into a discrete mapping (like flipping) by firing only when the rotation goes beyond a certain threshold within a time frame (e.g. more than 120° within 1 second).

Rotation as Manipulation

The second case of rotation applied by the pupils on pieces of paper is part of their manipulation. For example, in Quads, when the tool card was brought close to a shape to display its characteristics, the card could also be rotated. Indeed, when manipulating a physical element, it is hard to separate translations (i.e. changing the position) from rotations. It is then natural that pupils rotate pieces of paper in a fashion similar to how they translate them.

Again, we observed that shapes are most likely to be rotated, because their orientation do not have meaning. Cards have an intrinsic orientation (due to the text printed on it), but pupils changed it occasionally. In contrast, sheets are as stable in rotations as they are in position. This was observed in both Quads and SpaceJunk. This is understandable for Quads, because the sheet had printed text on it, which provides this intrinsic orientation. This is more surprising in SpaceJunk, where the drawing could be seen in any orientation, and it would have actually helped the measurement to rotate the sheet. In this case, the pupils were (wrongly) afraid to have to redo their measurements if the system lost track of the sheet.

The stability of the orientation of sheets was also observed in the Kaleidoscope activity where pupils had to find the symmetry axis of paper shapes. While the Kaleidoscope sheet was a stable referential, the pupils intensively rotated the shape to match its outline with its symmetric projection.

9.1.5 Edges

Visual Guidance

The edges of pieces of paper are a very salient property, which visually guided the pupils. This had both positive and negative influences, depending on the activities. For example, in the study where pupils used Kaleidoscope to create cut-outs, the visual aspect of the result was easy to associate with the draft outlined with a pencil. This resulted in a wide variety of creations. In contrast, in Quads, the fact that two random quadrilaterals had complementary shapes misled the pupils into thinking that this was a characteristic to use to classify quadrilaterals. However, Quads is also an example of how the edges of a paper artefact are more natural than an outline printed on a sheet: no pupil mistook a square for a rhombus, for example, because a cut-out shape has no stereotypical orientation. Indeed, if a shape is printed on a sheet, the shape inherits the intrinsic orientation of the sheet.

The activities based on StarrySheets further illustrate how the visual guidance offered by the edges can be more or less positive. In Messangles, knowing that the triangles were to be assembled into a rectangular A4 sheet was used as a heuristic by the pupils, which replaced the measurements of the angles. In AngleHunt, the borders of the pieces of treasure map was a helpful guide, because the pupils knew that the pieces of paper where contiguous, and could focus on measuring the angle to position them. In Sympliage, the edges of the resulting folded sheet could be compared to a template, in order to quickly check, without the Metroscope, whether the solution of a pupil matched the expected solution. Finally, in Triangram, when two cardboard triangles were positioned contiguously, the task became harder for the pupil who had to recreate the arrangement, because the drawing looked like a quadrilateral rather than two triangles.

Neutral Edges

Strictly speaking, cards and sheets have edges, but they can be considered visually neutral. The rectangle is a common shape for paper artefacts, and it is practical to produce. For example, cards can be printed on thick paper and cut out with a guillotine. Thus, if a paper artefact is not meant to interfere visually, it can have a rectangular shape.

However, even if they are visually neutral, the edges of pieces of paper influence the manipulation of a paper interface. The most obvious example comes from the studies of the two versions of SpaceJunk. In the first version, the laser control card was used in conjunction with a sheet, to which the origin of the laser was attached. We saw that the movements of the control card avoided the borders of this sheet. In contrast, the movements were more circular when the sheet was removed, in the second version of SpaceJunk.

9.1.6 Folds

Our studies were not conclusive on the use of folding. In fact, several prototypes of interaction based on folding were rejected at the early stage of their development, because they were simply not usable. We explain this by the fact that only a limited set of folds make sense. For example, folding a sheet on one of its axes of symmetry makes sense. So does folding a border onto another one, or a corner onto another one. However, arbitrary axes of folding are not practical. Another limitation for using folding is technical: it is hard not to occlude too much paper, such that the detection of the paper remains reliable. Even with markers on both sides, the piece of paper becomes to small for a precise estimation. This prevents the use of the folds as a way to specify a line on the sheet whose equation parameters can be used to control a value, in a fashion similar to how the position defines two values (x, y) to the system.

Another obstacle we encountered when designing an interaction based on folding was related to the topic of the activity. It seemed natural to use folding as a way to enact symmetries. However, folding makes the augmentation redundant, since folding physically implements the rules of axial symmetries. This was illustrated with Kaleidoscope, when pupils created cut-outs. They used the augmentations to draft the shape, but for cutting out the symmetric edges, folding was enough, and it could be done without the Metroscope.

In Sympliage, folding was the origin of another usability caveat. We used the resulting shape of a sheet of paper whose corners were folded on specific points to check that these points were correct. However, for a given set of folding axes, there can be several resulting folds, depending on the sequence of the folds. This was disturbing for pupils who found the right point, but folded the sheet in the wrong order, and received negative feedback from the Metroscope.

9.1.7 Ink

Paper interfaces make use of three kinds of inks. A document can be printed with *mechanical* ink. The paper can be annotated by users with *manual* ink. The projection of the Metroscope is a third layer made of *virtual* ink. The three types are linked with each other in a cycle: humans read printed content, and annotate accordingly. Metroscope detects the annotations and projects virtual ink. In Kaleidoscope, we made a full cycle when we printed the mandalas designed by the pupils shown by the Metroscope.

In SpaceJunk, the pupils drew construction lines on the satellite view to measure angles. Then, they wrote their measurements to pass them on to their classmate. In all the activities based on StarrySheets, the pupils were encouraged to annotate the various elements of the interface as they progressed in the task. For example, in Messangles, the measures of the angles were written down so that the triangles could be ordered later. In AngleHunt, the construction lines of the angle helped place the pieces of map relative to each other. In Triangram, the arrangement of triangles was made persistent using a sheet. In Sympliage, the iterative reflection of the starting point was marked before finding the endpoint onto which a corner had to be folded.

These examples of annotation are not specific to Metroscope. In contrast, Kaleidoscope showed an example of the integration of ink between paper and technology. Ink was taken into account by the Metroscope, and projected back in real time to illustrate a symmetric transformation. In this case, it is the variation of the ink property that drives the interaction, much like the variations of the position property that were used to interact in Angoli, for example. This is not possible without the Metroscope, where only the end result of a drawing process can be printed, and not its progressive construction.

9.1.8 Ephemeral and Persistent Properties

We just reviewed one dimension of our framework in the previous paragraph: the properties of paper that can be used to interact with the computer. We gave positive and negative examples of when to use some properties of paper, and how. However, these examples are relatively sparse in the design space of paper interface. Our intention was to base the interaction of paper with Metroscope on the interaction with paper without Metroscope. Hereafter, we cover the second dimension, i.e. the duration of the interaction.

In Chapter 3, we defined three ways to modify the properties of paper artefacts: permanently, persistently, and ephemerally. We gave examples of permanent modifications of the paper artefacts: the design of each activity defines these properties and can not be modified. However, our activities do not make an even use of the properties in the two other types of interactions.

Some properties are hardly modified ephemerally. Folding leaves marks on the paper. Ink is mostly persistent, except for pencils and special erasable pens, but even then, they leave a trace. Even in Kaleidoscope, where we described the apparition of the drawing as an ephemeral use of the ink, a persistent trace was left. Similarly, we mostly changed the other properties ephemerally: presence, side, position, and rotation are easily used during the interaction with the Metroscope, in a fashion similar to a Tangible User Interface. The edges are somewhat intermediate: cutting is a persistent change, but assembling two shapes by placing them contiguously is an ephemeral transformation.

This dichotomy can explain why the pupils preferred the drawing mode of the controlled study using Kaleidoscope to the taping mode: the taping mode showed that it was possible to persistently change the presence, position, orientation, and side of pieces of paper, but it was awkward.

Interestingly, we mostly used the ephemeral properties with artefacts made of thick paper (e.g. the position of cards and cardboard shapes), and the persistent properties with regular paper. This is simply a consequence of the fact that in the former case, paper is used for its tangible properties, and its manipulation is easier with thick paper because it can be grabbed more easily. In the latter case, paper is used for its ease of production and low cost, which allows for disposable interface elements.

The resulting rule of thumb is that interfaces making use of the ephemeral group of properties (classes III and IV in Chapter 1) are best implemented by small pieces of cardboard or thick paper, while other interfaces are best implemented by regular sheets of paper, or sheets that have undergone a simple transformation, so that they remain easily producible and can be altered without second thoughts.

9.1.9 Creative Appropriation

The goal of using properties in a familiar fashion is to obtain what we called *creative appropriation*, i.e. uses of the paper interface that were not designed for, but facilitated the use or the learning. The simplest example is derived from the manipulation of the paper interfaces. For example, in Quads, pupils displayed several characteristics of one cardboard shape by combining the control cards into a test bench.

StarrySheets was another domain where manipulations led to interactions not intended at the time of the design of the activity. In AngleHunt, several pupils manipulated the pieces of map with multiple hands, and a pupil improved the drawing technique shown originally by the experimenter (see Section 8.2.2 for more details). In Sympliage, pupils used the transparency of paper to check the alignment of starting points in exercises that they designed. In Triangram, the pupils overcame the difficulty of carrying the arrangement of triangles by placing them on a board.

Pupils were also able to change the pedagogical script itself. For example, in Quads, one pupil held the validation card above the exercise sheet to bypass the measure that aimed at preventing the shape to be moved while the validation was displayed. This pupil effectively transformed the discrete feedback into a continuous one. In contrast, most pupils using Angoli transformed the continuous feedback into a discrete one, by flipping the measure card only when they were ready to check their answer.

Finally, the most impressive creative appropriation happened during the development of the controlled study using Kaleidoscope. We had invited a teacher and her children to pre-test the activity. After a first run, during the debriefing with the teacher, her children went on playing with Kaleidoscope. One of them, Mia, figured out by herself that she could change the orientation of the axis by simply cutting out the fiducial marker of a vertical-axis sheet, and placing it next to a another sheet where a horizontal axis was expected (see Figure 9.1). This showed how the user of a paper interface can appropriate not only the manipulation or the interaction, but also the interface itself.

Figure 9.1: The triangles were placed assuming a vertical axis and the sheet was configured for a horizontal reflection. To fix this, Mia cut out the fiducial markers to align the expected and projected axes.

9.2 How do Pupils Learn with a Paper Interface?

We observed that paper interfaces were effective learning tools based on two main points: the exploration of a problem, and the collaboration between pupils. We explain how, and conclude by discussing a principle for the pedagogical design of activities based on paper interfaces.

9.2.1 Collaboration

Collaboration on Tangible Tabletop Systems

A paper interface shares two intrinsic benefits for collaboration common to tangible tabletop systems. First, as noted by Dillenbourg and Evans (2011), since the display surface is horizontal, the projected information is more *public* than if it was on a vertical screen. Indeed, the surface can be observed from any angle around it, while screens are typically limited to the few users directly in front of it.

Second, as the interface elements are tangible, they can easily be integrated in the argumentation between pupils. For example, in Angoli, we observed several examples of pupils explaining the difference between the orange control card (for clockwise angles) and the blue control card (for counter-clockwise angles) to the rest of the group. The explanations were supported by the pupils pointing at the graduations on the protractor. Supporting such argumentations is important, according to Collaborative Learning theories, because the effort needed to explain something is often internalized into a better understanding of the problem.

Scattered Interface

Tangible interfaces are often *scattered* (see (Cuendet et al., 2011)), which means that the controls and the displays are attached to multiple objects that are physically separated from each other. This is especially true for paper interfaces. Indeed, a piece of paper can represent an object of interest; another piece of paper can represent a function; another one represents an global option; etc. Table 9.1 shows how many elements of different types composed our activities. It does not mention non-augmented elements of the activities, such as pens, protractors, or tape.

Table 9.1: Scattering of the interfaces for the various activities.

Activity	Interface Elements
Quads	1 validation card, 3 tool cards, 20 cardboard shapes, 6 exercise sheets
Angoli	2 control cards, 10 measurement cards
SpaceJunk 1	1 satellite sheet, 20 ammunitions, 3 laser sheets, 1 control card
SpaceJunk 2	1 satellite sheet, 1 laser selection card, 1 laser control card, 1 laser trigger card, 1 communication sheet
Kaleidoscope 1	3-5 page booklet, 1 blank Kaleidoscope sheet, 1 cardboard shape
Kaleidoscope 2	4 exercise sheets, 6 cut-out paper shapes, 3 blank Kaleidoscope sheets
Kaleidoscope 3	4 page booklet, 5-15 paper triangles
Kaleidoscope 4	3 blank Kaleidoscope sheets
AngleHunt	8 pieces of a map, 1 instruction sheet
Sympliage	1 StarrySheet, 4 instruction sheets
Triangram	6 cardboard triangles, 1 support sheet, 1 instruction sheet
Messangles	5-10 pieces of a Kaleidoscope sheet, 1 instruction sheet

Interfaces are not only scattered in space but also in time. In addition to being divided into different pieces, the pieces are not always used all the time, or at the same time. Some elements, like the exercise sheets of Quads, are mutually exclusive. Others, like the feedback triggers, often remain in the interaction area, but are actioned less often.

One consequence of the scattering of the interface is that groups have a visualisation of their work. For example, in the first version of SpaceJunk, the pieces of paper representing ammunitions showed how many attempts had been made. When the observers, who gave the ammunitions to the controllers, thought that this resource was disappearing too fast, they told the controllers to pay more attention. This is an example of a public visualisation of resources of the group. AngleHunt gave an example of how a scattered interface materializes progress: the more pieces of map that were stuck together, the more advanced in the exercise. Angoli gave an example of the visualisation of the roles in the group: the pupils built the angles one after another, so the manipulating pupil was the one with the control card.

Furthermore, in Angoli, another pupil usually assisted the manipulating pupil, by giving or flipping the measurement card. This exchange of pieces of the interface was particularly interesting in Quads: the physical layout was such that the horizontal movements of the interface pieces showed how many exchanges between pupils happened. The layout of the exercise sheets was such that the vertical movements of the pieces of interface showed the progress in the task. Both of these features were correlated to the performance of the groups.

Interface Scattered in the Classroom

Paper interfaces can be used without augmentation, which supports their scattering. While the interest of a general tangible interface is limited without the augmenting device, the annotability of paper can be exploited to continue the activity without the Metroscope. Thus, the interface can be scattered outside of the interaction area, in our case, in the classroom.

The pupils could work individually at their desk, but we saw that scattered interfaces also triggered collaboration at the class level. Even though the first activity using Kaleidoscope was designed as individual exercises, the pupils gave feedback to each other when using the Metroscope. Again, this triggered argumentation between peers that are beneficial for the understanding of the problem.

The impromptu collaboration around the Metroscope happened because of the queue formed by the pupils waiting for their feedback. This is another advantage of sharing the Metroscope among the class, as shown in the studies for Kaleidoscope, Sympliage, Triangram, and Mesangles. The only drawback would be that pupils would wait too long for their feedback and lose time. The congestion was linked to the difficulty of the task. Only Sympliage was too hard, and resulted in an unproductively long queue. In the other cases, the pupils demonstrated a self-enforced, fair use of the Metroscope.

Finally, it is worth noting that paper interfaces can be scattered beyond the classroom, e.g. by bringing home an exercise to show it to the parents. However, we did not investigate this aspect.

9.2.2 Exploration

We saw that paper supports the exploration of a problem: in the pilot study for Quads, the groups came up with various groupings of quadrilaterals and different justifications very quickly. SpaceJunk showed how pupils could use seven degrees of freedom to describe a simple angle.

We see two reasons for paper interfaces to support exploration, which are in themselves defined by paper interfaces. First, they allow for the control of various computer-generated visualisations. Second, this control happens with intuitive manipulations, which are then continued naturally by an appropriation by the pupils. We detail these three points (feedback, manipulations, and appropriations) in this section.

Feedback

The activities demonstrated a variety of feedback. Quads and SpaceJunk had the simplest form: a binary validation, i.e. correct or not. In the case of Quads, an explanation accompanied the feedback, but the experimenters had to remind the pupils to read it; pupils cared more about the correctness of the feedback. Similarly, SpaceJunk offered a short visualisation of the shot, but pupils did not really pay attention to it or notice it. Again, they simply cared about the result. This is symptomatic of the trial and error strategy, which is the most popular among the pupils. They may think and discuss a priori about what to do, and want to know if it worked, without necessarily trying to understand why. An activity like SpaceJunk is thus more adapted to this kind of feedback, because the feedback is not directly given to the pupils who just acted. In this case, trial and error strategies fail, because they cannot see their errors. Eventually, the pupils are forced to reflect and discuss what their error was.

In Angoli, in addition to telling whether a solution was right or wrong, the feedback told whether it was exactly or approximately correct, and could display, in real time, the exact value of the correct solution. This had an impact on the engagement of the pupils, as they were often tuning an approximately correct solution into an exact one by using the continuous feedback. However, the biggest gain in engagement came from the discrete feedback: once again, the pupils preferred focusing on constructing a solution (in this case: building an angle) and checking it, rather than planning and reflecting. However, in our study, we saw that the feedback had a positive effect on the learning gain only for the class where the pupils were already familiar with the protractor.

Engagement resulting from the feedback is not necessarily playful. In fact, Quads showed that the use of the augmentations triggered by the tool cards were not random; the pupils were requesting helpful feedback rather than requesting just any feedback for fun. Indeed, the tool card revealing the characteristics useful for the current exercise (e.g. the tool card showing angles for the rectangles) was used more than the other tool cards. Similarly, the controlled study involving Kaleidoscope showed that the Metroscope was used more often for tougher exercises than for easier ones.

In Kaleidoscope, the feedback acted more like a mirror literally, but also figuratively: the feedback did not tell what is right or wrong, rather this was left to the observer to judge for himself. However, as noted by the teachers, the feedback was more precise than the one obtained with a regular mirror: it let the pupils explore the problem of symmetry at a finer granularity than what is possible without the Metroscope. Both validations are not mutually exclusive: we saw pupils using mirrors to quickly test their solution, before validating it at the Metroscope.

StarrySheet used abstract feedback. The interpretation of the feedback is left to the observers, but the feedback itself is not figurative. While some pupils expressed the wish to have more precise feedback instead (i.e. correct or not), most of them did not have problems deciding whether the feedback was positive or negative. Furthermore, as with Kaleidoscope, StarrySheets supports preliminary checks not involving the Metroscope. For example, In Messangles, the pupils could check whether the overall shape of the reconstructed sheet was not aberrant, and in AngleHunt, they could check that the paths drawn on the map matched.

Manipulation

An interesting feature of paper interfaces is that the feedback can be decoupled from the actions. For example, in SpaceJunk, the actions of the controllers are only seen by the observers. This mechanism can be the basis of so called jigsaw pedagogical scripts, where the information necessary to solve a task is scattered among the members of a group. The group then has to collaborate to recompose the original information.

Moreover, in Quads, the objects of interest (the cardboard shapes) are separate from the validation, even though it would be possible to continuously display whether a classification is correct or not. This highlights the tension between designing a user interface and a learning interface. A user interface aims at facilitating a task. However, a learning interface may be cumbersome by design, because eliminating difficulty from a task might be detrimental for learning. SpaceJunk is another example: while it would be possible to allow for continuous shots from the lasers for increased effectiveness, the activity would simply turn into a game where communicating angles, the central pedagogical idea, becomes useless.

It is also interesting to note that a given geometrical concept is not necessarily enacted by the most similar manipulation. For example, angles are more easily enacted by pupils by moving a control card in a polar referential rather than by rotating it. Symmetry can be enacted by folding, but this takes away gains from the Metroscope; instead, it is more interesting to use drawings. On a related note, manipulation is not necessarily more intuitive than writing, as we concluded from the controlled study on Kaleidoscope.

However, manipulation is an interesting feature to gain insights on learning, as seen in Quads. The layout of the exercise sheets allows for the interpretation of horizontal movements that have meaning for collaboration and vertical movements that have meaning in the solution space. More transitions between one area of the sheet and another correspond to more attempts. More generally, we observed increased manipulations when the tasks became more difficult. This implies the ability of paper interfaces to capture the manipulation of elements of the physical world with which the pupils are learning.

The manipulation of paper interfaces is also compatible with regular tools, as seen with the protractor in Angoli, SpaceJunk, AngleHunt, and Messangles. This integration is a real advantage in the learning context, as the whole complexity of the tool can be addressed: for example, in Angoli and SpaceJunk, the double graduation was confusing, and Messangles highlighted the need for a high precision in the measurement of the angles.

9.2.3 Splitting as Pedagogical Design Principle

Dillenbourg and Hong (2008) explained how splitting the collaboration between pupils allowed for the control of the focus the interaction. We implemented this principle in SpaceJunk. The scattered nature of paper interfaces makes it very easy to split the collaboration, but this property goes even further. We hereafter discuss how paper interfaces can be split in many more aspects, and how they allow a fine control of pedagogical activities.

For example, splitting a paper interface can support division of labour. Splitting a paper interface into several components allows pupils to use the separate pieces individually to support their explanations to the rest of the group by showing a precise object (e.g. the measure in Angoli). This also allows two pupils to share the manipulation in tricky processes (e.g. drawing on two pieces of paper in AngleHunt). The various elements of the interface allow the pupils to visualise the state of a resource (e.g. the ammunitions in SpaceJunk), or the progress (e.g. the exercise sheet in Quads). It also allows the pupils to measure the coordination or the difficulty that the groups experienced (e.g. the indicators in Quads) a posteriori. It also allows them to structure the activity in space (e.g. the two areas of the sheets in Quads).

The components of the interface can be used at separate times, which allows for the implementation of a scenario (e.g. the sequence of exercise sheets in Quads). This also allows for the control of the frequency with which an element can be used (e.g. the cool-off period in SpaceJunk) or the duration (e.g. the expiration time in Kaleidoscope). The use of a scarce resource can be distributed among a classroom (e.g. in Kaleidoscope). This temporal segregation can also have the effect of spatially grouping pupils at the same stage of the activity.

The paper interface can be split in a way that it separates a feedback and its triggering event. For example, in Quads, the feedback could be displayed all the time, but a separate card triggered it to limit trial and error. Other cards can determine the level of details given to the feedback. This can also control the engagement, where feedback can be more or less playful. The feedback can also be decoupled from the action for more natural enactments (e.g. circular movements in Angoli rather than rotations).

Finally, if the interface is split into simple, modular components, these pieces can be recombined and adopted, as shown with the several examples of creative appropriation.

To put it in a nutshell: any pedagogical object can be controlled by splitting it in a paper interface.

9.3 How do Teachers Use a Paper Interface?

Teachers can be involved with a paper interface during the design of the activity, and during its use in a classroom. We distinguish five phases in this involvement:

1. Definition of the features, e.g. a sheet that reflects one part over another, or a card that controls an angle,
2. Design of the activity, i.e. the specifics of how the pupils will use the features defined in the previous step,
3. Last minute adjustments, i.e. changing the activity in order to adapt to timely constraints,
4. Management of the activity, which may include a direct intervention from the teacher,
5. Follow-up of the activity, to integrate it into the normal course of the teaching.

We think that paper has the potential to boost the integration of digital technologies in schools. In this perspective, it is not satisfying to have the teachers simply use the interface once during a study, or simply give input in a co-design session. Instead, we aim for a deeper adoption and appropriation. In this section, we show how this happened for each phase. Kaleidoscope will be the most discussed activity, because it is the activity in which the teachers were the most involved. Other activities will, however, provide additional data points.

9.3.1 Appropriation of the Functionalities

We did not investigate extensively how teachers could appropriate the feature set of an activity. The main reason is that we needed to control our exploration of the paper interface. In addition to the pedagogical goals that the teachers were clearly able to address, we were also bound to interaction constraints, such as which property of paper to use (e.g. folding or the ink content), or what characteristic to highlight (e.g. the scattering of the interface among the class).

However, the teachers did not lack inspiration as to how to use the paper interface, as revealed by the two design sessions of the activities based on Kaleidoscope. The teacher of School 4 imagined using pre-printed shapes on paper cut-outs with a Kaleidoscope sheet (see Figure 9.2). We, on the other hand, only intended to use content printed or drawn directly on the Kaleidoscope sheet. This allowed the pupils to manipulate the shapes. More precisely, the pupils could rotate the shapes and see them in many orientations, not only the stereotypical ones (e.g. a square having sides perpendicular to the edges of the sheet). This rotation helped the pupils uncover all the axes of symmetry of the shapes; otherwise, pupils tended to see mostly vertical or horizontal symmetries.

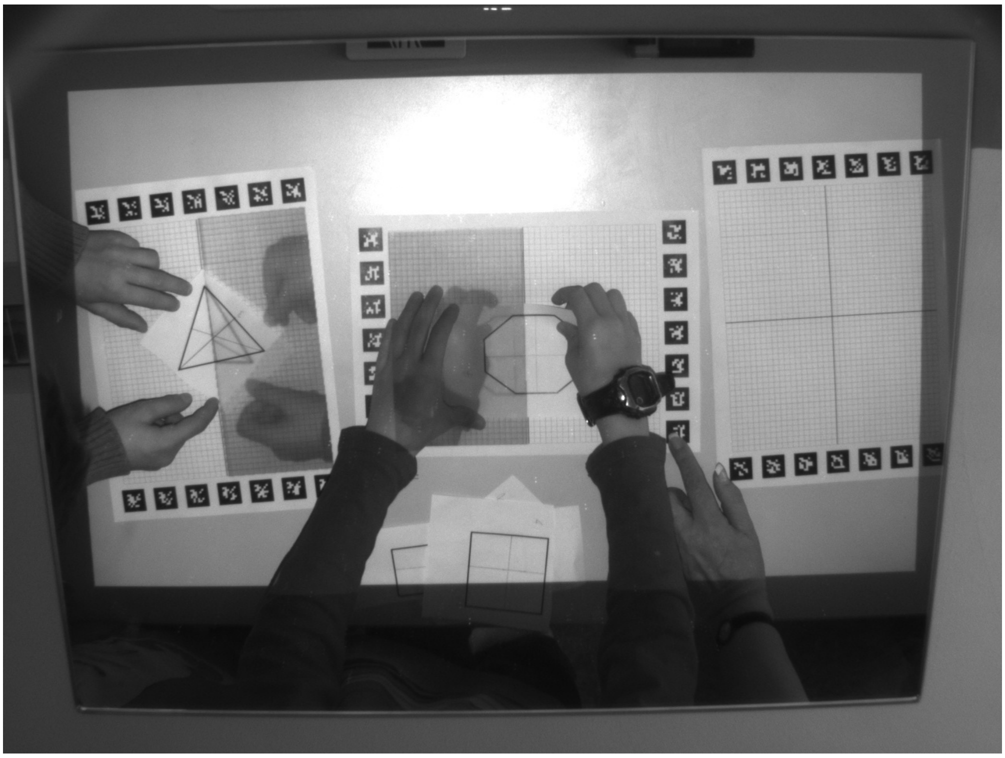

Figure 9.2: An overview of pupils manipulating outlines placed on paper cut-outs. They align the symmetry axis they find on the outlines with the symmetry axis of the symmetry sheets. If the projection of the outline matches their outline, their axis is correct.

Additionally, in a correspondence following the Kaleidoscope study, the teacher proposed a mechanism to limit the number of students able to access the Metroscope as an improvement to the activity. This is similar to the last activity using Kaleidoscope (for prototyping folded cut-outs), where we set an expiration time for the augmentation of each sheet. We did so for practical reasons (sharing one Metroscope among the class), but the teacher expected such a limitation to help the pupils "structure" their work, i.e. she saw the pedagogical reason first.

In the design session in School 3, one of the two teachers requested an animation to show the construction of the symmetric projection of a shape. This is indeed a good way to introduce symmetry which makes use of the dynamic display offered by a computer. Actually, one of his comments when he first saw the Metroscope was that such a system could help with the visualisation of dynamic concepts, such as a geometric construction.

During the same design session, the other teacher naturally assumed that a paper interface was part of the classical cycle of paper: printing, annotating, scanning, processing, and printing again. He asked whether it would be possible to print the mandala created by the augmentation of what the pupils had drawn, but having used normal paper before, he already knew this was possible.

It is important to note however that these suggestions happened during the design session for Kaleidoscope. In other words, we came up with a core functionality that could be instantiated and extended. In contrast, Angoli and SpaceJunk are the result of a design session of the core functionality itself. The resulting activities were very different from what was discussed with the teachers, because the discussion remained too abstract and hypothetical.

Our conclusion is that it is more productive to bring a working, adaptable paper interface for experienced users to improve, rather than try to co-design it from scratch. Paper is naturally adapted for this process: it is a concrete object, and thus requires lower cognitive effort to visualize its use. Furthermore, modifications of the paper interface can be drafted even without technical knowledge as to how they can be implemented. Paper interfaces are thus a natural middle ground to support a collaboration on their own design.

9.3.2 Instantiation of the Activity

Once the functionalities of the paper interface are set, they need to be instantiated into an activity for a specific class. In this perspective, Kaleidoscope showed the importance of an editable component. It allows for the accommodation of the idiosyncrasies of the teachers and their class.

Obviously, not all teachers have the same preferences in terms of instructional design. We already illustrated this in Kaleidoscope: one teacher used the Metroscope only for validation, while the other teachers had the pupils draw with continuous feedback. One teacher preferred full oral explanations for more flexibility, while the others had the instructions written on each sheet, as a visual support to their oral explanations. One teacher reused official content from the curriculum (for the symmetric patterns to draw), while the others defined their own patterns.

Every class has its own set of characteristics (the general level of performance, the subjects that have been addressed in the previous grades, the discipline, the focus, the social ambiance, the heterogeneity of the pupils, etc.). It is thus important to instantiate an activity for a class. Of course, this task is best done by teachers who know their class.

The teacher from School 4 gave an explicit example of such an instantiation for Kaleidoscope. She said that her class had a heterogeneous level of performance, which led her to create two versions of each exercise, so that the pupils could choose whether they wanted the more challenging one or not. We do not doubt that many more customizations were implicit. For example, the fact that the teachers from School 3 could have each pupil draw directly under the Metroscope is due to the schedule of the class had a slot allocated for group activities; one of the classes was of mixed level students, so not all pupils could perform exactly the same exercises.

More importantly, such instantiation was doable without any technical knowledge about the Metroscope. In fact, we found the fact that the researcher did not disturb the process very profitable. The teachers from School 3 designed the whole activity by themselves; we only transcribed the exercises into digital versions. This involvement from us was not even necessary: the teachers might as well have edited blank Kaleidoscope Sheets with a pen, and photocopied them. In fact, this is what we did with AngleHunt: we created the map manually. Teachers in general are obviously not computer illiterate either, as the teacher from School 4 demonstrated: she produced the whole activity by herself.

Our conclusion is that paper interfaces have the strength of being highly customizable (in the sense of tweaking final parameters) in addition to being easily designable (in the sense of core features, as discussed previously). This is clearly an advantage that allows for the easy integration of paper interfaces into classrooms, as the specific needs of both the teacher and the pupils are respected.

9.3.3 Just In Time Adaptations

The ease of adaptation is not only useful when planning an activity. It also allows for last minute changes that address unplanned concerns, which are bound to happen in real conditions. We involuntarily experienced such concerns on two occasions: one set of cards was missing for the controlled study of Angoli, and a set of exercise sheets was missing for the controlled study of Kaleidoscope (in both cases, we miscalculated the number of necessary items). Since we had no access to a printer, we simply used the photocopy machine at the school to duplicate the missing items.

Another example of last minute adaptation happened in the study related to Triangram. During the debriefing, the teacher advised us to give an example of a self-created exercise to the third graders before letting them create their own exercise. Again, this was simply addressed with the photocopy machine of the school, after we created a simple exercise ourselves. The teacher also advised us to reduce the number of triangles for the third graders, which was also an easy last minute preparation.

9.3.4 Orchestration

We saw different kinds of involvement by the teacher during the activity itself. Unsurprisingly, the teacher had the most ownership of the activity in the case where they designed it themselves: during the studies involving Kaleidoscope, the researchers only sat back and observed. The strongest ownership was demonstrated during the study based on Kaleidoscope in School 4: the teacher monitored the use of the Metroscope, roamed the classroom to help students at their desk, and eventually brought pupils from their desk to the Metroscope to incorporate a live demonstration into the explanations.

The ownership of the activity is a continuous process that improves at each run. The clearest example of this is the second run by the fourth grade teacher of School 3 of Kaleidoscope: the teacher directed the pupils to write their names on the symmetric sheet to see what happens, because the last group of the first run did so, and this was an interesting detail to discuss with the whole class.

The teachers also appropriated activities that they did not design, in several ways. The controlled study based on Kaleidoscope was the only other example where the teacher used the Metroscope to help the pupils. This may be linked to the fact that the Metroscope gave another representation of the topic of the lesson which could be used as an illustration by the teacher. In contrast, the other activities used the feedback of the Metroscope as a correction, which can be done by the teacher.

In the other activities which were run with the teacher present in the room, it was very natural for the teacher to help the pupils, even without the Metroscope. In Sympliage and Triangram, the teachers helped the pupils at their desk. In Sympliage, the excessive difficulty made the intervention of the teacher mandatory (and intensive). In Sympliage, the teachers explained what each pupil was doing wrong before the pupil checked his/her answer at the Metroscope. In Messangles, the teacher actually took part in the activity herself, as she found it “formative and playful at the same time”.

We also observed the appropriation of the activity with teachers with whom we had not previously collaborated. The AngleHunt and the SpaceJunk studies, which happened in the context of the open door day workshops, involved classes and teachers who had never seen the Metroscope. Even in this case, it was only natural for the teachers to help the pupils with the manipulation of the protractor without us even inviting them to do so.

We think that the easy appropriation during the activity is mainly due to the fact that we designed the activities to be *augmented by* the Metroscope, not *based on* it. None of these paper interfaces is completely meaningless without the Metroscope, in the sense that there is always a way to use it in a degraded mode. In the worst case, the validation can be done by the teacher. Conversely, the teacher can always validate the work of the pupil, and still maintains control over the activity.

9.3.5 Integration

The final step of appropriation of the paper interface by a teacher is its integration in the teaching beyond its use with the Metroscope, i.e. before and after the studies.

In a sense, an integration *during* the studies also happened. It consisted of finding a way to keep part of the class busy, so that the other part could use the Metroscope more intensively. We saw three strategies: in one case, the class was already split into two levels, in another case, half of the class was doing another activity with a different teacher, and in the last case, the pupils were distributed a drill exercise sheet as a time filler when the Metroscope was not available.

Before the study, the activity needed to be introduced. For new concepts like symmetry, the teacher needed a few teaching periods. Since at most two periods per week are allocated to geometry, this implied a few weeks between the design and the study. For a lesson about the use of the protractor, since angles had already been handled previously, one period was enough to introduce the use of a protractor, especially if the activity involving the Metroscope aimed at completing the discovery of the tool. However, these introductions were part of the normal course of the class. This can be contrasted, for example, with DGS, where a special introduction of the tool is needed.

Kaleidoscope was also integrated in the teaching after the activity in several ways. First, as designed, we processed the mandalas and sent them back per electronic mail for the teacher of School 3 to print and redistribute them. Second, the same teacher used the Metroscope without us to finish the activity that that been designed too optimistically regarding time constraints. Finally, the patterns were used by the teacher of School 2in an unrelated activity (a question about symmetry in a blank examination).

Finally, the design of the teachers of School 3 was integrated beyond the follow-up of the activity. Since the activity needed to be developed for three grades (third, fourth, and fifth), it was actually designed as a progression over the years.

9.3.6 Continuum of Appropriation

We conclude about the use of the paper interfaces by the teachers by discussing the importance of the appropriation. It is necessary that the teachers are able to use and manage the activity if it is to exist as something other than a research prototype. Another reason is that if the teachers cannot use the paper interface, it is likely that the pupils cannot either. In contrast, appropriation of the design is not a goal. Not every teacher has as much time or energy to invest in the design of paper interfaces as the teachers we had the chance to work with. Most of the time, teachers reuse exercises found in the official text books. Similarly, while it is interesting to let teachers appropriate the design of an activity, it is also interesting to propose activities that are ready to use.

Moreover, the appropriation of the design is not a binary variable. We already explained that design is twofold: the feature set of the interface is independent of the instantiation of the activity for a class. As for the latter part, the activities we developed range from almost no customizability (in SpaceJunk, only the order of the satellites to hit can be easily changed), to little customizability (in Quads, other quadrilaterals can be coded), to a modular activity (in Angoli, the cards can be used to implement other exercises based on the construction of angles), to an instantiatable activity (Kaleidoscope allowed the teachers to design a wide range of exercises on axial symmetry), to genericity (StarrySheets can be used to design activities about a wide range of topics related to 2D-geometry).

In our case, as illustrated in this section, Kaleidoscope seems to be the right level of appropriation for teachers to design activities. It was the favorite activity of the teachers with whom we collaborated. However, SpaceJunk was also appreciated by teachers who only used it in the context of the open door workshop. We think that StarrySheets has a big potential that could be exploited with some technical knowledge in order to generate new activities whose design is easily appropriated by the teachers. Our point is that there is no right or wrong level of appropriation, and paper allows us to use whichever fits the best for the specific situation.

9.4 Design Principles

To summarize, we propose several principles that emerged while developing our interface.

- *Make the interface as autonomous as possible.* The computer is a tool that can assist a user in a task. It should not become the task itself by demanding too much attention. Activities based on paper interfaces should thus use the computer to augment the paper, but should not be based on it. Practically, it means that the interface should gracefully degrade as the computer becomes less available to augment the paper. For example, if the computer is used to compute a feedback on the work of a pupil, it is interesting to let the pupil work with a limited access to the computer (as it is the case when one augmenting device is share among the whole class), and design exercises that teachers are still able to validate themselves, so that technical issues remain manageable.
- *Make the interface flexible.* It is a well known software design principle that simple, independent, and interoperable components can be flexibly combined into a tool for a precise use. This is easy to do with a scattered interface like paper: functionalities can be isolated, which allows users to add or remove them on the fly (e.g. a teacher taking away a tool-card that is over-used). The goal is to favour real-time adaptation of system, as it is possible for example by tearing away a part of the interface or annotating it.
- *Use artefact-linked practices.* It is not necessary, and probably confusing, to invent new interaction techniques. We are already trained in using a rich set of interaction techniques involving paper: many cultural practices, such as active reading, drawing schema with tools, doing origami (folding) or kirigami (cutting), can be used as base for

an interaction. This is by far preferable to inventing new techniques, such as defining a point by folding a corner on it.

- *Use paper as a memory.* Hutchins (1995) described how a cockpit “remembers” its state. The state of the activity can be stored with paper: with the dispositions of the various elements of the interface, the series of physical transformations, the annotations, etc. This makes the interaction more robust, because it obeys the rules of the physical world, which are easier to understand and use. This is even more interesting when taking the mobility of paper into account: the work done on an interface can be continued on another augmenting device, or even “offline”, just like paper allows a spatially distributed work flow (e.g. annotating a text book in a lecture hall, revising it in public transportation, and using it for working on exercises at home).

- *Design for the user, the group, and beyond.* We saw many examples of individual activities that turned out to be full of impromptu collaboration, with the paper interface drawing from its tangible aspect to make the work of one publicly visible. For example, the pupils gave each other feedbacks while they were waiting for a validation at the lamp, or passing on ad hoc fixes of technical glitches. Beyond the group, as we already mentioned, it is thus important to take more than the group into account, and even more than the interaction surface, as paper can be used autonomously from a computer, outside, etc.

- *Use the right facet for the right task.* A paper interface is a triplet: the computer, the paper as tangible, and the paper as support for content. Each facet complements the other with complementary sets of strength. Computers can take care of complicated task to facilitate the exploration of a problem, provide a simple control to a complex object, collect precise data for later analysis, give multiple, dynamic representations of a single object, etc. Paper is a support for content that is “deployed” on many “platforms”, “compatible” with a wide range of “interoperable” tools. Moreover, the content has very little constraint but remains easily usable by humans (e.g. it is much easier to work on a paper drawing than on a digital one). Paper can also be used as a tangible interface, leveraging most of their advantages, such as an access to intuitive interaction techniques

- *Use the continuum of paper interfaces.* The three components of a paper interface are not discrete: it is possible to use the continuum between them to precisely tune a system. For example, if paper exercises are boring and computer games more entertaining, a paper interface can set the right level of engagement by adding more or less features of one or another (e.g. more or less animations or text). Similarly, a designer can use the content-tangible continuum of paper to make an interface more or less abstract, and more or less generic.

- *Paper brings information in the real world.* Another advantage of the tangible facet of paper is that it does not oversimplify reality. For example, there is not need to binarize a diagram on a sheet or round positions to align pieces of paper on a grid. This is especially interesting in education, to let pupils confront the real constraints of a problem, rather than a sterilized environment. For example, thy can manipulate a real protractor rather than a virtual one that would be always correctly oriented, fully visible, used with angles matching exactly a graduation, etc.

- *Use the skills where they are.* As we showed previously, it is easy to transfer the responsibility of the design to users, or experts of a domain. This is because paper can be edited with familiar tools, and as a consequence, greatly facilitate the adoption and the usage. It is also very easy to iterate on a design to fix issue, be it ad hoc or post hoc.

Like any guideline, these principles should be used and adapted as needed. Furthermore, these principles are voluntarily vague, in the hope that it will foster their *creative appropriation*. Indeed, paper offers too many possibilities to restrict them by a set of precise rules.

9.5 Conclusion

Section 9.2 showed that a paper interface can be a great tool for learning. It is well adapted to the collaboration between pupils at the group level but also at the class level. Paper interfaces are also a great tool for exploring complex problems in an engaging way. From the pedagogical design point of view, the paper also allows a precise control on these aspects.

Without the adoption by the teachers and the pupils of paper interfaces, these strengths cannot be realized and are just potential. We saw the creative appropriation as the ultimate form of adoption: the users (teacher or pupil) use the interface for even more than what it was created for, in the sense that they enrich it with practices that the designer did not intend.

We think that these examples of creative appropriation happen on the basis of an artefact that is already familiar. Paper is associated with a rich set of practices, often taught in primary school. Thus, for an interaction technique to be accepted by pupils, the best is to use a property of paper interfaces in ways that they are typically used: ephemeral changes in regards to the presence, side, position, and rotation; and persistent changes on edges, folds, and ink.

As for the adoption by teachers, we found that the activities were best integrated in the work flow of the class if the activity was run by the teachers, which in turn was favored by their involvement in the design of the activity. This ownership is not surprising. What is more novel, however, is that teachers were strongly involved in the design of the Kaleidoscope activities, which followed a preparation phase independent of us. In other words, we did not organize brainstorming sessions or user group studies to involve the teachers from the very beginning but instead prepared an editable component that exploited the possibilities of paper interfaces on a given topic. This process makes the best use of the experience of each actor: researchers for highlighting possibilities, and teachers for exploiting them.

Finally, these findings were only possible because of the many cycles including trials in real conditions. Each study, even of small scale, was richer than laboratory-controlled usability tests. It is the testing in the field that makes a technology pertinent.

10 Conclusion

The previous chapter summarized our findings obtained during the studies. This chapter concludes the dissertation by giving a more general overview of our contribution, discussing its limitations, and bringing up possible continuations.

10.1 Overview

We explored how the power of computers could be integrated with paper. Paper is a cultural object associated with precise and effective uses, and at the same time is a malleable material that can be used in many other ways. We conducted our exploration in the domain of geometry education at the primary school level, because paper interfaces fit this field very well: paper is omnipresent in schools, drawing and manipulating are central in geometry, and computers are naturally adapted to mathematics.

We easily leveraged the possibilities offered by the powerful combination of paper and computers. Quads, Angoli, SpaceJunk, Kaleidoscope, and StarrySheets all cover different possibilities: manipulating, drawing, folding, cutting, etc. Furthermore, Kaleidoscope has been customized into variants, fitting the needs of specific teachers and their classes. StarrySheets, on the other hand, is the base component to generate activities where the paper undergoes transformations described by geometrical concepts.

Paper also helped us convince the teachers to host experiments in their classrooms. This gave us the chance to test our activities with hundreds of pupils, and in an ecologically valid setting. The field studies were instrumental to the validation of our work, providing us with all the complexity to reveal many strengths and weaknesses of our prototypes.

The computer allowed us to record the interaction with high precision. We could then study the activities based on many aspects, because the manipulation of an object in the physical world is full of subtleties that can be revealed with the right measure or visualisation. The ability to replay the studies was central in the creation of the methods to extract these measures and visualisations.

Guided by the studies and their analyses, we could follow a progression from abstract models of a paper interface to concrete tools adopted by users. We started by laying out the structure of the exploration as a framework derived from a theoretical model of paper as an input device. We gleaned observations on the convergence between paper practices and augmentations in increasingly valid settings: from a spare room isolated from the classroom, to a classroom recreated in our laboratory (during the open day workshops), to real classrooms where teachers eventually handled everything autonomously, from the creation of the activity to its usage.

In addition to the design implications gathered from this progression, the process itself is an important finding. Innovations are often hindered by a tension between researchers trying to push a foreign practice into an established environment[1], and users limiting creativity through their habits[2]. Paper interfaces offer a sound compromise between these two pitfalls. They do not require all the stakeholders to be present from the start to orient a design. Instead, in a preliminary preparation, a researcher can crystalize possibilities into augmented paper artefacts. These artefacts are freely adaptable by users, as long as they have the skills to work with paper. This ensures that the artefacts will be the result of the best of each actor: a researchers knows the possibilities best, and a teacher knows their use best. We referred to this situation as creative appropriation: a paper interface remains a rich support for creation at any time, even if teachers are restrained to its customization, or pupils to its use.

Creative appropriation, an intrinsic property of paper interfaces, is our instrument to involve teachers in the creation of pedagogical technologies and to integrate these technologies in classrooms.

10.2 Limitations

10.2.1 Experimental Methods

The complexity of the classroom was a rich opportunity to evaluate the ecological validity of a solution, especially regarding the orchestration aspects. However, it has a major drawback from the scientific point of view: it is impossible to control everything but one (or a few) variable(s), in order to study its effect. The main consequence is that while we have many observations related to learning, we do not have a robust proof of learning outcomes. This is mostly due to the study design.

There are several ways to design studies to evaluate a new technology for learning: pilot tests to uncover usability issues, formative evaluations (with semi-structured interviews and informal testing) to reveal efficacy problems, or learning research to understand learning impact. Our research consisted mostly of pilot tests and formative evaluations because paper interfaces are still not a well defined design space. It was thus very fruitful in terms of usability.

---

[1] "I suppose it is tempting, if the only tool you have is a hammer, to treat everything as if it were a nail." – Abraham Maslow

[2] "If I had asked people what they wanted, they would have said faster horses." – attributed to Henry Ford

Learning research, on the other hand, allows us to measure learning outcomes (what the pupils should know after using the system) or transfer (the ability to use gained knowledge). We tried to measure the learning outcomes by comparing post-tests and pre-tests, but these tests were not at an appropriate difficulty level. In order to avoid floor effects (the measure outcome cannot be measured because both scores in the tests are too low) or ceiling effects (when the scores are too high), the tests need to be refined in parallel to address the issues revealed by the pilot studies about the interface. As a result, even if we investigate the effects of paper interfaces on learning, there is a choice to make between studies that focus on drawing precise conclusions about learning, and studies that gather data about the learning environment. We made the latter choice.

Had we preferred the former choice, i.e. depth over breadth, we would have probably repeated the same activity many times – Quads – until we knew precisely which components were effective to learn the classification of quadrilaterals. Instead, we showed that paper allows us to design and set up learning activities. In the case of Quads, we showed, for example, that cardboard quadrilaterals allow for the visualization of the shapes.

On a related note, we did not compare our paper interfaces to alternatives, such as plain paper or only a computer. The consequence of this is that we did not draw conclusions on whether the technology benefited or lowered learning. Obviously, a technology is not intrinsically good or bad for learning: it depends on how it is used, and many factors can bias the conclusion: whether the difficulty level was adapted, what type of background the class has, whether the teacher was in a good mood, etc. Again, we preferred gathering more general observations that are less context specific. More pragmatically, our time with pupils was limited, because our studies integrated the regular schedules of the classes. Thus, using half of this time on activities that did not make use of paper interfaces also halved the exploration opportunities.

We thus preferred to compare variants of paper interfaces between each other. This can also be criticised, because it is not trivial to design alternatives that are similar enough to be compared. For example, in the controlled study based on Kaleidoscope, the differences observed between the taping and drawing modes are biased by practical issues: we did not measure how much “fat fingers” affect the placement of paper triangles. We only informally tried bigger triangles and saw that they did not clearly influence the manipulation. Again, this renders the conclusions on learning weak, but was an interesting comparison on the usability aspect.

As a result, our research was a series of case studies on paper interfaces applied to learning geometry, and not a longitudinal study on learning (a specific topic of) geometry with paper interfaces.

10.2.2 Generalisation

We used a camera/projector system in the context of primary school geometry. This is a rather specific way to explore the possibilities of paper interfaces compared to the alleged possibilities of paper interfaces. In this section we discuss how the generality of our approach remains to be explored in other educational domains, more general applications, and other augmentation devices.

Applications

We strived to concretely implement our designs to real situations, which explains why we focused on geometry education in primary schools. This gave us a rich field for exploration because of the links between geometrical concepts and manipulations (e.g. roto-translations). However, it is a relatively narrow topic compared to increasingly wider application domains:

- Geometry education in general. Our work focused on 2D geometry, and it is not clear how a paper interface can address the third dimension, which is an important topic of geometry, especially in higher levels.
- Mathematics education in general. Some mathematical concepts are more abstract than geometry, which has an influence on how a paper interface can be used to represent and embody these concepts.
- Education in general. Computers are well adapted for mathematics because of their computational nature. They also have powerful memories and communication systems, which can be helpful in other educational domains, but we did not investigate how paper interfaces can be leveraged for this purpose.

However, judging from the ease with which we could create activities on different topics with the generic StarrySheets, we remain confident that paper interfaces are well adapted to those we investigated.

Genericity

In education, it is easy to create a micro-world to explore a problem, i.e. an environment where all the rules are explicitly defined, and the possibilities are thus controlled. This helps the creation of an interface because the application logic can be set in advance. Being able to define a micro-world is a pedagogic feature, but it is not always possible or expected in other domains.

In many applications, there is a very large choice of commands and functionalities. These commands can be combined with many objects, which results in a very complex set of possibilities. As an illustration, Graphical User Interfaces often address this using menus, a convenient way to access an exponential number of functionalities. In this case, a user can point to an icon representing an object of interest, and select within menus and sub-menus the command to apply to it.

In the activities that we developed, each piece of the interface had a direct effect or directly represented one object. To obtain the same amount of possibilities as the one allowed by a WIMP interface, we would have to create an exponential number of paper artefacts, which would result in a very cluttered interface. Alternatively, we could have generic components (e.g. a card representing a variable fraction) whose meanings are defined by applying other components (e.g. rotating a card changes the numerator of the variable fraction; another card controls the denominator). In this case, we trade simplicity for expressivity. We did not investigate as to what the right compromise between a complex interface (with a few generic elements and their meta-controllers) and complicated interface (with many simple elements) would be.

Augmenting Technology

Another limitation of our work is that it is based on a camera-projector system to augment paper. We claimed that paper has the potential to integrate computers in the classroom, but if schools refuse to acquire even one Metroscope, the problem is the same. However, there are other technologies that allow interaction similar to that of the Metroscope.

In the simplest case, a camera and a regular screen can already be used for some of the activities. For example, the Kaleidoscope activity developed by the teacher of School 4 used the lamp only to check answers. In this case, the vertical surface of the screen suffices to display the image of the sheet and the augmentation; there is no point in displaying the augmentation directly on the physical sheet. In contrast, the activity developed by the teachers of School 3 made an interactive use of the lamp. The pupils could not see the symmetry being constructed in real time with a vertical screen.

Alternatively, tablet computers can also mix augmentations on the sheet. More precisely, they can act like a lens, letting reality pass through, only adding virtual elements to it. This would be one way to watch how a symmetry is constructed in real time, for example, but it has other drawbacks. Mainly, the tablet has to be held above the augmented area, which immobilizes one hand and is tiresome.

There is no perfect technology to augment paper (yet), and we do not pretend that camera and projector systems are superior. Paper is flexible enough that interaction can be designed with other alternatives, keeping in mind their strengths and weaknesses.

10.3 Perspectives

These limitations are not unsurmountable and can very well be addressed in future work. We see this dissertation as another step towards an *ideal paper interface*. This hypothetical device would allow us to control the capabilities of computers without sacrificing any advantages of paper: cheap production cost, ease of annotation and manipulation, etc. It would be equivalent to regular paper, but with dynamic displays, computational power, and access to large amounts of possibly remote data.

We speculate that computer interfaces are in a transition toward an ideal paper interface. The first addition to the original computing machines was the ability to input and output text, in order to store knowledge, like the main use of paper. Next, mainstream computer interfaces were based on the mouse and GUI, aiming at direct manipulation, closer to the manipulability of paper. More recently, computers started to have the form of tablets or smart phones, achieving a mobility closer to the ease with which books or sheets can be transported.

Our contribution towards these goals has been to precise which aspects of the manipulation of an ideal paper interface could be used as input. We used cards as flat tangibles that can be moved and flipped. We explored how sheets of paper extended the possibilities of manipulations with folding or cutting. We used a framework to reveal properties of paper that could be changed through manipulation, which was used as input. The framework, used for output, reveals other possibilities of interaction.

The output of our system covers only one small unit of this framework: ephemeral ink. The projector virtually printed with light. However, the rest of the framework is also relevant for output. For example, Amano and Yamamoto (2012) let a computer ephemerally control the position of pieces of paper on an electro-static surface. Hashida et al. (2012) let a computer persistently add and remove ink from a sheet with ultra-violet light. Koizumi et al. (2010) used shape memory alloys to control the curvature of a sheet and animate paper foldings ephemerally. Controlling the edges of a paper-like object seems more hypothetical, but full of promises. Imagination is the only limitation to the use of paper as an interface.

A Plan d’Étude Romand

This appendix contains the translation of the expectations defined by the Plan d’Étude Romand in the *space* and *measures* categories.

A.1 Second Cycle

At the end of the 2^{nd} cycle of primary school (6-9 year old), pupils are expected to be able to:

- name a circle, a square, and a rectangle and recognize them in various situations,
- use a ruler to finish the construction of a figure from a model,
- recognize and name a rectangled triangle, an equilateral triangle, an isosceles triangle, a rhombus, a parallelogram, and a circle,
- describe an equilateral triangle, a square, and a rectangle by the number of sides, the number of right angles, the sides having equal length, the parallelism of the sides, and the internal symmetries,
- construct an equilateral triangle, a square, a rectangle, parallel and perpendicular lines using geometrical tools (graduated ruler, set square, compass, protractor),
- recognize an axial symmetry and a translation,
- spot the symmetry axes of a planar figure,
- continue the construction of a frieze or a tiling pattern,
- reproduce a planar figure by a translation or an axial symmetry using equipment,
- draw a path on a map following instructions,
- locate the relative positions of objects on a map,

- use an orthonormal referential to place a point or communicate its position,
- compare lengths and areas of simple surfaces by manipulation and drawing,
- use a graduated ruler to measure or draw a segment,
- measure a length (segments, distance between two points) using adapted equipment and express the result in a conventional, adequate unit,
- compare angles by manipulation,
- compute lengths: lines, perimeter of polygons (regular or not),
- compute the area of a square or of a rectangle (integer measures),
- determine the area of a parallelogram, rectangled triangle, from the area of a rectangle (integer measures).

A.2 Third Cycle

At the end of the 3^{rd} cycle of primary school (9-11 year old), pupils are expected to be able to:

- recognize, name, describe and construct: parallel lines, perpendicular lines, altitudes, angles, triangles, quadrilaterals, circle, perpendicular bisector, bissector of an angle, incircle and excircles, tangent of a circle, median of a triangle, regular polygons,
- use the sum of the angles of a triangle,
- use the geometry tools in an appropriate way (ruler, set square, compass, protractor),
- sketch to support reflection, memorization, or unambiguous communication of information
- recognize, name and describe: cube, rectangular cuboid, right prism, pyramid, cylinder, cone, sphere,
- sketch the representation of a solid: by development, in perspective,
- sketches the development: of a cube, rectangular cuboid, right prism, cylinder,
- recognize, name and describe: an isometry, homothecy,
- construct the image of a planar figure by: an isometry, homothecy,
- construct: the axes of symmetry of a planar figure, the center of symmetry of a planar figure,
- scale a planar figure up and down,

- recognize similar figures,
- construct and use a planar referential to place points, communicate their position, or describe itineraries,
- read a map,
- express a value using conventional units,
- use the adapted instrument to measure a length, an angle,
- express a given value in various usual units,
- compute the perimeter and the area of: polygons, disks, compound figures, circular intersections,
- compute the area of rectangular cuboids, right prisms, cylinders and pyramids,
- compute the volume of right prisms and cylinders, pyramids and cones,
- compute a missing value out of the given ones,
- use Pythagoras theorem in the plane and in space,
- use the intercept theorem in the plane.

B Planar Calibration

Geometry requires a very high precision: if a point is projected more than one millimetre off its expected position, the system is not usable by pupils who are taught to be precise with their tools. We report on how we adapted the existing calibration method to achieve a sub-millimetre precision, under the condition of working on one plane only. We first describe the previous approach, our replacement, and evaluate the improvement.

Calibration corresponds to the mapping between the camera and the projector. More precisely, a typical augmentation workflow consists of three steps. First, a marker is detected in the image captured by the camera, which gives coordinates in pixels. Second, these coordinates are translated into a position in the real world (e.g. in millimetres). For example, this allows to compute the distance between two objects. Third, the physical position is translated into coordinates in the projected image, which are in pixels, but in a different referential than the camera. For example, this allows to project a text one centimetre right from a marker.

Previous Approach

The previous approach was two-fold. It consisted of determining the distortion coming from the lens of the camera, and estimating the plane of the interaction surface.

Regarding the camera, several tools allow the definition of a mathematical model of the camera (See e.g. the toolbox of Bouguet (2010)). This model includes the focal length and central point. It is obtained by taking multiple pictures of a chessboard pattern. This pattern allows the tool to extract the coordinates in the camera image of points that are known to be coplanar and equidistant in reality. With enough points, the model of the deformation of the pattern can be inferred. With this model, it is possible to compute the position in space of a point, if the distance between the camera and the point is known. In our case, this distance corresponds to the height of the camera, i.e. the distance between the interactive surface and the mirror, plus the distance between the mirror and the camera.

The second part of the approach described here consists of estimating the equation of the plane of the interactive surface. This is done by showing the system a sheet where the relative position of markers is known, and reversing the equation of the previous paragraph. A similar process estimates the plane 10 centimetre above the interaction surface, which correspond to the height of the tangible objects used historically. Finally, the parameters calculated at this step are applied to the three-dimensional graphics environment, which allows the developer to simply issue graphic command in the coordinates of the physical world.

Overall, this approach allows an acceptably precise calibration of the system in space. However, the first step involves some “black magic” when taking the pictures of the chessboard pattern: some picture sets are better than others, which influences the quality of the calibration.

Planar Optimization

Since geometry in primary school is mostly planar, we can devise a simpler, two-step calibration procedure. The first step consists in showing the camera a chessboard pattern that covers the whole interaction area. The chessboard pattern allows to find the coordinates of each square with a sub-pixel accuracy. Since the squares are conveniently printed to be 10 × 10 millimetre, the mapping from camera pixel to physical world millimetres is trivial for the positions of the corner squares, and can be interpolated otherwise.

The second step is the same, in the other direction: the projector shows the image of a chessboard, and the previously obtained mapping is used to map projector pixels to physical world millimetres, by detecting the corners of the squares with the camera. The newly obtained mapping is then inverted to associate physical wolrd coordinates to pixels in the projection image.

Evaluation

On the one hand, the planar calibration loses the ability to map points in space; it is restricted to the surface of interaction. On the other hand, the procedure is as simple as placing a poster-sized piece of paper, and removing it, which is a big improvement compared to the erratic previous procedure of the spatial calibration. However, the main objective was the improvement in precision. To measure it, we did the following experiment: we displayed a single plain circle with the projector, and detected it again with the camera. Detecting such a blob is another way to determine a precise position: its center can be computed with a sub-pixel accuracy. If the calibration is perfect, the circle will be detected exactly were it was expected to be projected. We measure the quality of a calibration as the distance between the expected position and the actually detected position.

We displayed a circle until it was detected three times, and placed circles one after another on a grid, spaced by 50 pixels on the camera image. We used the calibration parameters used during experiments using each of the calibration approach. Figure B.1 helps visualizin this procedure. It showing the displacement between the expected position and the average of the detected positions in each of the condition.

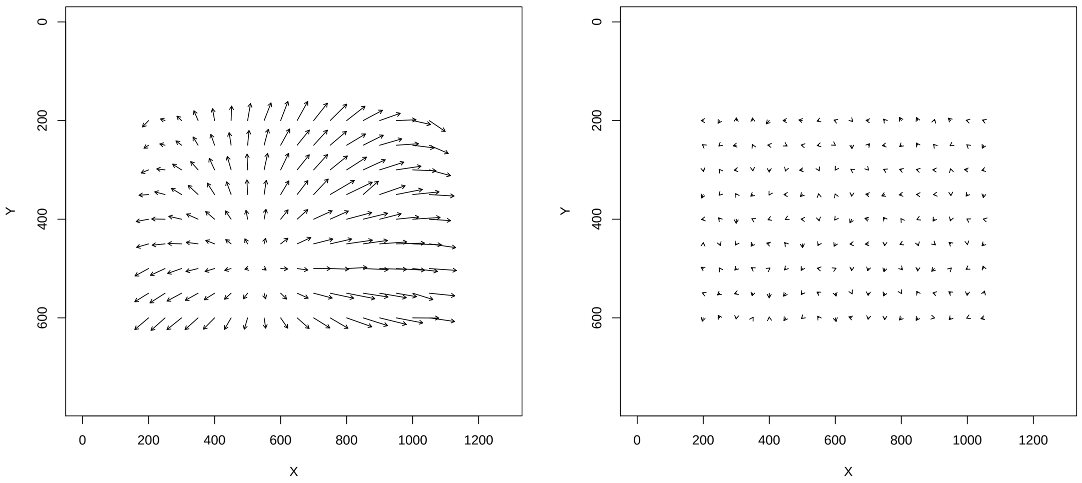

Figure B.1: The displacement between the expected position and the average of the detected positions, with the planar calibration (left) and the spatial calibration (right). The point on the tip of the arrow was expected at its base. For readability, the length of the arrows are scaled up by 10.

Figure B.2 shows the distribution of the errors, i.e. the distances between the expected position of the center of the circle, and the actually detected position. The distances are indicated in millimetre, by simply dividing the distance in pixels by 1.7, which is the pixel per millimetre ratio of the system used for this test. The average error with the calibration of the previous approach is 2.845 mm, while the average error with the planar calibration is 0.29 mm, which is a significant improvement, and passes the criteria of a sub-millimetre accuracy.

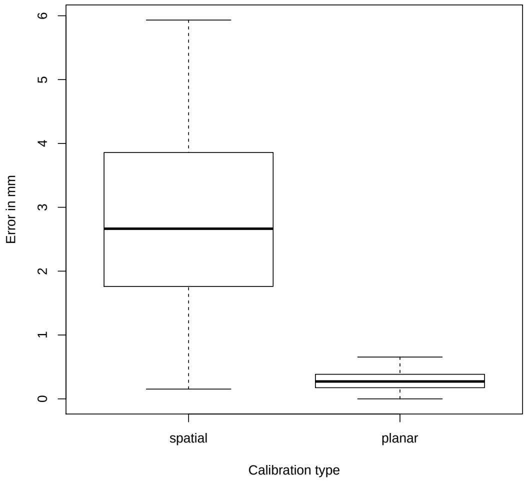

Figure B.2: Distribution of the distance (in millimetre) between the expected position of the circles, and the actually detected position, for the spatial calibration (left) and the planar calibration (right).

Bibliography

J. Ainley and R. Goldstein. *Making Logo work: a guide for teachers.* Basil Blackwell, 1988.

S. Ainsworth. The functions of multiple representations. *Computers & Education*, 33(2-3): 131–152, 1999.

S. Ainsworth. DeFT: a conceptual framework for considering learning with multiple representations. *Learning and Instruction*, 16(3):183–198, 2006.

C. Alcoholado, M. Nussbaum, A. Tagle, F. Gomez, F. Denardin, H. Susaeta, M. Villalta, and K. Toyama. One mouse per child: interpersonal computer for individual arithmetic practice. *Journal of Computer Assisted Learning*, page no–no, 2011.

D. Aliakseyeu, J. Martens, and M. Rauterberg. A computer support tool for the early stages of architectural design. *Interacting with Computers*, 18(4):528–555, 2006.

K. Amano and A. Yamamoto. Tangible interactions on a flat panel display using actuated paper sheets. In *Proceedings of the 2012 ACM international conference on Interactive tabletops and surfaces*, page 351–354, 2012.

T. Arai, K. Machii, and S. Kuzunuki. Retrieving electronic documents with real-world objects on InteractiveDESK. In *Proceedings of the 8th annual ACM symposium on User interface and software technology*, pages 37–38, Pittsburgh, Pennsylvania, United States, 1995. ACM.

T. Arai, D. Aust, and S. E. Hudson. PaperLink: a technique for hyperlinking from real paper to electronic content. In *CHI '97: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 327–334, New York, NY, USA, 1997. ACM.

E. Aronson and S. Patnoe. *The jigsaw classroom: building cooperation in the classroom.* Longman, 1997.

M. Artigue. Analyse de processus d'enseignement en environnement informatique. *Petit x*, 26: 5–27, 1991.

K. Asai, H. Kobayashi, and T. Kondo. Augmented instructions - a fusion of augmented reality and printed learning materials. In *Advanced Learning Technologies, 2005. ICALT 2005. Fifth IEEE International Conference on*, pages 213–215, 2005.

M. Back, J. Cohen, R. Gold, S. Harrison, and S. Minneman. Listen reader: an electronically augmented paper-based book. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '01, page 23–29, New York, NY, USA, 2001. ACM.

S. Barab and K. Squire. Design-based research: Putting a stake in the ground. *Journal of the Learning Sciences*, 13(1):1–14, 2004.

M. Bartolini Bussi, M. Boni, F. Ferri, and R. Garuti. Early approach to theoretical thinking: Gears in primary school. *Educational Studies in Mathematics*, 39(1):67–87, 1999.

Y. Baulac. *Un micromonde de géométrie, Cabri-géomètre*. PhD thesis, Université Joseph Fournier - Grenoble I, February 1990.

A. Benbachir and M. Zaki. Production d'exemples et de contre-exemples en analyse: étude de cas en première d'université. *Educational Studies in Mathematics*, 47(3):273–295, 2001.

R. Berthelot and M. H. Salin. Section v. the role of pupils' spatial knowledge in the elementary teaching of geometry. *Perspectives on the teaching of geometry for the 21st century*, 5:71, 1998.

M. Billinghurst, H. Kato, and I. Poupyrev. Collaboration with tangible augmented reality interfaces. In *HCI International*, volume 1, page 5–10, 2001a.

M. Billinghurst, H. Kato, and I. Poupyrev. The MagicBook - moving seamlessly between reality and virtuality. *Computer Graphics and Applications, IEEE*, 21(3):6 – 8, June 2001b.

A. Blackwell, M. Stringer, E. Toye, and J. Rode. Tangible interface for collaborative information retrieval. In *CHI'04 extended abstracts on Human factors in computing systems*, page 1473–1476, New York, NY, USA, 2004. ACM Press.

F. Block, M. Haller, H. Gellersen, C. Gutwin, and M. Billinghurst. VoodooSketch: extending interactive surfaces with adaptable interface palettes. In *TEI '08: Proceedings of the 2nd international conference on Tangible and embedded interaction*, page 55–58, New York, NY, USA, 2008. ACM.

K. Blythe. Children learning with logo. *National Council for Educational Technology*, 1990.

J.-Y. Bouguet. Camera calibration toolbox for matlab, 2010. URL http://www.vision.caltech.edu/bouguetj/calib_doc.

G. Brousseau. Etude de questions d'enseignement, un exemple: la géométrie. *Séminaire de didactique des mathématiques et de l'informatique*, (45):183–226, 1983.

A. L. Brown. Design experiments: Theoretical and methodological challenges in creating complex interventions in classroom settings. *The journal of the learning sciences*, 2(2): 141–178, 1992.

J. Bruner and H. Kenney. Representation and mathematics learning. *Monographs of the Society for Research in Child Development*, 30(1):50–59, 1965.

L. Buechley, S. Hendrix, and M. Eisenberg. Paints, paper, and programs: first steps toward the computational sketchbook. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*, page 9–12, 2009.

B. Bähr, S. Kratz, and M. Rohs. A tabletop system for supporting paper prototyping of mobile interfaces. In *PaperComp 2010: 1st International Workshop on Paper Computing. Workshop at Ubicomp 2010*, 2010.

S. K. Card, J. D. Mackinlay, and G. G. Robertson. A morphological analysis of the design space of input devices. *ACM Transactions on Information Systems (TOIS)*, 9(2):99–122, 1991.

S. Carter, C. Liao, L. Denoue, G. Golovchinsky, and Q. Liu. Linking digital media to physical documents. *Pervasing Computing*, 9(2):46–55, June 2010.

F. Chehimi. Storybook augmentation through portable projectors, 3D augmented reality and anaglyph. In *PaperComp 2010: 1st International Workshop on Paper Computing. Workshop at Ubicomp 2010*, 2010.

Y. Chevallard. Intégration et viabilité des objets informatiques dans l'enseignement des mathématiques. *L'ordinateur pour enseigner les mathématiques, Nouvelle Encyclopédie Diderot, Paris: PUF*, page 183–203, 1992.

Y. Chevallard. Enseignement insensé, enseignement raisonné et créativité sociale. *Revue de l'APMEP*, (435), 2001.

H. Cho, J. Jung, K. Cho, Y.-H. Seo, and H. S. Yang. AR postcard: the augmented reality system with a postcard. In *Proceedings of the 10th International Conference on Virtual Reality Continuum and Its Applications in Industry*, VRCAI '11, page 453–454, New York, NY, USA, 2011. ACM.

M. Chong and F. Kawsar. Improving paper books: Searchable books. In *PaperComp 2010: 1st International Workshop on Paper Computing. Workshop at Ubicomp 2010*, 2010.

P. Clarou and C. Laborde. Regards sur L'Intégration de cabri-géomètre. 2000.

D.-B. R. Collective. Design-based research: An emerging paradigm for educational inquiry. *Educational Researcher*, page 5–8, 2003.

A. Collins. Toward a design science of education. *New directions in educational technology*, pages 15–22, 1992.

Conférence Intercantonale de l'Instruction Public de la Suisse romande et du Tessin. Plan d'étude romand, 2010. URL http://www.plandetudes.ch/.

E. Costanza, M. Giaccone, O. Kueng, S. Shelley, and J. Huang. Tangible interfaces for download: initial observations from users' everyday environments. In *Proceedings of the 28th of the International Conference on Human factors in Computing Systems*, page 2765–2774, New York, NY, USA, 2010. ACM.

A. Crossen, J. Budzik, M. Warner, L. Birnbaum, and K. J. Hammond. XLibris: an automated library research assistant. In *IUI '01: Proceedings of the 6th international conference on Intelligent user interfaces*, page 49–52, New York, NY, USA, 2001. ACM.

S. Cuendet, P. Jermann, and P. Dillenbourg. Tangible interfaces: when physical-virtual coupling may be detrimental to learning. In *BCS '12 Proceedings of the 25th BCS Interaction Specialist Group Conference*, 2012.

S. Cuendet, Q. Bonnard, F. Kaplan, and P. Dillenbourg. Paper interface design for classroom orchestration. In *Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems*, CHI EA '11, page 1993–1998, New York, NY, USA, 2011. ACM.

M. Diaz, M. Alencastre-Miranda, L. Munoz-Gomez, and I. Rudomin. Multi-user networked interactive augmented reality card game. In *Proceedings of the 2006 International Conference on Cyberworlds*, page 177–182, Washington, DC, USA, 2006. IEEE Computer Society.

P. Dillenbourg and M. Evans. Interactive tabletops in education. *International Journal of Computer-Supported Collaborative Learning*, page 1–24, 2011.

P. Dillenbourg. What do you mean by collaborative learning? In P. Dillenbourg, editor, *Collaborative-learning: Cognitive and Computational Approaches.*, pages 1–19. Oxford: Elsevier, 1999.

P. Dillenbourg. Overscripting CSCL: the risks of blending collaborative learning with instructional design. In P. A. Kirschner, editor, *Three worlds of CSCL. Can we support CSCL?*, page 61–91. Open Universiteit Nederland, 2002.

P. Dillenbourg and F. Hong. The mechanics of CSCL macro scripts. *International Journal of Computer-Supported Collaborative Learning*, 3(1):5–23, March 2008.

P. Dillenbourg, G. Zufferey, H. S. Alavi, P. Jermann, L. H. S. Do, Q. Bonnard, S. Cuendet, and F. Kaplan. Classroom orchestration : The third circle of usability. In *Connecting Computer-Supported Collaborative Learning to Policy and Practice: CSCL2011 Conference Proceedings. Volume I - Long Papers*, volume 1, page 510–517. International Society of the Learning Sciences, 2011.

A. Dimitrakopoulou, L. Bollen, Y. Dimitriadis, A. Harrer, P. Jermann, V. Kollias, J. A. Marcos, A. Martínez, and A. Petrou. State of the art of interaction analysis for metacognitive support & diagnosis, 2006.

H. S. Do Lenh. *Supporting Reflection and Classroom Orchestration with Tangible Tabletops.* PhD thesis, Lausanne, 2012.

S. Do-Lenh, F. Kaplan, and P. Dillenbourg. Paper-based concept map: the effects of tabletop on an expressive collaborative learning task. In *Proceedings of the 23rd British HCI Group Annual Conference on People and Computers: Celebrating People and Technology*, BCS-HCI '09, page 149–158, Swinton, UK, UK, 2009. British Computer Society.

T. Doering, S. Beckhaus, and A. Schmidt. Towards a sensible integration of paper-based tangible user interfaces into creative work processes. In *Proceedings of the 27th international conference extended abstracts on Human factors in computing systems*, page 4627–4632, 2009.

N. Douek. Argumentation and conceptualization in context: a case study on sunshadows in primary school. *Educational Studies in Mathematics*, 39(1):89–110, 1999.

R. Duval. Section II. GEOMETRY FROM a COGNITIVE POINT OF VIEW. *Perspectives on the Teaching of Geometry For the 21st Century: an ICMI study*, 5:37, 1998.

M. Dymetman and M. Copperman. Intelligent paper. In *Electronic Publishing, Artistic Imaging, and Digital Typography*, pages 392–406. 1998.

A. Dünser and E. Hornecker. Lessons from an AR book study. In *Proceedings of the 1st international conference on Tangible and embedded interaction*, TEI '07, page 179–182, New York, NY, USA, 2007. ACM.

A. N. Eisenberg. *An educational program for paper sculpture: a case study in the design of software to enhance children's spatial cognition.* PhD thesis, University of Colorado at Boulder, Boulder, CO, USA, 1999.

M. Eisenberg, A. Eisenberg, S. Hendrix, G. Blauvelt, D. Butter, J. Garcia, R. Lewis, and T. Nielsen. As we may print. page 31. ACM Press, 2003.

M. Eisenberg and A. Nishioka. Orihedra: Mathematical sculptures in paper. *International Journal of Computers for Mathematical Learning*, 1(3):225–261, January 1996.

B. Erol, E. Antúnez, and J. J. Hull. HOTPAPER: multimedia interaction with paper using mobile phones. In *MM '08: Proceeding of the 16th ACM international conference on Multimedia*, page 399–408, New York, NY, USA, 2008. ACM.

M. Evans, E. Ryon, A. Feenstra, and D. McNeill. Discourse coding from a multimodal approach: Distributed cognition and geometric problem solving in young children. *the CSCL Alpine-Rendezvous Workshop*, 2009.

R. Falcade, C. Laborde, and M. A. Mariotti. Approaching functions: Cabri tools as instruments of semiotic mediation. *Educational Studies in Mathematics*, 66(3):317–333, July 2007.

K. Fishkin. A taxonomy for and analysis of tangible interfaces. *Personal and Ubiquitous Computing*, 8(5):347–358, July 2004.

G. W. Fitzmaurice and W. Buxton. An empirical evaluation of graspable user interfaces: towards specialized, space-multiplexed input. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '97, page 43–50, New York, NY, USA, 1997. ACM.

G. W. Fitzmaurice, H. Ishii, and W. A. S. Buxton. Bricks: laying the foundations for graspable user interfaces. In *Proceedings of CHI 1995*, page 442–449. ACM Press/Addison-Wesley Publishing Co., 1995.

G. Fitzmaurice. *Graspable user interfaces*. PhD thesis, Citeseer, 1996.

M. Fjeld, J. Fredriksson, M. Ejdestig, F. Duca, K. Bötschi, B. Voegtli, and P. Juchli. Tangible user interface for chemistry education: comparative evaluation and re-design. *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, page 805–808, 2007.

M. Flagg and J. Rehg. Projector-guided painting. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*, page 235–244, 2006.

R. Freitas and P. Campos. SMART: a SysteM of augmented reality for teaching 2nd grade students. In *BCS HCI (2)*, pages 27–30, 2008.

J. Garcia, T. Tsandilas, C. Agon, and W. Mackay. Interactive paper substrates to support musical creation. In *CHI '12: Proceedings of the 30th international conference on Human factors in computing systems*, page to appear (4 pages). ACM, 2012.

A. Girouard, E. Solovey, L. Hirshfield, S. Ecott, O. Shaer, and R. Jacob. Smart blocks: a tangible mathematical manipulative. In *Proceedings of the 1st international conference on Tangible and embedded interaction*, page 183–186, 2007.

S. Goldin-Meadow. *Hearing gesture: How our hands help us think*. Harvard University Press, 2003.

A. Grasso, A. Karsenty, and M. Susani. Augmenting paper to enhance community information sharing. In *Proceedings of DARE 2000 on Designing augmented reality environments*, pages 51–62, Elsinore, Denmark, 2000. ACM.

F. Guimbretiere. Paper augmented digital documents. In *Proceedings of the 16th annual ACM symposium on User interface software and technology*, pages 51–60, Vancouver, Canada, 2003. ACM.

D. Guin and L. Trouche. The complex process of converting tools into mathematical instruments: The case of calculators. *International Journal of Computers for Mathematical Learning*, 3(3):195–227, 1998.

S. Gupta and C. Jaynes. The universal media book: tracking and augmenting moving surfaces with projected information. In *Proceedings of the 2006 Fifth IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR'06) - Volume 00*, pages 177–180. IEEE Computer Society, 2006.

M. Haller, J. Leitner, T. Seifried, J. R. Wallace, S. D. Scott, C. Richter, P. Brandl, A. Gokcezade, and S. Hunter. The NiCE discussion room: Integrating paper and digital media to support co-located group meetings. In *Proceedings of the 28th international conference on Human factors in computing systems*, CHI '10, page 609–618, New York, NY, USA, 2010. ACM.

B. L. Harrison, K. P. Fishkin, A. Gujar, C. Mochon, and R. Want. Squeeze me, hold me, tilt me! an exploration of manipulative user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '98, page 17–24, New York, NY, USA, 1998. ACM Press/Addison-Wesley Publishing Co.

T. Hashida, K. NIshimura, and T. Naemura. Hand-rewriting: automatic rewriting similar to natural handwriting. In *Proceedings of the 2012 ACM international conference on Interactive tabletops and surfaces*, page 153–162, 2012.

F. Heinrichs, J. Steimle, D. Schreiber, and M. Mühlhäuser. Letras: an architecture and framework for ubiquitous pen-and-paper interaction. In *Proceedings of the 2nd ACM SIGCHI symposium on Engineering interactive computing systems*, EICS '10, page 193–198, New York, NY, USA, 2010. ACM.

S. Hendrix. *Popup Workshop: Computationally enhanced paper engineering for children.* PhD thesis, University of Colorado, 2008.

G. Holland. *GEOLOG.* Dümmler, 1993.

D. Holman and R. Vertegaal. Organic user interfaces: designing computers in any way, shape, or form. *Commun. ACM*, 51(6):48–55, 2008.

L. Holmquist, J. Redström, P. Ljungstrand, and H.-W. Gellersen. Handheld and ubiquitous computing. volume 1707 of *Lecture Notes in Computer Science*, pages 234–245. Springer Berlin / Heidelberg, 1999.

J. Hong, M. N. Price, B. N. Schilit, and G. Golovchinsky. Printertainment: printing with interactive cover sheets. In *CHI '99 extended abstracts on Human factors in computing systems*, CHI EA '99, page 240–241, New York, NY, USA, 1999. ACM.

M. S. Horn, M. Tobiasz, and C. Shen. Visualizing biodiversity with voronoi treemaps. In *Voronoi Diagrams, 2009. ISVD'09. Sixth International Symposium on*, page 265–270, 2009a.

M. S. Horn, E. T. Solovey, R. J. Crouser, and R. J. Jacob. Comparing the use of tangible and graphical programming languages for informal science education. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems*, page 975–984, New York, NY, USA, 2009b. ACM.

M. S. Horn, R. J. Crouser, and M. U. Bers. Tangible interaction and learning: the case for a hybrid approach. *Personal and Ubiquitous Computing*, 16(4):379–389, June 2011.

E. Hornecker and A. Dünser. Of pages and paddles: Children's expectations and mistaken interactions with physical-digital tools. *Interacting with Computers*, 21(1-2):95–107, 2009.

E. Hornecker and T. Psik. Using ARToolKit markers to build tangible prototypes and simulate other technologies. *Human-Computer Interaction-INTERACT 2005*, page 30–42, 2005.

E. Hornecker and J. Buur. Getting a grip on tangible interaction: a framework on physical space and social interaction. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, page 437–446, New York, NY, USA, 2006. ACM Press.

E. Hornecker, P. Marshall, N. S. Dalton, and Y. Rogers. Collaboration, interference: awareness with mice or touch input. *CSCW '08: Proceedings of the ACM Conference on Computer supported cooperative work*, page 167–176, 2008.

C. Houdement and A. Kuzniak. Un exemple de cadre conceptuel pour letude de l'enseignement de la geometrie en formation des maitres. *Educational studies in mathematics*, 40(3): 283–312, 1999.

M. Hsieh and J. Lee. AR marker capacity increasing for kindergarten english learning. In *Proceedings of the International MultiConference of Engineers and Computer Scientists 2008*, page 663–666. International Association of Engineers, 2008.

S. Hudson and J. Mankoff. Rapid construction of functioning physical interfaces from cardboard, thumbtacks, tin foil and masking tape. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*, page 289–298, 2006.

E. Hutchins. How a cockpit remembers its speeds. *Cognitive science*, 19(3):265–288, 1995.

R. Hölzl. How does 'dragging' affect the learning of geometry. *International Journal of Computers for Mathematical Learning*, 1(2):169–187, 1996.

H. Ishii and B. Ullmer. Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '97, page 234–241, New York, NY, USA, 1997. ACM.

N. Ito, N. Fujita, H. Shimazu, N. Nakajima, and K. Yamada. TransWorld: paper world as avatar of electronic world. In *CHI '99 extended abstracts on Human factors in computing systems*, pages 206–207, Pittsburgh, Pennsylvania, 1999. ACM.

N. Jackiw. The geometer's sketchpad, v3. 0. *Key Curriculum, California*, 1995.

R. J. Jacob, A. Girouard, L. M. Hirshfield, M. S. Horn, O. Shaer, E. T. Solovey, and J. Zigelbaum. Reality-based interaction: a framework for post-WIMP interfaces. In *Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, page 201–210, 2008.

R. J. K. Jacob, H. Ishii, G. Pangaro, and J. Patten. A tangible interface for organizing information using a grid. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, page 339–346, New York, NY, USA, 2002. ACM.

G. Jacucci, A. Oulasvirta, A. Salovaara, T. Psik, and I. Wagner. Augmented reality painting and collage: Evaluating tangible interaction in a field study. *Human-Computer Interaction-INTERACT 2005*, page 43–56, 2005.

P. Jermann, G. Zufferey, B. Schneider, A. Lucci, S. Lépine, and P. Dillenbourg. Physical space and division of labor around a tabletop tangible simulation. In *Proceedings of the 9th international conference on Computer supported collaborative learning-Volume 1*, page 345–349, 2009.

P. Jermann and P. Dillenbourg. An analysis of learner arguments in a collective learning environment. In *Proceedings of the Third CSCL Conference*, 1999.

M. Jo and N. Kim. Delphi survey on the use of robot projector based augmented reality in dramatic activity for young children. 2011.

W. Johnson, H. Jellinek, L. Klotz,Jr., R. Rao, and S. K. Card. Bridging the paper and electronic worlds: the paper user interface. In *Proceedings of the INTERACT '93 and CHI '93 conference on Human factors in computing systems*, CHI '93, page 507–512, New York, NY, USA, 1993. ACM.

G. Kadunz and H. Kautschitsch. Thales. *Software zur experimentellen Geometrie. Stuttgart: Ernst Klett Schulbuchverlag*, 1993.

H. Kaufmann and A. Dünser. Summary of usability evaluations of an educational augmented reality application. In *Proceedings of the 2nd international conference on Virtual reality*, ICVR'07, page 660–669, Berlin, Heidelberg, 2007. Springer-Verlag.

T. Kawashima, K. Imamoto, H. Kato, K. Tachibana, and M. Billinghurst. Magic paddle: A tangible augmented reality interface for object manipulation. In *Proceedings of ISMR2001*, page 194–195, 2001.

J. Kim, S. M. Seitz, and M. Agrawala. Video-based document tracking: unifying your physical and electronic desktops. In *UIST '04: Proceedings of the 17th annual ACM symposium on User interface software and technology*, page 99–107, New York, NY, USA, 2004. ACM.

S. R. Klemmer, M. W. Newman, R. Farrell, M. Bilezikjian, and J. A. Landay. The designers' outpost: a tangible interface for collaborative web site. In *Proceedings of the 14th annual ACM symposium on User interface software and technology*, page 1–10, 2001.

S. Klemmer, K. Everitt, and J. Landay. Integrating physical and digital interactions on walls for fluid design collaboration. *Human-Computer Interaction*, 23(2):138–213, April 2008.

S. R. Klemmer and J. A. Landay. Toolkit support for integrating physical and digital interactions. *Human-Computer Interaction*, 24(3):315–366, 2009.

S. R. Klemmer, J. Graham, G. J. Wolff, and J. A. Landay. Books with voices: paper transcripts as a physical interface to oral histories. In *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems*, page 89–96, New York, NY, USA, 2003. ACM.

H. Koike, Y. Sato, Y. Kobayashi, H. Tobita, and M. Kobayashi. Interactive textbook and interactive venn diagram: natural and intuitive interfaces on augmented desk system. In *CHI '00: Proceedings of the SIGCHI conference on Human factors in computing systems*, page 121–128, New York, NY, USA, 2000. ACM.

N. Koizumi, K. Yasu, A. Liu, M. Sugimoto, and M. Inami. Animated paper: A toolkit for building moving toys. *Comput. Entertain.*, 8(2):7:1–7:16, December 2010.

B. Koleva, S. Benford, K. H. Ng, and T. Rodden. A framework for tangible user interfaces. In *Physical Interaction (PI03) Workshop on Real World User Interfaces*, pages 46–50, 2003.

I. Kollar, C. Wecker, S. Langer, and F. Fischer. Orchestrating web-based collaborative inquiry learning with small group and classroom scripts. In *TEI '09: Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*, page 77–84, 2011.

C. Laborde. Integration of technology in the design of geometry tasks with cabri-geometry. *International Journal of Computers for Mathematical Learning*, 6(3):283–317, 2002.

C. Laborde. Technology used as a tool for mediating knowledge in the teaching of mathematics: the case of cabri-geometry. In *Plenary speech delivered at the Asian Technology Conference in Mathematics*, 2003.

C. Laborde. Géométrie–période 2000 et apres. *One Hundred Years of L'Enseignement Mathématique, Proceedings of the EM–ICMI Symposium*, 20:133, 2000.

I. Lakatos, E. Zahar, and J. Worall. *Preuves et réfutations: essai sur la logique de la découverte mathématique*. 1984.

A. Lam, K. Chow, E. Yau, and M. Lyu. ART: augmented reality table for interactive trading card game. In *Proceedings of the 2006 ACM international conference on Virtual reality continuum and its applications*, VRCIA '06, page 357–360, New York, NY, USA, 2006. ACM.

B. M. Lange, M. A. Jones, and J. L. Meyers. Insight lab: an immersive team environment linking paper, displays, and data. In *CHI '98: Proceedings of the SIGCHI conference on Human factors in computing systems*, page 550–557, New York, NY, USA, 1998. ACM Press/Addison-Wesley Publishing Co.

J. H. Larkin and H. A. Simon. Why a diagram is (Sometimes) worth ten thousand words. *Cognitive Science*, 11(1):65–100, 1987.

G. Lee, C. Nelles, M. Billinghurst, and G. Kim. Immersive authoring of tangible augmented reality applications. In *Proceedings of the 3rd IEEE/ACM International Symposium on Mixed and Augmented Reality*, page 172–181, 2004.

W. Lee, R. de Silva, E. Peterson, R. Calfee, and T. Stahovich. Newton's pen: A pen-based tutoring system for statics. *Computers & Graphics*, 32(5):511–524, 2008.

C. Letondal and W. E. Mackay. L'écriture augmentée: enregistrer des explorations interactives avec une feuille de données scientifiques. In *IHM '09: Proceedings of the 21st International Conference on Association Francophone d'Interaction Homme-Machine*, page 363–366, New York, NY, USA, 2009. ACM.

J. Li and N. M. Allinson. A comprehensive review of current local features for computer vision. *Neurocomput.*, 71(10-12):1771–1787, June 2008.

C. Liao, F. Guimbretière, K. Hinckley, and J. Hollan. Papiercraft: A gesture-based command system for interactive paper. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 14(4):18, 2008.

C. Liao, F. Guimbretière, and C. E. Loeckenhoff. Pen-top feedback for paper-based interfaces. In *UIST '06: Proceedings of the 19th annual ACM symposium on User interface software and technology*, page 201–210, New York, NY, USA, 2006. ACM.

W. Liu, A. D. Cheok, C. L. Mei-Ling, and Y.-L. Theng. Mixed reality classroom: learning from entertainment. In *Proceedings of the 2nd international conference on Digital interactive media in entertainment and arts*, DIMEA '07, page 65–72, New York, NY, USA, 2007. ACM.

P. Ljungstrand, J. Redström, and L. Holmquist. WebStickers: using physical tokens to access, manage and share bookmarks to the web. In *Proceedings of DARE 2000 on Designing augmented reality environments*, DARE '00, page 23–31, New York, NY, USA, 2000. ACM.

R. Luckin. The learner centric ecology of resources: A framework for using technology to scaffold learning. *Computers and Education*, 50(2):449–462, 2008.

W. Mackay and D. Pagani. Video mosaic: laying out time in a physical space. In *Proceedings of the second ACM international conference on Multimedia*, MULTIMEDIA '94, page 165–172, New York, NY, USA, 1994. ACM.

W. Mackay and A. Fayard. Designing interactive paper: lessons from three augmented reality projects. In *Proceedings of IWAR*, volume 98, page 81–90, 1999.

W. Mackay, D. Pagani, L. Faber, B. Inwood, P. Launiainen, L. Brenta, and V. Pouzol. Ariel: augmenting paper engineering drawings. In *Conference companion on Human factors in computing systems*, page 421–422, 1995.

W. E. Mackay, A.-L. Fayard, L. Frobert, and L. Médini. Reinventing the familiar: exploring an augmented reality design space for air traffic control. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '98, page 558–565, New York, NY, USA, 1998. ACM Press/Addison-Wesley Publishing Co.

W. E. Mackay, G. Pothier, C. Letondal, K. Bøegh, and H. E. Sørensen. The missing link: augmenting biology laboratory notebooks. In *Proceedings of the 15th annual ACM symposium on User interface software and technology*, UIST '02, page 41–50, New York, NY, USA, 2002. ACM.

Y. Makino and Y. Kakehi. Metamorphic light: a tabletop tangible interface using deformation of plain paper. In *ACM SIGGRAPH 2011 Posters*, page 48, 2011.

S. Malacria, T. Pietrzak, A. Tabard, and E. Lecolinet. U-note: capture the class and access it everywhere. *Human-Computer Interaction–INTERACT 2011*, page 643–660, 2011.

L. Malmborg, B. Peterson, and M. Petersson. Augmenting pen and paper to support creative collaboration in design education. In *7th International Digital Arts and Culture Conference*, 2007.

M. Mariotti, R. Sutherland, and J. Mason. Images and concepts in geometrical reasoning. *NATO ASI SERIES F COMPUTER AND SYSTEMS SCIENCES*, 138:97–97, 1995.

P. Marshall, Y. Rogers, and E. Hornecker. Are tangible interfaces really any better than other kinds of interfaces? *THE WORKSHOP ON TANGIBLE USER INTERFACES IN CONTEXT AND THEORY AT CHI 2007*, 2007.

J. Martin-Gutiérrez, J. Luis Saorin, M. Contero, M. Alcaniz, D. Pérez-López, and M. Ortega. Design and validation of an augmented book for spatial abilities development in engineering students. *Computers & Graphics*, 34(1):77–91, 2010.

D. McGee, P. Cohen, R. Wesson, and S. Horman. Comparing paper and tangible, multimodal tools. In *Proceedings of the SIGCHI conference on Human factors in computing systems: Changing our world, changing ourselves*, page 407–414, 2002.

R. Mechling. Euklid. *Another geometry software*, 1994.

A. Millner and M. Resnick. Tools for creating custom physical computer interfaces. In *Interaction Design and Children*, 2005.

M. Minsky and S. Papert. Draft of a proposal to ARPA for research on artificial intelligence at MIT. *As cited in Dreyfus*, 1972.

P. Mistry, P. Maes, and C. Stephanidis. Universal access in human-computer interaction. intelligent and ubiquitous interaction environments. volume 5615 of *Lecture Notes in Computer Science*, pages 547–556. Springer Berlin / Heidelberg, 2009.

A. Moore and H. Regenbrecht. The tangible augmented street map. In *Proceedings of the 2005 international conference on Augmented tele-existence*, page 249–250, 2005.

T. P. Moran, E. Saund, W. Van Melle, A. U. Gujar, K. P. Fishkin, and B. L. Harrison. Design and technology for collaborage: collaborative collages of information on physical walls. In *Proceedings of the 12th annual ACM symposium on User interface software and technology*, UIST '99, page 197–206, New York, NY, USA, 1999. ACM.

T. Nakai, K. Kise, and M. Iwamura. Use of affine invariants in locally likely arrangement hashing for camera-based document image retrieval. *Document Analysis Systems VII*, 2006.

L. Nelson, S. Ichimura, E. R. Pedersen, and L. Adams. Palette: a paper interface for giving presentations. In *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems*, page 354–361, New York, NY, USA, 1999. ACM.

W. Newman and P. Wellner. A desk supporting computer-based interaction with paper documents. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 587–592, Monterey, California, United States, 1992. ACM.

M. Norrie and B. Signer. Overlaying paper maps with digital information services for tourists. *Information and Communication Technologies in Tourism 2005*, page 23–33, 2005.

M. Norrie, B. Signer, and N. Weibel. General framework for the rapid development of interactive paper applications. In *Proceedings CoPADD*, volume 6, page 9–12, 2006.

R. Noss and C. Hoyles. *Windows on mathematical meanings: Learning cultures and computers*, volume 17. Springer, 1996.

A. M. O'Donnell, D. D. F, H.-L. R, and M. N. *Scripted cooperation in student dyads: A method for analyzing and enhancing academic learning and performance.* Cambridge University Press, 1992.

S. Olberding and J. Steimle. Towards understanding erasing-based interactions: Adding erasing capabilities to anoto pens. In *PaperComp 2010: 1st International Workshop on Paper Computing. Workshop at Ubicomp 2010*, 2010.

C. O'Malley and S. Fraser. Literature review in learning with tangible technologies. Technical report, 2004.

S. Oviatt, A. Arthur, and J. Cohen. Quiet interfaces that help students think. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*, page 191–200, 2006.

S. Oviatt, A. Arthur, Y. Brock, and J. Cohen. Expressive pen-based interfaces for math education. In *Proceedings of the 8th iternational conference on Computer supported collaborative learning*, page 573–582, 2007.

S. Papert. Teaching children to be mathematicians versus teaching about mathematics. *International Journal of Mathematical Education in Science and Technology*, 3(3):249–262, 1972.

S. Papert. *Mindstorms: Children, computers, and powerful ideas.* Basic Books, Inc., 1980.

Y. Park and W. Woo. The ARTable: an AR-Based tangible user interface system, 2006.

J. Patten and H. Ishii. Mechanical constraints as computational constraints in tabletop tangible interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '07, page 809–818, New York, NY, USA, 2007. ACM.

E. R. Pedersen, T. Sokoler, and L. Nelson. PaperButtons: expanding a tangible user interface. In *Proceedings of the 3rd conference on Designing interactive systems: processes, practices, methods, and techniques*, pages 216–223, New York City, New York, United States, 2000. ACM.

R. Perlman. Using computer technology to provide a creative learning environment for preschool children. Technical report, M.I.T., May 1976.

T. Pietrzak, S. Malacria, A. Tabard, and E. Lecolinet. What do u-note? an augmented note taking system for the classroom. In *PaperComp 2010: 1st International Workshop on Paper Computing. Workshop at Ubicomp 2010*, 2010.

A. M. Piper, N. Weibel, and J. D. Hollan. A pen-based toolkit for authoring collaborative language activities. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work Companion*, CSCW '12, page 269–270, New York, NY, USA, 2012. ACM.

A.-M. Piper, N. Weibel, and J. Hollan. Audio-enhanced paper photos: Encouraging social interaction at age 105. In *Proceedings of ACM Conference on Computer-Supported Cooperative Work*, 2013.

B. Piper, C. Ratti, and H. Ishii. Illuminating clay: a 3-d tangible interface for landscape analysis. In *Proceedings of the SIGCHI conference on Human factors in computing systems: Changing our world, changing ourselves*, CHI '02, page 355–362, New York, NY, USA, 2002. ACM.

E. Portocarrero, D. Robert, S. Follmer, and M. Chung. The NeverEndingStorytellingMachine a platform for creative collaboration using a sketchbook and everyday objects. In *PaperComp 2010: 1st International Workshop on Paper Computing. Workshop at Ubicomp 2010*, 2010.

D. Pratt and J. Ainley. The construction of meanings for geometric construction: Two contrasting cases. *International Journal of Computers for Mathematical Learning*, 1(3):293–322, 1996.

S. Price. A representation approach to conceptualizing tangible learning environments. In *Proceedings of the 2nd international conference on Tangible and embedded interaction*, pages 151–158, Bonn, Germany, 2008. ACM.

S. Price and Y. Rogers. Let's get physical: the learning benefits of interacting in digitally augmented physical spaces. *Comput. Educ.*, 43(1-2):137–151, 2004.

K. Probst, T. Seifried, M. Haller, K. Yasu, M. Sugimoto, and M. Inami. Move-it: interactive sticky notes actuated by shape memory alloys. In *Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems*, page 1393–1398, 2011.

H. S. Raffle, A. J. Parkes, and H. Ishii. Topobo: a constructive assembly system with kinetic memory. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '04, page 647–654, New York, NY, USA, 2004. ACM.

R. Rao, S. K. Card, W. Johnson, L. Klotz, and R. H. Trigg. Protofoil: storing and finding the information worker's paper documents in an electronic file cabinet. In *Proceedings of the SIGCHI conference on Human factors in computing systems: celebrating interdependence*, pages 180–185, Boston, Massachusetts, United States, 1994. ACM.

D. Reid. Conjectures and refutations in grade 5 mathematics. *Journal for Research in Mathematics Education*, page 5–29, 2002.

E. Richard, V. Billaudeau, P. Richard, and G. Gaudin. Augmented reality for rehabilitation of cognitive disabled children: A preliminary study. In *Virtual Rehabilitation, 2007*, page 102–108, 2007a.

P. Richard, J. Fortuny, C. Experimentals, M. Hohenwarter, and M. Gagnon. geogebraTUTOR: une nouvelle approche pour la recherche sur l'apprentissage compétentiel et instrumenté de la géométrie à l'école secondaire. In *Proceedings of the World Conference on E-Learning in Corporate, Government, Healthcare, and Higher Education*, page 428–435, 2007b.

J. A. Robinson and C. Robertson. The LivePaper system: augmenting paper on an enhanced tabletop. *Computers & Graphics*, 25(5):731–743, October 2001.

Y. Rogers and J. Ellis. Distributed cognition: an alternative framework for analysing and explaining collaborative working. *Journal of Information Technology*, 9:119–128, 1994.

Y. Rogers and S. Lindley. Collaborating around vertical and horizontal large interactive displays: which way is best? *Interacting with Computers*, 16(6):1133–1152, 2004.

J. Roschelle and S. Teasley. The construction of shared knowledge in collaborative problem solving. *Computer-supported collaborative learning. In C. O'Malley (Ed.)*, page 69–197, 1995.

N. Rouche. La géométrie et la nature des choses. *Proceedings I, Histoire et épistémologie dans l'éducation mathématique, Belgium*, page 43–63, 1999.

K. Ruthven. Instrumenting mathematical activity: Reflections on key studies of the educational use of computer algebra systems. *International Journal of Computers for Mathematical Learning*, 7(3):275–291, 2002.

D. D. Salvucci and J. H. Goldberg. Identifying fixations and saccades in eye-tracking protocols. In *Proceedings of the 2000 symposium on Eye tracking research & applications*, ETRA '00, page 71–78, New York, NY, USA, 2000. ACM.

G. Saul, C. Xu, and M. D. Gross. Interactive paper devices. page 205. ACM Press, 2010.

C. Scherrer, J. Pilet, P. Fua, and V. Lepetit. The haunted book. In *Proceedings of the 7th IEEE/ACM International Symposium on Mixed and Augmented Reality*, ISMAR '08, page 163–164, Washington, DC, USA, 2008. IEEE Computer Society.

E. Schneider. La TI-92 dans l'enseignement des mathématiques, des enseignant (e) s découvrent la didactique des mathématiques. *Actes du congrès "Calculatrices symboliques et géométriques dans l'enseignement des mathématiques*, page 49–60, 1998.

R. Schorr, J. Kaput, M. Heid, and G. Blume. Changing representational infrastructure changes most everything: The case of SimCalc, algebra and calculus. *Research on Technology in the learning and teaching of mathematics: Syntheses and perspectives*, 2007.

A. J. Sellen and R. H. R. Harper. *The Myth of the Paperless Office.* MIT Press, Cambridge, MA, USA, 2003.

O. Shaer and E. Hornecker. Tangible user interfaces: Past, present, and future directions. *Foundations and Trends® in Human–Computer Interaction*, 3(1-2):23–27, 2009.

R. Sidharta, J. Oliver, and A. Sannier. Augmented reality tangible interface for distributed design review. In *Computer Graphics, Imaging and Visualisation, 2006 International Conference on*, page 464–470, 2006.

B. Signer and M. Norrie. Interactive paper: Past, present and future. In *PaperComp 2010: 1st International Workshop on Paper Computing. Workshop at Ubicomp 2010*, 2010.

B. Signer, U. Kurmann, and M. Norrie. iGesture: a general gesture recognition framework. In *Document Analysis and Recognition, 2007. ICDAR 2007. Ninth International Conference on*, volume 2, page 954–958, 2007.

B. Signer and M. C. Norrie. PaperPoint: a paper-based presentation and interactive paper prototyping tool. In *Proceedings of the 1st international conference on Tangible and embedded interaction*, TEI '07, page 57–64, New York, NY, USA, 2007. ACM.

A. Sin and H. Badioze Zaman. Tangible interaction in learning astronomy through augmented reality book-based educational tool. *Visual Informatics: Bridging Research and Practice*, page 302–313, 2009.

J. Smith, J. Long, T. Lung, M. M. Anwar, and S. Subramanian. PaperSpace: a system for managing digital and paper documents. In *CHI '06: CHI '06 extended abstracts on Human factors in computing systems*, pages 1343–1348, New York, NY, USA, 2006. ACM.

A. Soller, A. Martínez, P. Jermann, and M. Muehlenbrock. From mirroring to guiding: A review of state of the art technology for supporting collaborative learning. *International Journal of Artificial Intelligence in Education*, 15(4):261–290, 2005.

H. Song, F. Guimbretière, C. Hu, and H. Lipson. ModelCraft: capturing freehand annotations and edits on physical 3D models. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*, page 13–22, New York, NY, USA, 2006. ACM Press.

H. Song, T. Grossman, G. Fitzmaurice, F. Guimbretiere, A. Khan, R. Attar, and G. Kurtenbach. PenLight: combining a mobile projector and a digital pen for dynamic visual overlay. In

Proceedings of the 27th international conference on Human factors in computing systems, page 143–152, 2009.

H. Song, F. Guimbretière, M. Ambrose, and C. Lostritto. CubeExplorer: an evaluation of interaction techniques in architectural education. In *Human-Computer Interaction – INTERACT 2007*, page 43–56, 2007.

H. Song, F. Guimbretiere, T. Grossman, and G. Fitzmaurice. MouseLight: bimanual interactions on digital paper using a pen and a spatially-aware mobile projector. In *Proceedings of the 28th international conference on Human factors in computing systems*, CHI '10, page 2451–2460, New York, NY, USA, 2010. ACM.

E. Sowell. Effects of manipulative materials in mathematics instruction. *Journal for research in mathematics education*, page 498–505, 1989.

D. Spikol and J. Eliasson. Lessons from designing geometry learning activities that combine mobile and 3D tools. In *Wireless, Mobile and Ubiquitous Technologies in Education (WMUTE), 2010 6th IEEE International Conference on*, WMUTE '10, page 137–141, Washington, DC, USA, 2010. IEEE Computer Society.

D. Stanton, H. Neale, and V. Bayon. Interfaces to support children's co-present collaboration: Multiple mice and tangible technologies. In *CSCL '02: Conf.on Computer Support Collaborative Learning*, page 342–351, 2002.

J. Steimle. Designing pen-and-paper user interfaces for interaction with documents. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*, pages 197–204, Cambridge, United Kingdom, 2009. ACM.

J. Steimle, M. Khalilbeigi, and M. Mühlhäuser. Hybrid groups of printed and digital documents on tabletops: a study. In *Proceedings of the 28th of the international conference extended abstracts on Human factors in computing systems*, page 3271–3276, 2010.

J. Steimle, O. Brdiczka, and M. Muhlhauser. CoScribe: integrating paper and digital documents for collaborative knowledge work. *IEEE Transactions on Learning Technologies*, 2:174–188, 2009.

J. Steimle, O. Brdiczka, and M. Muhlhauser. Digital paper bookmarks: collaborative structuring, indexing and tagging of paper documents. In *CHI '08 extended abstracts on Human factors in computing systems*, pages 2895–2900, Florence, Italy, 2008. ACM.

L. Stifelman. Augmenting real-world objects: A paper-based audio notebook. In *Conference companion on Human factors in computing systems: common ground*, page 199–200, 1996.

R. Sutherland and N. Balacheff. Didactical complexity of computational environments for the learning of mathematics. *International Journal of Computers for Mathematical Learning*, 4 (1):1–26, 1999.

A. Tabard, W. Mackay, and E. Eastmond. From individual to collaborative: the evolution of prism, a hybrid laboratory notebook. In *Proceedings of the ACM 2008 conference on Computer supported cooperative work*, page 569–578, 2008.

T. Tachi. Origamizing polyhedral surfaces. *IEEE Transactions on Visualization and Computer Graphics*, 16(2):298–311, April 2010.

K. T. W. Tan, E. M. Lewis, N. J. Avis, and P. J. Withers. Using augmented reality to promote an understanding of materials science to school children. In *ACM SIGGRAPH ASIA 2008 educators programme*, SIGGRAPH Asia '08, page 2:1–2:8, New York, NY, USA, 2008. ACM.

J. Tang and S. Minneman. VideoDraw: a video interface for collaborative drawing. *ACM Transactions on Information Systems (TOIS)*, 9(2):170–184, 1991.

H. Tap. Interactional features of a paper-based monitoring system. *Personal and Ubiquitous Computing*, 8(3):241–246, July 2004.

M. Terry, J. Cheung, J. Lee, T. Park, and N. Williams. Jump: a system for interactive, tangible queries of paper. In *Proceedings of Graphics Interface 2007*, page 127–134, 2007.

P. Thompson. Mathematical microworlds and intelligent computer-assisted instruction. 1987.

T. Tsandilas, C. Letondal, and W. Mackay. Mus ink: composing music through augmented drawing. In *Proceedings of the 27th international conference on Human factors in computing systems*, page 819–828, 2009.

G. Ucelli, G. Conti, R. Amicis, and R. Servidio. Learning using augmented reality technology: Multiple means of interaction for teaching children the theory of colours. In D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, M. Maybury, O. Stock, and W. Wahlster, editors, *Intelligent Technologies for Interactive Entertainment*, volume 3814, pages 193–202. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.

H. Uchiyama and E. Marchand. Deformable random dot markers. *Mixed and Augmented Reality, IEEE / ACM International Symposium on*, 0:237–238, 2011.

B. Ullmer, R. Sankaran, S. Jandhyala, B. Tregre, C. Toole, K. Kallakuri, C. Laan, M. Hess, F. Harhad, U. Wiggins, et al. Tangible menus and interaction trays: core tangibles for common physical/digital activities. In *Proceedings of the 2nd international conference on Tangible and embedded interaction*, page 209–212, 2008.

B. Ullmer, H. Ishii, and R. J. K. Jacob. Token+constraint systems for tangible interaction with digital information. *ACM Trans. Comput.-Hum. Interact.*, 12(1):81–118, 2005.

J. Underkoffler and H. Ishii. Illuminating light: an optical design tool with a luminous-tangible interface. In *CHI '99 extended abstracts on Human factors in computing systems*, CHI EA '99, page 542–549, New York, NY, USA, 1999. ACM.

D. H. Uttal, K. V. Scudder, and J. S. DeLoache. Manipulatives as symbols: A new perspective on the use of concrete objects to teach mathematics. *Journal of Applied Developmental Psychology*, 18(1):37–54, 1997.

D. Van Krevelen and R. Poelman. A survey of augmented reality technologies, applications and limitations. *International Journal on Virtual Reality*, 9:1–20, 2010.

L. Vygotski. *Mind in society: The development of higher psychological processes.* Harvard Univ Pr, 1978.

M. Waldner, J. Hauber, J. Zauner, M. Haller, and M. Billinghurst. Tangible tiles: design and evaluation of a tangible user interface in a collaborative tabletop setup. In *Proceedings of the 18th Australia conference on Computer-Human Interaction: Design: Activities, Artefacts and Environments*, page 151–158, 2006.

F. Wang and M. Hannafin. Design-based research and technology-enhanced learning environments. *Educational Technology Research and Development*, 53(4):5–23, December 2005.

R. Want, K. P. Fishkin, A. Gujar, and B. L. Harrison. Bridging physical and virtual worlds with electronic tags. In *Proceedings of the SIGCHI conference on Human factors in computing systems: the CHI is the limit*, CHI '99, page 370–377, New York, NY, USA, 1999. ACM.

R. Watanabe, Y. Itoh, M. Asai, Y. Kitamura, F. Kishino, and H. Kikuchi. The soul of ActiveCube: implementing a flexible, multimodal, three-dimensional spatial tangible interface. *Comput. Entertain.*, 2(4):15–15, 2004.

N. Weibel. *A publishing infrastructure for interactive paper documents.* PhD thesis, SWISS FEDERAL INSTITUTE OF TECHNOLOGY ZURICH, 2009.

N. Weibel, A. Ispas, B. Signer, and M. C. Norrie. Paperproof: a paper-digital proof-editing system. In *CHI '08: CHI '08 extended abstracts on Human factors in computing systems*, page 2349–2354, New York, NY, USA, 2008. ACM.

M. Weiser and J. S. Brown. The coming age of calm technology. 1996.

P. Wellner. The DigitalDesk calculator: tangible manipulation on a desk top display. In *Proceedings of the 4th annual ACM symposium on User interface software and technology*, pages 27–33, Hilton Head, South Carolina, United States, 1991. ACM.

P. Wellner. Interacting with paper on the DigitalDesk. *Communications of the ACM*, 36(7): 87–96, 1993.

D. West, A. Quigley, and J. Kay. MEMENTO: a digital-physical scrapbook for memory sharing. *Personal Ubiquitous Comput.*, 11(4):313–328, 2007.

A. Wilson. PlayAnywhere: a compact interactive tabletop projection-vision system. In *Proceedings of the 18th annual ACM symposium on User interface software and technology*, page 83–92, 2005.

R. Yeh, C. Liao, S. Klemmer, F. Guimbretière, B. Lee, B. Kakaradov, J. Stamberger, and A. Paepcke. ButterflyNet: a mobile capture and access system for field biology research. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, CHI '06, page 571–580, New York, NY, USA, 2006. ACM.

R. B. Yeh, A. Paepcke, and S. R. Klemmer. Iterative design and evaluation of an event architecture for pen-and-paper interfaces. page 111. ACM Press, 2008.

S. Yuen, G. Yaoyuneyong, and E. Johnson. Augmented reality: An overview and five directions for AR in education. 2011.

Z. Zhang, Y. Wu, Y. Shan, and S. Shafer. Visual panel: virtual mouse, keyboard and 3D controller with an ordinary piece of paper. In *Proceedings of the 2001 workshop on Perceptive user interfaces*, page 1–8, 2001.

O. Zuckerman and M. Resnick. System blocks: A physical interface for system dynamics simulation. In *Proceedings of CHI '03*, 2003.

G. Zufferey. *The Complementarity of Tangible and Paper Interfaces in Tabletop Environments for Collaborative Learning*. PhD thesis, 2010.

G. Zufferey, P. Jermann, A. Lucchi, and P. Dillenbourg. TinkerSheets: using paper forms to control and visualize tangible simulations. In *TEI'09, Third International Conference on Embedded and Tangible Interaction*, 2009.

Quentin Bonnard

Rue de Lausanne 19
1020 Renens
Suisse
✉ quentin.bonnard@epfl.ch

Education

2008-2012 **Ph.D. candidate in Computer Science**, *"Paper Interfaces: an HCI Approach to Geometry Education"*.
EPFL: Swiss Federal Institute of Technology of Lausanne, Switzerland

2003–2008 **Master of Engineering**, *in Information Technology and Computer Engineering*, summa cum laude.
"Grande École" INSA: National Institute of Applied Sciences in Lyon, France

2007–2008 **Master of Science**, *"Detection of Objects in Pictures using Graphical Models"*, magna cum laude.
École Centrale Lyon, France

2006–2007 **Erasmus Exchange**, *GPA: 83/100*.
Trinity College Dublin, Ireland

2003–2005 **Vordiplom Informatik**, *two year degree in Information Technology and Computer Engineering*, magna cum laude.
University of Karlsruhe (TH), Germany

2000–2003 **French scientific baccalauréat and German Abitur**, *major in mathematics*, magna cum laude.
European High school Charles De Gaulle, Dijon, France

Work Experience

Since 09.2008 **Assistant at CRAFT**, *EPFL*, Switzerland.
Design and production of a new series of interactive tables, including a novel interaction method. Assistance in preparing and running a Java project course. Supervision of bachelor and master projects.

06.2007 09.2007 **Software Development Engineer in Test intern**, *Microsoft*, Redmond, USA.
Research and realisation with Spec# of a model to automatically generate automated test cases. As an introduction, design and implementation in C# of a simple test tool, realisation of a test plan, test cases, and manual test passes.

06.2006 08.2006 **Java developer intern**, *Editions Diagonal*, Sophia-Antipolis, France.
Design and implementation in Java of a reusable architecture for word games. Implementation of application, applet and server.

10.2004 08.2005 **Teaching Assistant in Computer Science**, *Faculty of Informatics*, University of Karlsruhe, Germany.
Weekly classes of 1½ hours theory and 1½ hours lab work. Weekly correction of the assignments of about 20 students.

Publications

Full Paper	**Q. Bonnard**, P. Jermann, A. Legge, F. Kaplan and P. Dillenbourg. (2012). *Tangible Paper Interfaces: Interpreting Pupils' Manipulations.* Interactive Tabletops and Surfaces (ITS), Cambridge, MA, USA. Accep. $< 30\%$, best paper award nominee.
Full Paper	**Q. Bonnard**, H. Verma, F. Kaplan and P. Dillenbourg. (2012). *Paper Interfaces for Learning Geometry*. 7th European Conference on Technology Enhanced Learning (EC-TEL), Saarbrücken, Germany. Accept. $< 20\%$, best paper award.
Workshop Paper	**Q. Bonnard**, F. Kaplan and P. Dillenbourg. (2012). *Paper Interfaces to Support Pupils and Teachers in Geometry.* Digital Ecosystems for Collaborative Learning (DECL). Workshop in the International Conference of the Learning Sciences (ICLS), Sydney, Australia.
Abstract	**Q. Bonnard**, A. Legge, F. Kaplan and P. Dillenbourg. (2012). *Hands-on Symmetry with Augmented Reality on Paper*. 9th International Conference on Hands-on Science (HSci), Antalya, Turkey.
Full Paper	P. Dillenbourg, G. Zufferey, H. S. Alavi, P. Jermann, S. Do-Lenh, **Q. Bonnard**, S. Cuendet, F. Kaplan. (2011). *Classroom orchestration : The third circle of usability.* 9th International Conference on Computer Supported Collaborative Learning (CSCL), Hong Kong, China.
Poster	S. Cuendet, **Q. Bonnard**, F. Kaplan and P. Dillenbourg. (2011). *Paper Interface Design for Classroom Orchestration.*, 29th ACM Conference on Human Factors in Computing Systems (CHI), Vancouver, BC, Canada.
Workshop Paper	**Q. Bonnard**, A. Mazzei, F. Kaplan and P. Dillenbourg. (2010). *Vocal Sticky Notes: Adding Audio Comments on Paper Documents*. 28th ACM Conference on Human Factors in Computing Systems (CHI), Atlanta, GA, USA.
Poster	**Q. Bonnard**, F. Kaplan and P. Dillenbourg. (2010). *A Paper Interface for Code Exploration*. 1st International Workshop on Paper Computing (PaperComp). Workshop in the 12th ACM International Conference on Ubiquitous Computing (UbiComp), Copenhagen, Denmark.

Interests

Languages	French (mother tongue), English (fluent, TOEIC: 950/990), German (fluent), Japanese (beginner), Arabic (basic), Latin (basic)
Music	8 years oboe
Sport	Trampoline, swimming, underwater rugby...