Latent Social Information in Group Interactions with a Shared Workspace

THÈSE N^{O} 6616 (2015)

PRÉSENTÉE LE 28 AOÛT 2015

À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS
LABORATOIRE D'ERGONOMIE ÉDUCATIVE
PROGRAMME DOCTORAL EN INFORMATIQUE ET COMMUNICATIONS

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Himanshu VERMA

acceptée sur proposition du jury:

Prof. W. Gerstner, président du jury
Prof. P. Dillenbourg, directeur de thèse
Prof. Y. Rogers, rapporteuse
Prof. D. Lalanne, rapporteur
Prof. F. Kaplan, rapporteur

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Suisse
2015

To Béchamel Sauce, and its inherent sophistication and yet delicate humbleness,
that is the simplest ingredient for happiness.

Acknowledgements

This dissertation is one more step forward towards a better understanding of collaborative work, and it owes its own realization to a concerted collaborative endeavor, involving the efforts, contributions, and aspirations of many, who I had the honor to have worked with during my life as a PhD student. First and foremost, I wish to extend my sincere gratitude and utmost appreciation to my supervisor Pierre Dillenbourg, who ceaselessly provided me with academic support and advice during my studies. Pierre has the distinguished qualities of a great teacher through which he allowed me to explore and go beyond the conventional, while continuously encouraging, inspiring, and guiding me towards the successful conclusion of this adventure.

I would like to extend my earnest gratitude to the members of my review committee: Yvonne Rogers, Denis Lalanne, Frédéric Kaplan, and Wulfram Gerstner. Thank you for investing your time and effort to evaluate this research work, and your insightful comments and suggestions have helped me to make the manuscript better.

Special thanks goes to Patrick Jermann for his inexhaustible spirit for discussions and exchange of ideas, as well as his readiness to offer counsel and support through his knowledge of varied subjects, has significantly contributed to the fulfillment of my research objectives.

Research is a venture that continues on a winding road, full of twists and turns and numerous detours. I had the honor and opportunity to find excellent companions on this journey, who have helped me through choices to find the right direction. Many thanks to Hamed Alavi, Jan Blom, Guillaume Zufferey, and Łukasz Kidziński for your profound thinking and judgements that enabled me to conceptualize ideas and take the research objectives through all avenues to completion.

The different prototypes of *MeetHub* - the meeting technology introduced in this dissertation, would not have existed without the substantial assistance of of Cristián Alcoholado Moenne and Olivier Guédat. Also, thanks to Nan Li and Flaviu Roman for your help in extending the features and applications of our meeting technology.

The conducted user-studies wouldn't have been possible without the invaluable cooperation of professors Martin Vetterli, Paolo Prandoni, Marco Picasso, as well as their students. I would also like to extend my gratitude to the students of CSCW (2011, 2012, and 2013) course for their formidable enthusiasm and patience while participating in the studies.

To the members of CRAFT (now CHILI), I owe a substantial part of my admiration and gratitude, especially those who made a significant contribution towards the attainment of my PhD, and with whom I have had a pleasure of collaborating directly. Many thanks to Andrea Mazzei, Nan Li, Kshitij Sharma, Afroditi Skevi, Daniela Caballero Díaz, and Silvia Magrelli for your solemn dedication for the projects, and scintillating cheerfulness that resulted into a remarkably cheerful time spent while working.

I have been incredibly lucky to be blessed with amazing colleagues and friends in the lab, who brought joyfulness even to the most bleak days, through their charm and their propensity for the poetry called life, and I can proudly call you my "family away from home". My sincere admiration and gratefulness goes to Julia Fink, Quentin Bonnard, Hamed Alavi, Andrea Mazzei, Kshitij Sharma, Sébastien Cuendet, Valerie Bauwens, and Khaled Bachour for your friendship from the beginning of my PhD days till the end. You guys made my days in Lausanne a memorable and unforgettably happy days of my life. My heartfelt thanks goes to Guillaume Zufferey, Jessica Dehler, Huong-Ly Mai, Farzaneh Bahrami, Tabea Koll, Lorin Cuendet, Nisha Dalal, Son Do Lenh, Mirko & Sonja Raca, Daniela Caballero Díaz, Cristián Mansilla, Sophia Schwär, Guillaume & Jo'an Bardy, Lorenzo Lucignano and Luis Prieto for being my extended family in Lausanne, and for the crazily funny conversations shared over coffee, food, beer, games, hikes, and so many more occasions. Many thanks to you and all the other fellow colleagues from CHILI and CEDE: Mina Shirvani, Ayberk Ozgur, Séverin Lemaignan, David Bréchet, Beat Schwendimann, Gilles Raimond, Gwénaël Bocquet, David Meister, and Cibelle Avelino. I would also like to acknowledge the amazing CAPE team, especially Roland Tormey, Ingrid Le Duc, and Siara Isaac for many exhilarating discussions on numerous occasions.

Special thanks to Kshitij for being an awesome flatmate, amazing cook, and for your eternal optimism that brought me hope, clarity, and inspiration during the chaotically foggiest moments of my life. These acknowledgements are literally incomplete without mentioning the "secret ingredient" and the "backbone" of the lab, Florence Colomb, to whom I am sincerely grateful for her effervescent playfulness and dignified seriousness that catalyzes the fluid functioning of the group, and significantly increases its awesomeness index.

Last but not the least, to my late grandfather, my mother and father, and my sister, thank you for your blessings, wisdom, sacrifices, and most of all your everlasting love, that always gave me the strength to keep on walking forward, for your admiration and faith in my dreams, and your eternal support in letting me pursue them. My sincere gratitude to your untiring stamina, divine curiosity, and subtle patience, that I forever try to imbibe as my core values and virtues.

Lausanne, le 16 Juillet 2015 Himanshu Verma

Abstract

Shared artifacts, such as drawings and schemas on whiteboards, sticky-notes with ideas on walls, are often created and interacted with during meetings. These shared artifacts *a*) facilitate the expression of complex fleeting ideas, *b*) enable collaborators to establish a common ground and validate each others' understanding about the context, and *c*) extend the validity of shared information by making it permanent. By the end of a collaboration session, the shared content denotes the shared knowledge amongst collaborators, which emerged as a result of a recursive process of storage, transformation, and retrieval from an external memory such as a whiteboard. Although these interactions with the artifacts symbolize the important episodes in a group discussion, still the information contained within them has not been much leveraged in collaboration research.

Being well assimilated in the established work culture, collaborators do not heed the interactions with the shared artifacts, and therefore the nature of the social information contained within them is *latent*. However, from a research perspective this information is valuable and can offer insights into a few facets of ongoing group dynamics and processes. This thesis in particular *a*) identifies and examines the characteristics of the latent social information, *b*) studies the relationship of this information with different aspects of collaboration, and *c*) explores the practical utility of this information in collaboration assessment.

We start by designing a meeting technology - *MeetHub* that enables collaborators to share and interact with artifacts in an unconstrained manner over a shared workspace, and allow us to collect fine-grained interactional information. Then we present user studies, where we extract and comprehend the relevant social information from interactions with the artifacts, and analyze its relationship with collaborative processes. Our findings demonstrate that latent social information is significantly correlated with the task outcome, division-of-labor, and the quality of mutual understanding between collaborators. Finally, we present a prediction system based on this social information, capable of alerting the group members about the poorly grounded episodes in real-time, and thus enabling them to regulate their collaborative behavior. The final contribution of this work presents itself as implications towards the *dual* nature of shared workspace, supporting the creation and sharing of artifacts as well as an assessment tool.

Key words: *Computer-Supported Collaborative Work, Latent Social Information, Shared*

Workspaces, Artifacts, Interactions

Résumé

Les artefacts partagés, comme les dessins ou schémas sur un tableau blanc, ou les notes collées sur un mur, sont fréquemment créés et utilisés pendant les réunions. Ces artefacts partagés *a*) facilitent l'expression d'idées complexes, *b*) permettent aux collaborateurs d'établir une base commune et de valider la compréhension de chacun du contexte, et *c*) d'étendre la validité des informations partagées en les rendant permanentes. À la fin d'une session de collaboration, le contenu partagé dénote le savoir partagé par les collaborateurs, et émergeant du processus récursif de stockage, transformation et récupération d'une mémoire externe comme un tableau blanc. Bien que ces interactions avec les artefacts symbolisent des épisodes dans une discussion de groupe, les informations qu'elles contiennent n'ont pas été beaucoup exploitées dans la recherche sur la collaboration.

Comme les interactions avec les artefacts partagés sont bien assimilées par les professionnels, les collaborateurs n'y font pas attention, et les informations qu'elles contiennent sont *latentes*. Toutefois, du point de vue de la recherche, ces informations sont importantes et peuvent fournir des indications sur certaines facettes des processus et dynamique de groupe. En particulier, cette thèse *a*) identifie et examine les caractéristiques des informations sociales latentes, *b*) étudie la relation de ces informations avec différents aspects de la collaboration, et *c*) explore l'utilité pratique de ces informations dans l'évaluation de la collaboration.

Nous commençons par concevoir une technologie pour les réunions, *MeetHub*, qui permet aux collaborateurs de partager et interagir sans contrainte avec les artefacts dans un espace de travail partagé, et nous permet de récolter des informations précises sur l'interaction. Ensuite nous présentons des études utilisateurs, desquelles nous extrayons et saisissons les informations sociales pertinentes des interactions avec les artefacts, et analysons leur relation avec les processus collaboratifs. Nos résultats montrent que les informations sociales latentes sont corrélées significativement avec l'aboutissement de la tâche, la répartition du travail, et la qualité de compréhension réciproque entre les collaborateurs. Enfin, nous présentons un système de prédiction basé sur ces informations sociales, capable d'alerter en temps réel les membres du groupe sur les épisodes faiblement fondés, ce qui permet de réguler le comportement collaboratif. La contribution finale de ce travail se présente sous la forme de conséquences sur la nature duelle des espaces de travail partagés, qui supportent la création et le partage d'artefacts ainsi que d'un outil d'évaluation.

Mots clefs : *Collaboration assistée par ordinateur, Informations sociales latentes, Espaces de travail partagé, Artefacts, Interactions*

Zusammenfassung

In Meetings werden häufig geteilte Artefakte, wie Skizzen auf Whiteboards, oder Notizzettel auf Wänden, kreiert und gemeinsam bearbeitet. Diese geteilten Artefakte *a*) erleichtern, komplexe und flüchtige Ideen auszudrücken, *b*) helfen den Beteiligten, ein gemeinsames Verständnis aufzubauen und sich das Verständnis gegenseitig zu bestätigen, und *c*) machen die geteilten Informationen langfristig nutzbar. Am Ende einer Zusammenarbeitssitzung stellen die geteilten Inhalte das unter den Gruppenmitgliedern geteilte Wissen dar, welches aus einem rekursiven Prozess des Speicherns, Transformierens, und Abrufens aus einem externen Gedächtnis wie z.B. einem Whiteboard hervorgeht. Obwohl diese Interaktionen mit dem Artefakt wichtige Episoden in der Gruppendiskussion darstellen, wurde die im Artefakt enthaltene Information bisher in der Forschung nicht ausgenutzt.

Die soziale Information des Artefaktes ist von *latenter* Natur, da die Gruppenmitglieder die Interaktionen mit dem Artefakt nicht bewusst beachten, weil diese gut an die bestehende Arbeitskultur angepasst sind. Von einer Forschungsperspektive her haben diese Informationen jedoch hohen Wert und können Einsichten in einige Facetten der Gruppendynamik und -prozesse bieten. Diese Thesis *a*) identifiziert und untersucht die Eigenschaften der latenten sozialen Information, *b*) erforscht die Zusammenhänge zwischen dieser Information und verschiedenen Aspekten der Kollaboration, und *c*) erkundet die praktische Nützlichkeit dieser Information für das Assessment der Kollaboration.

Zunächst wird eine Meeting-Technologie entwickelt - *MeetHub*, die Gruppenmitgliedern erlaubt, ohne Beschränkungen Artefakte in einem gemeinsamen Arbeitsplatz zu teilen und zu bearbeiten. Dies erlaubt uns, detaillierte Informationen über die Interaktion zu erfassen. Danach werden Studien präsentiert, in denen wir die relevanten sozialen Informationen aus den Interaktionen mit den Artefakten extrahieren und deren Zusammenhänge mit den Gruppenprozessen analysieren. Die Resultate zeigen, dass latente soziale Informationen signifikant mit dem Aufgabenergebnis, der Arbeitsteilung und der Qualität des gegenseitigen Verständnisses zwischen den Gruppenmitgliedern korrelieren. Zuletzt präsentieren wir ein auf diesen sozialen Informationen basierendes Vorhersage-System, welches Gruppenmitglieder in Echtzeit in Episoden wenig geteilten Verständnisses warnt. Dies erlaubt den Gruppenmitgliedern, ihren Gruppenprozess zu regulieren. Die Implikationen dieser Arbeit richten sich auf die duale Natur von geteilten Arbeitsplätzen, die sowohl das Schaffen und Teilen von Artefakten als auch das Assessment unterstützen.

Stichwörter: *Computer-unterstütztes kollaboratives Arbeiten, Latente Soziale Information, geteilte Arbeitsplätze, Artefakte, Interaktionen*

Contents

List of Figures

List of Tables

1 Introduction

> "*L'essentiel est invisible pour les yeux.*"
> "What is essential, is invisible to the eye."
> LE PETIT PRINCE by *Antoine de Saint-Exupéry*

The invention of the *microscope* in the 17th century, was a significant step forward for the field of medical sciences, as it empowered scientists to study the human body from a completely new perspective, by making visible what happens underneath the surface, and see the smallest intricacies of human body. Since its inception, microscope has contributed significantly to the expansion of the field of modern medicine, and is even used in numerous fields other than medicine. Modern day clinical pathologists use the microscope to diagnose diseases through the analysis of tissues, cells, and body fluids. To study the cause of diseases, they use the analysis of the microscopic structures and organisms, as well as the thorough investigation of development mechanisms of pathogens, and the structural alterations of cells.

We use the example of a clinical pathologist and her microscope as an analogy to describe our approach of analyzing collocated collaboration, and the assessment of group processes. Pathology is the study of the properties and intricacies of human body that are invisible to the human eye, but still these intricacies significantly contribute to the well-being of an individual. Similarly in collocated collaboration environments, there is *latent social information* - the social information that is present but not evident, in the interactions between collaborators and artifacts. While the interactions between individuals has been exploited by researchers in the field of social signal processing (SSP), little is known about the collaborators' interactions with environmental objects such as documents and artifacts. In this dissertation, we emphasize on studying the interactions between collaborators and content artifacts.

Creating and sharing of content during meetings is an established work practice and part of the organizational culture. It facilitates better mutual understanding of the topic under discussion, and allows group members to consensually validate each others' ideas and understanding, which is crucial for a decision-making process. In addition, creation of drawings and sketches on a whiteboard enables an easy way to express complex ideas, and the permanence of artifacts on a physical medium renders the information valid for a longer duration, which won't be possible otherwise due to the ephemeral nature of verbal interactions. Collaborators' interactions with shared workspaces (such as whiteboard) also signify important episodes of collective information retrieval, processing, and storage, where the shared workspace plays the role of an external group memory. However, these interactions are not thoroughly explored and leveraged in CSCW (Computer-Supported Collaborative Work) research, even though they contain valuable information that is representative of ongoing group activity. The social information contained within these interactions is latent and the collaborators are not explicitly aware of this information because the practice of interacting with the artifacts is deeply rooted in the work culture, and individuals do not explicitly think about them.

Similar to a pathologist's microscope we also need a tool that can enable us to visualize and analyze the latent information in collaborators' interactions with artifacts. Our choice of an appropriate meeting technology will depend on the presence of a shared workspace where artifacts can be accessed and shared freely. The requirement of unconstrained access to the workspace is highly relevant, because the fact that individuals have to compete to gain access would restrain the group dynamics. Conversely, unconstrained access to the workspace may allow for varied group dynamics to emerge naturally depending on other factors such as task types, roles, etc., which enables us with a broad set of dynamics to analyze. Single Display Groupware (SDG) [Stewart et al., 1999] is an example of a meeting technology that satisfies our requirements. We design and develop a SDG in our research that serves a dual purpose of supporting group's content creation and sharing activity, and acts as a data collection tool for us to examine the hidden information in interactions with the shared workspace.

Finally, the methodology we use to analyze the group's interactions with the shared artifacts is thematically similar to the one that is used by a clinical pathologist to diagnose a disease which involves examination, identification and recognition, treatment, and prediction. Through our research methodology, we aim to *render the invisible, visible*, by using the hidden information to better understand the visible collaborative behavior, and comprises of the following steps:

- *Examination* of the different kinds of interactions with the shared workspace.
- *Identification* of the nature and properties of this hidden social information.
- *Recognition* of the relevant properties that might explain a few aspects of the ongoing activity and dynamics.
- Conducting user studies with different *treatments* to analyze the relationships of this latent information with ongoing collaborative behavior.
- Utilizing this latent social information to *predict* or *model* some aspects of collaboration.

The challenge manifests itself in obtaining an alternative set of variables that can be used to explain and model the collaborative behavior, which may also compliment the existing variables used in collaboration assessment such as the verbal interactions and gestures. In addition, as the meeting tool can analyze these variables in real-time, they can be used to render the process of collaboration assessment significantly quicker and less tiresome for the researchers, and pave the way for real-time automatic assessment. This is one of the main motivation behind our research work.

1.1 Organization of this Thesis

The remainder of the thesis is organized in the following way:

- We begin by presenting a brief review of the relevant literature in **Chapter 2** and **Chapter 3** in order to effectively situate our work. **Chapter 2** is dedicated to Computer Supported Collaborative Work (CSCW), Social Signal Processing, and Socially Shared Cognition. In **Chapter 3**, we review the role of Distributed Cognition in the design of collaborative systems, and examples of Single Display Groupware.

- In **Chapter 4**, we describe the context of this research work, as well as the general research questions that we aim to answer by means of this dissertation.

- The first version of our meeting technology - *MeetHub* is presented in **Chapter 5**. In addition, we present a user-study to analyze the role of different input modalities on the kind of shared representations produced during collaboration, as well as their influence on the collaborative behavior.

- Based on the design lessons learned from our first user study, we present the re-designed version of *MeetHub* in **Chapter 6**. In addition, we present our second user-study where we extract and identify the different characteristics of latent social information in group's interaction with the shared workspace, and also analyze their relationship with different aspects of collaboration.

- Our third user-study, where we examine the effects of latent social information on the collaboration outcome, is presented in **Chapter 7**.

- Group's interactions with shared artifacts of a different kind - MOOC video lectures, is studied in study group settings in **Chapter 8**. In addition, we investigate the role of different video watching configurations and groupware on the group's interactivity with the video lectures.

- **Chapter 9** discusses the practical utility of the latent social information as a prediction system that assesses the quality of collaborative episodes and presents the groups with an appropriate feedback.

- Finally, we conclude with a summary and general discussions about the contributions of this research work in **Chapter 10**. At the end, we discuss the limitations and the implications of this dissertation for future research.

2 Literature Review

2.1 Where do we stand?

Interdisciplinary research such as this one is situated at the intersection of several domains, as illustrated in Figure 2.1. It reports on the design and implementation of a meeting technology facilitating content creation and manipulation within small group collocated meetings. In addition, the meeting technology enables the extraction, interpretation and finally visualization of contextually implicit social information; as a group mirror providing awareness of the

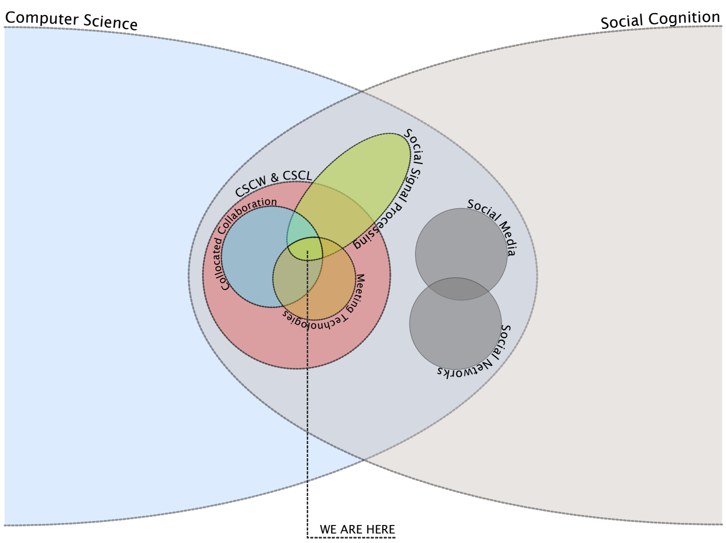

Figure 2.1 – *Where do we stand?* - Graphically situating this thesis. Besides *CSCW*, *Social Media* and *Social Networks* are shown as an example of other domains which lie at the intersection of Computer Science and Social Cognition. However, this thesis work is not related to Social Media and Networks.

underlying collaborative processes to the group.

The questions that we wish to raise and answer through this thesis lie under the umbrella of **Computer Supported Collaborative Work (CSCW) & Computer Supported Collaborative Learning (CSCL)** in the context of face-to-face (or collocated) collaboration. In addition, we take inspiration from the principles and methodologies in **Social Signal Processing (SSP)** to understand collaborative behavior and visualize relevant information to facilitate regulation of group's behavior. The principles of social signal processing will allow us to comprehend collaborators' non-verbal interactions, and use it to say something about the collaborative behavior of the group. Therefore, I will review the research works from these fields in this chapter, which has pertinence to this work. I will also delineate the theories and frameworks in **Social Cognition**, which are required to position this work.

2.2 Computer Supported Collaborative Work

Thirty years ago, Iren Greif and Paul Cashman coined the term **Computer Supported Collaborative Work (CSCW)** in an interdisciplinary workshop, to investigate technology's role in how people collaborate [Grudin, 1994a]. Interchangeably, CSCW has also been referred to as **Collaborative Computing** by Borenstein [1992]; Schooler [1996], and defines the use of computers to support coordination and cooperation within a group, performing a task together.

So, what does CSCW encompass? This question is answered in an extensive review on characterization of collaboration and collaborative systems by Bafoutsou and Mentzas [2002]. The authors broadly position CSCW research into three classes: *a*) *time* and *space* of collaboration; *b*) collaborative task type; and *c*) group characteristics . Independently or in a combination, these factors influence the attributes and functionality of the technology supporting collaboration [Tang, 1991]. I will next describe the research contributions in each of these dimensions.

Time and Location of Collaboration: The taxonomy presented by DeSanctis and Gallupe [1987] and Ellis et al. [1991] classifies collaboration based on *when* (synchronous or asynchronous) and *where* (collocated or distributed) the collaborators interact as shown in Figure 2.2. DeSanctis and Gallupe [1987] use this taxonomy to position various features of GDSS (Group decision support system) technology in each of the four quadrants (visualizing time and location along the horizontal and vertical axes). Grudin [1994a] further extends this categorization by including the factor of predictability; i.e. whether the collaborators are aware of others' location and time of collaboration.

Collaborative Task Type: The purpose of collaboration can be to *generate* (ideas or alternatives), *choose* (alternatives), *negotiate* or *execute* [McGrath and Hollingshead, 1994]. These collaborative intentions were mapped onto six different task types by DeSanctis and Gallupe [1987]; which are - planning, creativity, intellective, preference, cognitive conflict and mixed-

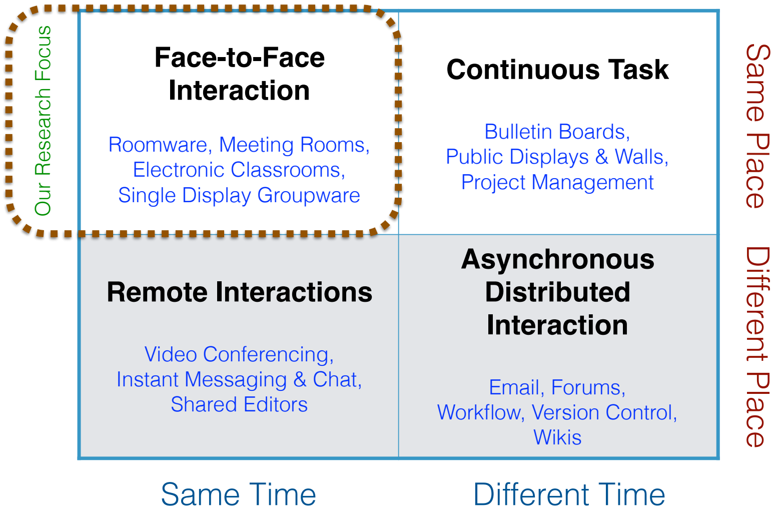

Figure 2.2 – *Time-Space Taxonomy* - Ellis et al. [1991] classified the collaboration based on the location and time. We also present some examples of groupware and technologies employed in a specific class of collaboration.

motive. In addition, Malone and Crowston [1994] used coordination theory as a basis for classifying task types into six categories concerned with the management of shared resources, producer/consumer activity, simultaneity constraints, task/subtask relationship, decision making and communication.

Group: Next, regarding the group itself as the subject of analyses, researchers have studied various factors concerning the group. In their taxonomy of collaborative systems, DeSanctis and Gallupe [1987] elaborate the role of group size (smaller versus larger groups as two distinct categories) on the properties and functionalities of GDSS; besides collaborators' proximity and task. In addition, the authors integrated these three factors (group size, task and proximity of collaborators) in a contingency framework which can be used to implicate the design of collaborative systems. Also, Mullen [1991] conducted a meta-analyses on perception of belongingness to a group; i.e. the sense of belonging to "we" instead of "they" in a group. His findings suggest that group composition, salience and cognitive representations may contribute to various group processes such as cohesiveness, participation-leadership effect, social projection, and ingroup bias. Furthermore, groups collaborating in the presence of a human facilitator were observed to perform better and experienced a higher sense of cohesion within the group, as compared to the groups without a facilitator [Anson et al., 1995]. The authors also identified that in groups with both a facilitator and a group support system (GSS), both the facilitator and GSS enhanced each other's influence on cohesion and group dynamics.

The complex interaction between these aforementioned factors define the role and functionality of the technology designed to support collaboration [Tang, 1991]. However, the task of designing technology supporting effective collaboration and their evaluation is a very difficult and complex one because of the multiplicity, diversity and significance of these factors [Neale et al., 2004]. The variables that need to be regarded while designing a groupware involve individual cognitive factors, cooperative and collaborative factors, usability issues for groups and individual, social, and organizational impact of the technology. This also accounts for the relatively low penetration of groupware in organizations when compared to single-user systems. Grudin [1989, 1994b] studied the challenges faced by the groupware designers from an organizational perspective. The author identified that the additional work required to effectively use these systems, and the perceived benefits to only a minority of organizational population are the main reasons for the failure of groupware. In addition, the author suggested that groupware designers and researchers should focus on incorporating user-centered design methodology (refer to Norman and Draper [1986]) by considering the group and organization as part of the design process. Taking into account the limitations and complexity involved in the design process of groupware, Mandviwalla and Olfman [1994] have specified seven generic guidelines which consider providing supporting for varied task types, multiple collaboration methods, and fluid interchangeability of interaction techniques. In addition to these requirements, the groupware should be flexible enough to adjust to group's context and sustain multiple behavioral characteristics.

The multiplicity and diversity of these factors also demonstrate the vastness of CSCW as a whole. Therefore in order to further position this thesis among these variables, I will use the classification by DeSanctis and Gallupe [1987]. Figure 2.1 shows that this thesis work examines real-time collaboration with collaborators in close proximity and technology supporting such collaboration. Besides these factors, this thesis work focuses on collaboration between small groups performing short tasks in the presence of an experimenter. This distinction is necessary to distinguish this research work from recurrent project management based collaboration consisting of multiple teams distributed in both space and time. Therefore, the next sections will focus on relevant research on collocated collaboration (see Section 2.2.1) and technologies (see Section 2.2.2) supporting it.

2.2.1 Collocated Collaboration

The close physical proximity among the group members during face-to-face interaction allows for peripheral awareness of the other collaborators and their actions, and this is used as a key resource by the group members to mediate their interaction [Tang, 1991]. On the contrary, in distributed (or distant) collaboration, Cummings and Kiesler [2008] identified that the lack of peripheral awareness and increased cost of coordination makes it harder for collaborators to sustain a strong working relationship. In addition to peripheral awareness, the access to numerous other behavioral cues such as gaze, gestures and posture put collocated (or face-to-face) collaboration at an advantage over distributed collaboration. In this section, I will review the previous research works concerning collocated collaboration, with few implications to distributed collaboration. In the rest of this thesis, I will use the terms "collocated", "copresent" and "face-to-face" interchangeably when referring to collaboration. Also, the term "distributed" collaboration will always indicate to distant or not collocated collaboration.

Awareness and Coordination

Dourish and Bellotti [1992] defined awareness as "*an understanding of the activities of others, which provides a context for your own activity*". This context enables collaborators to evaluate their actions and contributions so that they are well aligned with the group's objectives and dynamics. The primary role of awareness in collaboration is to bring forth an individual's activity to the group [Dourish, 1997]. Dourish and Bellotti [1992] have highlighted the relevance of passive mutual monitoring of others and the need for coordination towards successful collaboration. The authors provide a detailed account of two ways in which CSCW systems provide awareness to groups: *informational* and *role-restrictive*. Informational awareness refers to collaborators' actions that are meant to inform others of their activity; for example a changelog written by a collaborator on a version control system. Unlike informational awareness that concerns with the content of collaboration, role restrictive awareness concerns with the nature of the activity and refers to the information about the support for various roles within collaboration.

The criticality of workspace awareness is emphasized in the observational study of collaboration in the line control room for London's underground system, by Heath and Luff [1992]. They observed that the personnels in the control room explicitly make other collaborators aware of their activities, while also monitoring the actions of others. This is done for the purpose of better coordination among collaborators and to reduce any chances of error or misjudgment. Carroll et al. [2003] classified awareness into three types of awareness: *social awareness* refers to knowledge about who is present during interaction and their current state, *action awareness* refers to what others are doing, and *activity awareness* refers to the global task activity and knowledge of its progress.

In collocated collaboration, this awareness information arises automatically from each individual's activity, and group members are not required to explicitly seek for this information. On the contrary, in distributed collaboration, this information is managed explicitly by the collaborative awareness tools [Dourish and Bellotti, 1992; Dourish, 1997]. In collocated collaboration, the ease of access to awareness information does not necessarily imply that this information is effectively utilized by the groups to their benefits. Researches on awareness tools in collocated collaboration such as Meeting Mediator [Kim et al., 2008] and Reflect Table [Bachour et al., 2010] have demonstrated the importance of bringing out the invisible information (such as participation balance, speaking time, etc.) to group's notice, and their benefits on collaboration. I will discuss awareness tools in more detail in "Workspace Awareness Tools and Group Mirrors" in Section 2.2.2, as well as the role of social signals in awareness tools in Section 2.3.2.

Dourish and Bellotti [1992] have further examined the role of granularity of this awareness information on collaboration processes by classifying this information into high- or low-level awareness. High level awareness of other's actions is identified as a means to structure group activities and assists in avoidance of work duplication. Moreover the low level awareness of the content of other's actions allows for synergistic group behavior. For example, in shared editors an awareness that a collaborator is editing a document refers to a high level awareness, whereas the position of the editor in the text represents the lower level awareness information. In addition, tools and systems providing such a shared feedback to the groups enable groups to easily coordinate their activities [Dourish and Bellotti, 1992].

Dourish [1997] drew an important design guideline for presenting awareness information to the group in a way that the interpretation of awareness is moved from cognitive to perceptual systems of collaborators. In other words, the nature of awareness information should be passive, because an explicit nature of such information might increase collaborator's cognitive load and thus disrupt the actual task.

Tang et al. [2006] studied another aspect of workspace awareness concerning a collaborator's need or desire to work closely or independently of each other, and called it *collaborative coupling*. A tightly coupled group is the one whose actions are well coordinated because the goals and intentions of each member are known to the others, and the group is working as a single entity towards the attainment of task objectives. Conversely, in a loosely coupled group, collaborators focus on the attainment of individual objectives (for example, during completion of a sub-task) while detaching themselves from the group activity. In such a case, groups have to rely on explicit social protocols to coordinate and resolve conflicts. In addition, during the course of the collaboration, the group transitions fluidly from one form of coupling to another; i.e. from tightly coupled to loosely coupled phase and vice versa. The choice of coupling style and the transitions from one form of coupling to another are influenced by various factors such as the task structure, inter-personal relationships, and roles played by collaborators.

Further, Fisher and Dourish [2004] explored the potential of supporting coordination by providing access to collaborators with social and temporal structures of ongoing collaboration. The social structures are the patterns of interaction among the collaborating partners. The temporal structures describe the evolution of these patterns of interaction over time. The authors mention that these structures can be identified and extracted from the electronic traces of group activity while using a collaborative system such as a shared whiteboard or a shared text editor. Later, these extracted structures can be used to build contextualized awareness tools for the group. These contextualized awareness tools can be designed to present to the end-user, via a suitable selection of information, the activity in the form of a summary instead of the structure of the activity information. This implication is relevant for this thesis, as we will see in the later chapters.

Gestures

Tang [1991] conducted an observational study of small sized collocated groups performing a shared drawing task. He observed that a significant proportion of communicative actions (approximately one-third) performed by group members were hand gestures, thus suggesting that gestures are a prominent part of group activity. Likewise, Suthers et al. [2003] observed that expressing complicated ideas is easier when the component ideas are mentioned with the help of simple gestures. According to Dillenbourg and Traum [2006], collaborators utilize multiple modalities while communicating. Where speech can be used as the primary channel for expressing ideas, other more visual channels such as facial expressions, posture and hand gestures act as *backchannel markers* of grounding (for a detailed description refer to "Grounding Theory" in Section 2.4.2), attitudinal reactions and management of turn taking.

Hindmarsh and Heath [2000] and Cherubini et al. [2008] also observed that gestures and gaze (visual conduct) are interrelated processes that are used in tandem with speech while communicating; i.e. people generally look or point at objects they are talking about. In face-to-face interaction, gestures and gaze are available as complementary channels to convey information about the object or context of conversation. However, in distributed collaboration over video conferencing or audio channel, this information is missing, and therefore collaborators frequently formalize the conversation by explicitly describing the attributes of the objects they are talking about [Dillenbourg and Traum, 2006]. Fussell et al. [2000] demonstrated that groups completed the task (manipulation of a 3-dimensional object) much more quickly and with greater efficiency in collocated settings, than when the group members were collaborating remotely. This task required collaborators to effectively communicate ways of manipulating a 3D object, and therefore gestures constituted a significant part of the communication. The findings suggest the need for supporting gestures in distributed collaboration.

What kind of gestures are frequently performed by collaborators? To answer this questions, Bekker et al. [1995] defined a gesture coding scheme for design collaboration and established the dependence of a specific gesture on the task sub-type (designing, meeting management, etc.). Their coding scheme considers *four* types of gestures: *a*) kinetic, *b*) spatial, *c*) pointing, and *d*) other . The *kinetic* gestures enable the simulation of an action involved in a design process, such as movement of a specific machine part. *Spatial* gestures denote the location, size or distance of an object in question. Whereas the *pointing* gestures are used when pointing or indicating any object, person or place. All other gestures, not classified into the above three categories, fall under *other* class of gestures.

The significance of gestures in collaboration and CSCW has led to perennial research in this domain; mostly on the ability to support gestures in distributed collaboration using remote- or tele-pointers. Fussell et al. [2004] and Kirk et al. [2007] have investigated the influence of tele-pointers and remote gestures in collaborative physical tasks. A dyad consisting of a worker (individual responsible for manipulating objects) and a remote helper (individual instructing the worker on how to manipulate the object) were required to accomplish the task, with the support of video conferencing and remote gesture technology. In their study on distributed collaboration, Fussell et al. [2004] demonstrated that representational as well as pointing gestures can be easily implemented using simple tools like a mouse cursor and a pen-based annotation tools over video streams. Further, Kirk et al. [2007] formulated the factors influencing how remote gesture technology can be applicable to varied collaborative scenarios based on the novelty and urgency of the task and the expertise of the collaborators. However, as this thesis is concerned with collocated collaboration, and there is no immediate need for the implementation of remote-gestures in a workspace (such as a whiteboard). Still implementation of these features could prove to be beneficial in scenarios where two collocated and yet geographically distant teams are simultaneously collaborating on a single task.

Territories and Proximity

Collocated collaboration affords an easy access to peripheral actions and activities of the collaborators. Peripheral actions are generally perceived by the peripheral vision of an individual, such as the awareness about who is currently taking notes and who is currently typing on a laptop without gazing in their direction. This capability to perceive other's actions depends on the spatial positioning of collaborators (for example, around a table), as well as orientation of group members with the shared workspace. Furthermore, the ability to fluidly orient workspace items (documents, drawings, plans, etc.) among collaborators functions as a means to mediate the interaction within the group [Tang, 1991]. Kruger et al. [2003] studied the orientation and spatial positioning of objects on a tabletop during collaboration and reported on the partitioning of the space between *personal* and *group* spaces. In the personal space, the items were oriented towards the owner of that space, whereas in the group space, the items were positioned in a way that the group (or a sub-group) had an easy access and visibility to the items. In addition, the authors demonstrated that the orientation of tabletop items serves three important roles during collocated collaboration: *a*) comprehension of information; *b*) coordination of activities; and *c*) communication among group members .

Scott et al. [2004] analyzed the spatial interaction of collaborators around traditional tables while performing various collaborative activities, and further refined the findings of Kruger et al. [2003] by providing a categorization of *territories* used during collaboration. The authors identified three kinds of territories used by collaborators to coordinate group activity: *personal*, *group* and *storage*. Personal territory affords for individual activity and is related to the loosely coupled phase (see [Tang et al., 2006] or refer to "Awareness and Coordination" subsection in Section 2.2.1) of the collaboration. Personal territories correspond to the dedicated workspace of each collaborator on the table, and the items in this territory are oriented towards its owner. On the other hand, the group territory belongs to all the collaborators allowing the group to perform the main task activity, and is oriented in a way to provide easy access and best visibility to the group. Also, the group territory is actively used during the tightly coupled phases of the task. Finally, the storage territory acts as a permanent memory and is used by the group to store relevant information which is not frequently accessed. The location of storage territory depends on the proximity with the individual(s) responsible to store and retrieve items from this territory. The authors further indicate that these territories are transient, and they change as collaborators change their orientation around the table. Also the boundaries between the territories were found to be flexible.

Hall [1969] studied the role of distance and proximity in inter-personal relationships, and coined the term *proxemics* to describe the socio-cognitive use of space in physical world and personal space. The personal space is defined as an area surrounding an individual with invisible boundaries, which serves as a comfort zone in inter-personal communications. Further, Hall et al. [1968] classified the distance between individuals into four categories (with a focus on American population)- *intimate*, *personal*, *social*, and *public*, and also examined the role of these distances in communication, and the kinds of interactions they afford. The intimate and personal distances are the minimum and are generally observed between romantic partners and close friends, and the distance ranges between 0.5 - 1.25 meters. The social distances range between 1.25 - 3.5 meters, and are observed between colleagues in organizational settings. Finally, the social distances are greater than 3.5 meters and afford for public presentations, lectures, and speeches while addressing a group or crowd.

Kiesler and Cummings [2002] analyzed the role of proximity within work groups and its influence on collaboration. They observed that close proximity between collaborators is related to numerous emotional, cognitive, and behavioral changes that positively affect the collaborative processes. In addition, Allen et al. [1977] identified that the probability of two individuals communicating in a work environment is a "decreasing hyperbolic function of the inter-personal distance". Kraut et al. [2002] presented a framework that provides mechanisms to make collaboration easier by addressing the role of close proximity. In their framework, emphasis is placed on three effects of close proximity in collocated collaboration and how the collaborative processes are fostered by *a*) increasing the likelihood of opportunistic encounters that might initiate conversations amongst individuals; *b*) the presence of non-verbal signals and verbal cues, increased coordination in turn-taking, and increased chances of attaining a state of common ground; and *c*) the benefits of awareness about others' activity and actions. Close proximity between individuals was also identified to facilitate the formation of friendships, persuasion, and perception of others' expertise [Latane, 1981]. The effects of close proximity on collaborator's mental coherence was studied by Tang et al. [2006]. They observed that during the tightly coupled phase of collaboration, collaborators preferred to be in close proximity to each other, whereas in loosely coupled phase, collaborators tended to stand further apart.

Another aspect of physical proximity that is related to the subjective sense of "being together" in any environment is the idea of *co-presence*, which was defined by Lombard and Ditton [1997]. Collocated situations were identified to induce the most intense sense of co-presence by Zhao [2003], as compared to distributed collaborative scenarios supported by chat or video conferencing systems. Therefore, virtual multi-user environments aim to achieve a higher sense of co-presence as a primary design objective.

The role of proximity between individuals and artifacts was investigated by Krauss and Weinheimer [1966] and Tversky and Lee [1998]. Referential communication (when a speaker points towards an object as a reference while talking) was observed to have an important relationship with the development of a mutual belief in the identification of artifacts during discussions by Krauss and Weinheimer [1966]. In addition, this establishment of mutual belief was observed to be more efficient in the collocated interaction scenarios where the artifacts are in close proximity to communicators. Further, spatial features of artifacts such as proximity, salience, and permanence were observed to contribute to the easy identification of artifacts during discussions by Tversky and Lee [1998]. Finally, Hawkey et al. [2005] examined the impact of proximity of a group member with a shared workspace (large wall mounted display or a whiteboard), and the nature of interaction with the workspace (direct or indirect). Their findings suggest that a close proximity with another collaborator (standing or sitting adjacent to each other and a distance of arm length) and the shared workspace (standing close enough to be able to write directly on the whiteboard) was perceived to be more effective, more engaging, and less stressful, due to the ease of communication and reduced chances of communication breakdowns.

2.2.2 Meeting Technologies in Collocated Collaboration

Among many existing definitions of the term *groupware*, Ellis et al. [1991] provide a succinct description as "*computer-based systems that support groups of people engaged in a common task (or goal) and that provide an interface to a shared environment.*" The primary role of groupware, as emphasized by the authors, is to assist a group in a "common task", and this is achieved by providing a "shared interface". A shared interface allows group members to gain peripheral awareness of other's actions, which is an important ingredient for better coordination. However, this definition incorporates no distinction between same-time (synchronous) and different-time (asynchronous) collaboration. Therefore, computer-support for collocated same-time collaboration falls under *real-time groupware*.

In this section, I will review some examples of the meeting technologies supporting collocated collaboration (real-time groupware) and which hold relevance to the research work in this thesis.

Meeting Rooms and Roomware

Almost three decades ago at Xerox PARC, Stefik et al. [1987] observed that although computers were widely used to assist individual needs and purposes within organizations, during collaborative scenarios they were replaced with passive media such as flip-charts and chalkboards. These resources often serve as an external memory and communication mediators amongst collaborators. However, manipulating content on chalkboards is cumbersome, and chalk-

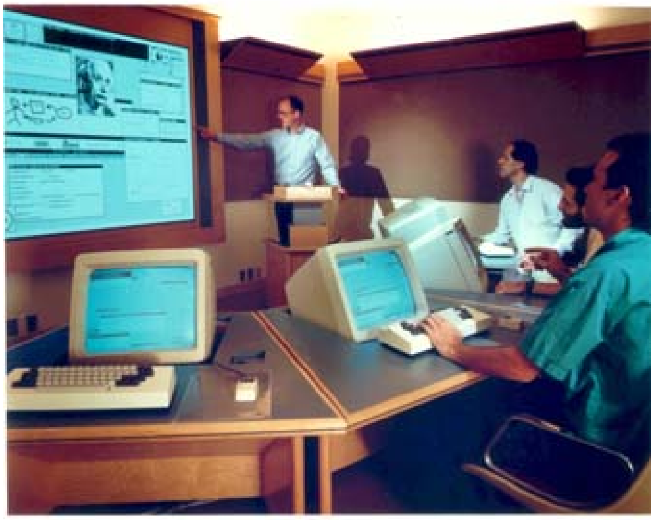

Figure 2.3 – The *Colab* Meeting Room [Stefik et al., 1987]: The meeting room at Xerox PARC designed for small collocated groups of 2 to 6 participants using a networked workstation. Besides individual workstations, the room is equipped with a large wall-mounted touch-sensitive display, and a stand-up keyboard.

boards cannot be used as a permanent storage for information. Therefore, in order to study the role of computers in meetings, and to effectively exploit the potential of computers over the limitations of conventional meeting tools, the authors are credited with the development of the first meeting room called *Colab* as shown in Figure 2.3. Colab was designed to support small group (2 to 6 participants) collocated meetings during problem-solving tasks for computer scientists in Xerox PARC. The meeting room was equipped with workstations for each meeting participant, and a large touch-sensitive wall-mounted display with a stand-up keyboard. The wall-mounted display and other individual workstations were connected to a local area network and a database. These hardware components were equipped with collaborative software meeting tools such as *Boardnoter*, *Cognoter* and *Argnoter*. Boardnoter is a free-hand drawing application with a large WYSIWIS (what you see is what I see) canvas and supporting tools to enable free-hand drawings. Further, Cognoter and Argnoter are respectively applications to generate and organize ideas while brainstorming; and a tool for considering and evaluating alternate proposals.

Providing each meeting participant with a workstation might prove to be disruptive to ongoing communicative and meeting processes in face-to-face interaction, because computer screens can obstruct gaze and gestures. Therefore, many researchers in HCI (Human-Computer Interaction) and CSCW have tried to utilize non-obtrusive and natural interaction modalities such as pen and paper to interact with digital content. One such example is the NiCE Discussion Room by Haller et al. [2010] as shown in Figure 2.4.

The NiCE Discussion Room [Haller et al., 2010] has three major components: *whiteboard*, *paper* and *laptop input*. The whiteboard is a large wall-mounted display on which content can be projected using three projectors as shown in Figure 2.4a. The whiteboard is completely covered with an over-sized Anoto[1] pattern under a transparent layer of acrylic laminate acting as a projection surface and enabling the simultaneous interaction with multiple digital pens. The whiteboard is also supplemented with a sketching application which supports drawing new content, annotating and manipulating content projected from a laptop. Users can change the attributes of their digital pen, such as pen color or switching to a different tool (eraser, highlighter etc.) by using the digital pie menu (displayed over the whiteboard) or using the supplied tangible tool palette. Similar to the whiteboard, the paper interface has Anoto dotted pattern printed on paper and users use the digital pens equipped with ball-point pen and ink. Therefore the paper can be used to create artifacts in a natural way and facilitates the real-time and deferred streaming of the content to the whiteboard. The deferred streaming of content enables the differentiation between the public and the private workspace (also refer to "Public & Private Workspaces" in Section 2.2.2). The real-time streaming from a paper and the whiteboard form the public workspace where shared content is immediately visible to others. On the contrary, as part of the private workspace, users can prepare content on paper and when convinced can share their content over the whiteboard by tapping over the share icon on the paper which automatically creates a paper layer on the whiteboard displaying the paper content (see Figure 2.4b). Furthermore, the laptop input enables the users to connect their laptop and share the screen over the whiteboard by invoking the screen capture layer

[1] Anoto Writing Solutions: http://www.anoto.com/

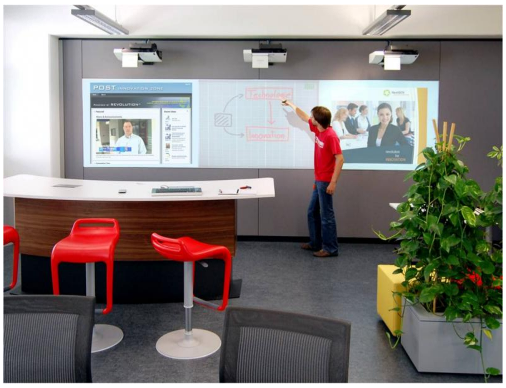

(a) The discussion room is equipped with a large interactive wall-mounted display. Users can interact with the content during meetings by using a digital stylus. In addition, the table enables meeting participants to connect their laptops and share content to the public display.

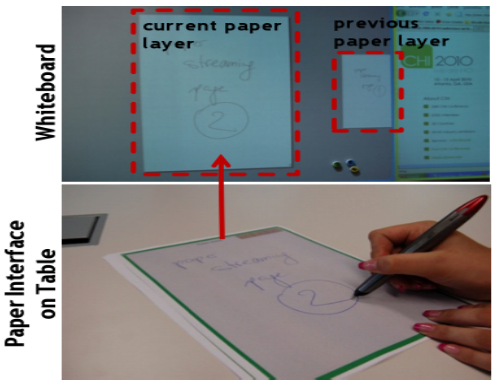

(b) The pen-based interaction allows participants to create content on a paper, and then upload this content onto the public display for others to interact with it. Such a functionality lets users to work on their individual ideas, and when ready share them with the whole group.

Figure 2.4 – NiCE Discussion Room [Haller et al., 2010] empowers meeting participants to use natural and conventional ways to create, manipulate and interact with digital content. The collaborative system uses pen and paper based interaction, similar to conventional collocated meetings.

over the whiteboard. The results of the system evaluation performed by the authors, show that users interacted two-third of the time over the whiteboard and the rest over paper. Also, analyses of system interaction logs revealed the emergence of territories on the whiteboard similar to the one observed by Scott et al. [2004].

The application of ubiquitous computing methodology [Weiser, 1993] to meeting rooms might render an improved collaboration experience while mitigating the disruptive side-effects of technology in meetings. One such implication of the amalgamation of ubiquitous computing with meeting technology design can be *roomware*. Roomware can be defined as the computer-augmented objects resulting from the integration of room elements, such as walls, doors, or furniture with computer based information devices [Streitz et al., 2001]. In the context of supporting collaboration, the authors refer to specifically tailored roomware components capable of facilitating flexible and dynamic **cooperation landscapes**. These cooperation landscapes encompass an ecology of collaborative tools which serve multiple purposes such as team rooms, presentation suites, learning environments, information foyers, and so on. As an example, I will describe the *i-Land* workspace environment [Streitz et al., 1999] along with its various constituents.

The *i-Land* environment as shown in Figure 2.5 is a conceptual framework which was envisioned and designed by Streitz et al. [1999] as their vision of the workspace of the future supporting the changing needs of dynamic collocated teams. Its design is inspired from the need to support newly emerging work practices among creative teams through the integration of architectural spaces and technology. The environment consists of three major components: *a*) an interactive whiteboard (*DynaWall*), *b*) an interactive electronic table (*InteracTable*), and *c*) mobile and networked chairs with integrated interactive devices (*CommChairs*) . Similar to the large whiteboard in the NiCE Discussion Room [Haller et al., 2010], DynaWall is a

Figure 2.5 – The *i-Land* Environment [Streitz et al., 1999] and its four components: DynaWall, InteracTable, CommChair and ConnecTable.

large interactive and touch-sensitive wall (see the wall-mounted display in Figure 2.5), which provides collaborators with a shared perspective to large volumes of information, besides creating and organizing new information. DynaWall supports two content sharing metaphors to transfer artifacts from one position to another on the wall. Using a *pick-and-drop* metaphor, one can pick an artifact from one position and drop it at another position by touching the wall again at that position. Secondly, using the *shuffle* metaphor, the artifacts can be thrown by one collaborator from one side of the wall, and caught by another individual. Next, the InteracTable (see on the right-side of Figure 2.5) is a bottom-up projected, touch-sensitive table, which enables a small group to create, draw, display and annotate information objects on a horizontal surface. Users can use their fingers or a pen to interact with the table. In addition, a wireless keyboard allows for extensive text entry on the table. Finally, the CommChair (see at the center of Figure 2.5) allows users to share artifacts with others sitting on a CommChair, standing in front of the DynaWall, or even with the InteracTable. Users can plug-in their laptops or use an integrated stylus-based tablet to generate and share content. Whereas, the CommChair affords for individual work where users work on ideas and when ready share it with others; DynaWall and InteracTable are designed for group centered activities.

The *ConnecTable* [Tandler et al., 2001] (see on the left-side of Figure 2.5) was added as a later addition to the iLand environment [Streitz et al., 1999], and also consists of a touch-sensitive display on top of it. The height of these tables can be adjusted to suit the individual working on it, and allows for varied working positions such as standing or sitting. In addition, the display can be tilted to suit the appropriate working view of the content displayed on it. Further, the integrated proximity sensors around the display detect the distance between two ConnecTables, and when sufficiently close to one another, the individual displays of the tables merge into a bigger display enabling for fluid collaborative operation around a bigger display.

Shared Workspaces

Independently of the task at hand, group members often have a need to share and store information, during the course of the collaboration. Conventional resources such as whiteboards and flip-charts often serve the purpose of shared workspaces in collocated collaboration, and enable group members to store and share information that is necessary towards the completion of the task. In addition, these resources assist group members to offload their cognitive load onto the shared workspace, and also provide awareness information regarding the information sharing activity in the group. In a user study, Dourish and Bellotti [1992] identified that awareness information provided and exploited passively through the shared workspace, allows groups to smoothly adjust the dynamics of the group by facilitating the transitions between loosely- and tightly-coupled collaboration phases (similar to the collaborative coupling suggested by Tang et al. [2006]). Further, regarding the role of shared workspaces in specific task types, Whittaker et al. [1993] observed that presence of a shared workspace led to increased communication efficiency and influenced the nature of communication in a graphical design and complex text-based tasks. This increased communication efficiency was a result of spatial relationships between the shared artifacts over the shared workspace, which also facilitated gesturing, direct monitoring, and coordination of activity by direct visual inspection.

In distributed collaboration, shared workspaces are designed into the collaborative environments to present the collaborators with aforementioned benefits of shared workspaces. For example, ClearBoard by Ishii et al. [1993] allowed distributed collaborators to draw using digital pens while enabling them to maintain eye-contact and perform natural gestures; as if they were standing face-to-face with a big glass wall in between them. However, it is generally not the case that distributed collaborative environments are equipped with shared workspaces. Colab (as shown in Figure 2.3) by Stefik et al. [1987] was one of the first collocated collaborative environment that supplied groups with a shared workspace. Similar to Colab, DynaWall [Streitz et al., 1999] and ConnecTable [Tandler et al., 2001] in the iLand environment (as shown in Figure 2.5) are also examples of shared workspaces in collocated collaborative scenarios. For more details on these environments, please refer to "Meeting Rooms and Roomware" in Section 2.2.2. Another class of groupware called the *Single Display Groupware* as defined by Stewart et al. [1999], empowers the shared workspaces with the capability of simultaneous interaction ability, making it more probable for group members to share. Single display groupware is the meeting technology central to this thesis work, and therefore I will discuss about them in detail in the next chapter (refer to Section 3.2 in Chapter 3).

Public & Private Workspaces

In line with the idea of territories in collaboration [Scott et al., 2004]; shared workspaces in collaborative environments can be partitioned into shared (or public) and individual workspaces. Tandler et al. [2001] presented a further classification of the individual workspaces into *personal* and *private* in their roomware - ConnecTable. While shared workspaces belong to the group and the group members have access to the content shared over it and rights to manipulate the content; personal and private workspaces are meant for individual interaction only. The only difference in between personal and private workspaces is in the visibility of their content to others, the content of the personal workspace can be visible to other collaborators, whereas an individual has complete ownership over a private workspace and others do not have visibility to the contents of a private workspace. These workspaces have been leveraged in groupware to support varied group dynamics and task demands. Single display groupware [Stewart et al., 1999] (refer to Section 3.2 in Chapter 3) provides a single shared workspace to the group. On the contrary, multi-display groupware incorporate the provision for both private and shared workspaces over different displays. For example, Colab [Stefik et al., 1987] and Dolphin [Streitz et al., 1994] combined these two workspaces, where the wall-mounted display served as a shared workspace, and the individual workstations supplied users with a private workspace. Collaborators can create artifacts while using their individual private workspace and once content they can share the artifact with the rest of the group on the shared workspace.

Several studies have evaluated the effects of different kinds of workspaces on collaborative processes. Haué and Dillenbourg [2009] investigated the role of multiple laptops (displays) in a collocated team by varying the number of displays from 2-4. They observed that laptop owners looked for the most amount of time on their respective laptop screens, rather than looking at their peers. In addition, presence of individual displays inhibited the coordination between team members, reduced the chance of emergence of leadership, and lead to poor quality strategies and performance. However, individual displays favored the parallel task performance of the individuals. Therefore, the authors recommended that groups should be supplied with fewer displays than the number of group members to facilitate better social interaction. Mandryk et al. [2002] examined the role of privacy in collocated collaboration supplied with a shared workspace. They observed that shared workspaces increased the collaborators' awareness of other's activities. However, this increased awareness might distract collaborators from the primary task. Further, Rogers and Rodden [2003] proposed that private workspaces should be integrated in collaborative environments, as they allow individuals to work on their ideas before sharing them, and thus preventing them from any social embarrassment which might happen if raw ideas are immediately shared on a shared workspace. Regarding group participation on a shared workspace, Rogers et al. [2009] examined the influence of input entry points or the number of input channels. Their qualitative findings showed that participation of group members was balanced in terms of verbal utterances in a condition with least restriction on the input entry points.

Looi et al. [2008] demonstrated that groupware should have a provision for both the private and shared workspaces so that their respective advantages are scaffolded for a balanced and effective collaboration, via their collaboration tool - *Group Scribbles*. Further, Streng [2009] investigated the role of public and private workspaces in scripted collaborative learning scenarios. Streng observed that the presence of both private and public workspace resulted in an increase of sequenced arguments produced by the learners, which can be considered as a desired result in the collaborative learning scenario. Further, the presence of a private workspace along with a public workspace reduced the instances of modification of the arguments once they were shared; as the learners worked on their respective arguments and shared them once they were satisfied.

Workspace Awareness Tools and Group Mirrors

Workspace awareness encompasses information regarding the status of the collaborators (such as geographical location in distributed collaboration), what are they doing and what can be the next set of actions they might take [Gutwin and Greenberg, 2002]. This information is often presented to collaborators via awareness tools. Jermann [2004] described the effects of feedback about collaborators' interactions as well as their problem-solving state on the collaborative processes. The interaction and participation data of the collaborators as well as the state of the task they are performing, which is collected in real-time by the collaborative system or groupware, can be fed back to the collaborators (*learners* in case of the author's research) in an aggregated and easy-to-interpret manner. This feedback was demonstrated to be important as it allowed for the planning and structuring of coordination as well as defines the group's well-being function; i.e. how well are the group members faring in the task they are required to accomplish. In other words, awareness tools implement the idea of outsourcing regulation.

Awareness tool is an umbrella term consisting of all kinds of awareness tools (even the one's which are meant for individual awareness). In collaborative scenarios, they are often referred to as *group mirrors*. By definition, a *group mirror* is a family of tools meant to guide collaborators towards a more desirable behavior, such as achieving balanced participation [Jermann et al., 2001]. During collocated collaboration, collaborators have a constant access to various awareness senses such as gaze, gesture, speech and peripheral awareness of others' activity, which groupware designer attempt to provide to an online group [Greenberg et al., 1996]. This does not imply that collocated groups have no need for awareness tools, and their real utility is limited to distributed collaboration. However, there are aspects of collaboration, which are invisible and not directly evident to group members in face-to-face setting (for example, how much is an individual speaking), and visualizing these aspects might influence the collaborative process. Even when the awareness information is available and visible, awareness tools attempt to offload this information off the collaborators, and the collaborators can refer to this information when needed [Bachour, 2010]. In a way, awareness tools in collocated collaboration reduce the cognitive workload of collaborators responsible for monitoring others' activities and actions, and therefore leave more cognitive space for performing the task.

The group mirrors can be subdivided into three types depending upon their degree of active involvement in the regulation process [Jermann et al., 2001]. The mirroring tools simply reflect each group member's actions to the group, and in turn supplementing the group members with a common representation of how the collaboration is unfolding. Next are the metacognitive tools, that provide a summary to the group of the state of ongoing interactions via a set of key indicators, without providing any advice on how to interpret this information and how to act on it. Finally, coaching systems constantly observe and interpret the ongoing collaborative processes and provide advice to the collaborators (learners in case of the author).

The way the awareness information is presented and visualized is also an important factor influencing the efficacy of the group mirror. Streng et al. [2009] observed that showing raw awareness information in the form of graphs and diagrams is less effective as compared to displaying the same information in a metaphorical way. In a user study, the authors observed that metaphorical way of information presentation was perceived to be more popular than charts and graphs. In addition, the group members regulated their behavior significantly faster on receiving a negative feedback from the metaphorical group mirror, than when the same information was presented as charts and graphs. Furthermore, the authors also observed that the presence of metaphorical as well as non-metaphorical group mirror was overall positive as compared to a control condition without any group mirror. The Reflect Table [Bachour et al., 2010] incorporated this finding in the design of their group mirror (refer to Section 2.3.2 where

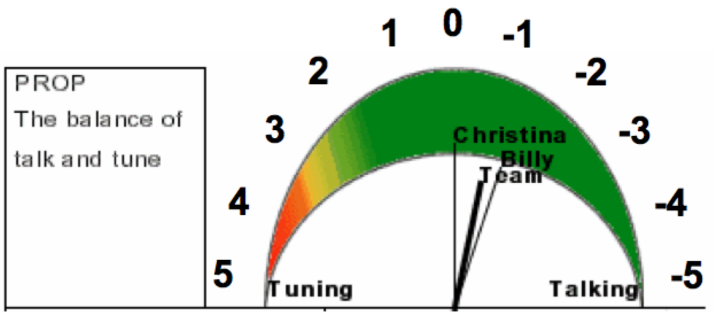

Figure 2.6 – Group Mirror [Jermann, 2004]: The group mirror visualizes the ratio of number of tuning actions performed by collaborators to the number of times the group members speak, in a form of a gauge. The actions of each user (Christina and Billy) are visualized along with the team average.

I describe few examples of group mirrors).

The group mirror designed by Jermann [2004] as shown in Figure 2.6 displays the ratio of actions performed to the quantity of speech of the collaborators. The task required a dyad (Christina and Billy in the figure) to adjust the parameters of a traffic-light system to minimize read congestion. The collaborative system recorded the number of times the collaborators were tuning the parameters and the times when they were talking, and then displays the ratio of these two values in form of a gauge. Besides keeping each group member aware of their respective actions, this group mirror also enables easy repair or avoidance of situations which might adversely affect the task performance.

I will also discuss few more examples of group mirrors and awareness tools which visualize social signals during collaboration in Section 2.3.2.

2.3 Social Signal Processing

Nonverbal communication forms a significant proportion of human communication process through the use of gestures, gaze, facial expressions, posture information and so forth. We as humans are constantly expressing our attitudes, feelings and personality through nonverbal social signals such as turn taking, agreement, politeness, and empathy. Also, we are particularly adept at sensing and interpreting these social signals [Knapp and Hall, 1972]. In spite of the advances in machine learning and computer vision techniques, and even with the immense processing power, modern day computer systems do not have this ability to make sense of the social signals. However, computer systems are likely to be considered as natural, efficacious and trustworthy, if they are capable of sensing social signals such as agreement, inattention, or dispute; and adequately adapt and respond to these signals in a non-intrusive manner [Vinciarelli et al., 2009a]. Therefore, the algorithms and techniques which empower computer systems with such an ability to make sense of social signals can be termed as *social signal processing* (SSP) or *socially aware computing*.

Alex Pentland used the term *social signal processing* for the first time as an umbrella term to describe his group's work at MIT on the extraction of voice based social signals in dyadic interactions (verbal communication); and later used these social signals to accurately predict the outcome of events such as salary negotiations, hiring interviews, and speed-dating conversations [Pentland, 2007]. Therefore by definition, SSP is not restricted to the sensing of only nonverbal social signals. [Vinciarelli et al., 2009a,b] and [Vinciarelli et al., 2012] have summarized most of the relevant research in the emerging domain of social signal processing (SSP) in their excellent research review. While, Vinciarelli et al. [2009b] have specifically focussed on the relevant research in nonverbal behaviors, Vinciarelli et al. [2012] have regarded the research on SSP from a perspective of bridging the social intelligence gap between humans and machines by considering the modeling, analysis and synthesis of human social behavior.

So, what is the difference between a "social signal" and a "behavioral cue"? In their review on the past research on SSP, Vinciarelli et al. [2009a] have answered this question. A social signal is a perceptual expression of an individual's reaction or attitude towards a social interactive situation (the context). These social signals are manifested through a multiplicity of non-verbal behavioral cues. More precisely, a behavioral cue "*encapsulates a set of temporal changes in an individual's neuromuscular and physiological activity that lasts for short intervals of time, as compared to behaviors*" [Vinciarelli et al., 2009a]. Examples of behavioral cues can be: gaze exchanges, blinks, smiles, head nods, crossed arms and so on. On the other hand, a careful interpretation of these behavioral cues within the appropriate context will give forth the social signal such as attention, empathy, politeness, flirting, disagreement, and so forth. Furthermore, a social signal and a social behavior are almost similar as both are demonstrated as recurring temporal patterns of nonverbal behavioral cues, except that social signals typically last for a small duration as compared to social behavior that lasts longer. Mostly, behavioral cues accompany verbal communications, and despite being invisible, they are sensed and interpreted outside of the conscious awareness of an individual; thus having a major influence on the perception of the verbal message and social situation [Knapp and Hall, 1972] (as cited in Vinciarelli et al. [2009a]).

In this section, I will review the relevant past works in SSP concerned with the automatic detection of some of the social signals during collaboration. Later, I will consider a few examples of collaborative and awareness tools applying SSP to support collaborative behavior.

2.3.1 Social Signals in Collaboration

In order to render computer systems more socially aware, researchers from varied domains and disciplines have converged their efforts. This requires the thorough understanding of human behavioral cues within a given context, in addition to the successful application of adequate computational techniques such as computer vision, speech processing and machine learning. Vinciarelli et al. [2009a,b, 2012] have conducted a comprehensive review on the various social signals (and behaviors) such as dominance, roles etc., which can be accurately identified and interpreted by using multiple behavioral cues and different computational techniques.

Next, I will discuss the social signals (and social behavior) in collaborative scenarios, and the behavioral cues which play a crucial role in the identification of these social signals. However, I will not provide a detailed account of all the computational techniques and the inferential statistics required in the processing of behavioral cues because it is not related to the methodology used by us in this research work.

Detection of Roles

The course of human interaction is shaped by the frequent involvement of behaviors with defined statuses and roles, which in turn also provide predictability to the interaction [Tischler, 1990]. Researchers in SSP have applied varied approaches to recognize roles within collaboration. Vinciarelli [2007] analyzed the turn-taking during conversations for the identification of roles. More precisely, the temporal proximity of the speakers was used to construct a social network of the group and to extract features which were fed to a Bayesian Classifier for recognition. Garg et al. [2008] and Salamin et al. [2009] combined the temporal proximity with the duration of intervention and the distribution of words in speech transcriptions for better role recognition accuracy. Another similar approach takes into account the probability that a group member starts speaking when everyone is silent or when someone is speaking, as the main feature for predicting the role of the individual [Laskowski et al., 2008]

It is evident that most research work focused on the identification of roles in collaboration, considers turn-taking (speakers' temporal proximity) as a crucial social signal. This social signal in combination with other vocal and movement-based features is mainly used to identify the role of a collaborator.

Dominance

A subset of researchers analyzing small group interactions have regarded the presence of a dominant person within the group. A dominant person is considered to have a higher influence on the development and outcome of an interaction [Levine and Moreland, 1998].

One of the successful methodologyies employed in the automatic detection of dominant person within the group regards the speaking activity of the speaker, which serves as the input feature for Support Vector Machines to classify dominance into three classes based on the degree of dominance: low, normal or high [Rienks et al., 2006; Rienks and Heylen, 2006]. The social signals influencing dominance are speaking times, number of turns, number of successful interruptions, and prosodic features such as pitch of the speaker. Speaking related activity has proved to be an effective predictor of dominance. However, when in combination with motion-based features such as the time during which a person moves and the frequency of the time interval when a person is moving, has also demonstrated to have increased the efficiency of prediction of the most dominant person [Jayagopi et al., 2009].

2.3.2 Visualizing Social Signals: Awareness Tools

In this section, I will review few examples of awareness tools that consider social signal indicators as a means to visualize the state of collaborative process.

Meeting Mediator

The *Meeting Mediator* (MM) by Kim et al. [2008] incorporates a sensor and a visualizer to mirror the ongoing group activity to the group (see Figure 2.7). The MM was designed to record multi-modal interaction data in real-time, followed by the assimilation of these multiple streams of data to produce a visualization of social signals such as participation balance and dominance of a participant (also refer to "Dominance" in Section 2.3.1 for more research in this domain).

Meeting Mediator uses a sociometric badge which is worn around the neck of a collaborator as shown in Figure 2.7a. The sociometric badge is an electronic sensing device capable of recording non-linguistic vocal features (speech energy, turn taking, etc.), body movement, distance to other badges and facial-orientation information to detect conversations among

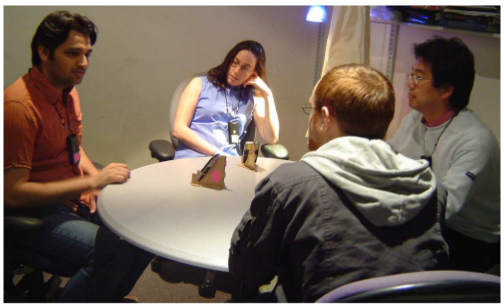

(a) The Meeting Mediator captures interactions among collaborators through the Sociometric Badges worn around each participant's neck, and a mobile phone to display this information in real time.

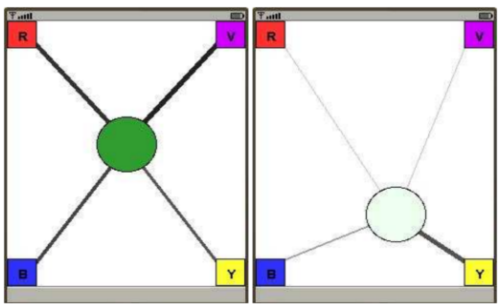

(b) Two contrasting cases of balanced- and highly interactive group (left) and unbalanced- and less interactive group (right), as shown in Meeting Mediator visualization. The position of circle denotes the balance in participation, and the color denotes the interactivity of the group. The thickness of the line represents speaking times of each participant (the four colored rectangles at the corners).

Figure 2.7 – Meeting Mediator [Kim et al., 2008] is a group mirror capable of recording and interpreting multi-modal and real-time interaction data, and visualizing this information for the group to regulate their behavior.

two group members. This recorded information is interpreted and analyzed in real time by a remote server and the resulting visualization is generated and displayed on hand-held displays like mobile phones placed in front of each participant in a persuasive and unintrusive manner. The visualization is shared for all the group members, but available on different displays. This facilitates the use of MM in collocated as well as distributed collaborative scenarios, as the recording of social signal and their visualization is done through individual devices, and the information between sensor and visualizer is communicated via a remote sensor.

The aspects of the group interaction that are emphasized by the Meeting Mediator visualization constitute dominance of a meeting participant, participation balance as well as the interactivity of the group as shown in Figure 2.7b. The four collaborators are shown as the four colored squared at the corners, and the thickness of the lines connecting participants to the circle in the center denotes the speaking time of each participant. The spatial position of the circle represents the participation balance of the group and the dominance of collaborator(s), whereas the color of the circle indicates the interactivity of the group measured as the speed of turn taking (green for a very interactive group, and white for a less interactive group). The visualization on the left-side in Figure 2.7b shows a balanced group (the circle is situated in the center) where the collaborators are speaking for almost equal amount of time (thickness of the lines), and are very interactive due to the frequent change of the speaker. On the other hand, the group in the visualization on the right-side visualization in Figure 2.7b represents a less-interactive group, where the participant on the bottom-right corner is dominating the discussion.

The authors conducted a study with the Meeting Mediator, and observed that it significantly reduced the behavioral difference between dominant and non-dominant group members, and enabled groups to regulate its behavior while promoting equality of participation.

Reflect Table

Reflect is a table and a roomware element which acts as a group mirror in collocated collaboration scenarios for small groups (4 - 6 participants), and was designed by Bachour et al. [2010] as shown in Figure 2.8. It was designed to analyze and visualize group members' spoken contribution in an unintrusive and semi-ambient way while performing a task in real-time. The table can effectively identify the direction of speech and the speaker (sitting around the table), with the help of a three microphone array situated at the center of the tabletop and a selective filtering technique called *beamforming*. Next, the table extracts the prosodic speech features (rhythm, pitch and relative emphasis) for each speaker separately, and then analyzes this integrated information to compute each individual's participation in the discussion. This participation information is visualized by the means of LEDs (light emitting diodes) under a frosted glass surface (see Figure 2.8) with a territorial metaphor; meaning that the territory grows for group members who participate more, and over-participating member's territory invades other's territories. In addition, the table surface can be normally used to place objects and documents during meeting. Further, Reflect Table does not require group members to

Figure 2.8 – Reflect Table [Bachour et al., 2010] is a group mirror, that visualizes the speech based participation information as territories of different color in front of the meeting participants in a semi-ambient way. The three microphone array situated at the center of the table in tandem with a selective filtering technique can accurately identify the direction of speech and the current speaker.

wear sensors like in case of Meeting Mediator [Kim et al., 2008].

The authors also conducted a formal study to investigate the influence of participation feedback on regulation of group behavior. They observed that the group members perceived an increased awareness in terms of their participation in the group. In addition, group members who believed the importance of balanced participation while collaborating were found to be significantly influenced by the awareness, and thus regulated their behavior towards a balanced group. Further, over-participators were observed to give away the floor to others and allowing them to speak after looking at their territory, and the under-participators started speaking more.

Conversation Clock and Conversation Votes

The *Conversation Clock* [Bergstrom and Karahalios, 2007a,c] is another group mirror that uses group members' speech as an input to visualize interaction patterns during conversation. In Conversation Clock, the authors aimed to display the persistent history of the conversation, while providing visual cues of various social signals such as dominance, interruption, turn taking and mimicry. Each group member is equipped with a microphone (as shown in Figure 2.9a), and the system visualizes speech as rectangular bars around a circular timeline following a clock metaphor (as shown in Figure 2.9b) in a shared physical space between the speakers through a projector. The color of the visualized speech is different for each speaker and matches the color code on the microphone. Furthermore, simultaneous speech is displayed as overlaid rectangular bars. The size of the rectangular bar is proportional to the speaker's volume or the degree of emphasis a speaker puts in her voice. A louder volume, and hence a larger visualized speech bar might indicate towards a desire to lead the conversation and consequently suggest dominance. Each concentric circle in the visualization (as shown in Figure 2.9b) is a representation of a minute of conversation; with the outermost circle denoting the most recent conversation. Silence during conversation is visualized as small dots.

In a user study conducted with the Conversation Clock, the authors observed that the visualization increased the awareness of the speakers' about their own conversations. In addition, a difference in perception and interpretation of the visualization was observed between the passive and active participants. On one hand, passive participants interpreted the visualization to see the overall dominance of others by commenting on the presence of one or two major colors in the visualization. On the other hand the active speakers focussed on the outer circle of the visualization and regulated the duration of their turns.

Further, *Conversation Votes* [Bergstrom and Karahalios, 2007b] extends Conversation Clock by incorporating user feedback into the group mirror. Each group member is equipped with an up-vote and a down-vote button, using which listeners can provide feedback to the current speaker, and this feedback is displayed anonymously as a positive or negative feedback over the rectangular bar as shown in Figure 2.10. In addition, an up-vote increases the length of the rectangular bar and the brightness of the color is increased as well. On the other hand, a down-vote does the opposite. The positive feedback could encourage current speaker to

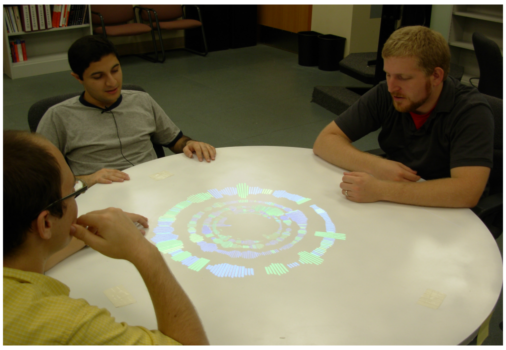

(a) Conversation Clock Setup: Each group member is equipped with a wearable microphone used to capture individual speech. The visualization is displayed in real-time in a shared physical space through a projector on the top.

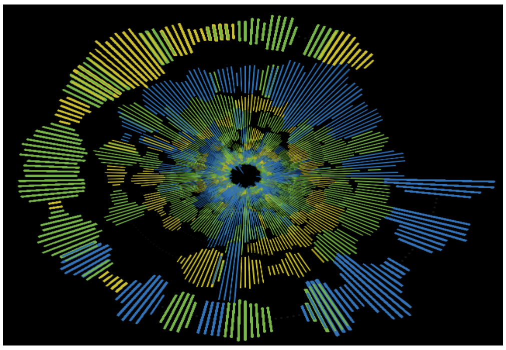

(b) Conversation Clock Visualization: Each rectangular bar around a circular timeline denotes speech by a group member at that moment. The color of the bar denotes the speaker, and the length of the bar indicates the loudness of the speech. Silence is represented as small dots along the time line. The outer circle represents the recent conversation, whereas the inner concentric circles denote history of conversation.

Figure 2.9 – The Conversation Clock [Bergstrom and Karahalios, 2007a,c] is a group mirror designed to display group members' conversational pattern in real-time along a circular timeline.

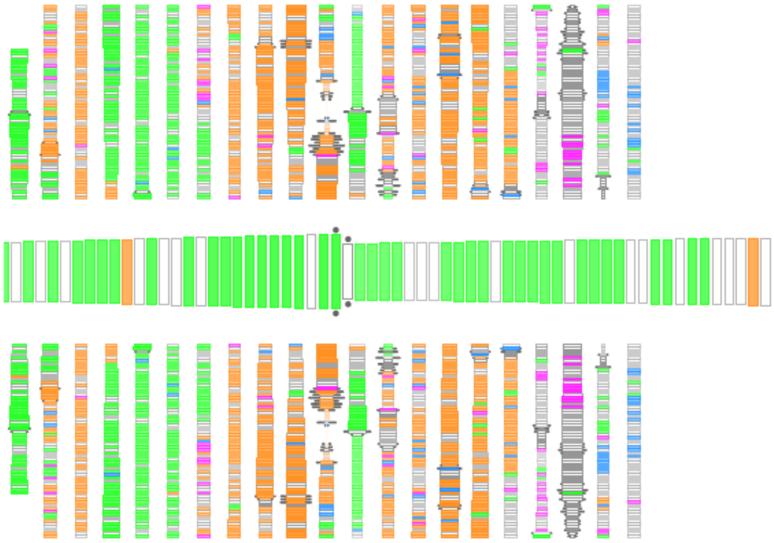

Figure 2.10 – Conversation Votes [Bergstrom and Karahalios, 2007b]: Listeners can provide feedback anonymously to the speaker by voting (positively or negatively), which are then displayed over the rectangular speech bars. The horizontal line represents the current ongoing conversation, whereas the vertical lines represent the history of past conversations.

continue with the discourse. However, the negative feedback could mean that listeners intend the speaker to yield the floor. Further, the capability to interact with group mirror in this way could lead to an efficient regulation of the conversation, in addition to the creation of flags of salient moments during conversation. However, voting in small group collocated scenarios can have some downsides due to the loss of anonymity.

2.4 Theoretical Framework: Social Cognition

The confluence of Computer Science and Social Cognition in CSCW empowers the researchers to microscopically study group processes from varied perspectives which are offered via numerous viewpoints (theories and frameworks) in Social Cognition, and thereupon supporting these processes with effective technologies. In the previous sections I reviewed pertinent researches from the Computer Science perspective, including works from CSCW, Meeting Technologies and Social Signal Processing. Next, I will review some of the theories and frameworks from Social Cognition, which are of direct relevance to the research work presented in this thesis, and provide us with a foundation to base our research methodology and results.

Philosophers and researchers studying human cognition have followed two distinct approaches namely *symbolic processing* and *situated cognition*. In the symbolic processing approach, the focus is entirely in the information processing structures of the brain and the symbolic representations of mind (also known as *knowledge representation problem*). On the other hand, in the situated cognition approach, emphasis is placed on the importance of historical influences, social interaction, culture and the environment [Thompson and Fine, 1999]. In their review article on Socially Shared Cognition, Thompson and Fine describe the motivations behind situated cognition as *the assumption that human knowledge and interaction cannot be divorced from the world; to do so is to study a disembodied intelligence, one that is artificial, unreal, and uncharacteristic of actual behavior*. Further, Levine et al. [1993] noted that more emphasis has been given to the symbolic processing approach (referred to as excessive mentalizing) to study human cognition, which diverts the attention from the understanding of behavior and interaction; which in turn requires the study of intentions, motivations, social interpretations, and cognitive functioning in interactions with others.

Thompson and Fine [1999] referred to the idea of *meaning*, and it's extension beyond an individual's cognition. They used the term *socially shared meaning* to denote the creation and utilization of interpersonal understanding within groups and other larger collectives. Further, the term *shared* in socially shared meaning can be interpreted in three different ways. Firstly, the term *shared* could denote "dividing something into portions", such as distributed knowledge, division of roles and responsibilities (in the case of transactive memory systems [Wegner, 1987]), and division of labor for knowledge (for example, distributed cognition [Hutchins, 1995]). The second possible meaning of *shared* could refer to the notion of common knowledge and maintaining overlapped cognitive representations of task requirements, procedures and role responsibilities (for example shared mental models [Converse, 1993]). Finally, the third interpretation of *shared* could indicate towards the notion of consensus and acceptance in conversations. It considers the idea of taking others' perspectives (for example intersubjectivity [Rommetveit, 1974], role taking, and tuning in conversations [Higgins and Rholes, 1978]), establishment of agreement about what is being said or understood (common ground [Clark, 1985]), shared recognition of shared meaning and so on. Thompson and Fine's idea of socially shared meaning is also closely related to the concept of *collective* or *interpersonal meaning* by Zajonc and Adelmann [1987]. This viewpoint considers the assumption that social cognition cannot be restricted to individual thought about social objects, but emerges through the social interchange constructed, maintained, shared and distributed among group of individuals. Consequently, it cannot be wise to assume the individuality of cognition detached from the external social process that significantly influence it [Levine et al., 1993] (as cited in Thompson and Fine [1999]).

In their review, Thompson and Fine [1999] outlined four themes that can be used to distinguish individual cognition from socially shared cognition: group behavior as the main unit of study, more emphasis on social activity rather than cognitive activity, coordination and synchrony between interacting individuals and the evolution of interaction with time. Besides these four themes, the authors also describe the four conceptual-empirical perspectives within the domain of social cognition: *a*) the supraindividual model, *b*) the information processing perspective, *c*) the communication perspective, and *d*) the social interaction perspective . The supraindividual approach suggests that individuals' behavior cannot be completely explained by their internal motivations and cognition, rather the force driving their social behavior is external. The information processing approach treats group interactions from the perspective of information processing units incorporating processes such as encoding, storage, processing, retrieval, response and feedback. Communication approach takes into account the creation of social reality among interacting individuals through communication. Finally, the social interaction perspective regards interaction as the core unit of study.

The frameworks and models from the information-processing and communication approaches that offer insights into the understanding of social cognition are highly relevant for the research work presented in this thesis. Therefore, in the next sections, I will review the relevant works from within these two approaches.

2.4.1 Information Processing Approaches

This viewpoint regards group as an information processor, where group members possess separate and independent memory structures situated within their own cognitive architecture, and others' memory structures can be accessed and utilized through communication. The access to others' memory expands the storage and retrieval channels available to any individual within the group, thus effectively extending the overall group memory structure [Thompson and Fine, 1999]. Next, I will provide a detailed account of three different models/theories aligned with the information processing view of socially shared cognition.

Transactive Memory

The theory of transactive memory is particularly inspired from the principles of cognitive information processing and symbolic representation of knowledge. Wegner [1987] coined the term *transactive memory* and defined it as a shared system for encoding, storing and retrieving information within groups. It comprises of the individual memories corresponding to each member of the group as well as the communication that takes place between the individuals [Wegner et al., 1985]. As the idea of transactive memory goes beyond the individual's memory, it expands the range of information that can be encoded, stored and retrieved from the transactive memory. In addition, each individual develops a mental model of others in the group and the group as a whole over time, but is not aware of how others perceive the system [Wegner, 1987].

Regarding at a finer level of granularity, there are four key components of a transactive memory system in addition to the communication between these components. These components are the individual memory, external memory, transactive memory and the meta-memory [Thompson and Fine, 1999]. As the name suggests, the individual memory is internal to a group member and represents the basic information processing model including the short- and long-term memory structures, and the external memory refers to the external storage of information. Meta-memory refers to the beliefs and knowledge that group members have regarding their own memory. Finally, the transactive memory refers to the knowledge that an individual has about the knowledge and expertise of other group members. Further, a successful retrieval of a memory item or information requires a prior encoding of two additional pieces of information in addition to the target information meant to be retrieved; i.e. a label or retrieval cue for the information to be retrieved and the knowledge of the location of the information. Other members in a group or community can act as the external storage for the information to be retrieved. However, a successful retrieval of this memory item requires that one person within the group has access to this information stored in other's memory structure due to the knowledge that others have a memory of this information with the correct label.

The development of transactive memory requires personal expertise and circumstantial knowledge in order to decide who stores a specific information item (i.e. the storage location of this item). In organizational groups, recognized experts in certain domains are responsible for encoding, storage and retrieval of information belonging to their domain [Thompson and Fine, 1999]. However, in intimate couples, one of the partner is responsible to hold certain type of information, so that the information required by the couple is available with one of the partners [Wegner, 1987; Wegner et al., 1991].

Various studies investigating the effects of transactive memory have identified a few advantages. In the context of work groups, transactive memory was observed to result into superior task performance [Moreland, 1999] and reduced worker fatality [Goodman and Shah, 1992]. In addition, transactive memory was observed to foster creativity among collaborating team members in the design and development of novel products, particularly in situations where the knowledge is differentiated, in a sense that group members do not have access to common shared information [Wegner, 1987]. However, Hill [1982] observed that transactive memory in groups resulted in below average performance as compared to the level of the best group member in tasks requiring the employment of logic, judgement, or problem solving skills. Also, Wegner et al. [1991] observed that tasks that prove disruptive to the transactive memory by imposing a different memory structure, put the groups and dyads (or couples) at risk of lower recall.

Shared Mental Models

Mental models are the "mechanisms by means of which humans create understanding and descriptions of system's purpose and form, its functioning and observed states, and likely predictions of its future states" [Rouse and Morris, 1986] (as cited in Thompson and Fine [1999]). Similar to individual's mental models, shared (or team) mental models refer to the development of common or overlapping cognitive representations of task requirements, procedures, and role responsibilities [Converse, 1993]. The concept of shared mental models was developed mainly as a descriptive and prescriptive tool for team performance suggesting that successful teams maintain a high degree of overlapped representations about the task requirements.

Converse [1993] described *shared mental models* as knowledge structures maintained by the group members that allow them to form accurate explanations and expectations about the task, and enable efficient coordination of their actions and adapting their behavior depending on the changes in task and team demands. The author draws out two behaviors: *taskwork* and *teamwork*, that are necessary for groups to coordinate their actions and adapt to changing task requirements. Taskwork refers to the skills required for an efficient execution of the task; whereas teamwork refers to the functions that allow the group members to coordinate and interact successfully.

Tindale et al. [1996] argued that social influence processes within group interactions are influenced by the degree of shared representations of the task. Therefore, within decision-making groups, minorities that favor normative incorrect positions are highly likely to win if their position is aligned with the shared cognitive models of the group. In addition, consensus within groups can be assumed to be a result of shared mental models, and groups reach consensus either due to shared meanings prior to the discussion or the development of this shared meaning during the discussion.

Distributed Cognition

Distributed Cognition is a theoretical framework by Hutchins [1995] that regards a distributed, socio-technical system as a primary unit of analysis instead of an individual mind. This framework is aligned to the information processing perspective of socially shared cognition, as it is concerned with the representation of information, and the ways in which information is transformed and propagated during task performance. In other words, distributed cognition regards cognitive processes, regardless of where they may be situated, based on the functional relationships of the elements that participate together in a process [Hollan, Hutchins, and Kirsh, 2000]. The authors argued that "*a cognitive process is not cognitive simply because it happens in a brain, nor is process non-cognitive simply because it happens in the interactions among many brains*". Further, in distributed cognition a system is expected to dynamically adapt based on the interaction and coordination among different subsystems, in order to accomplish the task. In addition, a cognitive process is delimited by the functional relationships among the participating elements, rather than by the spatial colocation of the elements.

Hutchins [1995] used the cockpit of a commercial airliner as an example of a socio-technical system, and the unit of analysis to elucidate the concept of distributed cognition. He primarily observed the interactions between the two pilots, as well as the cognitive task to compute and remember a set of correspondences between airspeed and wing configuration during the course of a flight, especially while landing. The focus of his study was the many representations which are inside the cockpit system (for example the speed card booklet, airspeed indicators with internal and external bugs, etc.), and yet outside the heads of the pilots. Besides these observed media elements (various representations of airspeed and wing configurations), the memories of the pilots form a memory system. Further, the process of continuous interaction among these elements can be called as the cockpit's memory; because the whole process involves the creation of a representational state that is saved and used to organize subsequent activities within the system.

2.4.2 Communication Approaches

The communication perspective to the understanding of social cognition focuses on how individuals create social reality through communication, and this perspective embodies three central ideas or themes of communication [Thompson and Fine, 1999]. The first idea concerns with how people define the situation; also termed as *intersubjectivity* by Rommetveit [1974]. In order for a communication to succeed, the communicating individuals should respond to the same stimulus; i.e. they establish a common definition of "here and now" of the situation. The second idea looks into the interpersonal relationships and the contextual usage of appropriate language; which are the assumptions, norms and rules that govern and shape the process of communication. Finally, the third idea regards tuning or perspective taking where communicators take the perspective of others.

According to Rommetveit [1974], "Intersubjectivity regards language as a purely social phenomenon and communication aims at transcendence of the private worlds of individuals". Intersubjectivity regards communication with an individualistic approach aiming at the establishment of a common social reality. However, another conceptually similar idea to intersubjectivity is the establishment of a common ground which is an interactionist view of looking at communication [Clark, 1985]. In the next section, I will discuss more about the research on common ground and grounding theory.

Grounding Theory

During a conversation, in order to facilitate coordination, the participating individuals often shape and guide their communication by taking into account the recipient and her perspective. This process of taking other's perspective while communicating, and taking appropriate reparatory action in case of misunderstanding is referred to as the establishment of a *common ground*, and its aim is to reduce the thematic distance between what the speaker means and what a listener understands [Clark, 1985]. The establishment of a common ground is facilitated in situations where individuals share mutual knowledge and suppositions. Further, the development, maintenance and usage of common ground is achieved through a process of *grounding*, where participants make sure that what has been said is understood by others [Clark and Brennan, 1991]. According to Clark and Brennan [1991], two main factors govern the grounding process: *purpose* and *medium*. The purpose defines what the participants are trying to achieve through their conversation; and medium refers to the techniques available to accomplish the purpose along with the cost of using these techniques. Further, speakers often exploit the common ground they share with the listeners in creating referring expressions (such as "the red chair on your right-hand side"), which in turn improves listener's comprehension and requires a degree of perspective taking on the speaker's part.

Fussell and Krauss [1991] suggested that conversations among individuals are understandable because communicators formulate a shared communicative environment. However, this formulation could be flawed because speaker's estimates about listener's knowledge could be biased in the direction of speaker's own knowledge. Furthermore, the degree of difficulty or easiness experienced by communicators in establishing a common ground is influenced by two kinds of social knowledge, which helps communicators to devise a tentative hypothesis about the shared communicative environment. These (two kinds of) social knowledge are: theories and intuitions about the listener's beliefs and background knowledge, and the knowledge of interaction rules and conversational resources such as verbal and non-verbal feedback.

Cooperated social activity such as consensually validated roles and relationships are required for the mutual creation, monitoring and sustaining of the social reality or common ground; and this serves as the primary norm in efficient communication [Thompson and Fine, 1999]. In addition, Higgins [1981] developed the "communication game" as an alternative to the individualistic information transmission approach, where he determined the extensive list of general rules governing communication for both the speaker and the listener. The central idea employed by Higgins [1981] to develop the communication game was to study how people use communication rules in social interactions to optimize the process of establishment of common ground.

Finally, the idea of perspective taking or *cognitive tuning* in communication, requires that communicators should also take into account the attitudes in addition to the background knowledge about their audience [Higgins and Rholes, 1978] (as cited in Thompson and Fine [1999]). Further, Wilke and Meertens [1994] studied the task groups and extended the idea of tuning to develop a theory of group performance based on four kinds of tuning: cognitive, reflective, communicative and structural tuning (as cited in Thompson and Fine [1999]). The authors suggest that in order to succeed in a task, the group members should build a collective representation of the task, as well as the construction of consensual notion that they agree to engage in the task and the organization of the group during the task. Amongst the four kinds of tuning identified by the authors; *cognitive tuning* refers to the collective representation of the task to be performed and the consequent organization of the group to perform the task. The *reflective tuning* regards the evaluation of various task processes such as task engagement, task continuation, and task abandonment; and their consequences. The *communicative tuning* concerns how group members' individual cognition about themselves, other group members and the task environment; are introduced within the group and are aligned with those of other group members. Finally, the *structural tuning* refers to the cognitive and normative social representation of how group members and tasks are interrelated.

2.5 Summary and Our Approach

The context of the research work presented in this thesis regards the analyses of collaborative processes within small-group collocated meetings. Collaboration analysis is considered to be a difficult and a time-consuming task, partly because of the complex interaction and influence of multiple and diverse processes [Tang, 1991; Neale et al., 2004]. However, a better understanding of these collaborative processes is crucial because of their role in the design of technology to support collaborative work. Further, collaboration analysis methods such as the one suggested by Meier et al. [2007], is highly time-consuming as it requires extensive coding and rating of multi-modal collaboration data such as speech, language, gestures, gaze, and so on. In addition to these issues related to collaboration analysis, most contemporary group research until now has not paid much attention to the analyses of dynamic group processes, but rather the focus is more on the task and performance outcomes [Worchel, 1994]. On a similar note, Dillenbourg et al. [1995] defined the "interaction paradigm" of studying learning in the context of computer supported collaborative learning(CSCL). Unlike judging collaborative learning based on the learning outcome and performance of groups, the interaction paradigm seeks to identify learners' interactions which can lead to better learning outcomes, and try to elicit these interactions by means of technology.

In our research, we focus on the analysis of dynamic group processes through the study of group's interaction with the artifacts created and shared among the group members. Further, as creation of a shared reality during collaboration involves communication and creation of mutual knowledge through sharing of artifacts and content. Our approach regards the study of group's interaction with a shared workspace during the course of performing a task collaboratively. We believe that many ongoing visible collaborative processes such as conversation, turn-taking, etc., may be tightly coupled with how group members interact with shared artifacts. In other words, group's interaction with a shared workspace can act as a proxy for the actual collaborative processes, and their analysis can have implications towards automating the tedious process of analyzing collaboration and making it real-time, and towards the design of informed awareness tools and group mirrors.

In order to record and analyze group member's interaction with the shared workspace, we employed an iterative design methodology to design and develop a shared workspace with simultaneous interaction ability for the group members. We designed a single display groupware (SDG) [Stewart et al., 1999] as a meeting technology to allow us to analyze group member's unconstrained interactions at a finer level of granularity. Therefore, in the next chapter, I will review the research work on the role and influence of SDGs in collaboration, and analyze the link between relevant theories from social cognition and the design of groupware capable of answering the research questions that we wish to ask. Finally, the research work presented in this thesis is concerned with collaboration work, and does not include collaborative learning as the primary focus. However, implications of this research are not restricted to collaborative work only, and might have implications towards collaborative learning scenarios as well.

3 Single Display Groupware

Designing a meeting technology that facilitates the creation, manipulation and sharing of artifacts in collocated meetings, is a primary requirement for this research work that aims to study group's interaction with a shared workspace. We aim for a design that satisfies both the functional requirements as well as the conceptual (or theoretical) requirements. By functional requirements, we refer to the features that directly address the way the system will be used. The conceptual design requirements inform the design of the system from the perspective of the cognitive processes that the users engage in and the resources that they use, which in turn defines the meaningfulness of their actions and experiences. We believe that the design and appropriate choice of the meeting technology is not just driven by the desired functionalities of the system. It is equally important to have a fine-grained understanding of how individuals create meaning through interactions with others as well as with objects situated within the environment, including technologies and representations. In collocated meetings, interactions don't just happen among participants; individuals frequently interact with artifacts and various kinds of external representations that are created, modified, maintained, and shared during the meeting. In addition, the interactions with others as well as different media and technology defines the properties of the whole collaboration scenario. Therefore, we consider the collocated meeting ecosystem as the socio-technical system and a primary unit of our analysis, with an emphasis on the human interactions with artifacts and external representations. In that regard, we employ distributed cognition [Hutchins, 1995] as a theoretical framework to comprehend, examine and augment our research context of collocated small-group meetings.

Next in this chapter, we will discuss how distributed cognition as a theoretical framework is used to inform our conceptual design requirements. Later, we will describe and review *Single Display Groupware (SDG)* as the meeting technology that satisfies both our functional and conceptual design requirements. Further, we will discuss a few examples of SDG used in educational and organizational collaboration scenarios, and their influence on collaboration. Finally, we will conclude the chapter with the design implications for the meeting technology that we developed and used in our studies.

3.1 Distributed Cognition as a Theoretical Framework

Distributed cognition analyzes the whole environment along with its constituent actors and their interactions, and how these actors go about to coordinate their actions [Hutchins, 1995]. In addition, Hollan, Hutchins, and Kirsh [2000] argued that "*distributed cognition provides a radical reorientation of how to think about designing and supporting human-computer interaction (HCI)*". It differs from traditional views of studying cognition not only by expanding the boundaries of the unit of analysis by including environmental artifacts, and accordingly extending the range of mechanisms that are assumed to contribute to the cognitive processes. In distributed cognition, a system is considered adaptive if it can dynamically reconfigure itself so as to bring the constituent elements into coordination while accomplishing different functions. Also, the cognitive processes are marked by the functional relationships among the participating constituent elements, rather than by the spatial collocation of the elements [Hollan et al., 2000].

Hutchins [1995] delineate the three apparent ways in which cognitive process can be distributed:

1. Distribution across the members of the group.
2. Coordination between internal and external representations and structures.
3. Distribution over time, in a way that the outcome of an earlier event consequently influences the nature of later events.

3.1.1 Three Tenets of Distributed Cognition

To establish a well grounded understanding of the various interactions undergoing in collocated collaboration, in this section, we structure a discussion that employs the three tenets of distributed cognition as defined by Hollan et al. [2000]:

- Social organization
- Embodied cognition
- The relationship between culture and cognition

The first tenet is the idea that "*social organization in itself is a form of cognitive architecture*" [Hollan et al., 2000]. The context of an activity reinforces this architecture by defining the trajectories through which information is transmitted and transformed within a group. The distributed cognition framework includes this idea through the study of interactions between cognitive processes at three levels, described via the following three questions:

1. How are the cognitive processes that are generally associated to an individual's mind, implemented within a group of individuals?
2. What are the differences in the cognitive properties of groups and the cognitive properties of the participating group members?

3. How does participation in group activities influence the cognitive properties of the participating group members?

The second tenet of distributed cognition is related to the idea of embodied cognition [Clark, 1997], arguing that the human body and the surrounding environment cannot be treated as separate entities in the study of cognitive processes. Therefore, minds cannot be considered as "passive representational engines" responsible only for creating internal models of the external world. In fact, there exists a complex relationship between the internal mental processes and the external processes. It involves coordination at different time scales between the internal mental resources (such as memory, attention, and executive functions) and external resources (such as objects, artifacts and elements of the surrounding environment). The principle of embodied cognition is highly relevant for the design of collaborative environments, where the various collaboration tools play a role as extensions of individuals, rather than a simple means of causing stimuli in a disembodied cognitive system. In this way, collaboration tools can foster new ways in which collaborators can think, visualize and effectively steer around a shared activity.

Finally, the third tenet regards the role of culture on the way human cognitive processes are shaped and governed. Hollan et al. [2000] discussed the role of culture as the means of shaping the cognitive processes through the influence of history of usage of artifacts and established work practices; and thus suggesting that cognition cannot be isolated from the culture. While embodied cognition refers to the usage of tools as an extension of one's cognition, culture takes into account the history of socially agreed work practices and interactions with artifacts (such as whiteboards, data visualizations, paper notes, etc.), which are developed over a prolonged period of time. Culture can be regarded as a process of accumulation of partial solutions to frequently encountered problems. In the absence of this residue of experience and knowledge from previously encountered problems, individuals would have had to start from scratch to search for solutions. We believe that culture heavily influences the design of meeting technologies and collaborative environments, because the adoption of new interactions or work processes which are disparate or non-conventional to the already existing ones might lead to non-synergistic behavior within groups. On the other hand, there can be a possible restrictive effect of culture on cognition in a sense that the way individuals think about certain things is rendered non-malleable because of culture, where a new way of thinking could be more effective.

The main consequence of these three tenets (social distribution of cognition, embodiment, and cultural immersion) was a call for a new kind of cognitive ethnography, which considers the environmental materials and the social construction of meanings through actions [Hollan et al., 2000]. This is in contrast to prior methods that focus only on the study of individual cognition. The main methodological focus of Hutchins, therefore, is on events or actions within the system. He argues that, it is not just important to study the internal mental processes, it is equally necessary to regard how the information being processed is positioned within the material and social world.

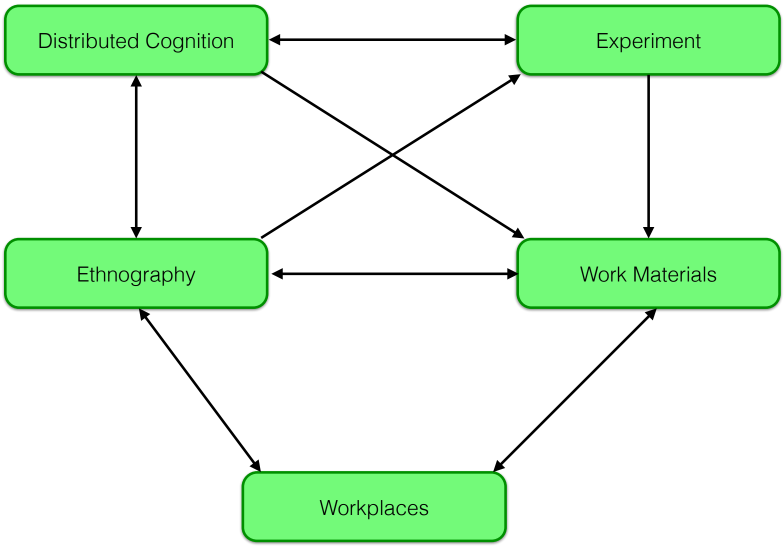

Figure 3.1 – The integrated research framework of studying distributed cognitive systems, presented by Hollan et al. [2000].

Hollan et al. [2000] also proposed an integrated research framework (as shown in Figure 3.1) that consists of a loop from observation to theory to design and back to new ethnographic observations, in order to design and test work environments that effectively support the execution of various collaborative functions during an activity. In this framework, distributed cognition theory identifies the individual actors, information trajectories, and various external resources incorporated during the information processing activity. In addition, cognitive ethnography is employed to observe, document, and analyze the ongoing phenomena, with a particular emphasis on the direction of information flows, cognitive properties of the system, social organization, and cultural processes. Finally, the impact of any variable or a process can be tested through experimentation.

3.1.2 Distributed Cognition in Collocated Meetings

We consider a face-to-face meeting as a distributed cognitive system, and employ the observational methodology of cognitive ethnography. We describe various relevant interactional processes occurring during meetings. Furthermore, we show how these processes influence the design of our meeting technology.

The most prevalent interactional processes occurring during meetings are conversations among group members, the use of gestures while referencing and establishment of common ground, and the use of gaze to signal attention during conversation. The established dominance of these processes in CSCW and CSCL research is made evident by the fact that they are the primary resources that are used in the assessment of collaboration quality. Many researchers (for example Baeza-Yates and Pino [1997], Meier et al. [2007], Kahrimanis et al. [2009], etc.) have developed and employed methods to assess collaborative behavior through the analysis of conversations and gestures in audio- and video-recorded meeting sessions. Other than conversations and gestures that are ephemeral in nature, there are other processes that play a role in collaboration, for example when collaborators leave traces of crucial episodes in the form of written content and external representations [Snyder, 2014]. The nature of them can be both private and public, thus also defining the accessibility of information for the group members.

Whiteboards, flip-charts, and presentation slides are some of the resources often used by collaborators as publicly accessible information storage for the group. The interaction with these resources generally implies the broadcasting of information meant for the whole group. For example, when one of the group members explicitly approaches the whiteboard to present an idea, it is considered as a sign for the group to converge their attention on the whiteboard. In addition, the interaction with these resources may also define the role of the user, which can also be dynamically acquired based on who uses the whiteboard as suggested by Dourish and Bellotti [1992]. Furthermore, whiteboards and flip-charts are considered as permanent storage media for crucial facts and knowledge discussed during the course of a meeting, as they allow for offloading the group knowledge such that it can be made permanent [Dillenbourg and Traum, 2006]. Similarly in Hutchin's example [Hutchins, 1995], a speed-correspondence booklet and speed-cards are used by a team of pilots in cockpit as publicly accessible (to the group) external representations, to coordinate and regulate airplane speed and wing configuration during the landing phase.

By the end of a collaborative session, the shared information upon the whiteboard can be assumed to be mutually agreed upon and well grounded by the group, as it can be assumed to have undergone a process of negotiations to establish a common meaning. Moreover, collaboration over whiteboard can denote negotiations and discussions within the group, which can also be reflected by how the content on the whiteboard is manipulated, and by whom. Finally, these resources can also be regarded as shared workspaces, and the shared content over them can be considered to be owned by the whole group.

Besides the publicly accessible work materials such as whiteboards, group members also spend a significant amount of time interacting with private work materials (artifacts) such as notebooks and documents (digital or on paper). These artifacts are generally used in an individualistic manner, and can be regarded as part of the individual information storage and processing corresponding to private territories as identified by Scott et al. [2004]. Also, unlike the tightly-coupled collaboration phase when the group interacts with the whiteboard, interaction with the individual notes corresponds to a loosely-coupled phase (refer to "collaboration coupling" as defined by Tang et al. [2006]), because the individual is cognitively decoupled from the collective for some time. The use of these resources can also be assumed to be a consequence of task-based division-of-labor. Segments of the information contained in individual notes can be made public by sharing them with the whole group. This can be achieved by replicating this information on the public medium such as a whiteboard, or passing the notebook around the group for everyone to see. These individual artifacts thus serve as private workspaces to work out on an idea before sharing it with the group, or serve as recording medium for the trace and history of the ongoing meeting by noting the crucial episodes and the decisions made during the meeting.

At the core of this research work are the physical media used to create, store, and share information during collaboration, as well as the representations created on them. Depending on the context, the term "representation" can mean different things. Scaife and Rogers [1996] described the distinction between representation as a process, and representation as a product: the transformations and preservations that lead to the attainment of representations as a final product, defines the process part of representation. In our research, we are interested in the process of creating representations more than the end-product, and the influence of this process on the dynamics and outcome of the collaboration. Therefore, we intend to use a meeting technology that equips group members with a shared workspace as a medium to create and manipulate digital artifacts. The shared workspace, similar to the one available on whiteboards and flip-charts, will assist the group members to create and share artifacts, and enable us to record and analyze the group interactions in real-time. Besides providing a shared workspace, we would like to empower the group members with simultaneous interaction ability. Tang [1991] emphasized on the need for simultaneous (or concurrent) interaction ability with the shared workspace, as it enables group members to work on the same part of the problem at the same time. In addition, he indicated that concurrent activity over a shared workspace reduces the competition for conversational turn-taking, and mitigates the problem of floor monopolization where one group member explicitly acquires the workspace and barely allows others to express their thoughts. However, concurrent activity is difficult and troublesome over physical whiteboards and flip-charts because in order to interact with them, the individual has to be in close proximity to them. This requires collaborators to stand in front of the whiteboard, and also obstructs the content shared over them for other group members. Therefore, we choose to use a *Single Display Groupware* (SDG) [Stewart et al., 1999] as the medium for our research, since it provides a shared workspace that is accessible from a few meters distance as well as simultaneous interaction ability for the group members. In the next section, we will describe Single Display Groupware as a meeting technology with its different form factors, and review the relevant research conducted in this area.

3.2 Single Display Groupware

The desktop metaphor in computing was initially designed for individuals with no thought for collaborative usage. Since the inception of the first personal computer, modern day computers haven't changed much in the way users interact with them. A single mouse and a keyboard are still the input means, and the interface allows only for one mouse pointer to interact with various windows, menus and to set the input focus for text entry. Connecting multiple mice to a single desktop doesn't change the situation, as the operating system merges the multiple input streams (corresponding to multiple mice) into a single input stream represented by a single mouse pointer. Further, out of multiple mice or keyboards connected to a computer, only one is considered active at any time.

Stewart, Bederson, and Druin [1999] challenged this desktop paradigm by proposing a model for collocated collaboration called *Single Display Groupware* (SDG) or programs that enable collocated group members to work together while using a single shared display and simultaneous use of multiple input devices - one for each group member. While designing technologies for collaborative learning activities involving elementary school children, Druin et al. [1997] observed that children in small groups crowded in front of the computer screens during learning activities. In addition, the learning experience was observed to be more enjoyable when a child could control the application while interacting with a mouse. This was the main motivation behind the development of the SDG framework by Stewart et al. [1999]. Besides this example of children collaborating during learning activities, there are numerous situations where there is a need to support collaboration on a single shared display and an ability to freely interact with the computer application using an input device for each collaborator, such as pair-programming, product design reviews, and so on. Such an interaction technique is contrary to the individualistic desktop metaphor, as it concerns the design of applications and other GUI (graphical user interface) elements that are responsive to multiple inputs.

SDG can be regarded as a subset of groupware that specifically focuses on collocated collaboration with multiple individuals at the same time and place. Stewart et al. [1999] differentiated

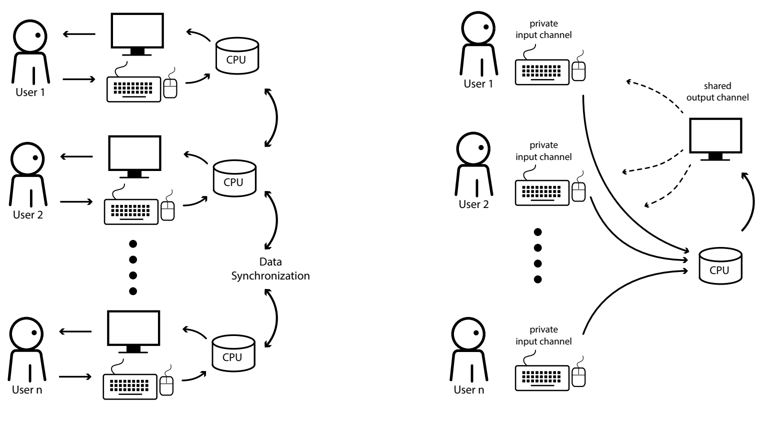

Figure 3.2 – Conventional Groupware (on the left) versus Single Display Groupware (on the right). This figure differentiating the two kinds of groupware is taken from Stewart et al. [1999].

SDG from a conventional groupware by comparing the difference in input and output channels. An input channel enables a user to interact with the computer application through the use of an input device. An output channel is a means through which the computer application communicates back to the user through a display, speech or haptic feedback. A conventional groupware equips each group member with an input and output channel of their own (similar to multiple networked computers), and at the background the data models (or databases) corresponding to different users are synchronized (as shown in the left-side of Figure 3.2). In single display groupware there is no need to synchronize multiple data models because there is just one data model present on a single computer with a single shared output channel and multiple input channels connected to this computer (as shown in the right-side of Figure 3.2). Further, the SDG framework is not strictly restricted to mouse and keyboard as the only input channels; multi-touch tabletop applications, and collaborative motion sensing games with Microsoft Kinect[1] are also considered as SDG. Based on the idea of input and output channels, Stewart et al. [1999] identified three main characteristics which can be used to discriminate SDG from conventional groupware:

1. *Shared User Interface*: In SDG, the user interface and its elements such as buttons, menus, etc. should be designed to respond to multiple simultaneous input devices. However, in a conventional groupware the interface is designed to respond only to a single input channel, similar to conventional GUI applications.
2. *Shared Feedback*: The users interacting with a SDG are presented with the same interface due to the presence of a single display. Therefore, any changes made in the state of the interface by one user are reflected to other group members immediately because the group shares a single view.
3. *Coupled Navigation*: As users share a single interface in SDG, a user navigating to a different part of the data model will affect others as well. In case of tight coupling, the whole group navigates together. However, if loose coupling is implemented, one user's navigation might partially obscure with another person's view.

Enabling multiple equivalent input channels over a single shared display has numerous advantages as discussed by Stewart et al. [1999]. Explicit provision of a shared working space coupled with individual input channels reduces existing proximity-based social barriers; which inhibits collaboration in the absence of SDG. In addition, SDG enables parallel work which was observed to enhance the collaboration experience by making it more enjoyable for the collaborators [Druin et al., 1997; Stewart et al., 1998] (see Figure 3.3), and reduces the potential conflicts associated with turn-taking to take control of the input device in order to interact with the application. SDG can also prove beneficial for peer-tutoring and peer-learning scenarios, as the roles associated with different input channels can be scripted in the SDG application development. Finally, SDG can mitigate floor monopolization within groups where one dominant group member takes control of an input device and monopolizes the task. In such a case, even less dominant group members can contribute to the activity

[1]Microsoft Kinect URL: https://www.microsoft.com/en-us/kinectforwindows/ (visited on 04-May-2015)

Figure 3.3 – Kidpad [Druin et al., 1997] is a drawing application that allowed a group of two children to draw simultaneously using the two mice provided to each child. In their study with Kidpad, Druin et al. [1997] observed an increase in engagement and enjoyment of the drawing activity as compared to the single-mouse condition.

by interacting with the workspace through their own input channel. However, there can also be some disadvantages associated with the new form of interaction enabled by SDG, such as conflicts arising during concurrent access to the same part of the work area by different collaborators. Also, the lack of support for the access of multiple input channels in almost all modern operating systems makes it significantly harder for application developers and researchers to develop SDG applications. In the next section, we will discuss the issues with the development of SDG applications, and how researchers have found varied solutions to tackle with these issues.

Single Display Groupware as a model for co-present collaboration is highly significant for our research work as it suffices our functional design requirements. The provision of multiple equivalent input channels corresponding to each group member enables unconstrained and concurrent interaction with the workspace. In addition, a single shared display provides a single common perspective for the whole group to the state of the activity in the form of shared content and artifacts. Further, concurrent access to shared artifacts might facilitate discussions and negotiations among the group members over the shared content, and consequent establishment of common ground. Finally, as each group member is equipped with a separate input channel; this enables us to uniquely identify, record and analyze each group member's interaction with the shared workspace.

In the next sections, we will discuss previous research works investigating the influence of SDG in various collaborative scenarios, and the computer science perspective into the issues related with the development of SDG applications. In addition, I will also perform a brief review of different SDG systems that were employed in educational and organizational scenarios.

3.2.1 Support for the Development of SDG

The peculiarity of SDG's characteristics as compared to conventional single-user GUIs makes it difficult to rapidly design and develop SDG applications [Tse and Greenberg, 2004]. In addition, the lack of support in modern operating systems as well as the absence of effective cross-platform software frameworks and APIs (application program interface) to ease the development of SDG applications exacerbates the complexity involved in the development of SDG. Consequently, most SDG applications are developed from scratch, and require application developers to handle low-level input events to enable multiple simultaneous interactions. This further reduces the reusability of SDG applications and frameworks because the design and development of these applications is tightly bound to the context and specific to a platform. One such example is MMM (Multi-Device Multi-User Multi-Editor) by Bier and Freeman [1991], which enables multiple users to enter text into different text editors as different windows with simultaneous input focus on a single computer system. Each user was equipped by their own mouse and keyboard. However, MMM was never made available to the SDG research community, and therefore its functionalities were not extended to other tasks, and no user study was conducted with it to study its influence on co-present collaboration [Stewart et al., 1999].

However, some researchers have made significant attempts in making it easier to develop and test SDG applications by developing software frameworks and also embedding support for SDG at the operating system level. Bederson and Hourcade [1999] developed a software package in Java called MID (Multiple Input Devices) that extended the available mouse events in Java by providing additional mouse events corresponding to different devices. MID enabled SDG developers to handle events from different USB (Universal Serial Bus) mouse devices connected to a single computer, and thus significantly reduced the application development time. Kidpad [Druin et al., 1997] (see Figure 3.3) was a SDG application that enabled a group of children to draw in a single application window using multiple mice; was developed using the MID framework. However, MID only worked with the Windows 98 operating system, and was never updated to work with other operating systems. In addition, MID didn't provide support for handling events from multiple keyboards connected to a single computer; thus limiting the functionalities of the SDG application.

Tse and Greenberg [2004] performed an exhaustive requirement elicitation for the design of a SDG software framework, and presented *SDGToolkit* which can be used by application developers to rapidly prototype SDG applications without concerning themselves with low-level handling of input events from multiple devices. SDGToolkit was developed in C# programming language, and is supported on Windows operating systems, starting from Windows XP. Further, SDGToolkit was the first attempt in uniquely identifying and gaining multiple input streams from both mice and keyboards, and it provided support for multiple mouse cursors - one for

(a) Third-grade students in a class in Chile solving exercises on geometrical properties of triangles collaboratively. The problems are presented on the public display in front of the class, and each quadrant of the display corresponds to one student group.

(b) Each child is equipped with a mouse device of their own, through which the kids interact with the SDG application.

Figure 3.4 – *One Mouse Per Child* Project: The project aims to develop affordable interpersonal computers for classrooms in developing countries. These images were taken from the study of Caballero et al. [2014] focusing on the learning of geometrical properties of triangles for third-grade students.

each user. Besides the capability to handle input events from multiple devices, SDGToolkit also supplies application developers with simple interface elements and widgets to significantly reduce the development time. In addition, SDGToolkit facilitates the development of SDG applications with different display orientation; i.e. for both vertical displays and horizontal tabletop displays.

Microsoft Research and Miguel Nussbaum of Pontificia Universidad Catòlica de Chile are collaborating on an educational project for schools in developing countries called "*One Mouse Per Child*"[2] project [Alcoholado et al., 2012] as shown in Figure 3.4. The project aims to use SDG and multiple mice in immersive and collaborative learning activities through gamification. In this context, Microsoft developed an API similar to the SDGToolkit [Tse and Greenberg, 2004] to foster the development of SDG applications using multiple mice devices. The API is called *MultiPoint Mouse SDK*[3] and enables application developers to easily handle input events from multiple mice connected to the same computer. The API can effectively handle up to 25 connected USB mouse devices, however there is no support for multiple keyboards. In addition to the ability to handle events from multiple mice, the API also provides basic widgets and provision for the development of SDG enabled interface elements.

Hutterer and Thomas [2007] went a step further to enable multiple input support for legacy GUI applications, by embedding SDG support into the X windowing system on UNIX based operating systems (see Figure 3.5). The authors presented *Multi-Pointer X Server* (MPX) as a first Groupware Windowing System (GWWS) together with a *Multi-Pointer Window Manager*

[2] One Mouse Per Child Project Webpage: http://research.microsoft.com/en-us/projects/onemouseperchild/ (visited on 24-February-2015)

[3] MultiPoint Mouse SDK: http://www.microsoft.com/multipoint/mouse-sdk (visited on 24-February-2015)

(MPWM). MPX and MPWM enable support for multiple true system cursors corresponding to each mouse connected with the computer, as well as sophisticated floor control for input-focus related conflicts and window overlays, facilitating annotations. The main advantage of embedding SDG support in the windowing system is the availability of collaboration support within plethora of legacy applications without any additional change. In addition, software developers have access to conventional APIs and toolkits while developing novel features in the application without relying on a single toolkit that singularly supports SDG functionalities. Further, MPX also supports the simultaneous execution of multiple SDG and legacy applications as different windows on a single display.

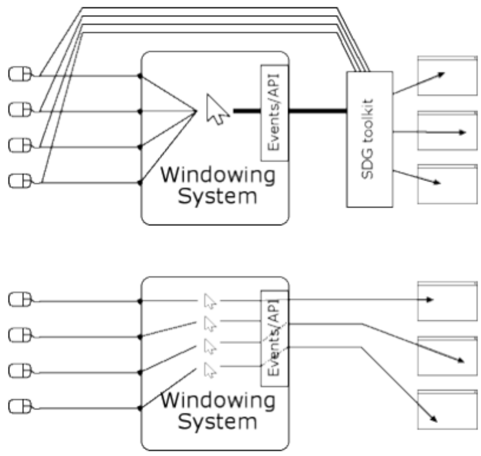

Figure 3.5 – A comparison between a typical SDG API such as SDGToolkit [Tse and Greenberg, 2004] (the figure at the top), and the implementation of SDG support in the operating system's windowing system by Hutterer and Thomas [2007] (the bottom figure) as presented in their article.

Finally, regarding the development of SDG applications for tabletop based interaction, Shen et al. [2004] presented the *DiamondSpin* API to support multi-user, concurrent and touch-enabled interaction for tabletops. DiamondSpin was developed in Java and was designed to support application development for varied tabletop shapes; mainly polygonal shapes (rectangular, circular, etc.). Further, the API enabled developers to incorporate separate work areas and also allows for arbitrary orientation and positioning of documents and windows to support multi-directional viewing angles during collaboration around a tabletop. In addition to supporting concurrent handling of multi-touch events from different users, DiamondSpin also supports multi-device events from connected mice. Similar to SDGToolkit [Tse and Greenberg, 2004], DiamondSpin also provides application developers with a selection of widgets and interface elements specifically modified and attuned for tabletop interaction.

3.3 Use of Single Display Groupware

In their attempt to distinguish personal computers from interpersonal computers, Kaplan et al. [2009] indicated that interpersonal computers such as SDG that allow for collective manipulation and group awareness (through simultaneous interaction with a common shared workspace) have a positive influence on collocated collaboration. SDG introduced a new form of interaction with shared workspace, which was previously attained through networked personal computers. Therefore, many researchers have studied the influence of sharing a workspace in collocated settings with SDG. The studies have focused on varied collaborating audience and settings such as children solving collaborative puzzles [Inkpen et al., 1999; Scott et al., 2003], learning activities in classrooms [Moraveji et al., 2009; Alcoholado et al., 2012; Caballero et al., 2014], and organizational meetings [Izadi et al., 2003; Wigdor et al., 2009]. In this section, we will discuss the findings from some of the studies focusing on the effects of SDG.

Druin et al. [1997] and Stewart et al. [1998] have reported the first observational studies concerning the use of a SDG application with children. Both these studies compared Kidpad - a collaborative drawing application running on a single computer with two mice devices (see Figure 3.3) with a similar drawing application using just a single mouse. Their observations showed an enhancement in the overall collaboration experience and an increase in enjoyment of the activity while using Kidpad as compared to the condition with only a single mouse.

In another set of similar studies with children solving a puzzle collaboratively with a SDG, Inkpen et al. [1999] and Scott et al. [2003] also observed a significantly higher level of engagement with the activity, increased activeness and preference for a SDG as compared to a system with a single mouse. Both these studies analyzed the collaboration between two kids in three conditions: solving a puzzle on a paper, solving the same puzzle on a computer equipped with a single mouse and the third condition provided a pair of kids with a computer with two mice. The first condition where the children solved a puzzle on paper was similar in nature to the collaborative drawing activity in the observational study of Tang [1991]. Both studies reported that off-task behavior (frustration, boredom, distraction, etc.) was exhibited significantly more in the condition with a single mouse device, as compared to the other two conditions. The authors observed that the kids were frequently active at the same time in the two-mice and the paper condition, however this was not the case in the single mouse condition as only one child could interact with the puzzle; thus the other child displayed off-task behavior as boredom and frustration due to the lack of interaction with the activity. Further, Scott et al. [2003] observed that the paper-based condition was similar to the SDG condition with two mice because the kids could simultaneously engage in the activity; similar to playing a video game in the multiplayer mode with two joysticks. Finally, Scott et al. [2003] conducted a second study where the kids solved the puzzle collaboratively, but were asked to use different computers connected over a network. Their findings suggested that children had difficulty establishing a common ground and mutual understanding when using individual and separate displays as compared to the SDG condition. In addition, the puzzle was perceived relatively easier in the SDG condition as compared to the condition where kids used separate displays to collaborate. These findings of Scott et al. [2003] are in-line with the observations made by Wallace et al. [2009] in their study investigating the role of SDG and MDG (Multi-Display Groupware) on teamwork and taskwork within collaborating teams. In their study, Wallace et al. [2009] have established that SDG is beneficial for coordinating over shared resources (teamwork), whereas multi-display groupware (MDG) is advantageous for individual task duties (taskwork). These findings clearly show the benefits of using a SDG in collocated settings. In the next two sub-sections, I will specifically regard the usage of SDG in educational settings and organizational meetings.

3.3.1 In Educational Scenarios

Education was envisioned as one of the possible usage scenarios of single display groupware by Stewart et al. [1999]. The presence of multiple input channels enable learners to collaborate on the same problem space, and coupled navigation associated with the shared workspace might help learners to guide each other within the problem space. Thus SDG could have positive implications towards the design and testing of peer-learning and peer-tutoring scenarios.

Previous studies by Inkpen et al. [1999] and Scott et al. [2003] focused on collaboration between a group of children, instead of a classroom. Moraveji et al. [2009] conducted one of the first user studies in a classroom full of students, and focused on the usability issues related to the use of SDG, rather than learning itself. The authors investigated the influence of group size (ranging from 1 participant to a group of 32 participants) on children's task performance while using a SDG. An increase in the size of the group might impose constraints on the size of the display and the content that can fit over it. In addition, presence of numerous active mouse cursors (corresponding to each member of the group) might make it harder for individuals to identify their respective cursors, cause clutter of mouse pointers, might lead to visual distraction, and might also occlude the content underneath the cursors; this in turn might affect the performance of the group. In the study of Moraveji et al. [2009], each child was provided with her own mouse, and the group of children aged 10-12 years were seated in a classroom facing a large wall-mounted projected display. The groups were asked to complete two simple tasks of reciprocal pointing at an object (the color of the object different from the color of the user's cursor), and entering a five-lettered word while using an on-screen keyboard. Their findings suggest that SDG can be used for classroom level interactions; provided that the target size is not too small and the cursors are designed carefully to minimize occlusion. Further, the effect of the group size was observed to be more evident when pointing was distributed spatially and temporally (such as different users using on-screen keyboard to enter different words). However, group size had no significant effect when everyone was trying to aim at the same target concurrently. These implications for the design of the learning activity with a SDG are that one should distribute the different learning tasks in a way that the problems related to occlusion and distraction can be minimized.

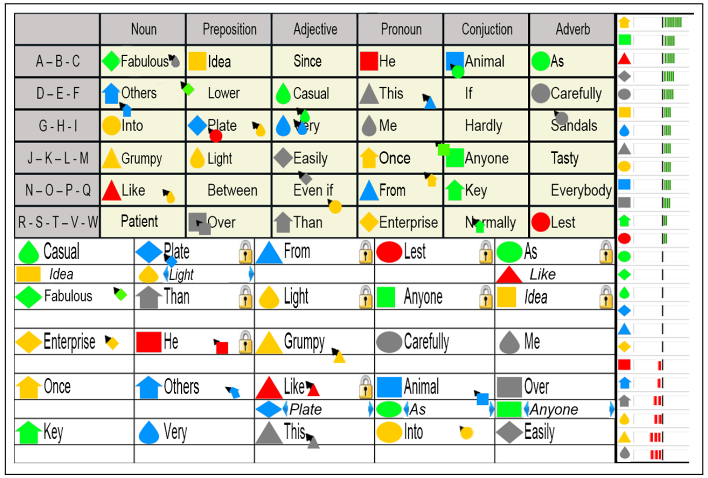

(a) Screenshot of the Spanish word-class association activity by Szewkis et al. [2011]. The students (represented by each cell in the bottom-half of the screen) are required to place the assigned word in correct order in the matrix at the upper-half of the screen.

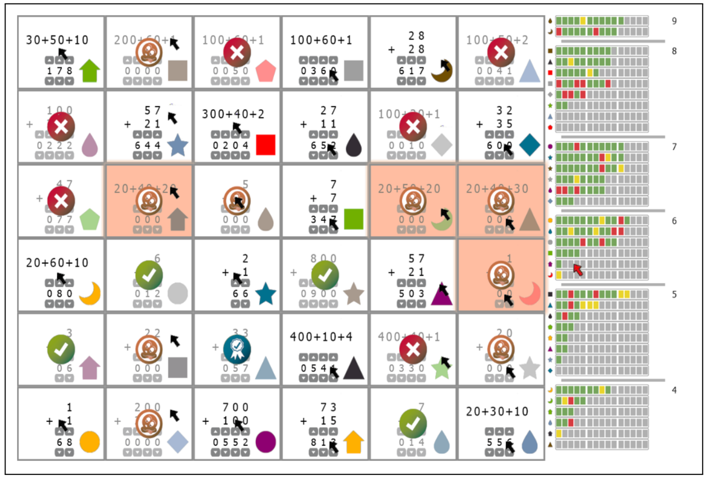

(b) Screenshot of the arithmetic activity by Alcoholado et al. [2012]. Each cell corresponds to the exercise for one student, and the column on the right shows the activity status of each student.

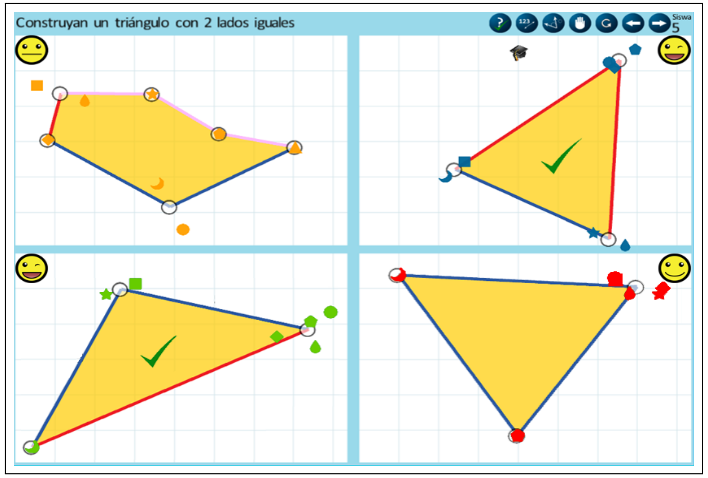

(c) Screenshot of the triangle activity by Caballero et al. [2014]. Each quadrant corresponds to the exercise for a group of students.

Figure 3.6 – Screenshots of the various SDG applications within the *One Mouse Per Child* project.

Next, in the framework of *One Mouse Per Child* (OMPC) project several studies have been conducted to examine the effectiveness of learning with a SDG in real classrooms (see Figure 3.4). Unlike the *One Laptop Per Child*[4] project, OMPC aims to design scripted learning scenarios and equip schools with affordable technological solutions in developing countries such as Chile and India. Since SDG requires a single computer per classroom, and a mouse for each student, the cost is significantly reduced compared to providing each student with a laptop [Alcoholado et al., 2012]. Besides the affordability, SDG has been shown to be beneficial for collaboration among children, with increased activeness, engagement and enjoyability of the activity [Inkpen et al., 1999]. The user studies conducted within the framework of the OMPC project have focused on diverse sets of learning themes and subjects, such as spanish language learning [Szewkis et al., 2011] (see Figure 3.6a), simple arithmetic [Alcoholado et al., 2012] (see Figure 3.6b), and geometrical properties of triangles [Caballero et al., 2014] (see Figure 3.6c). In all these studies, a class of up to 25 (or more) students were equipped with a mouse, through which they can interact with exercises presented to them on a large projected display in front of the classroom.

The research findings in these three studies (Szewkis et al. [2011]; Alcoholado et al. [2012]; Caballero et al. [2014]) showed a significantly higher learning gain while using a SDG (with and without the intervention of the teacher), as compared to the control condition where students learned the same concepts in a typical classroom environment with a teacher. Alcoholado et al. [2012] also observed that the use of SDG was specifically more beneficial (higher learning gain) for weaker students (with lower pre-test scores), mainly because the SDG adapted to the level of the student and presented suitable arithmetic problems (mainly addition exercises for third grade students). They also examined the usability issues related to the use of SDG in schools, in two completely different socio-economic cultures (schools in Chile and India). They demonstrated that the system was easy to use, even for the students with no prior experience with computers, and students had no difficulty in identifying and interacting with their respective private workspace on the shared display.

Finally, Szewkis et al. [2011] and Caballero et al. [2014] defined and studied the idea of "silent collaboration" in contrast to spoken collaboration while learning with a SDG in a classrooms. Unlike spoken collaboration, silent collaboration works efficiently in classroom scenario (with approximately 30 students) as the students collaborating with each other may not be sitting adjacent to each other. Besides, spoken collaboration might not be very feasible with many students and might easily lead to a ruckus within the classroom, and thus makes the situation harder for the teacher to manage and orchestrate. In their study on spanish accent rules and word-class associations, Szewkis et al. [2011] designed the learning activity to facilitate silent collaboration in a way that students were required to place the word assigned to them in a correct position and class within a matrix (see Figure 3.6a). In order to successfully complete the activity, the student was required to exchange the word with her peer by clicking on the desired word to be exchanged (belonging to another student). However the other student

[4] One Laptop Per Child Project Page: http://one.laptop.org (Visited on 25-February-2015)

has a choice to accept or reject the exchange request of the former student; in which case the former student has to find an alternative word to exchange. Silent collaboration is rendered possible because each student has access to the workspace of all her peers in the classroom over a common shared workspace making it easier to compare one's solution with other's, and the additional ability to interact with other's workspace with the equivalent input channel. Further, the results of the study revealed that silent collaboration positively influenced the learning gain of the students.

In another study concerning the geometrical properties of triangles, Caballero et al. [2014] observed silent collaboration of a different nature where upon finishing the assigned task, the student openly assists other students to complete their activity. Such a collaboration where students openly intervene and interact with their peers is made possible because of the use of SDG.

3.3.2 In Meetings

Unlike the education scenario, the effect of collaborating with a SDG is not very well investigated in organizational scenarios, such as the role of a SDG in group's performance during a task. One possible explanation for the lack of studies assessing collaboration with a SDG is that unlike in educational contexts organizational collaboration is complex to assess and analyze due to the change of dynamics from one group to another, interpersonal relationships among collaborators, the degree of flexibility ingrained in the nature of the desired solution, and the lack of a baseline condition to compare one collaboration to another. However, emphasis has been given to the design and development of SDG technologies that enable collaborators to fluidly share content, visualizations and applications from a personal computer onto a public display such as *PointRight* [Johanson et al., 2002b] (see Figure 3.7a), *Dynamo* [Izadi et al., 2003] (see Figure 3.7b) and *WeSpace* [Wigdor et al., 2009] (see Figure 3.7c).

Dynamo [Izadi et al., 2003] (see Figure 3.7b) was designed as a public wall-mounted surface intended to interactively access, share and exchange digital media (documents, pictures, etc.) spontaneously by connecting external storage, laptops or PDAs among multiple users. In addition, users can also remotely connect to their personal desktops to access digital media, as well as the surface allowed for simultaneous interaction through multiple mice and keyboards. Further, Dynamo was also intended to be used in public spaces (such as a corridor or a foyer of a conference hall) by users as an extension of their personal physical space. In addition to the provision of accessing, viewing and exchanging digital media, Dynamo also enabled users to annotate information. Upon connecting with Dynamo, the users can select a part of the surface as their personal workspace where they can share information. The users had the choice to make their workspace public for others to access or they can grant access to a selected few who can access the information shared within this workspace. Dynamo also allowed for the deferred exchange of information among users by dropping a package with access rights only for a specific person or a group. Izadi et al. [2003] conducted a semi-formal

(a) *PointRight* [Johanson et al., 2002b] was developed as a mouse and keyboard input redirection software for multi-display, multi-user and multi-computer environments such as the iRoom [Johanson et al., 2002a] environment as shown here.

(b) *Dynamo* [Izadi et al., 2003] is a communal display designed for the purpose of exchanging digital media spontaneously. Users can select an area on the display as their private workspace, and can also provide access rights to others over their workspace and shared content.

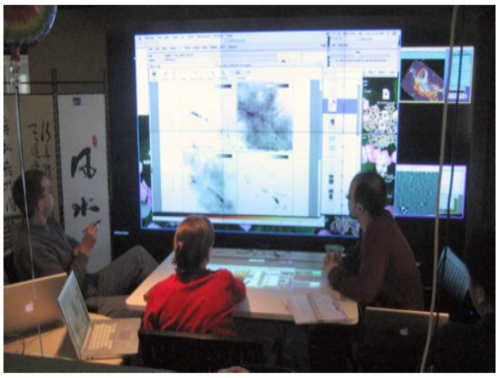

(c) *WeSpace* [Wigdor et al., 2009] was designed for the purpose of sharing and visualizing large data in research oriented meetings. It consists of a large wall-mounted display and a multi-touch table. Users can also connect their individual laptops to share content.

Figure 3.7 – Examples of various SDG systems implemented for the usage in meetings.

usability evaluation of Dynamo with 65 conference participants, and the participants who used the systems reported on the benefits of simultaneous interaction during information sharing and found it easy to use.

Unlike Dynamo [Izadi et al., 2003] which was designed as a communal workspace, WeSpace [Wigdor et al., 2009] (see Figure 3.7c) was designed for the purpose of large data visualization and exploration for scientists and researchers meeting in collocated small groups. WeSpace comprises a large wall-mounted display and a multi-touch table. In addition, WeSpace allows multiple users to connect their laptops and share data, visualizations and applications simultaneously over the data wall, and the multi-touch table enables the group members to explore and interact with the shared data. The users can choose to share the renderings of the data they wish to share with others without revealing the underlying data. Later, the users can annotate and overlay several visualizations to compare, and save the overall product of the collaboration. Wigdor et al. [2009] conducted no formal evaluation of WeSpace but performed an observational study with a group of three researchers, and observed a balanced participation among the group members during the meeting session.

Finally, PointRight [Johanson et al., 2002b] (see Figure 3.7a) is a software framework which allows multiple users to use multiple mice and keyboards to interact with multiple display environments and machines connected together. Originally, PointRight was developed as a mouse pointer and keyboard redirection system for the four wall-mounted displays (SMART Boards[5]) and a bottom-projected tabletop present in the *iRoom* environment [Johanson et al., 2002a]. The novelty associated with PointRight was the implementation of the geometric model capable of redirecting mouse and keyboard input across displays driven by multiple independent machines and running different operating systems. Users can connect their laptops to one of the displays in the environment and use PointRight to share information and applications from the laptops to the public displays. However, in PointRight as only one mouse pointer can be active on a single display or a machine, it greatly inhibits the collaboration that can happen over the shared content.

3.4 Single Display Groupware as an Analytics Tool

Methods of rating and assessing collaboration (such as the one presented by Meier et al. [2007] and Kahrimanis et al. [2009]) rely on video and audio recordings of collaboration sessions to assign various codes to events and behaviors before assessing them. This is considered to be a difficult and extremely time-consuming process as it requires multiple raters and coders to navigate through the video recordings and audio transcriptions to code the desired behavior. The transcription of a collaboration session is in itself a rigorous and time demanding task. Thus it is greatly desired for this assessment process to be automated. In their article on the implications of distributed cognition on human-computer interaction (HCI) research; Hollan et al. [2000] have cited Hill and Hollan [1994] on the criticality of automated recording of histories of interaction with technology as an increasingly important source of data for analysis. Jermann [2004] also emphasized the need to collect the interaction and participation data of the collaborators in real-time via meeting technologies or the collaborative environment. This interaction data can then be integrated (in case of multi-modal data), analyzed, and presented back to the group as an easy-to-interpret awareness information. Further, Jermann [2004] demonstrated that this awareness information about group's interaction and participation has a positive influence on collaboration, and enables group members to coordinate and regulate the future set of events (for more discussion on awareness tools, refer to "Workspace Awareness Tools and Group Mirrors" in Section 2.2.2).

[5]SMART Boards Webpage: http://education.smarttech.com/ (visited on 26-February-2015)

The deployment of interactive meeting technologies can prove to be of great assistance in the process of collaboration assessment; as the interactions made by the users with the technology can be recorded and analyzed by the system. Later, analyzed interactions can be interpreted by the group members, and the researchers to better understand the collaboration, and make informed decisions regarding ways to improve the collaboration. Further, the consistent increase in computational power and advances in sensor technology in tandem with smart implementation of machine learning algorithms might lead to the automation of the process of collaboration assessment. Our research work regards the process of creation, manipulation and sharing of artifacts during collaboration while using a SDG. Therefore, in this section I will discuss the role of SDG as an analytics tool; capable of recording interactions of the users, analyzing them and presenting the group with appropriate feedback of their participation.

3.4.1 Data Collection

Concurrent input channels in a SDG, one for each collaborating member, can be used to uniquely identify as well as record interactions of different group members. Besides logging the knowledge about who is interacting, SDG enables researchers to collect data about the kind of interaction as well, such as the input device being used (mice, keyboard, touch, etc.), the type of the event being made (drag, click, key press, etc.), and so on. However, the kind of input device being used or the type of event are not very informative by themselves. Therefore, combining this information with the object, artifact or the tool that is being used by an individual and the knowledge of the task, might help researchers to ground the event within the appropriate context and defines the semantics of the interaction. For example, when a user opens a powerpoint presentation or a product design image on to the shared workspace, it tells a lot about the role of the user as well informs the researchers about the kind of task. Also, brainstorming tasks can be identified by the groupware, when group members initially create textual ideas concurrently, and during the later phase of the task try to assimilate this initially shared information.

In addition, to the interaction with the workspace, SDG researchers can also record and analyze multi-modal data such as speech and gaze. Prosodic features such as speech rhythm, stress and pitch can be extracted in real-time and has already been achieved by Bachour et al. [2010] in their roomware *Reflect*. Recording gaze in an accurate manner is a challenging task as it requires group members to wear mobile eye-trackers, and also requires additional post-hoc analysis on the gaze data to achieve meaningful results. Further, the social information concerning the proximity and associations of artifacts created by different group members might also provide valuable information about the collaboration. For example, if a group member places her text in a close proximity to an already existing text created by someone else, this might indicate a semantic relationship between these two objects. Such a combination of multi-modal data and the interactional logs collected with the groupware might inform the researchers in social signal processing (SSP) to easily discern the context of the activity, and thus facilitate the automatic analysis of varied collaboration aspects.

The nature of data that can be recorded and analyzed via a SDG is quite diverse and rich, and can be very informative about the ongoing collaborative processes. One example of the richness of interactional data while using a SDG was observed during the use of WeSpace [Wigdor et al., 2009] by a group. The initial observation of the interactional data by Wigdor et al. [2009] showed that all the three participants had relatively equal participation during the task. However, analyzing the interaction data in 5 minute interval over the whole activity showed turn-taking among the participants. During each of these 5 minute intervals, one of the three participants would take the floor and dominate the discussion by taking control of the input with the shared workspace, as the majority of the input during this interval was made by one participant.

3.4.2 Group Process Feedback

In their study with primary school children on an arithmetic learning activity in a classroom with a SDG, Alcoholado et al. [2012] showed different kinds of feedbacks meant for single students, the teacher and the whole classroom as shown in Figure 3.6b. In the right-hand side column the progress of each student during the activity was displayed for the whole classroom, and was intended to induce competitive behavior among students. Further, each student got feedback about the correctness of the solution upon completing each exercise, in their own workspace. Also, lack of activity during the exercise was shown by changing the background color of the workspace to red; this enabled the teacher to intervene and investigate about the problem being faced by the student. Caballero et al. [2014] also incorporated the functionality of providing each group with a feedback of their activity based on how the group members were interacting with the exercise (geometrical properties of triangles) as different smiling faces as shown in Figure 3.6c. The decision about the smiling-face feedback took into account if all the group members were participating in the activity, and the conceptual distance between the current state of the solution and the desired solution.

These examples demonstrate that groups' interactions with a SDG can be interpreted in the context of the task, and appropriate feedback can be provided to the groups for effective regulation of the group processes. The process of providing feedback completes the process loop of interaction analysis and has been observed to positively influence the group processes as suggested by Jermann [2004]. In our research work, we aim to emphasize on the use of SDG as an analytics and collaboration assessment tool by augmenting selective collaborative processes, as we will describe in later chapters.

3.5 Summary

This chapter regarded Distributed Cognition [Hutchins, 1995] as a theoretical framework, and a basis for comprehending various ongoing collaborative processes during conventional meetings. The primary aim of this thesis concerns the thorough examination of group members' interaction with artifacts that are often created, manipulated, and shared during meetings.

Moreover, considering the whole collaboration eco-system as a single cognitive system can help us examine these interactions with artifacts microscopically, and more importantly to identify their inter-dependence on the interaction between the group members (conversations, gestures, etc.). In this regard, distributed cognition theory along with cognitive ethnography (as described by Hollan et al. [2000]) enabled us to identify the theoretical design requirements for a suitable meeting technology, which is capable of supporting interaction with content and artifacts in a non-intrusive way. We recognized that Single Display Groupware (SDG) [Stewart et al., 1999] is the most suitable meeting technology that enables us to support group's interaction with shared artifacts in collocated settings.

Single display groupware belongs to the class of groupware designed to support co-present collaboration, and provides the group with a functionality of a single shared workspace and concurrent interaction via multiple input channels. In this chapter, I reviewed the previous research work regarding the usage of SDG in different collaborative scenarios. SDG has been observed to be specifically beneficial for learning scenarios within classrooms and collaboration among children. The influence of SDG in meetings and organizational collaboration is not very well investigated, the emphasis was so far given towards the design and development of technologies facilitating sharing and exchange of information and media (refer to Johanson et al. [2002b], Izadi et al. [2003], and Wigdor et al. [2009]). The lack of studies in the domain of meetings can be attributed to various factors such as the diverse group dynamics, interpersonal relationship between collaborators, the degree of openness ingrained in the nature of desired solution, and the complex nature of the collaboration assessment process. The research work presented in this thesis approaches the problem of collaboration assessment from the perspective of analyses of group members' interaction with the artifacts and shared content. In the next chapters, we will present and discuss the research questions that we wish to answer through this thesis.

4 General Research Questions

During conversations, what motivates an individual to pick up a pen and start sketching the idea on a napkin or a whiteboard? From a cognitive science perspective this need to externalize information on a persistent medium can be determined by the relevance of information being communicated, and the expected need for this information to ground future conversations via references to this particular piece of information. Besides, as human cognition is predominantly visual in nature, the permanency of information on an external media such as a paper or a whiteboard might enable collaborators to effectively communicate and understand the context; which might not be so evident due to the ephemeral nature of verbal utterances and gestures.

This question of expressing an idea by sharing it on a persistent medium has been studied in parts by different researchers, as summarized in the following list.

- Snyder [2013, 2014] studied the motivations behind collaborators to create a visual artifact (drawings) during conversations, from a socio-linguistic perspective. In her qualitative study, she identified the episodes of communications that result into creation of drawings in face-to-face interactions, and the role of this drawing activity in dialogue management and coordination.

- The role of persistent mediums for sharing information such as whiteboards, as well as their impact on the mutual understanding amongst collaborators during tasks was examined by Dillenbourg and Traum [2006]. They compared the role of a whiteboard (persistent medium) and a chat system (ephemeral medium) in group communications and the attainment of different levels of common ground.

- Scaife and Rogers [1996] have approached this question in a much broader sense by examining the cognitive effects of varied external representations (static diagrams, animations, and virtual reality), and the nature of the relationship between internal cognitive structures and their external counterparts.

At their essence, these research works have stressed on the coupling between the artifacts as an information container, and the human-artifact interactions. Snyder et al. [2014] summarized this relationship between the artifact and the process of visual representation as: "*Process often involves a tight coupling of observation, interpretation, and communication. The visual artifact is simultaneously analytic tool, interpretive framework, and presentation vehicle. This coupling can give the impression of being closer to the origin of ideas or phenomena, especially when documenting activities and interactions that are ephemeral or fleeting*". Combined together, the process and the artifact might also represent crucial episodes in collaboration which have not been thoroughly analyzed in previous research works, and therefore it forms the crux of our research work. Also, a detailed examination of these episodes might highlight some insights into the ongoing collaborative process.

In the next sections, I will state the general research questions that we will attempt to answer by means of this thesis work. The research questions presented in this chapter are broad in nature and scope, which will be refined for each specific user-study, presented in the upcoming chapters.

4.1 Role of different Input Modalities

Single Display Groupware (SDG) (as described in Chapter 3) equips each group member with a separate input channel (input device), and the capability to interact with the shared workspace simultaneously. Input devices of different kinds (mouse & keyboard, digital pen, multi-touch tables, etc.) have been utilized with the SDG. However, it is still unclear which kind of input device effectively and naturally supports the collaboration scenarios that are relevant for this research work. For example, a pen can be considered as an intrinsic means for quick expression of an idea by drawing or sketching. On the contrary, a mouse can be effective to spatially arrange or order multiple artifacts on a canvas to facilitate an interpretation. The choice of an effective input device is also crucial for the design of our meeting technology, which can facilitate group members to share and interact with shared artifacts simultaneously. Therefore, in Chapter 5 we will analyze the role of different input modalities in the creation of varied kinds of shared representations, and their influence on the group behavior.

The research questions concerning the role played by different input modalities can be formulated as following:

Question 1

What are the effects of using different input devices on the collaboration with a SDG? These effects might comprise the differences in task-completion times, performance, subjective perceptions about the task and other group processes, and so on.

Question 2

How does the interactional dynamics of a group vary with the change of the input device? By *interactional dynamics* we refer to the group members' interaction with the shared content, as well as the usage pattern of different input devices. Different interactional dynamics might itself be influenced by the affordances of the input device, and might also contribute towards the way information is represented while using different input devices.

4.2 Social Information in Interactions with the Shared Workspace

In face-to-face interactions between individuals, the interaction is not limited to conversations. Often individuals externalize information in order to express themselves better, or to offload their working memory on a physical medium for later reference. Gesturing and gaze are often used as supplementary signals to ground the information being communicated, as well as to refer to objects and artifacts. We believe that in the presence of a SDG (a shared workspace) these interactional processes might be coupled with the group members' interaction with the shared artifacts. In other words, the interactions with the shared workspace might serve as a proxy for the ongoing face-to-face interactions. An analyses of these interactions with the workspace might enable us to understand the ongoing collaboration processes. In addition, assessment of collaboration is regarded as a tedious and time-consuming process that requires extensive coding of dialogues, gestures, and gaze from the video data. A thorough understanding of group's interactions with the shared workspace might highlight some crucial relationships with various aspects of collaboration, which can be modeled or predicted in real-time. We also believe that the varied actions that group members perform over the shared artifacts while creating or modifying them might contain some social information that is representative of the ongoing collaborative processes. This social information can also be assumed to be similar in nature to the actual conversations and gesturing between the collaborators especially while using a SDG because group members can interact with each others' artifacts over a shared workspace.

Therefore, in Chapters 6, 7, and 9 we will focus on the extraction, interpretation, and analyses of social information within the group's interactions with the shared workspace. The research questions that we will attempt to answer regarding this social information can be formulated as follows:

Question 3

What are the relevant episodes in group members' interaction with the shared workspace? And what is the nature of social information that is embedded in these interactions?

Question 4

How does this social information vary across different collaborative settings with different kinds of tasks, presence of roles, familiarity between collaborators, etc.?

Question 5

How does the social information relate to different aspects of collaboration such as communication, coordination, task division, etc.?

Question 6

How can one utilize this social information in ecologically valid situations to assess, predict, or provide feedback to groups about the state of their interactions and the possible outcome?

4.3 Conclusion

In this chapter, we presented the motivations behind our pursuit to investigate the role of group members' interactions with the shared content and artifacts. In addition, we formulated the research questions that we will attempt to answer in the rest of the thesis. We will take an experimental approach to study these research questions. We will begin with the design and development of a meeting technology and various features to support collaborative content creation and sharing. Besides supporting the group members to create and share artifacts, we aim to use the meeting technology to assist us in collecting fine-grained behavioral data for later analyses. We will also conduct user-studies to explore and investigate the role of social information, and its relationship with different aspects of collaboration.

5 Complementarity of Input Devices

During meetings, the creation of artifacts is a distributed activity where group members have their own copy of notes, and they are not aware of their peers' actions and thought process. Moreover, to communicate an idea to the whole group or to express an opinion, group members share these artifacts among each other. Sharing can be achieved either by the transfer of artifacts to others, or explicit acquisition of shared resources like whiteboard or a flip-chart. Physical transfer of artifacts can be inconvenient because of the reduced visibility of the shared information to only a subset of group members who are in close proximity to the artifact. Further, the problem with the acquisition of shared resources is that it might lead to floor monopolization where a dominant member takes the floor and then does not yield the floor to others. In addition, there are certain limitations of sharing information on physical medium such as whiteboards because the content shared over it is not permanent, cannot be replicated and is hard to modify. Also, the product of collaboration in form of knowledge on the whiteboard cannot be saved for later references or carried away by collaborators after the meeting is over [Stefik et al., 1987].

As we discussed in Chapter 3 - Section 3.2, a Single Display Groupware (SDG) [Stewart et al., 1999] might mitigate the problems and issues faced while creation and sharing of artifacts during collocated meetings. The shared workspace supplied by a SDG can prove to be an effective medium for group members to share and manipulate artifacts (drawings, notes, etc.), and the concurrent input channels (one for each user) might make sharing more probable and assist in the avoidance of situations where a group member claims ownership over the shared resources. Furthermore, we believe that the affordances of an input device being used in the SDG play a crucial role in the choice of activity that can be best supported by the device. For example, a keyboard is a convenient way of inputting text to a computer, whereas a pen is more effective for spontaneous expression of ideas by drawing schemas and sketches. Therefore, it is crucial to leverage these factors as well in the design of the SDG that we plan to develop to support the auxiliary activity of content creation and sharing.

In this chapter, we will present the design and development of the first iteration of *MeetHub* - a SDG meant to support content creation and sharing in small group collocated meetings.

We will start with the discussion of design rationale that we employed in our design process followed by the description of the development process and the various characteristics of the system. We also present a user study comparing different input configurations, and its influence on the collaboration and the outcome of the task.

5.1 MeetHub 1.0

MeetHub 1.0 is the first iteration of our attempt to design and develop a SDG that facilitates creation and sharing of content during meetings, as well as enable us to register the interactions of each group member with the shared content. At the beginning of the development process, we were confronted by the question: "*Which input device to use?*". We hypothesized that the affordances of the different input devices (pen, mouse, keyboard) might influence the way the users interact with the system, the nature of the artifacts being created, and the task outcome. Therefore, we designed MeetHub with two separate input configurations - Mouse & Keyboard and Pen & Paper.

Next, we will discuss in detail the various design decisions that led us to the development of MeetHub. Later, a description of the technical specifications and the characteristics of MeetHub will be provided.

5.1.1 Design Rationale

Unlike single-user software, groupware have found it difficult to become a mainstream product during the last three decades of research in CSCW. In addition, the significant contribution in the field of groupware is observed in distributed collaboration scenarios. In the context of collocated meetings, participants frequently use widespread conventional meeting resources such as whiteboards, paper charts, pen and markers. Laptops or tablets are predominantly used as a private resource for information seeking or note-taking activities. Even in the case of presentations meant for a group, laptops serve as a personal storage of ideas and information being presented.

Grudin [1989, 1994b] elaborated the many reasons for the low prevalence of groupware in organizations as well as laid out the various challenges faced by groupware designers. Grudin suggested that few reasons for the failure of groupware are the additional amount of work required to make it usable in proportion to a small collective benefit, disruption of traditional work practices, and support for infrequent group processes. Therefore, in the design of our groupware, we decided to provide groups with a minimalistic interface, and without any additional provision of connecting computers and entering login information. In addition, the choice of the input device was considered carefully by conducting a user study as presented later in this chapter.

Shared Workspace

The orientation of the shared workspace over the public display is an important design decision to make. Both horizontal and vertical displays have their advantages and limitations as studied by Mandryk et al. [2002] with regard to the influence of privacy and user's proximity to the display. Mandryk et al. [2002] suggested that based on their affordances, horizontal displays are perceived to be more natural and comfortable for collocated collaboration, while on the contrary vertical displays provide a better perspective of the shared content. The shared content on horizontal workspaces such as the ones provided by multi-touch tables has reduced visibility among the group members primarily due to the orientation of users around the table, irrespective of the shape of the table [Shen et al., 2004]. In addition, the space over multi-touch tables is meant for the purpose of interaction and cannot be used to lay physical objects such as documents [Hinrichs et al., 2007], and the users sitting around multi-touch tables can occlude the displayed information by their hands and physical objects [Shen et al., 2009]. Further, the prolonged usage of tabletops has also been observed by several studies to result into fatigue as referred by [Wallace et al., 2009].

In our design of MeetHub, we decided to go ahead with the vertical display as it offers the same perspective of the shared information; thus keeping the table aside for the purpose for storing and holding of documents and input devices.

Mouse & Keyboard Setup

Mouse and keyboard as input devices for each meeting participant was our preferred choice during the initial development phase of our prototype. Mice offer a higher degree of pointing accuracy and keyboards are regarded as a faster, convenient, and robust means of entering text as compared to touch-based or on-screen keyboards [Hartmann et al., 2009]. Besides, computer literate users can be assumed to be quite familiar and comfortable interacting with these devices. Further, Hansen and Hourcade [2010] conducted a user study to compare a multi-touch table with a multi-mouse SDG during a collaborative task focusing on coordination. Their findings suggested that groups were more efficient while using the multi-mouse SDG in terms of task completion times, and fewer manipulations were performed to complete the task. However, the multi-touch table was the preferred condition amongst the participants, and groups using the tables were observed to be more communicative with twice more words per minute as compared to multi-mouse condition. In addition, the users reported on performing big hand gestures to interact with the surface of the table, whereas this was much more convenient in the multi-mouse SDG.

As part of the iterative development process of our prototype mouse & keyboard SDG, we tested it in several meetings and brainstormings (mostly concerned with the progress of ongoing research projects) in our lab. We observed that mouse was a very efficient device to locate, move, edit and delete objects (mostly text, images, and drawings) on the shared workspace. However, the mouse was regarded as tiresome and inefficient when a user wanted to quickly express an idea by means of drawing a sketch or a schema, because the users had to switch between different tools to achieve this goal (such as switching from line drawing tool to ellipse). We chose to provide a very minimal set of tools to the users that were required to create and share content; still mouse was not found to be easily compatible with the real-time dynamics of the meeting. In addition, the users reported (in informal qualitative feedbacks after meeting sessions) on an increased extraneous cognitive load while performing the drawing activity with a mouse, and this made it hard for them to focus on the ongoing collaboration simultaneously. Therefore, we decided as well to implement support for pen-based interaction in MeetHub as discussed next.

Pen & Paper Setup

Pen is considered to be more efficient than a mouse during the creation of drawings and schemas [Wolf et al., 1992]. In addition, Wolf et al. [1992] observed that pen based interfaces are unobtrusive and do not interfere with the interpersonal communication among the group members during collaboration. This might suggest that an obtrusive input device (like the use of a mouse to produce drawings, during our observations) increases the extraneous cognitive load of the user resulting in the attentional shift from the collaborative task to the device usage. Also, Terrenghi et al. [2007] studied the affordances of physical and digital artifacts (over interactive surfaces), and indicated that the tangibility of an artifact and its surrounding environment allows for easier manipulation of the artifact. Therefore, the tangibility of pen and paper as compared to digital artifacts on multi-touch table might render it more natural for users to create artifacts during meetings. Besides, providing each group member with a pen might prevent collaborators from competing each other to acquire the floor by the explicit acquisition of the pen. Another motivation for us to rely on the design of tangible paper-based interface was the success behind a few projects that were designed in our lab. For example, Do-Lenh et al. [2009] and Bonnard et al. [2012] designed various paper-based learning activities for concept-map task and primary school geometry learning methods respectively.

Despite the presence of varied systems supporting meetings, pen and paper remain popular in meetings, probably because they are faster and effective means of expressing an idea or making a point. Consequently, perennial research has been done to leverage the pen based interaction in meeting technologies. Few examples are *We-Met* [Wolf et al., 1992], *NoteLook* [Chiu et al., 1999], and *NotePals* [Davis et al., 1999]. We-Met [Wolf et al., 1992] was a prototype system designed to support communication and information retrieval needs of collaborators engaged in small group meetings (both collocated and distributed). The system worked with stylus and LCD-digitizing tablets connected over a network together with a facility of a telephone conference call. We-Met provided each meeting participant with a shared workspace to draw using a stylus. The drawing and sketching action was synchronized among all the group members, thus enabling group members to see the real-time actions of their peers. Next, NotePals [Davis et al., 1999] was designed as a note-sharing system among group members that used a shared note repository. NoteLook [Chiu et al., 1999] facilitated the production of meeting minutes in collocated meetings, where participants can directly annotate and write over the images captured during meetings. These meeting images corresponded to the crucial meeting episodes and collaboration activity, and were automatically extracted through the use of multiple video cameras installed in the room and the use of real-time meeting segmentation algorithms. However, both NoteLook and NotePals used stylus based interaction over PDAs, which restricted the interaction due to the small screen sizes of the PDAs. Also, Haller et al. [2010] leveraged the pen-based interaction in their design of NiCE Discussion Room, where meeting participants can use digital pen to interact with whiteboards as well as paper documents (for more details refer to "Meeting Rooms and Roomwares" in Section 2.2.2 in Chapter 2).

5.1.2 Technical Setup

Our design of MeetHub 1.0 incorporated a front projected display of 2m×2m size that served as the shared workspace. In addition, we designed a meeting table capable of accommodating up to 5 participants situated in front of the public display as shown in Figure 5.1. The meeting table was designed to ergonomically facilitate face-to-face communication among the group members, as well as interaction with the shared workspace as shown in Figure 5.2. Unlike rectangular- and circular-shaped tables, the undulations at the edge of our table design allowed meeting participants with appropriate degree of freedom to shift their focus from the speaker to the display and vice versa, without leading to fatigue in the neck region due to prolonged usage of the system. In order to compare the influence of various input devices on collaboration with a SDG, we designed two input configurations in MeetHub: in the first

Figure 5.1 – The Mouse & Keyboard setup of MeetHub 1.0: Each group member was provided with a wireless mouse and a keyboard that were color coded. The users can interact with various tools and artifacts over the shared workspace (the vertical wall-mounted display) via their mouse pointer in the color of the device color code.

Figure 5.2 – The Pen & Paper setup of MeetHub 1.0: Each group member was provided with a stack of A4 size sheets and a digital pen. The shared workspace was displayed similar to the mouse & keyboard setup.

setup each user was equipped with a wireless mouse and a wireless keyboard (MK setup) as shown in Figure 5.1; while in the other configuration each participant was provided with a digital pen and paper (PP setup) as shown in Figure 5.2. We color coded each input device so that the meeting participants can identify their corresponding mouse cursors (in MK setup), and their content created by them over the shared workspace (in PP setup, and only graphical content in MK setup). The users were not shown any cursors in the PP setup, because the users were meant to interact directly with the paper. Finally, the input devices and the public display were connected to a single computer, similar to any single display groupware.

5.1.3 Implementation Details

In this section, I will briefly discuss the implementation details of both the mouse & keyboard (MK) and the pen & paper (PP) setup of MeetHub 1.0.

Mouse & Keyboard Setup

Readily available and updated software frameworks that allow the handling of multiple input streams, such as Microsoft's Multipoint Mouse SDK[1] only allow the developers to handle events from multiple mice only. However, in order to develop the MK setup of MeetHub, we required the effective handling of events from both multiple mice and keyboards. The *SDGToolkit*[2], which was developed by Tse and Greenberg [2004], enabled us to handle these events after small modifications in the framework itself. The prototyping was done in C#

[1] MultiPoint Mouse SDK: http://www.microsoft.com/multipoint/mouse-sdk (visited on 24-February-2015)

[2] SDGToolkit: http://grouplab.cpsc.ucalgary.ca/cookbook/index.php/Toolkits/SDGToolkit (visited 05-March-2015)

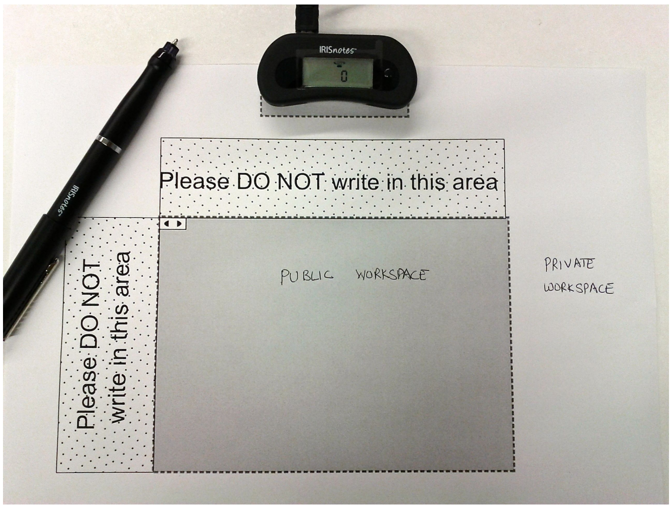

Figure 5.3 – The digital pen and the tracker used in the PP setup of MeetHub 1.0. The tracker device in the top-center of the figure tracks the position of the pen within the gray rectangular area on the A4 sheet (used in the landscape orientation). The arrows in the top-left corner of the gray rectangle allow users to browse between the pages on the shared workspace, simply by tapping their pen tip in these boxes.

programming language and the Windows Forms GUI library that is supplied with the Microsoft .Net Framework[3].

The MK setup of MeetHub allowed for the simultaneous handling of eight connected mice and keyboards. Each mouse was paired together with a keyboard to represent a user. Also, each input device was color coded and corresponding to each mouse, a cursor was displayed in the same color as the device color code. In order to keep the design minimalistic, we decided not to display the user's name or picture next to the mouse cursor as this might have required users to enter their login information. In addition, we intended to use MeetHub with small groups (maximum 5 participants) where the information from the peripheral awareness about the person(s) interacting with the shared workspace is readily available in collocated settings.

Pen & Paper Setup

The pen-based configuration of MeetHub used four position-based digital pens[4] (one for each user) and stack of A4 sized sheets as shown in Figure 5.2. All these four pens were connected to a single computer system, that was connected to the shared workspace over the public display. The position of the pen-tip on paper while writing was tracked by a tracker device (as shown in the top-center of Figure 5.3) at the rate of 60 Hz and is relayed in real-time to the computer. The computer regards the events from a pen similar to a mouse; i.e. the idle movement of pen over paper is treated as mouse move events, and when in use the pen strokes

[3]Microsoft .NET Framework: http://www.microsoft.com/net (visited on 05-March-2015)
[4]IRISNotes: http://www.irisnotes.de/ (visited on 06-March-2015)

on paper are considered similar to mouse move event with the left-button down. Similar to the MK setup of MeetHub, each pen was also color coded, which enabled users to identify their content over the shared workspace. The tracker device was also used to clamp together multiple A4-sized sheets with a gray rectangular area that mapped to the shared workspace (as shown in Figure 5.3). When a user produced a pen stroke within this gray area, it was concurrently reflected over the shared workspace in the color of the device color code.

For the purpose of development we used the Multi-Pointer X (MPX) server in the Groupware Windowing System (GWWS) developed by Hutterer and Thomas [2007]. MPX extends the existing behavior of X11[5] by providing an independent pointer to each mouse device, where all these cursors act as true system cursors. Further, a specific keyboard device can also be paired with a particular mouse cursor, thus equipping the windowing system with a true SDG functionality. Furthermore, every event generated by each device also contains a *Device ID* field, assisting in the easier identification of the actual device which generated the event.

Our SDG application incorporated an architecture comprising of two layers. The top-most layer in the application is a transparent X11 window overlay that registered the multiple pens connected to the computer. The second layer underneath was a Java Swing window that behaved as an actual shared workspace, which was responsible for interpreting the incoming events from multiple pens and rendering the pen strokes over the window. In order to communicate the input event information from the X11 window (transparent overlay that registered the multiple pens), we used Java Native Interface (JNI) framework to implement the callback functionality of relaying the event information from the X11 window to the Java application. Upon receiving the event information, the Java window processes the event to identify the pen that generated it and plots the stroke in appropriate color corresponding to the device color code.

5.1.4 System Features and Functionalities

In this section, I will describe the various tools provided to the users to create and share content, as well as other features of MeetHub 1.0.

Shared Workspace over the Public Display

Similar to conventional paper-based flip-charts, we used a page metaphor for the shared workspace; i.e. once the meeting participants have used up the available space on one sheet of paper, they can flip to the next page, while the content on the previous page can still be accessed by flipping the previous page back. Unlike the shared workspace in NiCE Discussion Room [Haller et al., 2010] that enables a span-and-zoom metaphor for content on the digital whiteboard, we have chosen the page metaphor as the pages are arranged sequentially and allow for quick browsing. In addition, we believe that the sequential arrangement of pages

[5]X11: http://www.x.org/wiki/ (visited on 06-March-2015)

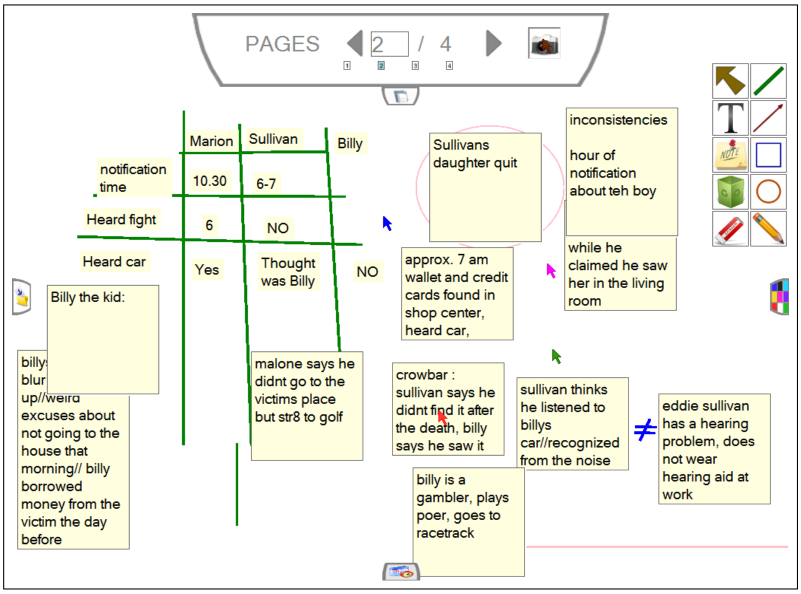

(a) Mouse & Keyboard Setup

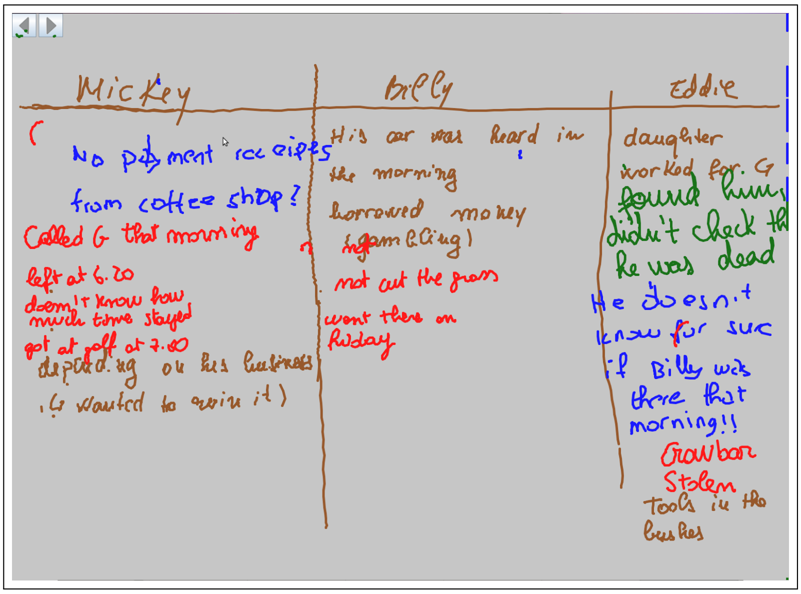

(b) Pen & Paper Setup

Figure 5.4 – Screenshots of the interface in the two configurations of MeetHub 1.0

affords for the storage of temporal dynamics of the content shared during a meeting session; i.e. every succeeding page would contain artifacts that were shared in the increasing progression of time during the collaborative activity. Therefore, there would be a temporal ordering to the shared content, and it won't be spatially cluttered, like on a big canvas. Figure 5.4 shows an example of how the shared workspaces looked like in both the MK (mouse & keyboard, Figure 5.4a) and PP (pen & paper, Figure 5.4b) setup of MeetHub.

In the mouse & keyboard setup of MeetHub, the shared workspace allows multiple users to concurrently add and manipulate artifacts such as graphics, text and images. In case of shortage of space on the current page, users can create a new blank page while concurrently saving the previous work. The groups are also provided with a *Page Manager* widget (as shown in the top-center of Figure 5.4a), which was designed to enable the group to browse through previously created pages as well as to create new ones.

In the pen setup, the users are provided with normal A4 sized sheets in landscape mode (see Figure 5.3). Upon each sheet is marked the area that is mapped to the entire public display (the gray rectangle in Figure 5.3 maps to the public display). Anything written within this area is simultaneously reflected on the public display in the color of the device color code as shown in Figure 5.4b. Changing of page on the shared workspace is accomplished by tapping the pen over one of the two squares printed on the top left corner of the marked area as shown in Figure 5.4b, indicating the direction of the page change. The users can change the A4 sheets simply by removing the topmost sheet and continue working on the one below. However, there was no cause-and-effect relationship in the changing of page in physical world and its counterpart in the digital world over the shared workspace, which could lead to confusion among the users. Therefore, we explicitly informed the users before each meeting session about the page change process.

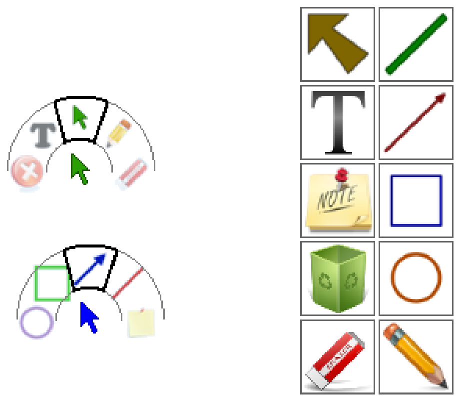

Figure 5.5 – The two versions of tool-collections provided to the users in the mouse & keyboard setup of MeetHub 1.0. The two semi-circular tool sets on the left-hand side are available to each user's mouse cursor and can be invoked by right-click on each mouse, and the users can browse through tools using the mouse wheel. The supplied set of tools contain shapes (line, arrow, rectangle, ellipse, and free-hand sketch tool), text (sticky-note and normal text boxes), eraser, and garbage bin (to delete the whole artifact). The same set of tools were also made available to the group publicly as shown in the right-hand side of the figure.

Tool Collection

Abundance of tools consume a lot of screen space as menu items, buttons, etc., and involve higher learning efforts and time on part of the users. This might also lead to an increase in extraneous cognitive load and would lead to less cognitive capacity for the execution of the task. Therefore, we provided users only with the minimal functionality that is necessary to support the task at hand.

In the mouse & keyboard setup, the basic set of tools consist of the ability to draw shapes (such as lines, arrows, ellipses, rectangles, and free-hand drawings), text (sticky-notes and text-boxes with transparent backgrounds), and images. The users were equipped with two identical tool collections containing the above mentioned functionalities as shown in Figure 5.5. One of the toolsets is personal to each group member, and can be invoked next to each mouse cursor by right-clicking on the respective mouse device. This toolset is shown next to the mouse cursor as a semi-circular context menu. The users can browse through the set of tools by using the mouse wheel and the selection of a tool can be made after pressing the left-button on the mouse. The second tool collection is public to the whole group and is displayed in the right hand side of the shared workspace as shown next to the right-hand boundary of shared workspace in Figure 5.4a. In addition, to these basic set of artifacts, we also added an annotation layer over the sticky-notes and images, so that users can also add annotations over these artifacts. The artifacts can also be moved by dragging them from one place to another over the shared workspace, as well as stacked by placing one artifact over another. Furthermore, any user can edit and annotate the artifacts created by others

Unlike the mouse & keyboard setup, in the pen & paper configuration no additional tools were provided to the users. We envisioned that the natural behavior of the users with the pen would be to write and draw directly over paper, and using a pen to select different tools would be troublesome. In addition, providing tools such as in the MK setup might result into disparity in the version of content created over paper and the shared workspace, which might result into usability issues with this configuration. The pen was intended to be used in the natural way as an individual would use it on a notebook or a flip-chart.

5.2 Research Hypotheses and Questions

The affordances of an input device render it appropriate for specific kinds of interaction, and this might contribute effectively to only a small set of task activities where the input device proves more useful. For example, a pen is more efficient while drawing schemas or quick expression of ideas through sketching, and mouse is designed for more precise referencing and manipulating tasks. Similarly, a keyboard is designed specifically for fast text entry. In addition, we believe that the use of an appropriate input device during a task might also affect various attributes of the task such as the task completion time, performance and participation by group members. On the other hand, during the execution of the task, the kind of input device used might also influence the type of information (facts, schematic representations, etc.) that is shared by the group. The user study presented in this chapter compares the collaborative task performed by groups in two conditions determined by the input configuration of MeetHub; i.e. *mouse & keyboard* (**MK**) condition and *pen & paper* (**PP**) condition. Based on the initial observations of the system usage during the iterative development process, and our beliefs we formulated the following hypotheses.

Hypothesis 1
: As group members can quickly express an idea by drawing or sketching it on paper, whereas sharing the same information while using a mouse and a keyboard requires users to select and switch between various set of tools, we expect that groups will complete the task faster in PP condition as compared to the MK condition.

Hypothesis 2
: Based on the affordances of input devices, as pen is a faster means of creating schemas and figures and it requires more effort to create the same schema using a mouse, we hypothesize that groups in PP condition will create more graphical elements as compared to MK condition (**Hypothesis 2a**). Similarly, a keyboard is a faster, convenient, and robust medium of entering text; therefore we suppose that more textual elements will be created in MK condition as compared to PP condition (**Hypothesis 2b**).

Hypothesis 3
: Considering the simultaneous interaction ability provided by MeetHub; in PP configuration as there is a lack of visual feedback of other's content creation activity on the paper, we expect that there will be more cases of content overlaps (or coordination breakdowns). This absence of visual feedback in PP condition can be attributed to the fact that individuals look directly at paper while writing, therefore they will not be aware of the spatial location of other's content to avoid overlaps. However, in the MK setup as users directly look at the shared workspace while interacting, there will be less overlap of content (due to location coordination) and coordination during the activity might be easier.

Hypothesis 4
: The use of keyboard in MK condition might influence group members to generate content in the *sentential form*. On the other hand, the use of pen in PP condition might facilitate the creation of *diagrammatic form*. Further, the diagrammatic forms are identified to be more effective in expressing and interpreting information as compared to sentential forms of representation [Scaife and Rogers, 1996]. In addition, it requires more cognitive capacity to process the sentential forms. Therefore, we hypothesize that groups in the PP condition might be better than the groups in MK condition in terms of task outcome, because the groups in the MK condition may be more prone to mistakes due to increased cognitive load.

5.3 Study I: Comparing Between Input Modalities

In order to verify our hypotheses, we designed the following exploratory user-study.

5.3.1 Participants

Sixty-six participants (23 females and 43 males) aged 17-35 (average age of 23) years were recruited for the study from our university. The participants were students studying various engineering domains. We formed groups of 4 participants each, except for two cases where the experiment was performed with a group of 3 participants as one participant failed to show up at the scheduled time. Both these groups with 3 participants belonged to the PP condition. In total, we finished the study with 17 groups.

The participants were randomly assigned to the groups across the two conditions. Further, there were 11 mixed-gender groups, 2 groups were all female groups, and 4 groups were all male groups. Regarding the familiarity among the group members, there were 7 (out of 17) groups where at least 2 participants were acquainted with each other. Also, before the start of each experiment session, we explicitly asked each group member to informally introduce themselves to increase familiarity among participants. Further, 2 experiment participants (out of 66) reported themselves to be colorblind, one for red-green pair and the other one for blue colors. Participants were paid 25 Swiss Francs (equivalent to 25 US Dollars) for their participation in the study, after the completion of each session.

5.3.2 Experimental Task

During the experiment session, each group belonging to either MK or PP condition was asked to complete the murder mystery task designed by Stasser and Stewart [1992], where each group was required to use the assigned input configuration of MeetHub (defined by the experimental condition) to create and share artifacts. Stasser and Stewart [1992] initially designed the task for a group of three participants. This task was later adapted for a group of four participants by Bachour et al. [2010]. We used the adapted version for the groups of 4 participants, and the original version of the task by Stasser and Stewart [1992] was used for the groups of three. Each group member was supplied with investigation logs that were designed for this task, and contained the transcripts of interviews conducted by the investigators with the suspects and victim's relatives, as well as maps and a snippet of a news article.

The murder mystery task is a hidden profile task where some crucial pieces of information about the suspects are not shared equally among the group members. Some important excerpts in the investigation logs that were made available to an individual were not available to others in the group. The correct identification of the suspect most likely to have committed the crime, requires the group members to combine the information by sharing and discussing as much as possible. Failing to share and discuss during hidden profile task might lead to inferior solutions and premature solution convergence as suggested by Mennecke [1997]. Furthermore, Mennecke [1997] also identified that bringing explicit awareness to the group members about the nature of the unshared information in hidden profile tasks leads to increased sharing of information by the group members during the task. Therefore, we informed each group about the existence of hidden pieces of information (but not the location) in the experiment booklets (investigation logs, maps, news snippet, etc.) at the start of each experiment session.

This task was chosen because it resembles the collaborative decision making strategy common to meetings where the group has some shared knowledge about the task, as well as some unshared information in the form of an individual's knowledge about a specific subject. The choice of the input device used to share information might influence the way groups create shared representations of the murder mystery task. For example, group members in the MK condition might decide to share facts about important events and later link them. On the other hand groups in the PP condition might come up with the timeline of the murder.

5.3.3 Procedure

A list of steps that were followed during each experiment session are presented as a list below.

1. Before the start of each session, the participants were asked to choose their preferred seat around the table.
2. Following this, they were asked to introduce themselves to their peers in the study.
3. Subsequently, the group members were asked to complete a pre-experiment questionnaire (refer to Appendix A.1) recording basic demographic data, familiarity with computers, and knowledge of collaborative tools.
4. In addition, the group members were also asked to mark their seating position corresponding to the public display and were requested not to change their seat during the course of the experiment.
5. Next, the group members were introduced to the appropriate configuration of MeetHub defined by the experimental condition (mouse & keyboard or pen & paper).
6. The participants were given some time to try out the system and familiarize themselves with various tools and features of MeetHub.
7. After the familiarization phase with the system, the experiment booklets (including investigation logs, maps, news snippet) were distributed to the group members.
8. The group was asked to complete the task by suggesting the name of the most likely suspect to have committed the crime. The task was not time bound, but an upper limit of two hours was set, and the groups were asked to finish the task within this time limit.
9. After the completion of the task, the group members were asked to complete a post-experiment questionnaire (refer to Appendix A.2) recording their perception of the groupware usage and the collaboration.
10. Finally, after disclosing the true identity of the murderer and communicating the actual motives behind the user study, all the subjects were given the compensation and were thanked for their participation.

5.3.4 Design

In our user study, we used a two input configuration (mouse & keyboard (MK) versus pen & paper (PP)), between-subjects design. Groups of 4 participants completed the experimental task in either the MK condition or the PP condition. The average time to complete the task by all the groups was 90 minutes. Finally out of 17 groups, 9 groups completed the experiment in MK condition and 8 groups in PP condition.

5.3.5 Data Collection

Before as well as after the experiment session, each participant was asked to complete two questionnaires (refer to Appendix A). The pre-experiment questionnaire collected basic demographic information as well as working preferences (if the individual prefers to work alone or in groups), weekly frequency and average duration of meetings attended, and the nature of meetings attended. Further, experiment participants were also asked if they suffered from colorblindness and of what kind, because the content over the shared workspace as well as different mouse cursors were represented in different colors corresponding to the device color code.

Next, the post-experiment questionnaire recorded participant's perception of the task, the emerging social structure during collaborating, the usage and usability aspects, and the purpose for which MeetHub was mainly used. Amongst the information that concerned the collaborative task, we recorded the participant's perception about the group's consensus on the attained solution, sufficiency of time to finish the task, and difference in opinion among the group members. Further, the statements related to the emerging social structure regarded the presence of a group leader, perception of homogeneity in group member's participation towards the task, and the sense of acceptance of an individual's contribution by the group. The part of the questionnaire that concerned with the usage of MeetHub, recorded the user's perception about the effort required to create artifacts (such as text, drawing, etc.), and the effort required to coordinate. Finally, the last part of the questionnaire recorded how the user's perceived the purpose and utility of the system while completing the task. For example, if the system was primarily used to collect facts, make hypotheses about the subjects, or to recreate the story line. We also provided participants with additional space for them to leave some open ended comments and suggestions. This questionnaire comprised of statements meant to gather information about various aspects of collaboration and tool, and the participants were asked to specify their agreement with these statements. The agreement with statements was recorded based on 5-point Likert scale, where 1 was regarded as strong disagreement and 5 as strong agreement.

Besides the questionnaire data, two observers were present during the experiment session, and recorded the number of deictic gestures and utterances for each participant in the group. An utterance was considered to be a complete spoken phrase without hesitation and a delay by a participant. Further, the utterances were either classified to be related to the task and the experimental material, or to the organization of activities within the group. Similarly, the deictic gestures were categorized so as to refer towards the task material (experiment booklets given to the participants), or to the content that was shared over the public display. Finally, all the interactions with the shared workspace were also recorded in system log files. However, the experiment sessions were not audio- or video-recorded.

5.4 Results and Analyses

	Experimental Condition			
	Mouse & Keyboard		Pen & Paper	
Dependent Variable	Mean	SD	Mean	SD
Task Completion Time (in minutes)	105.78	16.29	78.23	14.64
Number of words created	65.69	79.35	36.30	33.18
Number of drawing elements	7.92	12.18	11.16	13.98
Number of pages created	1.89	1.01	2.67	1.47
Number of utterances	65.0	30.89	57	25.85
Number of gestures towards the shared workspace	6.46	6.06	4.57	5.89
Number of gestures towards the experiment material	16.97	14.97	16.13	14.31

Table 5.1 – The mean and standard deviation values of dependent variables across the experimental condition (Mouse & Keyboard versus Pen & Paper).

5.4.1 Effects of Input Devices

Table 5.1 summarizes the mean and standard deviation values of different dependent variables across the conditions. An analysis of variance (ANOVA) performed with task completion time as the dependent variable and the experimental condition (MK versus PP) as the independent variable validates our first hypothesis (**Hypothesis 1**, see Section 5.2). The groups in the PP condition completed the task in significantly shorter time as compared to the MK condition ($F(1,15)=11.70$, $p<.01$). One possible explanation for this finding could be that the users were more efficient while using the pen and paper and this might have lead to a faster completion of the task.

Next, we will look at the differences in shared content across conditions. In order to do so, we categorize the content based on its nature; i.e. graphical and textual. The textual content is quantified as the number of words generated by each group member during the experiment. Similarly, the number of graphical elements correspond to the number of drawing strokes (line, ellipse, etc.). A repeated-measure ANOVA performed with the participating group as the grouping factor reveals that there was a significant difference in the amount of text generated across the two conditions. Participants in the MK condition generated more textual elements as compared to the PP condition (F (1,64)=4.08, p=.04). On the other hand, more graphical elements were shared in the PP condition than the MK condition. However, this difference was not statistically significant. Therefore, our **Hypothesis 2** (see Section 5.2) is partially validated where **Hypothesis 2b** is statistically validated, but we don't have enough statistical evidence to validate **Hypothesis 2a**. In addition, no statistical difference was observed in the number of pages that were created by groups across the two condition.

Considering the post-questionnaire data about participants' perception of the effort required to create and manipulate artifacts (both graphical and textual in nature), we observed that it required significantly higher effort to generate text in the PP condition (F(1,63)=4.30, p=.04). On the contrary, creating drawing elements was perceived as significantly easier in the PP condition (F(1,62)=5.71, p=.02). This finding complements our previous result that more text was shared in the MK condition. The perception based findings suggest that keyboards were considered as an efficient way of inputting text and pens are a quicker means of drawing. Thus, indicating that the affordances of an input device might influence the task depending on the differences in the nature of the content. We will analyze these difference in detail in Section 5.4.4.

5.4.2 Perceived Effort to Coordinate

Further, regarding the perceived effort required to coordinate content sharing activity over the shared workspace by avoiding content overlaps; the participants reported that it was easier to coordinate in the PP condition. However, this difference was not observed to be statistically significant, and this finding is also contrary to our **Hypothesis 3** (see Section 5.2). One possible explanation could be the lack of transparency of the sticky-notes widget in the MK setup of MeetHub. Also, the sticky-notes adjusted their size automatically depending on the quantity of text contained; as more text was added the widget grew in size and possibly occluded the artifact under it. This phenomenon was also observed by Zanella and Greenberg [2001], and recommended on the use of transparency as a means of avoiding interference and occlusion in SDG. Besides occlusion, we observed a change in participant's behavior in the PP condition while creating artifacts. Instead of looking at the paper, the users were looking directly at the shared workspace (over the public display) while writing on paper. Upon asking a few participants after the session about this change in behavior in an informal interview, the participants reported that it was solely to coordinate better and to avoid content overlaps. Therefore we observed that interaction with paper without visual feedback of other's activity in a SDG enforced a different and more indirect form of interaction, and emphasized the need for visual awareness on the paper itself.

We also conducted a factor analysis with the post-experiment questionnaire data, that revealed crucial aspects related to the coordination and the groupware usage. We found that the the effort required by the users to coordinate defined the purpose of the system. In other words, if group members perceived a higher effort to coordinate, the shared workspace was primarily used as a group memory to collect facts and information that can be referred later. On the other hand, the easier it was perceived to coordinate, the the content over the shared workspace contained conceptualized information such as a timeline that was used to create a narrative.

5.4.3 Task Outcome

Table 5.2 summarizes the differences in various dependent variables across the various instances of task outcome. The groups that correctly identified and suggested the name of the suspect who committed the crime in the murder mystery task, were considered successful groups. Now, studying the difference in success (or failure) of groups across the two conditions, we observed that there was no significant difference (χ^2=0.04, p>.05, Kruskal-Wallis Test). Therefore, our **Hypothesis 4** (see Section 5.2) is not statistically validated. Stasser and Stewart [1992] suggested that success in hidden-profile tasks is influenced by numerous factors such as information the group is aware of, effective sharing of unshared information, motivation of the group members, and the importance of information perceived by the collaborators. Different input techniques can only facilitate the process of effective sharing, while they might not have any influence on other aspects suggested by Stasser and Stewart [1992].

Tang [1991] and Mennecke [1997] also indicated that enabling group members to share information simultaneously reduces the chances of *production blocking* defined as the lack of opportunities for group members to share their ideas. Production blocking is more likely to happen in scenarios where collaborators have to compete for resources in order to share artifacts. For example, a whiteboard, where the group as a whole cannot collaborate simultaneously, and some group members might have to wait for their turn to grab the pen and explicitly acquire the floor. The SDG used in our experimental setup handles with this issue

	Task Outcome			
	Success		Failure	
Dependent Variable	Mean	SD	Mean	SD
Perception about similarity in opinions of group members (Likert Scale)	1.86	1.22	3.48	1.31
Number of utterances	65.94	30.45	56.19	26.29
Number of gestures towards the shared workspace	3.26	4.01	8.13	6.82
Number of gestures towards the investigation logs	13.17	12.67	20.45	15.78
Number of pages created	2.68	1.45	1.74	0.85
Number of textual elements	55.88	78.17	48.32	43.88
Number of drawing elements	10.94	15.12	17.64	10.14

Table 5.2 – The mean and standard deviation values of dependent variables across various instances of possible task outcome.

by providing each group member with their own input device. However, there can also be a downside associated with this parallelism in sharing artifacts, which might impede groups from discovering the true hidden profiles as the individuals might fail to assimilate the information shared by their peers; i.e. *consumption blocking*. This situation might present itself when the group focuses more on simultaneous sharing of information, and not stressing on the phases of pausing and processing the already shared information.

On the contrary, we observed some behavioral differences between successful and unsuccessful groups. The successful groups produced more utterances, but this difference was not statistically significant. In addition, successful groups made significantly fewer gestures towards the shared workspace (χ^2=12.86, p<.001, Kruskal-Wallis Test), as well as towards the experiment material (F(1,64)=4.32, p=.04). Further, the successful groups created significantly more pages as compared to the unsuccessful groups (χ^2=8.51, p<.01, Kruskal-Wallis Test). Also, successful groups shared more text and drawings than the unsuccessful groups, but these differences were not found to be statistically significant.

Finally, the perceived difference in opinions among the group members (recorded as 5-point Likert scale value in the post-experiment questionnaire) was also observed to be significantly related to the outcome of the task. The group members in successful groups differed in their opinions more than the unsuccessful groups (F(1,64)=27.29, p<.001). This finding might suggest that the difference in opinions among the group members might have resulted into dialogue, and while considering the different perspectives of collaborators, the group was more likely to discover the hidden profile and identify the correct suspect. This explanation is supported by our previous result indicating that successful groups had more utterances than the unsuccessful groups; however this difference was not found to be statistically significant.

5.4.4 Shared Knowledge Representation

The successful completion of the task requires groups to unify initially shared (common to all the group members) as well as initially unshared (specific to each user) information. Unifying these information demands simultaneous sharing of facts about the suspects, connecting and associating the shared facts to create a narrative. During the user study, we observed that similar kinds of shared representations were produced across the two conditions (MK and PP), which can be briefly classified into three categories: *facts* (Figures 5.6a & 5.6d), *timelines* (Figures 5.6b & 5.6e), and *syntheses* (Figures 5.6c & 5.6f).

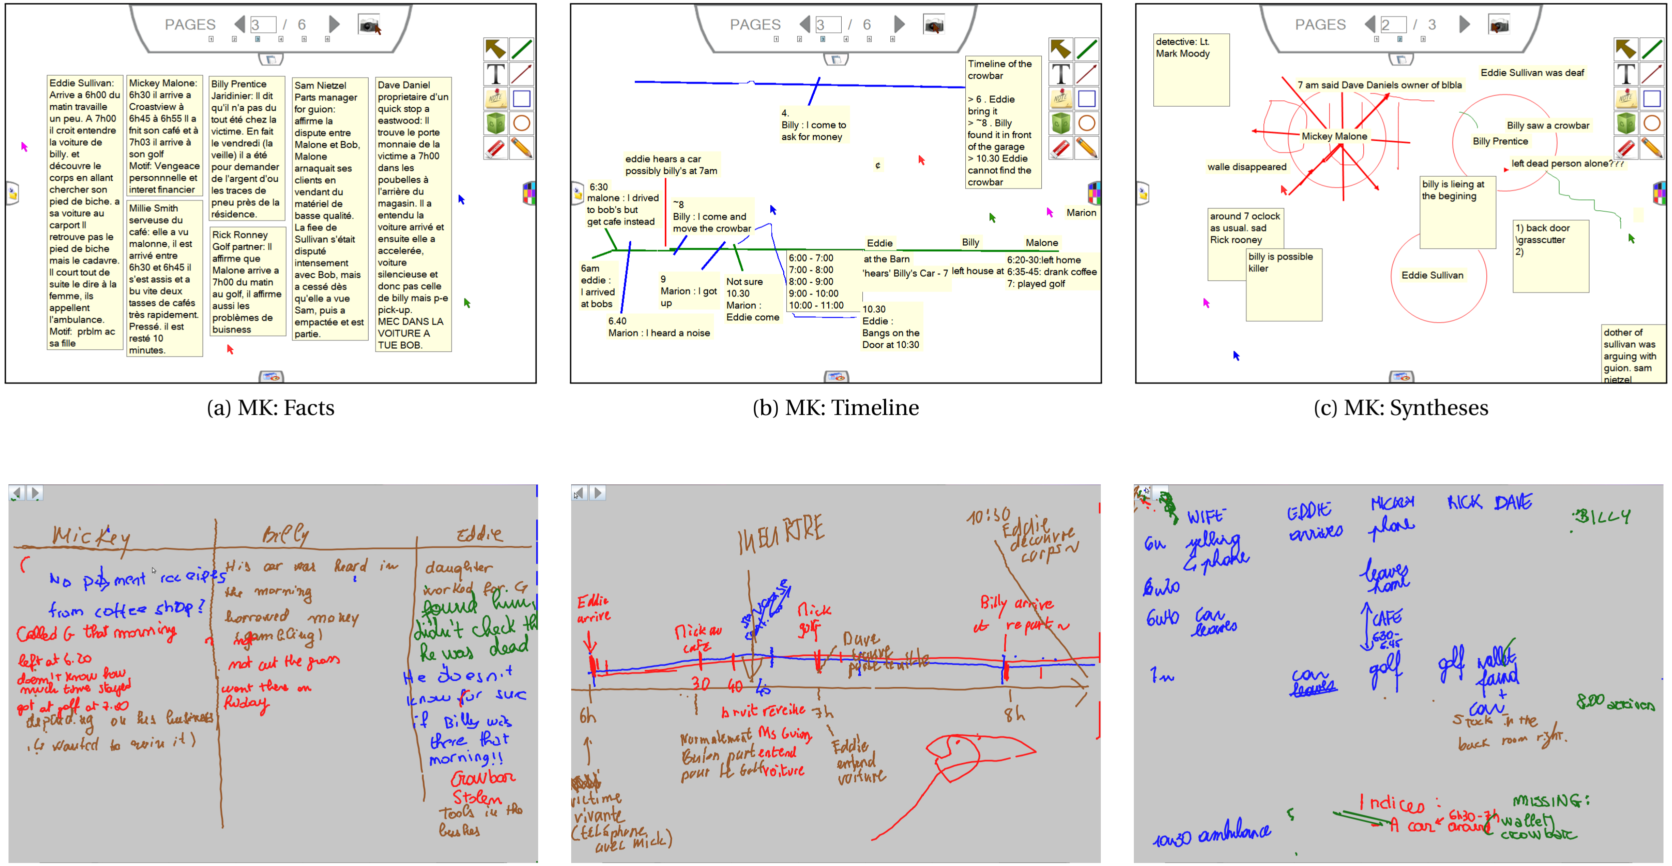

(a) MK: Facts

(b) MK: Timeline

(c) MK: Syntheses

(d) PP: Facts

(e) PP: Timeline

(f) PP: Syntheses

Figure 5.6 – The classification of content over the shared workspace into three classes: *facts*, *timelines*, and *syntheses*.

Facts were collected by the group members during the starting phase of the session while reading the investigation logs. The collaborators tended to partition the space over the workspace, where each partition corresponded to a suspect or another character as shown in Figures 5.6a and 5.6d. In the MK setup, a single sticky-note denoted the separation of one suspect from another, and over time facts about suspects were added to the appropriate sticky-note by the group members (see Figure 5.6a). However in one team in the PP setup, a group member took the initiative to manually partition the space as shown in Figure 5.6d. In some cases, the group members skipped the fact collection phase and went directly to the diagrammatic representation such as a storyline or a grid. In order to do so, the group members explicitly showed each other their investigation logs to visually spot the differences in the text, and therefore reducing their effort to share a lot of facts and sharing only crucial pieces of information.

Timelines were made towards the end of the reading phase when the group members started to create a narrative from the collected facts. The timelines in both the conditions looked visually similar as shown in Figures 5.6b and 5.6e. In the MK condition, groups created a timeline by spatially situating the textual episodes at appropriate position along the horizontal line, followed by linking these facts with lines to signify orderly occurrence of various events (see Figure 5.6b). Whereas, in the PP condition group members wrote directly at the exact spot (see Figure 5.6e). However, we observed that only one (out of nine) groups in MK condition created a timeline. On the other hand, many timelines were created by groups in the PP condition. This observation again reveals the difficulty faced by groups in the MK condition to create timelines, as it required users to switch between tools several times while drawing, dragging text from one position to another, and editing textual objects.

Syntheses were created during the convergent phase of the discussion, where users were considering the key episodes of the story and debating about the most likely suspect as shown in Figures 5.6c and 5.6f. In other words, this kind of representation served as a buffer to temporarily store relevant information needed to reach a decision. In the MK setup, the syntheses resembled a screen with randomly distributed textual facts that were later assigned to different clusters or concepts as the discussions progressed (see Figure 5.6c). The clusters were either explicitly created in the form of a circle to contain related facts, or were symbolically represented by placing sticky-notes over one another with slight overlap. Similarly, in the PP condition, participants browsed for relevant facts from the previous pages and created a hybrid representation of crucial episodes and links (see Figure 5.6f).

This categorization of shared representations during the task demonstrates a transition between individual sharing activity and collaborative information processing; similar to the one observed by Tang et al. [2009] in case of whiteboards. The group members start collecting facts and figures at the beginning, and by the time they reach the convergent phase of the meeting they have built a more globally shared knowledge about the task. This crucial group dynamics is probably supported and governed by the transitions between these shared representations.

Further, Scaife and Rogers [1996] conducted a review on external cognition in order to establish an empirically backed relationship between graphical representations and internal mental models. In their review, Scaife and Rogers [1996] discussed the study comparing the diagrammatic form of representation with sentential form by Larkin and Simon [1987]. The findings from this study suggested that diagrammatic form of representations greatly reduce the effort to search and recognize information, and also requires less inferencing as the “computational offloading” is considerably less as compared to sentential forms. In our study, we also observed an evolution from more sentential form of shared representations towards more diagrammatic forms during the convergent phase of the task. However, as more of the diagrammatic forms were produced in the PP condition as compared to the MK condition, we can assume that the pen facilitated the creation of diagrammatic forms as compared to mouse, and this might have lead the groups in PP condition to finish the task faster. These observations partially conform with our **Hypothesis 2** (see Section 5.2) as there is a difference in amount of sentential forms and diagrammatic forms across the two conditions, but both these representational forms exist in both the conditions. One possible explanation for this could be the nature of the experimental task, which might also determine the group’s dependence on one representational form than another.

5.4.5 Differences in Group Composition

Familiarity among the group members might affect the collaborative processes across the two conditions, and introduce a bias in our study. Therefore, we performed a Chi-square test to study if previously-acquainted subjects were distributed differently across the two conditions. The results demonstrate that there is no significant difference in distribution (χ^2(1,N=66)=1.25, p=.26). We can conclude that familiarity had no effect on other group processes as the groups were equally distributed with respect to familiarity.

	Experimental Condition	
Gender	Mouse & Keyboard	Pen & Paper
Females	8	15
Males	28	15

Table 5.3 – Distribution of recruited experiment participants based on gender across the experimental condition (Mouse & Keyboard versus Pen & Paper).

Next, the recruitment among EPFL students resulted in fewer female participants as compared to males as shown in Table 5.3, and this might introduce a gender bias and might also affect the dependent variables. Consequently, we conducted the Chi-square test, and the results revealed that there are significantly fewer than expected females assigned to the MK condition than to the PP condition (χ^2(1,N=66)=4.40, p=.03), which might appear as a bias. Therefore, we investigated if the amount of shared content produced (number of textual and graphical elements) and participation (utterances and gestures) varied with gender. Results of a mixed-effect ANOVA showed that there is no significant difference in the number of textual elements(F(1,64)=0.92, p=.34), and the number of drawings created (F(1,64)=0.04, p=.84) across gender. These results signify that the amount of shared content is not influenced by gender of the participant. Next, considering the number of gestures and utterances made by each individual, we found no significant gender difference in the gestures made towards the shared workspace (F(1,63)=0.08, p=.78), as well as towards the investigation logs (F(1,64)=0.46, p=.5). Similarly, the number of utterances made by the group members were not observed to be statistically different across gender (F(1,64)=0.28, p=.59). Consequently, we can conclude that gender did not affect other content creation and sharing processes.

5.4.6 Transition in Groupware Usage and Leadership

In Section 5.4.4 we demonstrated an evolution in the shared representations produced by the participating groups, from a primarily factual content towards a well assimilated conceptual representations. Similar to a transition in type of shared representations, we observed a transition in the way groups used MeetHub by the examination of the concurrent usage for all the experiment groups; i.e. the number of participants simultaneously interacting with the groupware. We observed that the group members concurrently interacted with the shared workspace during the divergent phase of the task (while collecting facts). On the contrary, MeetHub was used in a single-user mode during the convergent phase (while discussing and making a decision) of the task, where one group member took up the responsibility to interact with the shared workspace while others chose to be part of the ongoing discussion. A similar transition from an individual (or parallel) activity mode to a collective mode of operation was observed within groups of 4 participants while interacting with a tabletop, by Ryall et al. [2004]. Our observation of this transition can either imply that roles (especially that of a leader) emerged within the groups, or indicates towards the problems related to the ergonomics of the system that might have caused coordination breakdowns.

Considering the emergence of role within groups, we only analyzed the role of a *group leader*. In the post-experiment questionnaire (refer to Appendix A.2), each participant was asked if they perceived an emergence of leadership in their group. After each experiment session, the two observers were asked to rate each group member of displaying the qualities of a group leader, on a 5-point Likert scale (where 1 indicated absence of leadership qualities and 5 signified strong leadership quality). Later, the group member with the highest average rating was designated as the group leader. However, the difference in opinions and perception of the two observers might influence the perception of a group leader, and this might be a limitation with this method. Moreover, the result of the regression analysis showed that the leaders interacted 12.67% more with the shared workspace based on the total number of actions performed over the workspace (including creation, editing, and moving of textual and drawing elements), than other group members (t(48)=2.63, p=.01). Besides interaction with the shared workspace, the group leaders also had significantly higher number of utterances (19.28 utterance) than others (t(47)=3.15, p<.01). Also, group members exhibited more gestures towards the shared workspace (3 gestures, t(47)=2.73, p<.01) and the experiment booklets (7.49 gestures, t(48)=4.06, p<.001). These results indicate that group members were comparatively more interactive (gesturing, speaking and interactions with the shared workspace) than other group members, and also took responsibility of interacting with the shared workspace during the convergent phase of the task.

Next, we investigated if the problems related to the ergonomics of the system contributed towards one participant taking control of the shared workspace. In order to do so, we analyzed the relationship between the seating position (or orientation with respect to the shared workspace) and emergence of leadership. The results of the Chi-square test demonstrated that there was no significant relationship between the group leader and the seating position around the table (χ^2(4,N=44)=5.14, p=.27).

These findings indicate towards a crucial group dynamics of task validation and coordination observed in decision making teams. The convergent phase of the task (information assimilation and decision making) requires relatively higher coordination among the group members, as well as multiple perspectives to the shared knowledge. As a result, one group member explicitly acquires the control of the input device and decides to interact with the shared workspace. While interacting with the shared workspace, this group member is also taking

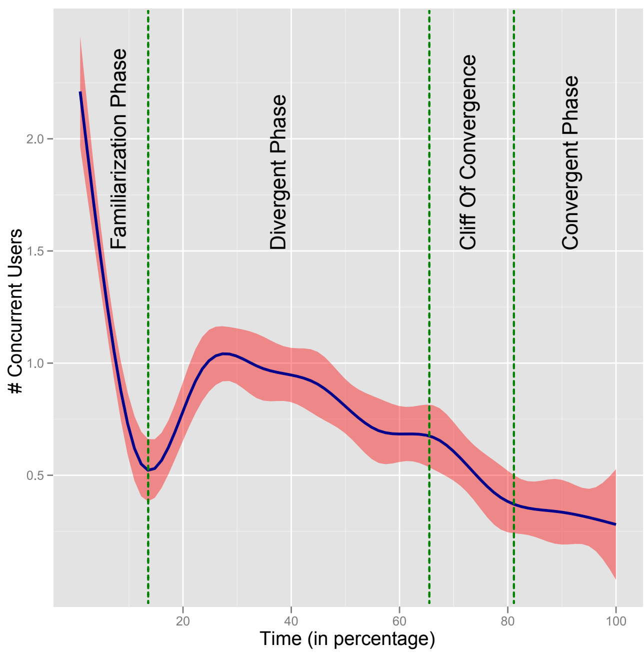

Figure 5.7 – A moving average plot with the confidence interval (shown as the red region) shows the number of concurrent users simultaneously using MeetHub, in all the participating groups (9 in MK condition, and 8 in PP condition. The X-axis represents the normalized time for all the groups, and the Y-axis represents the number of users within a group simultaneously using the system.

into account the consensus of others about search, annotation, and organization of shared knowledge so as to facilitate the decision making process. These observed phases in our study are consistent with the ones indicated by Arrow et al. [2000] and Tschan [2002].

This transition of groupware usage from multi-user mode to the single-user mode is visualized in Figure 5.7, where each groups transits into the single-user mode of operation by sliding down a *Cliff of Convergence*. The name of cliff of convergence denotes the group's transition from the divergent phase of the task (see at the middle of Figure 5.7) to the decision making or convergent phase (see at the right side of Figure 5.7). This phenomenon was observed for all the groups at roughly the same time. The region at the left side of Figure 5.7 represented the phase when the group members were familiarizing themselves with MeetHub before the start of the task. Most existing SDG systems like WeSpace [Wigdor et al., 2009] assist only in the divergent phases of tasks that primarily require sharing and creation of artifacts. However, this transition phenomenon emphasizes on the need to support the convergent phases of the task as well, by incorporating specific set of tools in the SDG.

5.5 Discussion

The user study presented in this chapter was intended to investigate the effects of different input devices on the content creation and sharing activity while using a SDG. During meetings, creation and sharing of artifacts is not the main activity, but rather an auxiliary or the supporting activity for the task at hand. During the task, group members might decide to create diverse representations of their shared knowledge depending on the task demands and the individual knowledge of the group members. Therefore considering that different tasks demand varied representations, it is impractical to restrict the interaction with the SDG by choosing only one kind of input mechanism. In other words, limiting group members with just one type of input device allows for only a narrow set of representations that can possibly be produced by the group. In that regard, our user study offered some important implications for the design of SDG.

The input devices in the two conditions did not influence the success or failure of groups in the hidden profile task, as this is influenced by other factors including the effective sharing of information, motivation of the group members, and the perceived importance of information [Stasser and Stewart, 1992]. However, the experiment results concerning the usage of various input devices clearly indicate that there exists a mapping between the input device affordances and the shared representation required by the task. Further, the availability of multiple kinds of input devices for each user might allow for an opportunity to select the input mechanism most appropriate to adequately support the task. In other words, a provision for an ecology of input devices capable of supporting various kinds of content representations might enable participants to choose the appropriate channel of interaction with the shared workspace. This choice may be governed by the mapping made by a user between the perceived affordance of the suitable input device and the representation demanded by the task.

We call this phenomenon *opportunistic sharing*, as the mapping between the input device and shared representation is made opportunistically by the user. However, it is still unclear if each group member should be provided with multiple input devices (for example, digital pen, mouse, and keyboard) of her own, or the distribution should be made based on user's preference over one input device than another. In the case where distribution of input devices is made on the basis of user's preference, the task of sharing a specific kind of content (such as drawing) can be delegated to the group member who has the adequate input device for creation of such content. However, the preference based distribution can have a downside if every group member prefers to choose the same input device.

An analysis of the shared artifacts created over the workspace by the groups were classified into three categories based on the time of creation, visual differences, attributes, and the immediate purpose served by the content. Our observations of the shared representations exhibit the variety in the shared representations produced by the groups. This classification also offers an insight into the role of the shared workspace (over the wall-mounted display), and its effects on the transitions in the collaborative coupling (similar to the one observed by Tang et al. [2006] around tabletops) of the group members. During the execution of the task, group members used the shared workspace as an external memory where the cognitive load of the group members was offloaded by sharing artifacts. At the later stage of the task, this information is processed, conceptualized, and comprehended by the group and acts more as the group knowledge which supports the decision making process.

The temporal evolution of shared representations also highlight the transition in group members' collaborative coupling from a loosely coupled group (individual task of sharing and surfacing facts) to a more tightly coupled group where the group members build a shared understanding. These observations about transitions are inline with the ones also observed by Tang et al. [2006] and Tang et al. [2009]. We observed that such a transition in collaborative coupling happened only once during the task, whereas one can expect such transitions to occur several times during the task. However, these transitions might depend on the kind of task, and certain tasks might have these transitions multiple times, once during each phase of the group activity.

The analysis of number of simultaneous users interacting with the shared workspace shows that MeetHub was used concurrently during the information collection phase, whereas during the decision making phase one group member was interacting with the shared workspace at any time. The transition in groupware usage from the divergent phase of the task to the convergent phase happens over a cliff of convergence as shown in Figure 5.7, after which only the emerged group leader interacts with the shared workspace, while others prefer to be part of the ongoing discussion and validation process. This observation also points out to the need for tools in SDG which can assist group members during the convergent phase of the task; for example, voting systems, text summarization, or agents enabling automatic organization of shared artifacts in formats that enable discussion about them. This is a crucial design guideline as most existing SDG are functionally positioned in the divergent phase of the group activity requiring creation and sharing of the content.

5.6 Limitations of the Study

Unlike a well controlled study, the nature of our user study was exploratory and was designed to examine the effects of different input devices on the task as well as the content sharing activity. In addition, we aimed to extract some valuable design guidelines for the next iteration of MeetHub. However, there were certain limitations with the system as well as with the study as I mention next.

In the pen & paper setup of MeetHub, the lack of visual feedback of others' activity was one of the limitation, as it altered the natural behavior of the group members while writing on paper. The group members focused their attention on the public display while writing on the paper in order to avoid overlaps. Next, the rating scheme used to identify the emergent group leader by the two observers was not an accurate way of identification, as the opinions of the two observers might differ on the definition and attributes of a leader. In addition, we did not ground these characteristics with the observers before the experiment. Also, the experiment sessions were not video recorded; thus we have no records of the ongoing conversations which might have contributed to the richness of the analyses.

Finally, in the user study, we did not explicitly test for the differences in the types of users and their behaviors across the two conditions. Also, participant's preferences for input devices were not taken into account for the analyses, as the keyboard & mouse and digital pen do not have the same prevalence in our sample population. This is because most of the subjects were not familiar with the digital pen, and consequently preference would not be informed by experience but rather but rather by novelty.

5.7 Conclusion

In this chapter, I presented MeetHub - the single-display groupware (SDG) that we designed to support the creation and sharing of artifacts in small-group collocated meetings. MeetHub is composed of a shared workspace over a wall-mounted display, and enables group members to interact with it concurrently while mitigating issues such as floor monopolization. The initial prototype of MeetHub was designed with two separate input configurations - Mouse & Keyboard (MK) and Pen & Paper (PP) along with an ergonomically designed table. The provision for the support of two different kinds of input mechanisms was done in order for us to understand the role played by different input devices in the creation and sharing of artifacts, as well as their effect on the group processes; therefore we conducted an exploratory user study.

Our findings demonstrate that the input devices do not play a significant role in the outcome of the task. However, the different input devices complement each other with regard to the shared representation that can be produced using them, where one kind of representation (graphical or textual) can be conveniently produced with one kind of input device. In addition, the evolution of shared representations during the task demonstrates the role of shared workspace as an external memory and how changes in the memory over time affect the shared understanding and knowledge held by the group. Further, the transition in the groupware usage highlights crucial aspects of group dynamics and how the need for better coordination and knowledge validation affects the groupware usage.

Our findings also provided us with some important design implications such as equipping group members with varied input devices so that they can opportunistically map the adequate input device to the representation demanded by the task. The lack of visual feedback in the PP condition was a limitation as it altered the natural behavior of users while interacting with a paper-based interface. Therefore, another crucial design implication is to provide visual feedback of others' activity with a pen-based SDG. We will incorporate these design lessons in the second iteration of MeetHub presented in the next chapter. In addition, the next chapter (Chapter 6) will focus on the identification and extraction of implicit social information from group members' interaction with artifacts over the shared workspace.

6 Latent Social Information

Other's Contribution and Mine

The user study presented in this chapter was a collaborative work with two more colleagues, and aimed to investigate different aspects of group behavior. Therefore, I will distinguish between my peers' contribution in contrast to my contribution. In this chapter, I will not go into details about my peers research, however I will briefly mention it in places where necessary.

My Contribution: I analyzed the group members' interaction with the artifacts that were created and shared over the shared workspace during the experiment sessions. In addition, I will extract the latent social information from these interactions and investigate their relationship with various aspects of ongoing collaborative behavior.

Other's Contribution: One of my peer was studying the effects of temporal awareness on task organization during meetings via a time management widget that was implemented as a component of MeetHub. Another peer analyzed the explicit collaborative search activity during meetings and compared it with automatically generated search recommendation produced after the analysis of shared text over the shared workspace.

The explicit visual representation of information on a shared workspace increases the accessibility of complex ideas that are not visible or immediately obvious in the ephemeral nature of conversations and gestures [Snyder et al., 2014]. In addition, according to the characteristics of distributed cognition presented by Hollan et al. [2000], the cognitive processes are not just distributed across the members of the group. There exists a strong coupling in between internal cognitive structures and interactions with the environment variables, and this process is also subjected to evolution with time. Consequently, we believe that ongoing interactions (conversations and gestures) between the group members can be effectively mapped on to their interactions with the artifacts over the shared workspace. In other words, group members' interaction with artifacts may act as a proxy for visible collaborative processes; and a thorough understanding of these interactions might assist us in gaining crucial insights into the ongoing collaboration without the need to code and transcribe video recordings of meetings, as well as without employing time-consuming methods to analyze collaborative processes.

In addition, from the perspective of social signal processing [Pentland, 2007], a specific kind of interaction with an artifact can be treated as a behavioral cue, which when interpreted in the context can provide us with social signals. These social signals can be later used to assess the ongoing collaboration or provide feedback to the group.

In this chapter, we will investigate the latent social information in group members' interactions with shared artifacts, which might hold relevance in our understanding of face-to-face interactions between group members. Further, the analysis of group members' interactions over time, with multiple artifacts in order to create, manipulate, and organized shared knowledge might contain valuable social information and patterns which are not as visible as the dialogues, gestures, and gaze of the meeting participants. Therefore we use the term *latent* to define the nature of this social information. In order to record and analyze these interactions, we also redesigned MeetHub (presented in Section 6.1) after extracting some valuable design guidelines and experiences from our first user study presented in Chapter 5.

6.1 MeetHub 2.0

Figure 6.1 – MeetHub 2.0: A group using the system during an experiment session. Each group member was supplied with a mouse and keyboard to interact with the artifacts over the wall-mounted display. In addition, each group member also had a tablet and stylus to create and manipulate content. The workspace was synchronized on all the tablets and the public display.

In our first user study that we presented in Chapter 5, we made a comparison between different input configurations of MeetHub 1.0 (Mouse & Keyboard versus Pen & Paper), as well as analyzed the role of these input devices on the content creation and sharing activity during collaboration. The findings suggested that the affordances of input devices render them complimentary for the content creation activity, and the choice of only one kind of input device for the users might adversely affect the groupware usage by limiting the varied kinds of representations that can be produced by the group [Verma et al., 2013]. Consequently, we redesigned MeetHub to allow users to interact with the shared workspace via either a mouse & keyboard (similar to the MK setup of MeetHub 1.0 presented in Section 5.1) or a stylus on a tablet. The choice of stylus based interaction on a tablet was made in order to provide the users with a visual feedback (which was absent in the PP setup of MeetHub, refer to Section 5.1) of other's activity over the shared workspace; thus facilitating better coordination in interactions with the shared content among the group members.

Next in this section, we will describe the technical specifications, implementation details, and various features of the redesigned version of MeetHub.

6.1.1 Technical Setup

Similar to the first prototype of MeetHub, MeetHub 2.0 also comprised of a front projected wall-mounted display of size 2m×2m, and a meeting table (same as the one in MeetHub 1.0) as shown in Figure 6.1. The shared workspace was displayed on the wall-mounted display. In addition, each group member was equipped with a mouse and a keyboard as well as a tablet[1] and a stylus[2] to enable interaction with the shared workspace. Each device was color coded to assist users in identifying their respective mouse cursor on the public display, as well as to distinguish one's content from others.

The workspaces on all the tablets were synchronized in real time with the one on the wall-mounted display via a server embedded in the MeetHub application; thus enabling group members to interact with the shared content either directly on the wall-mounted display (using mice and keyboards) or on their respective tablets (using stylus). The wall-mounted display as well as the input devices were connected to a single computer, that was also running a server to synchronize the content between all the tablets.

Further, as described by Stewart et al. [1999], a single display groupware consists of a single output channel and multiple input channels (refer to Section 3.2 in Chapter 3). However, in our setup of MeetHub 2.0 the presence of multiple tablets might characterize it as a multi-display groupware (MDG). But, as the workspace was synchronized across all the tablets, and the group members had the same view across all the displays because the users cannot navigate to a private working space; we can regard MeetHub 2.0 to possess the attributes of a single display groupware. The decision not to provide users with a private working space was made because this thesis aims to model the teamwork aspect of collaboration by means of analyzing

[1] Apple iPad 2nd Generation: http://www.apple.com/ipad/ (visited on 15-March-2015)
[2] Just Mobile AluPen^{TM}: http://www.just-mobile.eu/ (visited on 15-March-2015)

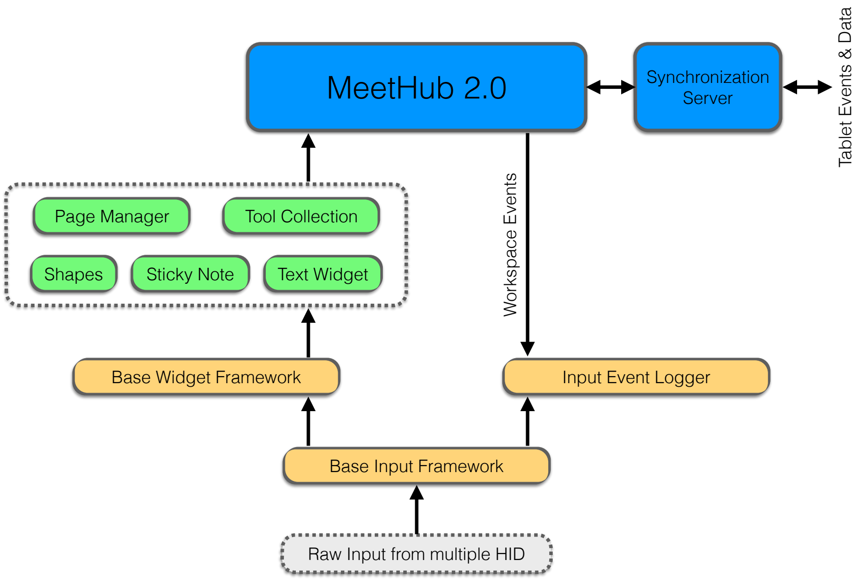

Figure 6.2 – *Base Input Framework* - An architecture of our software framework that was used for the implementation of MeetHub 2.0.

the interactions with a commonly shared workspace. Also, presence of a private workspaces for each group member might afford for a different work dynamics involving division-of-labor or task division based on individual's knowledge, expertise, or role; with less discussions and negotiations about other's viewpoints and ideas. Therefore, we chose to provide the groups with just one view of the workspace along with strongly-coupled navigation to avoid task division.

6.1.2 Implementation Details

Already discussed previously in Sections 3.2.1 and 5.1.3, developing a SDG application is a complicated and a tedious task due to the lack of support for handling events from multiple connected input devices. Many of the existing software frameworks such as SDGToolkit [Tse and Greenberg, 2004] are outdated, and others do not support the handling of variety of input devices (for example Microsoft MultiPointMouse SDK[3] only supports the handling of mice inputs). Furthermore, SDGToolkit employs an efficient software architecture that enables for the quick prototyping of the SDG, but as this API (Application Program Interface) relies on Windows Forms GUI library[4] it is not possible to develop widgets with transparent (or semi-transparent) background to prevent occlusion of content over the workspace. Therefore, we chose to develop our own software framework that can enable us to design SDG applications

[3]MultiPoint Mouse SDK: http://www.microsoft.com/multipoint/mouse-sdk (visited on 24-February-2015)
[4]WinForms GUI Library: http://en.wikipedia.org/wiki/Windows_Forms (visited on 16-March-2015)

and tools without a need to deal with the low-level complexity of handling events, pairing devices, logging events, and designing multi-input enabled controls and widgets (like buttons, text boxes, etc.).

Our software framework is called *Base Input Framework* (BIF) and was developed in C# programming language, with support for Windows Presentation Foundation[5] (WPF) - a graphical subsystem for developing user interfaces for Windows operating system. Figure 6.2 shows the architecture we used in the implementation of MeetHub 2.0. The handling of events from multiple mice devices is achieved via Microsoft MultiPoint Mouse SDK. Unlike the handling of input events from mice, the handling of input events from multiple keyboards is complicated because of the large number of keys present on the keyboards, and the variability of keyboard formats across different regions and languages. We used Microsoft's RawInput API[6] which is a generic framework to get access to any Human Interface Devices (HID) such as joysticks, microphone, and so on. The RawInput API enabled us to get access to the messages that are generated within the operating systems upon a key press in keyboards (or an equivalent input event). We parsed these event messages to gather information about which device generated the event (or which keyboard is used), and the information related to that event (i.e. which key was pressed). Next, we mapped the key press event with the appropriate language and region setting in the operating system so as to get the correct letter being typed.

We designed a class in BIF that encapsulates the handling of both mouse and keyboard events, as well as assigns a separate colored cursor to each mouse device, creates pairs of mouse and keyboard device and treats them as a single user. A simple instantiation of this class within any WPF window application makes it multi-input enabled and facilitates the quick prototyping of SDG applications on the Windows operating system. Besides multiple input device support for applications, we also designed controls (buttons, progress bars, etc.) and widgets (drawing canvas, sticky-notes, etc.) that are responsive to events from multiple input device. These widgets implemented the software interfaces that were developed within the BIF, and defined the appropriate behavior for different kinds of input events. In addition, we designed and developed a generic widget that can be used by developers to develop custom widgets with very specific behaviors such as a multi-input enabled document reader or file browser.

In MeetHub 2.0, the handling of input events from multiple mice and keyboards was thus achieved through the Base Input Framework (as shown in Figure 6.2). Besides events from multiple mice and keyboards, MeetHub 2.0 also enabled users to interact with the shared workspace via tablets and stylus. Therefore, we developed a similar client application for tablets in Objective-C programming language (used to develop applications for Apple devices), and we also implemented a server within MeetHub 2.0 that synchronized the content over all the tablets and the wall-mounted display. The tablet devices connected with the server via wireless network, and used event-oriented design to communicate all the interaction events over the tablet to the server (as soon as an event occurred), which later pushed this

[5]Windows Presentation Foundation: http://goo.gl/aj11f9 (visited on 16-March-2015)
[6]RawInput API: http://goo.gl/wLTwRM (visited on 16-March-2015)

interaction information to other tablets and also reflected the changes over the wall-mounted display. This enabled users to interact with the shared artifacts either over the wall-mounted workspace or over their respective tablets.

6.1.3 System Features and Functionalities

In this section, I will describe the various tools that were provided to the users to facilitate content creation and sharing, as well as other features of MeetHub 2.0.

Identification of Input Devices

Each mouse and keyboard was color coded, and corresponding to each mouse device a mouse cursor in the same color as the device color-code was displayed in the shared workspace over the wall-mounted display. These cursors were not displayed in the shared workspace over the tablets because the intended interaction medium was a stylus, and presence of multiple cursors might prove to be confusing and disruptive to the users.

In order to identify the tablet device corresponding to each user, the color code of the tablet device was displayed within the tablet application itself (for example, the device color of the tablet in Figure 6.3b is blue as shown in the top-right side). This color was chosen from among the color-codes of the mice and keyboard devices, and once assigned the tablet was considered to be paired with the mice and keyboard device. This restriction of pairing each mouse, keyboard and tablet was imposed for the purpose of our user study, so that we can analyze the difference in usage of different input devices during the task. However, in the later versions of MeetHub we removed this constraint so that the group members can choose the input device based on their preference or task demands, and without a need to provide all the devices to each group member.

Shared Workspace

Shared workspace as shown in Figure 6.3 allowed group members to collectively create and manipulate content. Similar to the first prototype of MeetHub, we considered a *page* as a functional unit of the workspace over which all the content can be created. This is analogous to the conventional flip-charts, with the only difference being that instead of providing the whole group with one pen, each user has her own input entry point. Further, in case of shortage of space on the current page, group members could create a new page while simultaneously saving the content on the current page. A page manager widget was made available on both the wall-mounted display (see the top-center of Figure 6.3a) as well as in the client application on all the tablets (see the top-left of Figure 6.3b), which enabled group members to browse through the previously saved content as well as create new pages. A change of page on any of the display was immediately reflected on all the other displays.

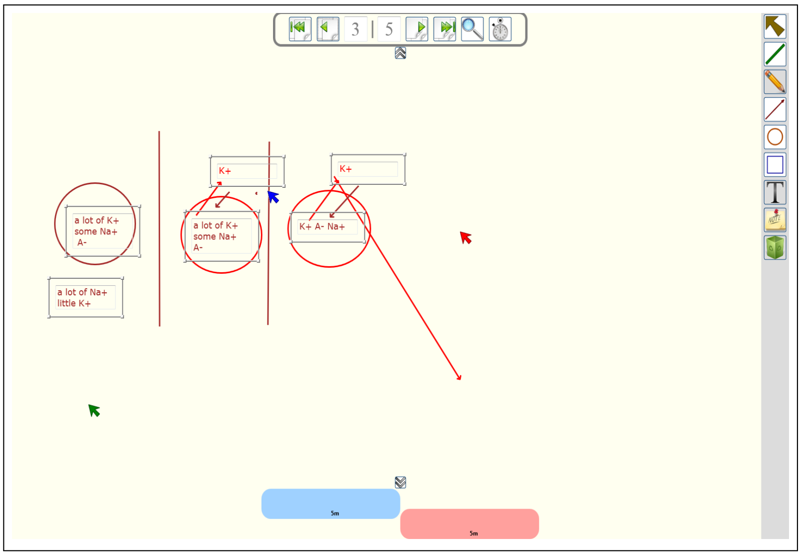

(a) Shared workspace on the wall-mounted display.

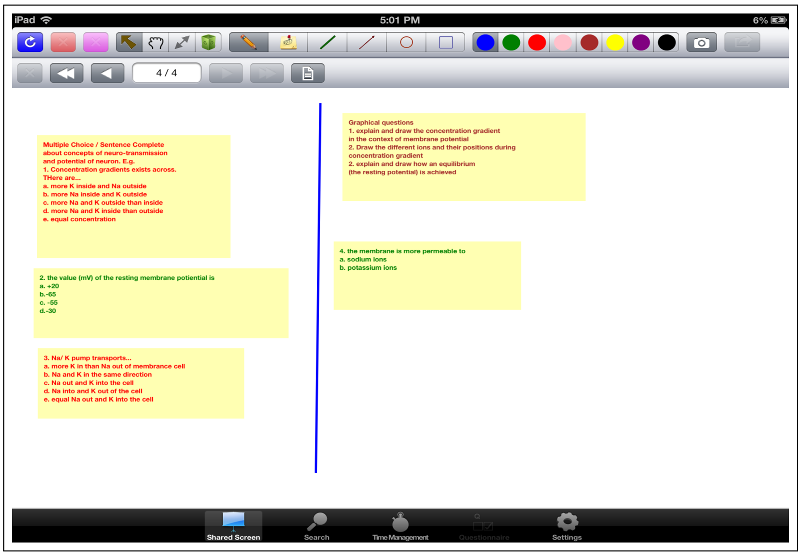

(b) Shared workspace on the tablet.

Figure 6.3 – The user interface of MeetHub 2.0 on the wall-mounted display as well as the tablets.

Changing of page can be regarded as a *global conflict* as defined by Morris et al. [2004] because of the strongly coupled navigation; i.e. if one group members changes a page, it might disrupt other users currently interacting with the shared workspace. In MeetHub, we did not implement any conflict resolution mechanism such as voting or ranking (also described by Morris et al. [2004]), instead we relied on social protocols between the group members to negotiate the change of a page. Therefore, before each experiment session the groups were reminded of this strongly coupled navigation, and were asked decide collectively when they wanted to change a page. Further, the conflicts related to the modification of any artifact (defined as *whole-element conflict* by Morris et al. [2004]) such as a sticky-note were handled by the system depending on the current input device that has an input focus on the artifact. In other words, if a group member explicitly wanted to edit an artifact, she would have to click on it so that she gains the input focus before manipulating that object.

Tool Collection

The tool collection provided in MeetHub 2.0 was similar to the one in the MK setup of the first prototype of MeetHub, as presented in "Tool Collection" in Section 5.1.4. This tool collection enabled group members to draw shapes such as lines, arrows, ellipses, rectangles. In addition, the users could also create free-hand drawing using a handwriting tool either by using their stylus on the tablets or mouse on the wall-mounted display. All these drawing elements were rendered as vector graphics enabling users to change the graphical properties of drawings such as position, size and rotation. Next, the users were supplied by two separate tools to create textual elements - sticky-notes and text boxes. The sticky-notes had a semi-transparent yellow background, and the text boxes had a completely transparent background. These text based objects automatically adjusted their sizes based on the amount of text contained within them.

The tool collection was displayed on the right-hand side of the shared workspace as shown in Figure 6.3a. On all the tablet devices, the same set of tools were displayed in the top-center part as shown in Figure 6.3b. Moreover, the content on the shared workspace was created in the color of the device color code, and the users cannot change the color of their content.

Search Tool & Search Recommendations

Besides the shared workspace that allowed group members to share content, a search tool was developed as a plugin to MeetHub by a fellow colleague, who aimed at investigating the role of explicit and opportunistic search (the search performed after viewing the recommended search keywords) in collaboration. A collaborative search tool was developed and integrated as part of the shared workspace on the wall-mounted display, which could be maximized and displayed when any user moved a mouse cursor in the bottom-left corner of the workspace. In order to minimize it, the user had to move their mouse to the bottom-right corner. The search tool was designed like a internet browser with a text field to type in the search queries. The users also had a capability to search for information on their respective tablets via a search tool that can be invoked by touching on the search button at the bottom of the shared workspace as shown in Figure 6.3b. The specific information searched over the tablets could also be brought to the group's attention by sharing the webpage from the tablet, which was immediately displayed over the wall-mounted display.

Besides explicitly searching for information by using the search tools, the search tool also analyzed the textual content shared by the group and provided the group member with search recommendations through an interactive visualization. However, the influence of search tools, recommendations, and collaborative search is the topic of analysis of a colleague, and therefore I will not analyze the aspects of search in this chapter.

Time Management Tool

A time management widget was also supplied to the group members, which provided groups with periodic temporal awareness of the elapsed time of a meeting as well as the visualization of the meeting progress in terms of the current phase of the meeting session. The widget was developed as a plugin of MeetHub by another colleague with an aim to study the influence of temporal awareness on meeting participants. Depending upon the task and the meeting goals, the group members could decide on the number of phases and their respective durations prior to the start of the meeting. Once the meeting started, the widget provided the group members with notifications over the wall-mounted display at critical times such as the completion of a single phase.

6.2 Research Hypotheses and Questions

Snyder [2012, 2013] identified that drawings exhibit a dual nature as an artifact and a communication activity during spontaneous face-to-face interactions. More specifically, Snyder regarded the activity of creating tangible artifacts spontaneously from the perspective of interactional sociolinguistics; i.e. the means by which collaborators create a shared and well-grounded meaning through the use of language, and the role played by artifacts in the establishment of this shared meaning. Snyder analyzed the communications within groups that lead to the creation of an artifact, and how group members engaged and coordinated during the drawing activity. The findings of her study suggested that the creation of artifacts is tightly bound to the ongoing conversations, and these two activities cannot be separated. In addition, the act of sharing information on a persistent medium such as a whiteboard or paper can be considered equivalent to “mode switching” where the persistent medium is used as an alternate supporting communication medium. This act of sharing artifacts also bears contextual signals about the potency of information under consideration, as well as the increased need for better grounding.

The decision to interact with a physical artifact (during creation or editing) at any point of time during collaboration holds significant relevance due to the ephemeral nature of face-to-face interactions, and marks these episodes as important because the information under consideration is key to the attainment of shared meaning (either by offering a different perspective to the problem or adding a key element to the final solution). Consequently, an examination of these episodes of collaborators’ interaction with artifacts can provide us with valuable contextual social information concerning the group dynamics (engagement, participation, roles, etc.), as well as the nature of task (decision making, brainstorming, etc.).

Unlike the collaboration scenario studied by Snyder [2012, 2013] where the collaborators focused their attention on a single informal physical artifact, meetings can encompass numerous kinds of artifacts such as meeting minutes, individual notes, information shared on the whiteboard, and slides. Moreover, in the presence of a SDG and the shared workspace, these numerous artifacts are commonly available and accessible to all the group members, and the interactions with these artifacts are complicated in nature, because similar to the dialogues between collaborators, group members can simultaneously contribute to the same part of the problem or on the same artifact. Many of these interactions with the shared workspace can be latent in nature (unlike the explicitness of dialogues and gestures, which are meant to acquire others' attention) as any group member can silently offload an idea over the shared workspace in a sticky-note, or spatially organize someone else's idea to bring it near similar ideas or visualize the problem from a different perspective.

Furthermore, Hollan et al. [2000] emphasized on the conjugation of individual cognitive processes and interactions with external representations and environment materials. Snyder [2012, 2013] also demonstrated this conjugation in her studies in the form of a duality in the ongoing conversations and the activity of creating artifacts. These observations highlight the inherent causality between the relationship between individual's internal cognitive processes, face-to-face interactions between the collaborators, and their interactions with external representations. We believe that exploration and establishment of the nature of this relationship, and leveraging it might equip CSCW researchers with a new microscope to look at collaboration.

Our motivation to analyze the group members' interactions with shared artifacts is backed by the lack of research emphasizing on this aspect of collaboration. Previous research in CSCW and collaboration assessment provided more emphasis on the individuals and the interactions between them, rather than the continuous interactions between individuals and artifacts, which principally lies within the domain of Distributed Cognition [Hutchins, 1995]. Examples of few previous researches that focused on the individuals's interactions and their role in collaboration assessment include the multi-dimensional coding scheme based on video analyses of conversations, gestures, etc. by Meier et al. [2007] and Kahrimanis et al. [2009]. Another approach called DISCOUNT by Pilkington [1999] focused on coding and analysis of dialogues in collaborative learning scenarios.

By means of this thesis we aim to emphasize the need for a paradigm shift from an individual oriented analysis of collaboration towards an artifact (or object) centered analysis. Modern groupware and SDG allow researchers to record and analyze the interactions with the shared workspaces in real-time, and this can consequently reduce the time required to assess collaboration; which is a time-consuming and a tedious process. In addition, the shared workspace of a SDG can store the traces of interactions and manipulations performed on all the shared artifacts. These histories of interactions can be utilized to construct a state-machine for each shared artifact in real-time, and enable us to examine and identify the crucial episodes of changes that resulted in the development of shared knowledge over the workspace. Now that we have established a well-grounded background for our motivations, we will formalize the research questions that we will examine later in this chapter.

Question 1

Collaborators perform different kinds of interactions while creating and editing artifacts (both graphical and textual in nature) over a shared workspace. The social information that is contained within these interactions (or actions) depends on *who* is performing an action, and what is the *nature* of this action. For example, an artifact which was originally created by Alice could be later modified by another collaborator Bob, or by Alice herself. The modifications by Alice and Bob provide information about different collaborative practices: the modification by Bob might indicate a dialogue between the two, whereas a modification by Alice might denote individual task responsibility. Also, the time when this action was performed plays a crucial role. Furthermore, moving an artifact over a shared workspace is different from editing the contents within an artifact such as a sticky-note.

Considering the different kinds of actions that might happen over artifacts, we can formalize our research question as: *What are the different characteristics of the social information that can be extracted from the varied actions on artifacts?*

Question 2

The varied kinds of tasks require group members to produce different shared representations. For example, the brainstorming task and a decision making task might differ in the kind of shared representations that allow for the attainment of the desired result. In addition, roles, expertise and inter-personal relationships might influence the collaborators' interactions over the shared workspace. For example, if Alice and Bob are not well acquainted with each other, then during collaboration, it is less probable that Alice will edit artifacts created by Bob freely. Also, in our previous study presented in Chapter 5, we observed that group leader performed more interactions with the shared workspace than others.

Regarding the possible effects of these aspects, we can ask our second question as: *How is the latent social information influenced by various collaboration aspects such as task type and presence of roles?*

Question 3

Finally, as an interaction with a shared artifact during verbal discussions can be supposed to be related to a need to establish a common ground. We can hypothesize that the latent social information might be related to the state of mutual understanding within collaborators, and this information can be used for modeling and prediction purposes.

Therefore, our third research question is: *How does latent social information relate to the state of mutual understanding between collaborators? How can we leverage this relationship?*

The research questions that we wish to examine and answer in this chapter are broad in nature and require an exploratory analyses to answer them precisely in the context of small-group collocated collaboration supported by a SDG - MeetHub. In addition, the analyses presented later in this chapter will follow a top-down methodology, where an action by a group member is further broken down in components and interpreted in the context of the task and the individual.

6.3 Study II: Latent Social Information in Interactions

In order to answer and examine the questions posed previously in Section 6.2, we designed and conducted a user study where we varied across different kinds of tasks and roles of group members.

6.3.1 Participants

Twenty-five participants (5 females, 20 males) aged 21-28 years (mean=23.6, SD=1.53) from our university, took part in the user study. The participants were all students belonging to computer science, communication systems, and management of technology domains. The participants were enrolled in a CSCW course (master level course) being taught at the university and their participation in the study was part of the course curriculum. Six groups of 4 to 5 participants (5 groups of 4 participants, 1 group of 5 participants) completed the experiments in a period of 4 weeks, where the first week was meant to familiarize the participants with MeetHub, and the rest of the three weeks corresponded to different kinds of tasks. There were 2 all male groups, and 4 mixed gender groups. Finally, as participating in the experiment was part of the CSCW course curriculum, the participants were not financially compensated for their participation in the study.

6.3.2 Experimental Tasks and Meeting Types

During the user study, each group was asked to complete three tasks corresponding to three different kinds of meetings: *brainstorming*, *decision making*, and *problem solving*. These meetings took place over the course of three weeks, where each group participated in a meeting session per week. The group members used the shared workspace in MeetHub 2.0 along with other tools like time-management and search tools during each week's meeting.

The following list describes the three tasks that were used during the study:

- **Brainstorming:** The first task (refer to Appendix B.1) was devised and used by Taylor et al. [1958], and required groups to brainstorm about the advantages and disadvantages if every human being born from the following year is born with an additional thumb in both the hands.

- **Decision Making:** In the second week of the user study, the groups participated in a decision making task to devise a sustainable solution to an impending energy crisis in a Chinese province. This task was designed by us (refer to Appendix B.2), and required groups to firstly estimate the power shortage in five year time based on the supplied statistics. Later, building upon the estimate, the groups were required to provide a plan to increase the energy production in the next 10 years to overcome this power shortage. The groups were asked to consider the geographical, logistical, environmental, and economical factors in mind while working on the plan for energy production.

- **Problem Solving:** Finally, the groups completed a problem solving task during the third and final week of the study. This task was designed by Sangin et al. [2011] (refer to Appendix B.3), and required group members to read and understand instructional text pertaining to the role of resting membrane potential in neurotransmission. After reading the instructional text, the group members were asked to conceptually visualize this phenomenon by illustrating the generation of resting membrane potential. Furthermore, the groups were asked to prepare a comprehensive assignment for students, assuming that meeting participants were a group of teaching assistants for the neuroscience course. In addition, as none of the subjects had a background in neuroscience (refer to Section 6.3.1 for the background of recruited participants), the task required them to understand and discuss within the group so as to have a consistent understanding of the topic.

The choice of these three tasks was made to address the variability in shared representations that can be produced, and differences in strategies and dynamics that can be adopted by groups. Such a variability might allow us to examine the latent social information in different task contexts, and might enable us to apply our findings to a wide range of tasks.

All these tasks were time bound and the groups were asked to complete all the tasks within the allocated time frame. Thirty (30) minutes were assigned for the brainstorming (additional thumb) task, whereas the group had 45 minutes each to complete the energy crisis (decision making) and the neuroscience (problem solving) task.

6.3.3 Experimental Condition: Presence & Absence of Roles

As hypothesized in Section 6.2, the interactions with shared artifacts might vary in nature and intensity based on the presence or absence of different roles. Therefore, in order to study the influence of roles on latent social information, we divided the participating groups across two conditions. In half of the groups (3 groups) roles were assigned to each group members (ROLE condition), and in the other half (3 groups) no roles were assigned (NOROLE condition).

The assigned roles were that of a *group leader*, *time manager*, *information searcher*, and *content organizer*. As the name suggests, the group leader was responsible for resolving conflicts and deadlocks, asking for others' opinions and briefing the group about the decisions being made.

The time manager was responsible for notifying others about the elapsed time periodically, as well as providing awareness about the status of the task so that the group is aware of the remaining tasks to be performed as well as the time at hand. The information searcher was required to search the internet for information demanded by the task, as well as monitoring the search recommendations provided by the search plugin in MeetHub, and notify the group of any useful information by sharing it with the group. Finally, the content organizer was in-charge of keeping the shared artifacts on the workspace uncluttered and organized at all time during the session. These roles were assigned based on the conceptual criteria used by Antunes and Costa [2002]. In addition, all the three roles except the group leader were functionally related to the three meeting tools in MeetHub; i.e. content organizer was related to the organization of the shared workspace, time manager to the time-management widget, and the information searcher to the collaborative search tool. Once assigned with a role, the individual played the same role during the three weeks of the user study corresponding to the three tasks.

6.3.4 Procedure

Familiarization Week

One week before the start of the user study, each group was asked to participate in a dummy collaborative task where the group was asked to prepare a good dinner for some guests who announced spontaneously about their arrival. The group members were supposed to decide on the meals to cook based on their expertise, and organize a list of items required to prepare the meal, followed by a recipe of these dishes. This was done in order to familiarize the group members with the features and tools of MeetHub, as well as to moderate the bias that might occur due to the novelty of the groupware. During this familiarization week, participants' interaction with the system were not recorded and hence it will not form a part of our analyses.

During the Study

During the user study, before the commencement of each weekly session, we used the following procedure:

1. The participants were welcomed and were asked to seat themselves around the table. During the first week the subjects were asked to choose their preferred seat. However, over the next two weeks the participants were asked to sit at the same place with respect to the shared workspace where they sat for the first week. Once seated, the participants were asked to complete a pre-experiment questionnaire recording the basic demographic data as well as the personality information about themselves (for more details about questionnaires, refer to Section 6.3.6).

2. Subsequently, the participants were provided with the instructional material describing

the task to be performed, and were asked to read the material.

3. Once all the group members had read the instructional material, they were asked to discuss about how they will structure the meeting session by deciding on the agenda of the meeting by estimating the number of phases as well as segmenting the total allocated time into these respective phases. This step was provided information required by the time-management tool, and was part of a colleague's research on the effects of temporal awareness in meetings.

4. Next, the group members started with the meeting task, where they were free to discuss and share content over the shared workspace.

5. Once finished, the participants were asked to fill-in the post-experiment questionnaire (refer to Appendix C) by recording their experiences of their interaction with MeetHub and the usage of various tools supplied by MeetHub.

6. Finally, the group members were thanked for their participation in the user study.

6.3.5 Design

We used a one-factor design for our user study, where each group completed three meeting tasks (measurements) over a period of three weeks (one task per week) as shown in Figure 6.4. The experimental condition was the presence (ROLE) or absence (NOROLE) of roles within

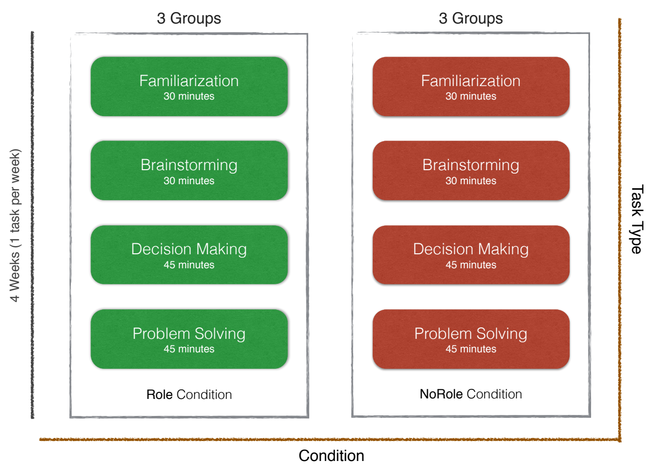

Figure 6.4 – The design of our second user study. We used a one-factor design where each group completed three tasks in a period of three weeks. In addition, the six groups were divided across a control condition, and the experimental condition where group members were assigned different roles.

the group. The NOROLE condition was the control condition, where the group members were not assigned to any role. On the contrary, in the ROLE condition, each group member was assigned a specific role to play during the meeting session. The participants were explicitly made aware of their duties while performing a certain role before the start of the first week's task. These roles were considered as additional responsibilities in addition to performing the experiment tasks.

6.3.6 Data Collection

Questionnaire Data

During the course of three weeks, each experiment participant was asked to complete two questionnaires: one before the experiment session and one after. The pre-experiment questionnaire collected basic information about the participants including the frequency and duration of collaborative work done by them through collocated meetings to discuss projects or assignments. In addition, the working preferences (if a participant preferred to work in groups or alone) as well as the familiarity of participants (rated on a 5-point Likert scale) with their peers was recorded in the pre-experiment questionnaire.

The post-experiment questionnaire (refer to Appendix C) recorded participant's perception of the task, group work, and the usage of different tools (shared workspace, time management widget, and collaborative search) supplied by MeetHub during collaboration. The questionnaire contained statements that the participants had to agree (or disagree) on by marking their ratings on a 5-point Likert scale with 1 as an indicator of strong disagreement, and 5 denoting strong agreement. The statements concerning the task and the group work recorded the participant's perception of the consensus on the task outcome, sufficiency of time to finish the task, differences in opinions among the group members, and acceptance of individual's contribution by the group. Furthermore, the usage preferences of the participants were collected in the questionnaire by asking them about their preferred input device for writing, drawing, and manipulating artifacts over the shared workspace. In addition, we asked participants to state their preferred display when interacting with the shared artifacts (wall-mounted display or tablets). The overall perceived usability of MeetHub was also recorded in the post-experiment questionnaire.

Interaction Logs & Video Recordings

All the interactions of the group members with the shared workspace were also recorded in the system log-files, as well as the sessions were video recorded. Later, the recorded video files were hand-coded to extract phases when a group member was speaking and the phases of no speech. Also, the speech in the videos was later classified as either *on-task* or *off-task* speech.

Group Performance Score

Finally, the groups were evaluated for task performance, and given a score by two experimenters based on the final representations produced over the shared workspace. The criteria used for the evaluation of different task-types is presented next.

- In the brainstorming task, the groups were graded based on the number of ideas (concerning both advantages and disadvantages of having an additional thumb) spooled over the shared workspace, weight of advantages as compared to disadvantages (measured as a ratio), and the proportion of initially spooled ideas that were available in the final solution.

- Next, the ranking of groups in the decision-making task was done based on their correct estimation of energy deficit in the next 5 years, cost-benefit analysis for power generation, focus on sustainable development, and consideration for environmental factors.

- The neuroscience (or the problem solving) task was rated based on the quality of illustrations and schemas of the resting membrane potential created by the groups, and the quality of assignment that was designed by the group members as part of the task.

6.4 Results and Analyses

In Section 6.2 we discussed about the shift in paradigm from an individual centered assessment of collaboration towards an artifact centered analyses in order to comprehend the latent social information embedded in group members' interactions over the shared workspace. In order to achieve this goal, we studied the various kinds of actions that were performed by different group members on artifacts. Consequently, the system log files were parsed in a way to create a state-model for each artifact created, modified, and deleted during the experiment sessions by the group members.

In this section, I will use a two-step methodology for the analyses of the experiment data, where the first step corresponds to the *exploratory* phase concerned with visualizing the different features and properties of the interactions to anticipate visual differences across tasks and roles. Subsequently, I will move to the second step that concerns the *inferential* statistics. For the inferential statistics we will use repeated-measures ANOVA to study the differences in various process variables (properties of the interactions) across condition (ROLE and NOROLE), and across different task types (brainstorming, decision making, and problem solving).

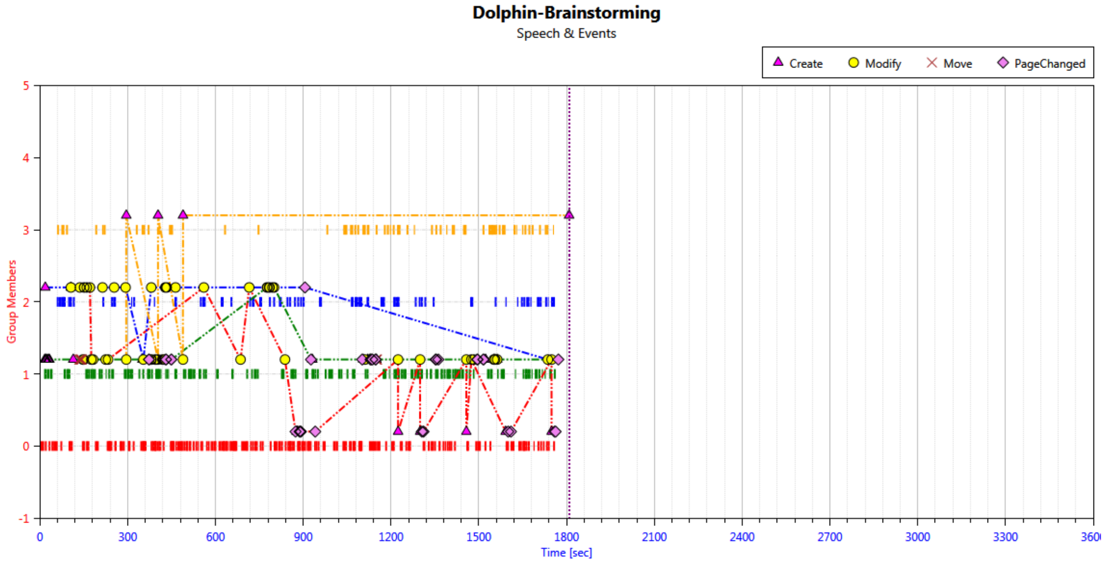

(a) Brainstorming: More Transacting Group

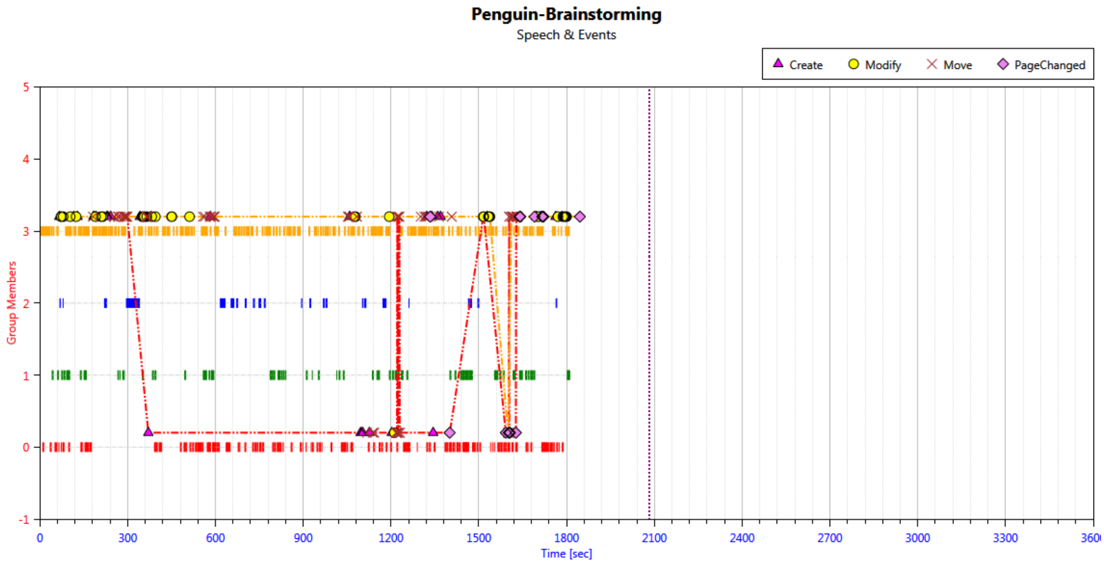

(b) Brainstorming: Less Transacting Group

Figure 6.5 – Visualizing each group member's speech and actions over the shared workspace. The X-axis represents session time in seconds, and the discreet Y-axis denotes various group members in different colors. The bars of different width denote the start and end of speech. On the other hand the user's actions (or interactions with shared artifacts) with the shared workspace are denoted by dotted lines in the user color code. A participant interacting with her own artifacts is represented by a dotted line, and a jump in vertical direction denotes an action where a group member is interacting with the artifacts that were created by another member. The different kinds of actions (create, modify, etc.) are denoted by various symbols as shown in the legend. Finally, the vertical purple dotted-line denotes the end of the experiment session.

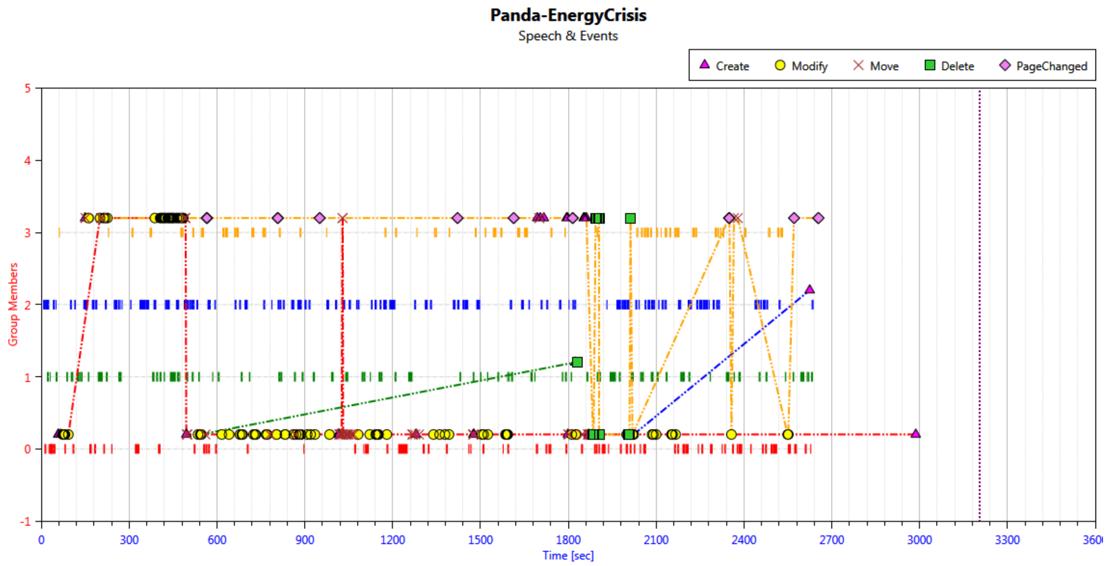

(a) Decision Making: More Transacting Group

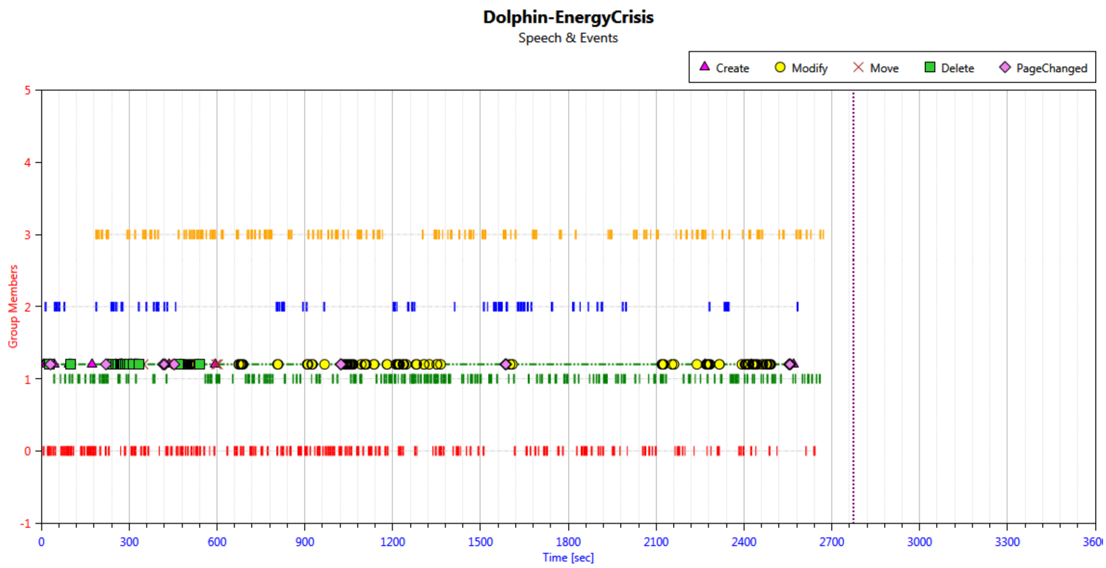

(b) Decision Making: Less Transacting Group

Figure 6.6 – Visualizing each group member's speech and actions (continued) over the shared workspace. The X-axis represents session time in seconds, and the discreet Y-axis denotes various group members in different colors. The bars of different width denote the start and end of speech. On the other hand the user's actions (or interactions with shared artifacts) with the shared workspace are denoted by dotted lines in the user color code. A participant interacting with her own artifacts is represented by a dotted line, and a jump in vertical direction denotes an action where a group member is interacting on the artifacts that were created by another member. The different kinds of actions (create, modify, etc.) are denoted by various symbols as shown in the legend. Finally, the vertical purple dotted-line denotes the end of the experiment session.

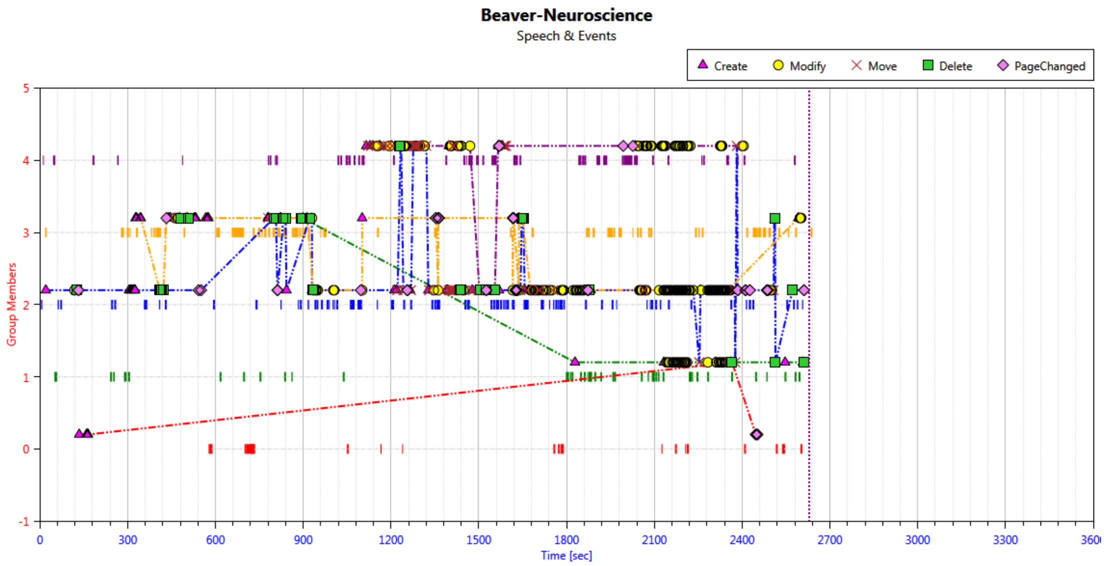

(a) Problem Solving: More Transacting Group

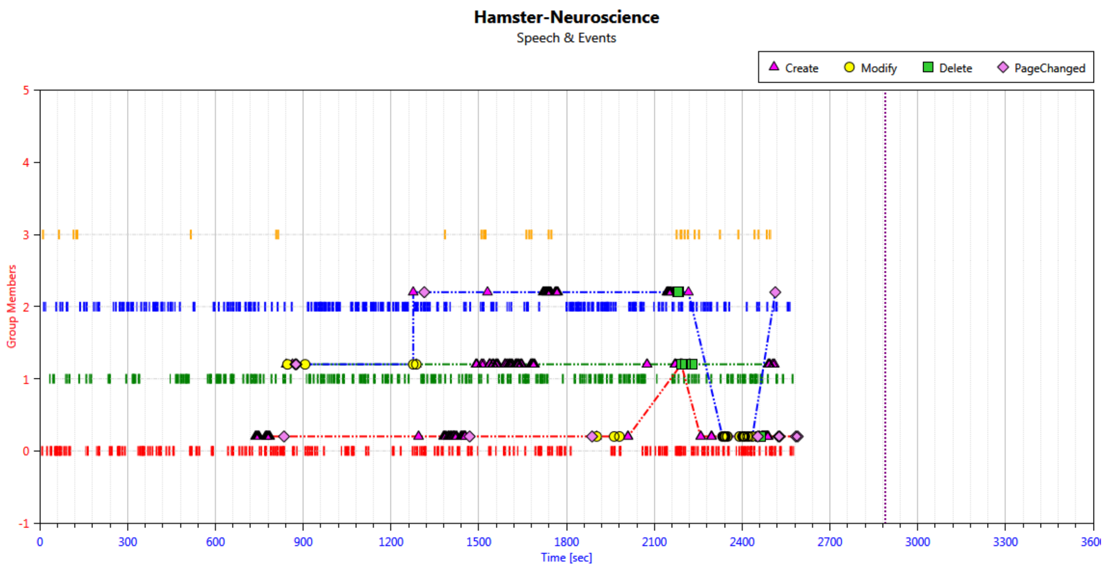

(b) Problem Solving: Less Transacting Group

Figure 6.7 – Visualizing each group member's speech and actions (continued) over the shared workspace. The X-axis represents session time in seconds, and the discreet Y-axis denotes various group members in different colors. The bars of different width denote the start and end of speech. On the other hand the user's actions (or interactions with shared artifacts) with the shared workspace are denoted by dotted lines in the user color code. A participant interacting with her own artifacts is represented by a dotted line, and a jump in vertical direction denotes an action where a group member is interacting with the artifacts that were created by another member. The different kinds of actions (create, modify, etc.) are denoted by various symbols as shown in the legend. Finally, the vertical purple dotted-line denotes the end of the experiment session.

6.4.1 Visualizing Actions and Speech

We visualized each group member's speech as well as actions over the shared workspace, on a timeline for each experiment session corresponding to different groups (see Figures 6.5, 6.6, and 6.7 for a few examples). Based on the type of interaction with the shared workspace and the artifacts, the interactions with the shared workspace were categorized as either of the five kinds: create, move, modify, delete, and page changed. The first four kinds of actions apply locally to a specific artifact (drawing or text), and the page changed action applied globally. In addition, the move, modify, and delete actions were further classified as either performed on the artifacts that were initially created by the same user or a different one. In the following list, we provide a brief description of these different actions.

- The *create* action represented the creation of a drawing or textual element over the shared workspace.
- *Move* action corresponded to the spatial displacement, when an artifact was dragged from one location to another, using either a mouse or a stylus over tablets.
- The *modify* action applied to the textual artifacts only and signified the change in text.
- Finally, the *delete* action corresponded to the permanent removal from the workspace.

In Figures 6.5, 6.6, and 6.7, each plot represents one experiment session where the horizontal axis corresponds to the session time in seconds, and the vertical purple dotted-line represents the end of session. The group members are represented on the discreet vertical axis (or levels) in different colors. The speech of each group member is represented as thick discontinuous line where each bar denotes the start and the end time of an individual's speech. Next, the different actions are represented over a dotted line in the same color as the color of individual's speech and is plotted right above the speech. The dotted line corresponding to actions of a group member is sometimes seen to connect to a different level corresponding to a different group member. This event represents a situation where an individual worked (or interacted with) on artifacts that were created by her peers. We visualized all the 18 experiment sessions corresponding to six groups and three tasks, and we observed some contrasting patterns as shown in the top- and the bottom-side plots in Figures 6.5, 6.6, and 6.7.

The visualizations shown in Figures 6.5, 6.6, and 6.7 provided us with an overview of the nature of interactions that the group members performed over the shared artifacts, as well as some clues regarding the ongoing group dynamics. We visually classified these plots into two separate categories based on the degree of interactivity with other's artifacts (i.e. are group members editing the artifacts that were created by others or not). The groups with more than half of collaborators interacting and editing the artifacts, which were not initially created by themselves, were classified as *more transacting* groups (see Figures 6.5a, 6.6a, and 6.7a). On the other hand, other groups were categorized as *less transacting* groups (see Figures 6.5b, 6.6b, and 6.7b).

It is important to note here that the purpose of these visualizations was primarily to inform us about the nature of the latent social information and observation of visible differentiable patterns in group's interaction with the workspace, rather than rigorous statistical differentiation and analyses, which forms the second stage where we identify individual features and study the differences across the independent variables (task and condition).

The transacting actions (when an individual manipulates the artifact created by another individual) performed by the group members might signify the negotiations happening over the content artifacts, between the group members. These negotiations might be concerned with the validity of information contained within an artifact. Based on the social protocols that are followed by the collaborating groups, these transacting actions can be preceded or succeeded by verbal discussions or acknowledgements between the group members to validate the action over the artifact. The discussions can be detailed in nature and can last for several seconds, or there can be an agreement by a simple gesture of acknowledging this action. In addition, these transacting actions can be peculiar to single display groupware as any group member can interact with the artifacts without competing for the floor, or the acquisition of the input device. Furthermore, we can regard these actions as one attribute of the social information that can be extracted from the interactions with the shared workspace.

These transacting actions can be considered similar to the “silent collaboration” phenomenon referred by Szewkis et al. [2011] and Caballero et al. [2014] (refer to Section 3.3.1 in Chapter 3). Silent Collaboration was defined as a phenomenon where students in a classroom, while using a SDG, assisted their peers in problems by requesting to interact with the problem space of the student they want to assist. Consequenntly, a student can help her friend without offering verbal explanations, simply by interacting with the problem space (or artifact) of the student in need for help. Similarly, the transacting actions refer to an individual's interactions with the artifacts that were created by her peer, in a collocated meeting. Therefore, these transacting actions refer to the silent collaboration, which is happening in meetings, over the shared workspace. However, we can expect a difference in the nature of negotiations and observed social protocols among the school students and the collaborators in meetings.

The visualizations also demonstrated some interesting group dynamics concerned with the usage of SDG. For example, in a group performing the decision making (energy crisis) task as shown in Figure 6.6b, only one group member (in green color) took up a role to interact with the shared workspace during the whole session, whereas others just engaged in the verbal conversations and discussions. In another example, two group members in a group performing the brainstorming task interacted with the shared workspace, and the other two decided to engage in conversation and did not interact with the workspace as shown in Figure 6.5b. Further, among the two group members who interacted with the workspace in this group, only one (in red color) performed the transacting actions (or transactions). Similar pattern can also be seen in the problem solving (neuroscience) task as shown in Figure 6.7b where only one group member (in blue color) performed transactions towards the end of the task. These few examples highlight the divergent strategies used by various collaborating groups such as role-based division of labor, which might provide sufficient evidence that MeetHub can be used not just as a tool that supports the creation and sharing of content, it can also be used to capture the group dynamics and later use this information as an awareness to the group, or predict a certain collaborative process or task outcome.

In the next section, we will list the various indicators of social information that we believe can be extracted from the group members' interactions with the shared workspace. We will also examine their influence on the tasks, condition (ROLE and NOROLE), and outcome of the task.

6.4.2 Attributes of Social Information

Based on the visualizations of speech and actions presented in Section 6.4.1, we extracted the following attributes of social information from the group's interactions with artifacts. It is worth emphasizing again that the following measures were calculated for each artifact that was shared over the workspace during the meetings, and these measures were later aggregated for statistical analyses. Next, we will define each of these attributes followed by their detailed analyses.

Transactions can be defined as asynchronous collective information processing, where group members take turns to interact with, and manipulate an artifact using their respective input devices. In other words, transactions can denote a turn-taking in interactions with artifacts, where an artifact can be assumed to be exchanged between collaborators. A single transaction is counted when a group member modifies an artifact (by moving, deleting, or modifying the content) that was previously manipulated by her peer, irrespective of who initially created the artifact.

Average Transaction Time is the average time in between two consecutive transactions. A lower value of average transaction time can refer to frequent transactions being performed on an artifact, or frequent exchanges of the artifact from one collaborator to another.

Self-Transactions can be defined as repeated individual manipulation of an artifact. Unlike transactions, which signifies a dialogue between two group members via the artifact under consideration, a self-transaction is counted when a group member edits an artifact which was previously also edited by herself.

Average Self-Transaction Time is the mean time in between two consecutive self-transactions. A lower value of average self-transaction time for the whole meeting, might represent individual task work, where collaborators are working on their sub-problems on their own set of artifacts.

Degree of Transaction is measured for each artifact shared during the session, and represents the number of group members who interacted with it. It is an integer value in a range of [0, N-1], where N is the number of collaborators within a group. A higher value of degree of transaction might denote that the artifact was transacted amongst more group members, and might also refer to the fact that the artifact was part of a collective effort, and could also be a part of the final solution. For example, a value of 0 denotes that the artifact was never transacted, and a value of 3 shows that 3 collaborators excluding the creator of the artifact, manipulated the artifact.

Nature of an Action can be defined as either *epistemic* or *cosmetic*. A cosmetic action concerns the visual aesthetics of the shared workspace by arranging them in a visually distinguishable manner, and preventing the workspace from cluttering. It also regards the change in visual features of the content to make it more discernible by changing font size, line thickness, font color, etc. (in this study these features were not supplied to the group members, but they were available in the next study in Chapter 7). On the other hand, as the name suggests, epistemic actions such as modification of text, creation or deletion of artifacts, are related to the state of the content itself and its evolution.

Ownership of an Artifact is a binary variable that is related to the degree of transactions. An artifact can either be *jointly* owned, or *strongly* owned. A strongly-owned artifact has a degree of transaction equal to zero indicating that no transactions were performed. Conversely, a jointly owned artifact has been exchanged (through manipulations) within the group (or a sub-group).

Both transactions and self-transactions are not complementary in a sense that if one increases the other one decreases. In addition, the difference of these two quantities can be understood by relating them to collective and individual taskwork. On the one hand transactions might refer to the episodes where an artifact is exchanged among the collaborators as part of the communication. On the other hand, self-transactions might refer to the individual work, where a group member is responsible to edit an artifact as a task taken up by herself or assigned by the group based on some division-of-labor.

Now, once we have identified some measures of the social information, we will study the effects of these quantities on the independent variables such as task type, and the experimental condition (ROLE and NOROLE). The inferential statistics presented in the next two sections (Effects of Task Type & Effects of Experimental Condition) uses the data consisting of each artifact (both graphical and textual) that was created during the course of the study. We used linear-mixed effect regression, to examine the relationship between the dependent and independent variables. Finally, as some groups finished the experimental task a few minutes before the allocated time duration, we normalized the variables by the length of the session.

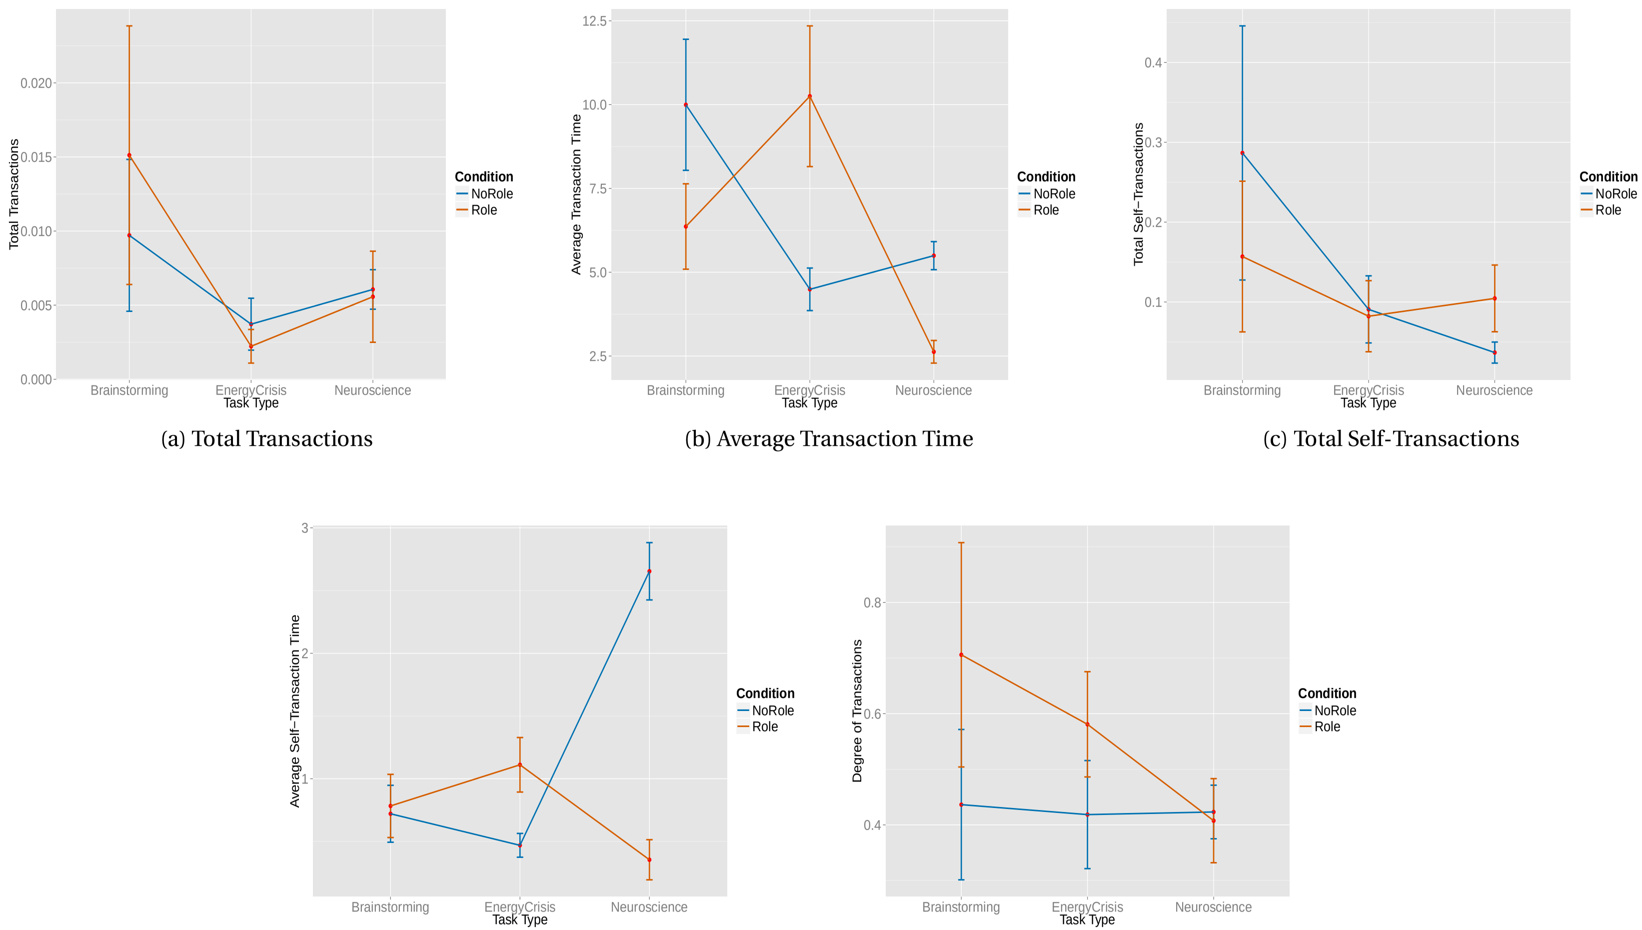

(a) Total Transactions

(b) Average Transaction Time

(c) Total Self-Transactions

(d) Average Self-Transaction Time

(e) Degree of Transactions

Figure 6.8 – Plots visualizing the means and confidence interval for various attributes of social information across the different kinds of task. The values were normalized by the length of the session.

	Task Type					
	Brainstorming		Decision Making		Problem Solving	
Dependent Variable	Mean	SD	Mean	SD	Mean	SD
Transactions	0.39	0.17	0.18	0.14	0.92	0.94
Average Transaction Time	265.69	189.23	271.19	290.15	148.21	104.84
Self-Transactions	7.48	4.79	5.07	7.79	8.98	5.22
Average Self-Transaction Time	23.70	12.70	42.34	27.88	45.58	67.56
Degree of Transactions	0.66	0.29	0.62	0.33	0.59	0.50
Epistemic Actions	9.14	3.53	8.45	3.99	15.60	7.38
Cosmetic Actions	7.58	8.06	3.27	3.44	7.53	8.53
Proportion of Speech	0.63	0.17	0.65	0.11	0.59	0.13
On-Task Speech Proportion	-	-	0.92	0.04	0.84	0.07
Off-Task Speech Proportion	-	-	0.08	0.05	0.16	0.07

Table 6.1 – The mean and standard deviations for various attributes of latent social information across the different task-types. The variables were normalized by the duration of an experiment session.

Effects of Task Type

The differences in the aforementioned attributes of social information across different task types are summarized in Table 6.1. The mean and confidence interval for each task type and for the different experimental conditions (i.e. ROLE and NOROLE) are displayed in Figure 6.8. The results are summarized in the following list:

Transactions

Significant statistical differences were observed in the number of transactions performed by the group members over the shared artifacts across different kinds of tasks ($F(2,1473)=11.56$, $p<.0001$) as shown in Figure 6.8a. More transactions were observed in the brainstorming task, and the least number of transactions were observed in the decision making task (energy crisis task). In addition, the groups in the ROLE condition performed more transactions in the brainstorming task as compared to the groups in the NOROLE condition. On the contrary, groups in ROLE condition had less transactions than NOROLE groups in the decision making and problem solving task.

Average Transaction Time

Considering the average transaction time, a significant difference was observed across the different kinds of tasks ($F(2,237)=4.63$, $p=.01$). The groups in the problem solving task demonstrated the least average transition time as shown in Figure 6.8b. However, no difference in average transaction time was observed in the brainstorming and the decision making task.

Self-Transactions

A significant difference was observed across the different tasks ($F(2,1473)=19.11$, $p<.0001$). A higher number of self-transactions were observed in the brainstorming task, as compared to the decision making and problem solving task as shown in Figure 6.8c. This observation is characteristic of the brainstorming task, where the group members shared their individual ideas over the shared workspace in the initial divergent phase of the task leading to relatively higher self-transactions.

Average Self-Transaction Time

We observed a significant difference in average self-transaction time across the task-types ($F(2,429)=11.71$, $p<.0001$). The lowest value of average self-transaction time was observed in the brainstorming task, which can be attributed to the divergent idea generation phase of the task, where group members share a lot of ideas without judging them. This might reduce the mean duration of self-transaction time for the whole group. Another interesting split regarding the mean self-transaction time was observed in the problem solving (neuroscience) task, where the groups in the ROLE condition had lower self-transaction time as compared to the groups in NOROLE condition as shown in Figure 6.8d, which lead to the overall mean in the problem solving task to be higher than other tasks.

Degree of Transactions

Next, we analyzed the differences in the degree of transaction across the various tasks, and we found a significant difference ($F(2,1473)=4.186$, $p=.02$). Groups in the brainstorming task shared artifacts with a comparatively higher degree of transaction, and the problem solving task had the least degree of transaction. Furthermore, the groups in the ROLE condition demonstrated a higher degree of transactions than the groups in the NOROLE condition as shown in Figure 6.8e. One possible explanation for this observation could be the presence of a leader in the ROLE condition whose role was to coordinate the task activities, participation of the group members, and asking for other's opinions on the shared ideas. However, in the NOROLE condition, an almost horizontal line in Figure 6.8e indicates that collaborators were not interacting with artifacts created by others. This could also be attributed to the lack of a group leader.

Epistemic & Cosmetic Actions

A significant difference in both the epistemic ($F(2,1468)=9.92$, $p<.0001$) and cosmetic actions ($F(2,1468)=12.35$, $p<.0001$) was observed across the different kinds of task, where the brainstorming task had the highest number of both kinds of actions, and the problem solving task had the least number of epistemic and cosmetic actions.

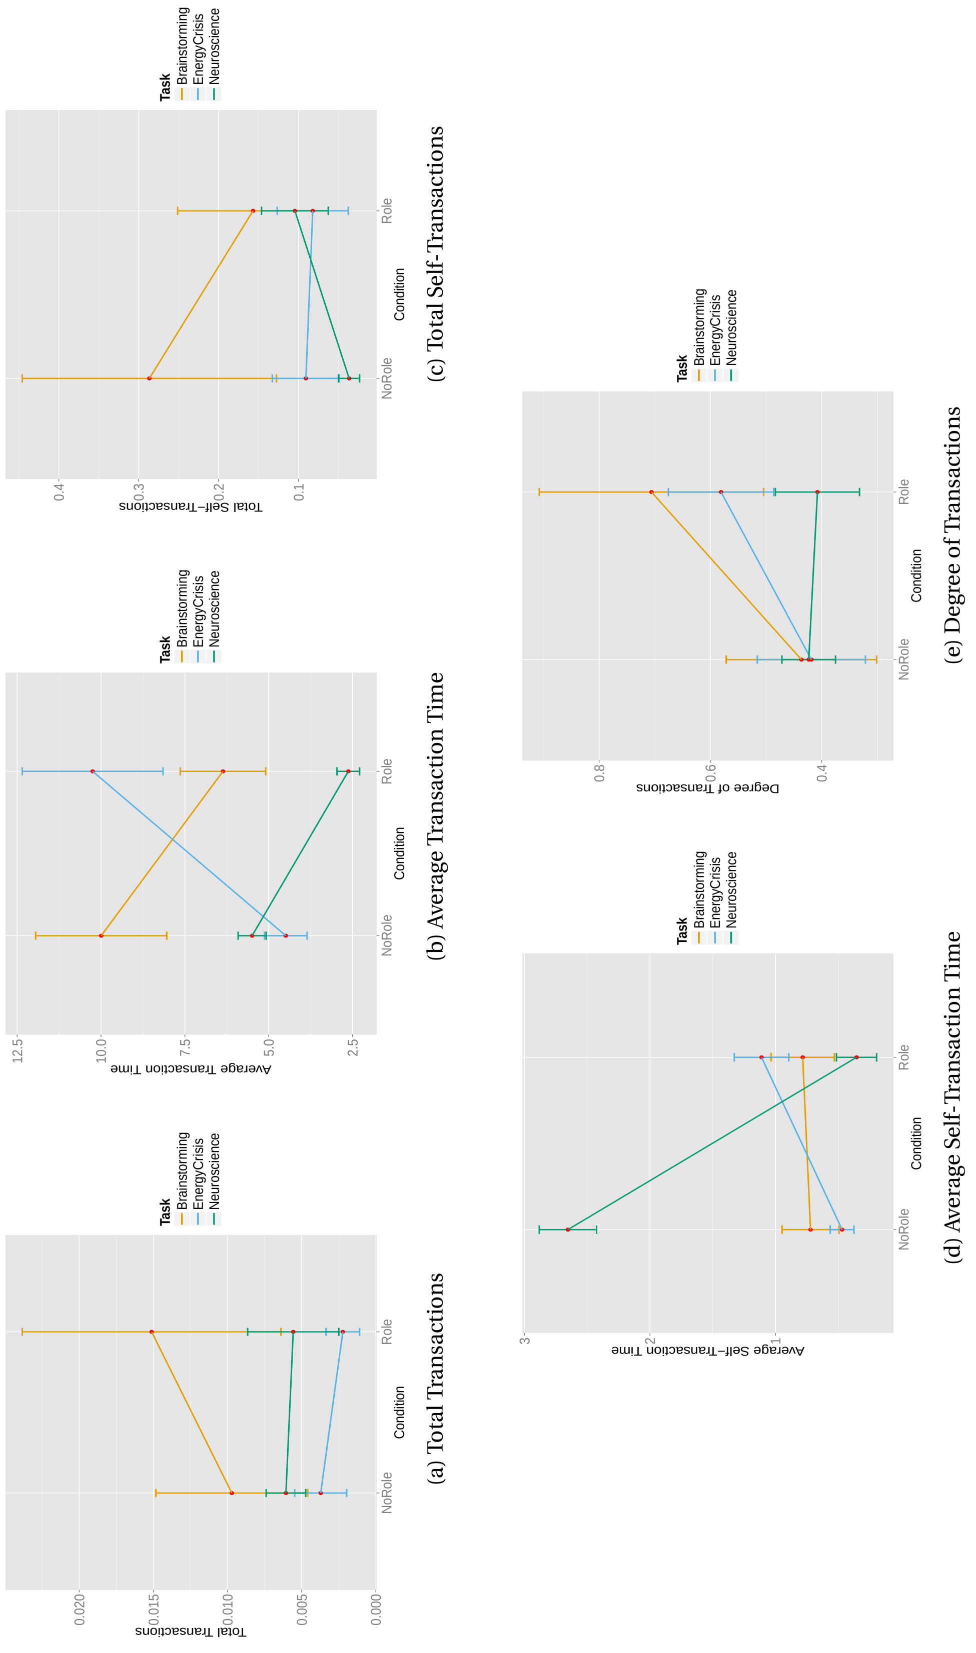

(a) Total Transactions

(b) Average Transaction Time

(c) Total Self-Transactions

(d) Average Self-Transaction Time

(e) Degree of Transactions

Figure 6.9 – Plots visualizing the means and confidence interval for various attributes of social information across the experiment condition. The values were normalized by the length of the session.

	Condition			
	NOROLE		ROLE	
Dependent Variable	Mean	SD	Mean	SD
Transactions	0.62	0.82	0.36	0.29
Average Transaction Time	238.74	170.69	211.33	236.17
Self-Transactions	7.95	4.48	6.40	3.99
Average Self-Transaction Time	41.69	54.74	32.72	26.74
Degree of Transactions	0.52	0.28	0.72	0.43
Epistemic Actions	13.27	7.01	8.84	3.86
Cosmetic Actions	6.21	7.10	6.06	3.44
Proportion of Speech	0.60	0.17	0.72	0.43
On-Task Speech Proportion	0.52	0.16	0.56	0.06
Off-Task Speech Proportion	0.08	0.03	0.06	0.05

Table 6.2 – The mean and standard deviations for various attributes of latent social information across the experiment condition. The variables were normalized by the duration of an experiment session.

Effects of Experimental Condition

Figure 6.9 shows the differences in the social information across the experimental condition; i.e. the presence or absence of roles (ROLE and NOROLE). The mean and the confidence interval values are further separated for different tasks to bring out the differences across the to independent variables. Also, Table 6.2 summarizes the results in the form of mean and standard deviation values, and the results are summarized in the following list:

Transactions
: No significant differences were observed in the number of transactions across the conditions (F(1,1474)=0, p=.9) as shown in Figure 6.9a.

Average Transaction Time
: No significant differences were observed in the average transaction time across conditions (F(1,238)=0.58, p=.4) as shown in Figure 6.9b. The groups that belonged to NOROLE condition, exhibited a higher average transaction time as compared to the groups in ROLE condition (only the brainstorming and problem solving task) as shown in Figure 6.9b, which might signify that group members in the ROLE condition were more frequently exchanging artifacts between themselves. Conversely, a completely reversed pattern was observed in the decision making task.

Self-Transactions
: No significant difference was observed in the total number of self-transactions performed by the groups across the conditions (F(1,1474)=2.26, p=.13) as shown in Figure 6.9c.

Average Self-Transaction Time

A significant difference across conditions was observed with the average self-transaction time ($F(1,430)=29.26$, $p<.0001$). The groups in the ROLE condition had a significantly less self-transaction time. However, if we split the data further based on the different tasks as shown in Figure 6.9d, we see that there is no significant difference in average self-transaction time, except in the problem solving task where the groups in NOROLE condition have the highest mean self-transaction time, which in turn increases the mean value for the NOROLE condition.

Degree of Transaction

A significant difference was observed in the degree of transaction across the condition ($F(1,1474)=5.36$, $p=.02$), where groups in ROLE condition had a higher degree of transaction. Further, taking into account the differences in the various task-types shows some interesting findings as shown in Figure 6.9e. The groups in the NOROLE condition had almost the same mean value of degree of transaction, which is equal to the groups in the ROLE condition and performing the problem solving task. Conversely, the groups in ROLE condition and performing the brainstorming and the decision making task showed a higher degree of transaction.

Epistemic & Cosmetic Actions

No significant differences were observed in epistemic and cosmetic actions across the two conditions.

Effects on Task Performance

The performance of the groups in the three tasks was evaluated by two experimenters based on different criteria as mentioned in Section 6.3.6. The three tasks that we used in our study were different, and the grading criteria employed by us varied with task. Therefore, the scores for each groups were later normalized to lie within the continuous range of [0,1]. A score of 0 denoted a failure to complete the task, and the score of 1 represented complete success.

Upon analyzing the effects of different attributes of social information (transactions, average transaction time, self-transactions, average self-transaction time, degree of transaction, and epistemic & cosmetic actions), we found no significant effect of these variables on group's performance in the task. This can be ascribed to factors related to the nature of the task, such as the open-endedness of the tasks, and the availability of multiple valid solutions. In one aspect, open-ended tasks can be regarded as important while assessing the collaboration because they closely model the kinds of tasks performed by groups in meetings in the wild. Conversely, the open-ended tasks are hard to evaluate due to the availability of several valid solutions, and the role of context and circumstances which also influences the group to consider one solution over another. Therefore, we conducted another user-study with a close-ended task and reported in Chapter 7.

6.4.3 Division of Labor

Jermann [2004] classified the different types of division of labor among dyads based on the actions that each collaborator performed on the various parts of the problem. He defined three categories: *task based*, *role based*, and *concurrent editing*. The task based division of labor was performed by splitting the task into sub-tasks and each group member works on her sub-task. In role based division of labor, one group member acquires a role, and thus performs more actions with the artifacts than others so as to organize content, and coordinate actions within the group. Finally, in the concurrent editing type of division of labor, all the group members equally interact with the problem or artifacts. Further, Jermann discriminated between the aforementioned three types of division of labor based on two dimensions: difference in types of actions, and asymmetry in activity performed by the group member. The former dimension regards the different types of actions that are performed by different group members (create, modify, move, etc.). The later dimension addresses the asymmetry in the collaborators' activity over the workspace in terms of the artifacts that are interacted and manipulated by an individual.

We visualized the division of labor in the groups that participated in the user study across different tasks, based on the classification of division of labor suggested by Jermann [2004]. We extended Jermann's categorization, which was based on dyads, to bigger groups (4-5 members). Due to this extension we expected observing a variety of different types of division of labor within a single group during a task, because of the possible existence of sub-groups as well as roles acquired by group members. We will show some examples of different kinds of division of labor observed in our user study.

Firstly, we investigated the *difference in type of actions* (create, modify, move, delete, or page changed), the varied actions performed by the group members were visualized as bipartite graphs as shown in Figure 6.10. The orange vertices in Figure 6.10 denote the different group members, and the blue vertices represent the different kinds of actions that were performed over the shared workspace. The group on the left-hand side demonstrates a *concurrent editing* type of division of labor, where the two group members perform all the varied actions over the workspace. The other two group members in this group (left-hand side graph of Figure 6.10) do not interact with the shared workspace. The groups in the middle and the right-side graphs in Figure 6.10 exhibit both *role-* and *task-based* division of labor. Regarding the group in the middle graph in Figure 6.10, we see that group member 0 is performing all kinds of actions, on the contrary group member 1 and 3 perform only a small set of actions. If we take into account group member 0 against the collaborators, we observe a role-based division of labor. However, if we regard group members 1 and 3, a task-based division of labor is observed. Similar conclusions can also be made for the group shown at the right-side of Figure 6.10.

Secondly, regarding the *asymmetry in activity* over the shared workspace to categorize the

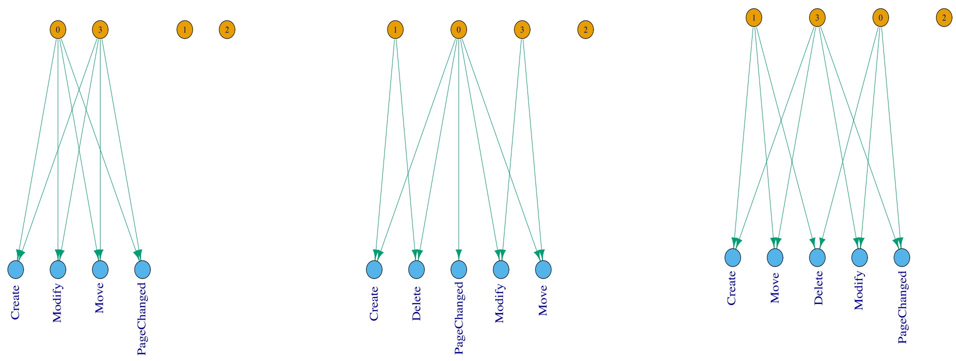

Figure 6.10 – An example of a group participating in the three tasks and exhibiting different types of division of labor based on the *difference in actions* as bipartite graphs. The orange colored vertices represent the group members, and the blue vertices denote the different kinds of actions performed by users. The arrows denote the actions that were performed by each group member. The left-side graph shows *concurrent* division of labor. The graphs in the center and the right-side exhibit the properties of both *task*- and *role-based* division of labor.

division of labor, we visualized, in a bipartite graph, the interactions between individuals and the artifacts that were interacted and manipulated by them as shown in Figure 6.11. Prior to visualizing, in order to obtain a reasonable and informative sample of artifacts over which several actions were performed, we removed all the artifacts that were created and never modified, or were modified only once. The bipartite graphs shown in Figure 6.11 represent the group members as blue vertices and the artifact (texts and drawings) as orange vertices. Further, a mix of different kinds of division of labor can also be observed in this case. Figure 6.11a shows an example of concurrent editing, where group members 0, 1, and 2 interacted with almost all the artifacts. Next, the role-based division of labor is exemplified in Figure 6.11b where group member 1 performed varied actions over all the artifacts. However, a task-based division can also be debated in Figure 6.11b, as group member 0 and 3 split the artifacts and interact with only one set of artifacts. Finally, Figure 6.11c shows a division of labor that resembles a task-based division where group members 0, 1, and 3 interact with their individual set of artifacts with minimal overlap in these sets.

These visualizations shown in Figures 6.10 and 6.11 demonstrate that one can utilize the interactions with shared artifacts to interpret the ongoing group dynamics. This further enhances the relevance of the information contained within the interactions with the artifacts, because the state of the shared workspace reflects the way the group is structured, and the nature of interactions between collaborators.

After regarding the different kinds of division of labor, we looked at the difference in group members' action on the artifacts created by them versus the artifacts created by others. Fig-

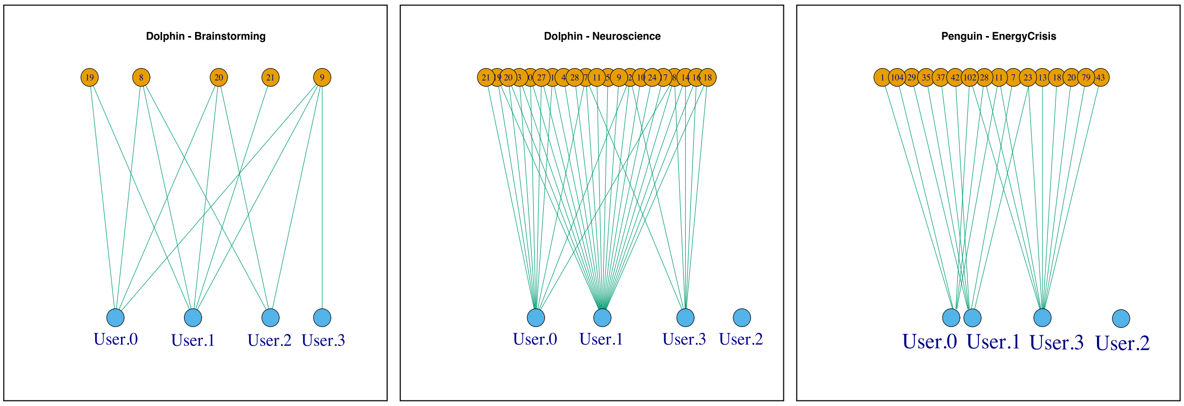

(a) Concurrent (b) Role Based (c) Task Based

Figure 6.11 – An example of groups exhibiting different kinds of division of labor based on the asymmetry of actions. These bipartite graphs show which group member interacted with which artifact during the course of the task. The blue vertices denote the group members, and the orange vertices represent artifacts with a unique ID.

ure 6.12 shows a few examples of group members' interactions with self and other's artifacts as a graph. A weighted directional edge between two group members represents the proportion of an individual's interactions with other's artifacts. For example, an edge from group member X to Y denotes that X worked on the artifacts that were created by Y. The weight of the edge represents the ratio of X's actions on Y's artifacts to the total actions performed by X. Further, a loop (or an edge to the node itself) represents the proportion of actions that a group member performed on the artifacts created by herself. In Figure 6.12a we see that ZERO interacted not just with her own artifacts, but with the artifacts that were created by others as well, mostly the artifacts that were created by THREE. On the contrary, the other three group members in this group (ONE, TWO, and THREE) only interacted with their own artifacts. This indicates a presence of a role (the actual role assigned to this participant was time manager, however she took a role of leader as well) where ZERO interacted with other's content to organize and integrate others' ideas in the brainstorming task. Another example that shows an individualistic group, where group members preferred working on their own set of artifacts is shown in Figure 6.12c, except for a very small proportion where TWO interacted with the artifacts created by ONE.

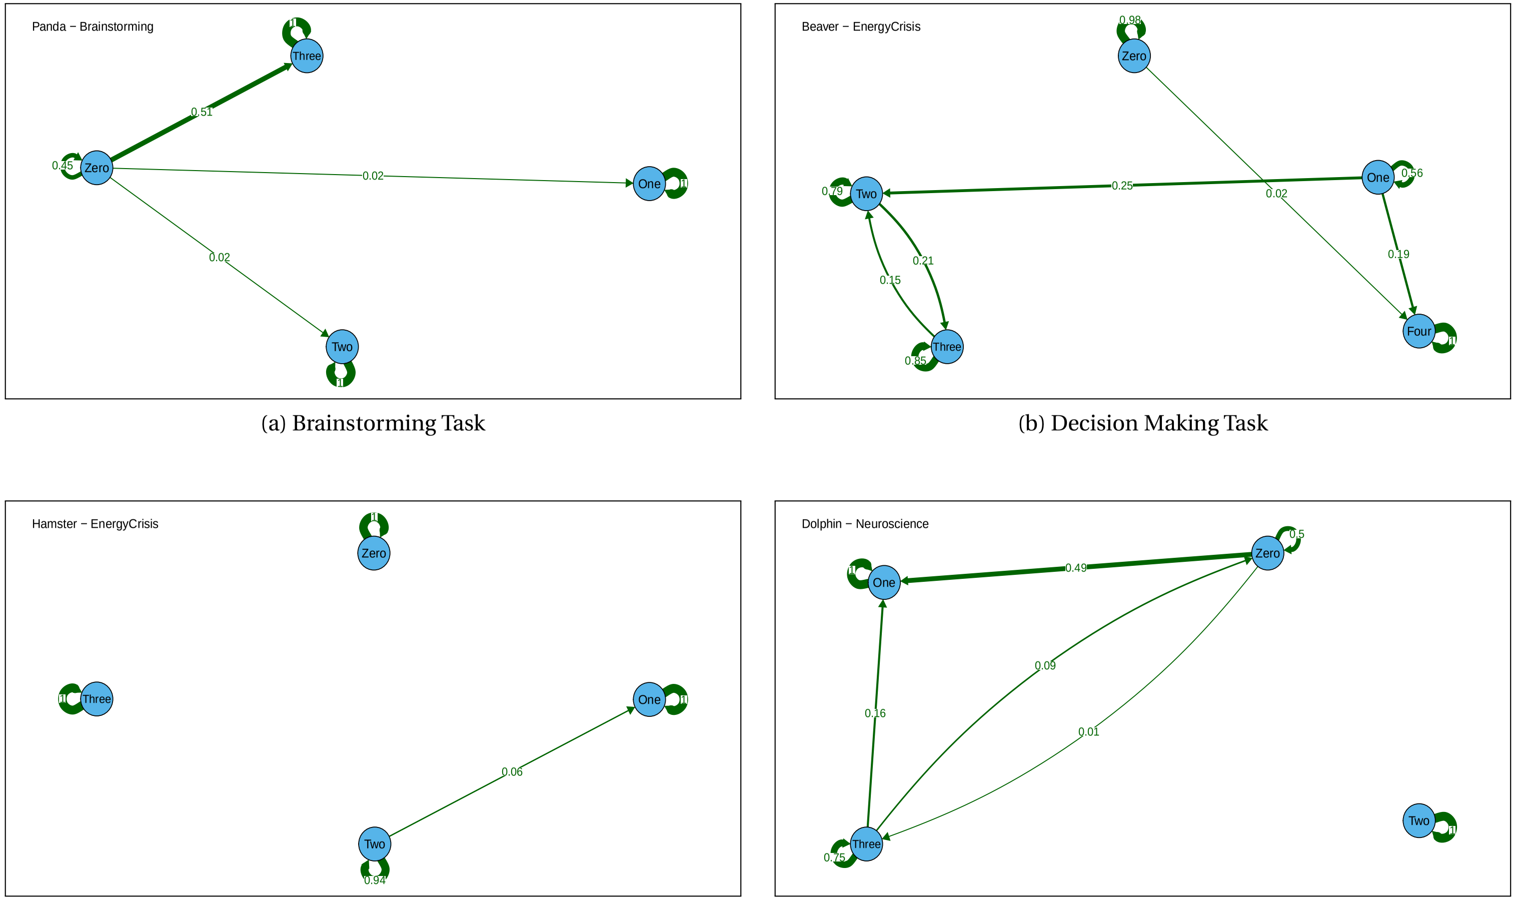

(a) Brainstorming Task

(b) Decision Making Task

(c) Decision Making Task

(d) Neuroscience Task

Figure 6.12 – *Creator-Editor Plots*: The plots represent the proportion values of the total actions a group member performed on the objects that were created by herself (represented by the loop on the vertices), and the proportion of actions when a group member worked on the artifacts that were created by others. The vertices denote the group member, and the directional edges denote the actions. For example, an arrow from member A to B denotes that A worked on the artifacts that were created by B.

These examples could indicate to collaborators' *sense of ownership* of their artifacts against others' artifacts. In some groups, collaborators might perceive to jointly own the shared artifacts, and freely modify the artifacts created by others. However, collaborators in other groups might have a strong sense of ownership, and therefore hesitate to manipulate the artifacts which were not created by them. In order to quantify this sense of ownership, and compute a single measure for the group, we used the *creator-editor* matrices that were used to generate the graphs shown in Figure 6.12. The rows in a matrix correspond to the individuals who created the artifacts, and the columns correspond to the editors (or modifiers). Thus a cell c_{ij} (where *i* denotes the creator and *j* denotes the editor) would correspond to the frequency of *j* modifying the artifacts created by *i*. The diagonal elements of the matrix represent the frequency of modifying one's own artifacts. Further, based on these matrices we formulated the group's sense of ownership through the *ownership index*, and is shown as the following formula:

$$Ownership\ Index = \frac{\sum_{i,j=1}^{N} c_{ij} - \sum_{i=1}^{N} c_{ii}}{\sum_{i,j=1}^{N} c_{ij}}$$

where, N is the size of a group, and c_{ij} is a cell in the creator-editor matrix. The formula generates a number in the continuous range of [0,1], where 0 denotes a strong sense of ownership to one's artifacts, and 1 signifies a joint ownership. In groups where all the groups members work on their individual artifacts all the time, the value of ownership index is zero (0). On the contrary, in a group where group members interact with the artifacts of others all the time, the ownership index is 1. For example, the group shown in Figure 6.12c will have a very low value of ownership index because the group members mostly interact with their own artifacts. On the other hand, the group in Figure 6.12d will have a relatively higher value of ownership index.

Next, we examined the relationships between different attributes of social information and the ownership index. The results demonstrated no significant correlations between the social information (transactions, self-transactions, degree of transactions, and epistemic & cosmetic actions) and the sense of ownership.

6.4.4 Effects on Group Speech

The video-recordings of the experiment sessions were coded to identify the phases of *speech* and *no speech*. In addition to this coding, the students in the CSCW course (who also participated in the study) transcribed the videos of two tasks (decision making and problem solving) as part of an assignment. These transcriptions were later used by two experimenters to further classify speech as either on-task speech or off-task speech. The *on-task* speech referred to the moments of conversations where group members talked about the problem at hand, information contained in the experiment material, and the ways to solve the problem. The *off-task* speech referred to utterances concerning coordination of activities (for example, asking someone else to move a specific artifact on the shared workspace), and technical issues faced by the group members. These speech durations were later normalized by the length of the experiment session in order to get the proportion values.

A linear mixed-effect regression with the different groups as the grouping factor was performed to examine the influence of independent variables (task type and condition), and group performance on speech. The results are summarized in the following list:

Effects of Task Types
: A significant difference in the speech duration (normalized by the session length) was observed across the different tasks (F(2,66)=4.01, p=.02). Also refer to Table 6.1 for the mean and standard deviation values across different tasks.
Group members spoke more in the brainstorming task and the decision making tasks, whereas the least amount of speech was observed in the problem solving task. Furthermore, a significant difference in the speech duration was observed in between the decision making and problem solving task (t(66)=-2.83, p=.006). However no difference was observed in between the decision making and the brainstorming task (t(66)=-1.45, p=.15).

Effects of Condition
: We found no significant difference in the amount of speech across the conditions (ROLE and NOROLE). However, the groups in the ROLE condition spoke more than the groups in the NOROLE condition. Refer to Table 6.2 for the mean and standard deviation values.

Effects on Group Performance
: No significant correlations were observed between the amount of speech and the performance in the task.

Next, we regarded the relationship in between speech and the attributes of social information. We observed no significant correlation between proportion of speech and transactions, degree of transactions, average transaction & self-transaction time, and epistemic & cosmetic actions. However, a significant correlation was observed between self-transactions and amount of speech (F(1,11)=5.45, p=.04).

On-Task & Off-Task Speech

The on-task and off-task speech was coded only for the decision making and problem solving task, and was normalized by the length of the session. The following list summarizes the effects of different variables on the on-task and off-task speech.

Effects of Task Types
: Table 6.1 summarizes the mean and standard deviation values across the two tasks (decision making and problem solving). We observed a significant difference in both the on-task speech (F(1,42)=50.98, p<.0001) and the off-task speed (F(1,42)=53.47, p<.0001) across these two tasks. The groups in decision making task exhibited relatively higher

on-task speech than the groups in problem solving task. Further, the groups in the problem solving task had higher off task speech.

Effects of Conditions

Table 6.2 summarizes the mean and standard deviation values for the two conditions. We observed that group members had a higher proportion of on-task speech in the ROLE condition than the NOROLE condition, but this difference was not found to be statistically significant.

Effects of Group Performance & Social Information

No significant correlations were observed in between the two kinds of speech (on- & off-task speech) and the social information and group performance.

6.5 Collaboration Quality Assessment

In Section 6.4 we discussed the different characteristics of the latent social information within the group member's interactions with the artifacts and content over the shared workspace. We also extensively described the differences in various variables across the two independent variables in our study: task type and presence or absence of role. Besides these differences, we also examined the relationship between social information and amount of speech during each session, as well as the group performance in the tasks. However, these results do not offer us insights on how to utilize the extracted social information to assess the quality of collaboration.

Assessment of the collaboration quality is regarded as a tedious and time-consuming process as it requires researchers to extensively code conversations in video recordings. In addition, the time consuming nature of collaboration assessment task renders it impossible to provide groups with real-time feedback of their collaboration quality, so that they can regulate their behavior in order to achieve higher quality collaboration. This problem of assessing collaboration quality has been addressed differently in the domain of CSCW and Social Signal Processing.

Researchers in CSCW have emphasized on analyzing conversations, gestures, gaze, and posture; most of which is extracted and hand-coded from video recordings of the collaboration sessions, such as the collaboration assessment scheme developed by Meier et al. [2007]. A subgroup within CSCW research domain have successfully applied machine learning techniques to identify different interactional episodes through segmentation (for example McCowan et al. [2003], Zhang et al. [2006], Dielmann and Renals [2007], etc.), by utilizing audio and video recordings of meetings. These research works use the interactions among individuals to identify episodes like monologues (or presentations), note-taking, whiteboard usage, and discussions between group members. However, these approaches provide researchers with a sequence of episodes from the meetings, but not the assessed ratings of the overall collaboration quality. Other research approaches that utilize conversational records (meeting transcripts) and machine learning techniques have implications for information retrieval of meeting records, but not necessarily collaboration assessment. As an example, Lalanne et al. [2003] and Mekhaldi et al. [2004] used machine learning approaches to thematically align the meeting conversations and the document content (what part of a document was referred or talked about at what time) followed by a thematic segmentation of documents to enable effective retrieval of meeting information.

Furthermore, researchers in the social signal processing domain have also emphasized on analyzing individuals' interactional features that are extracted from speech (for example Bachour [2010]), posture, and gaze (refer to Vinciarelli et al. [2009a] for an extensive review of the used features for analysis). Even though some of these approaches are thematically far away from our research, we emphasize them as examples to stress the fact that interactions between individuals is a dominant methodology used to analyze collaborative processes.

Through the analyses of group members' interaction with the shared workspace, we also aim to assess collaboration and to predict some (if not all) aspects of collaborative behavior. Our approach of examining and understanding the actions on the shared workspace is analogous to a microscope that enables pathologists to study micro-organisms and diagnose their effects on human health. Therefore, in this section we study the relationship in between the social information that we extracted and an already established collaboration assessment methodology. This serves two purposes of saving us from additional efforts to develop our own collaboration assessment scheme, and enable us to establish a ground truth by finding the relationship between extracted social information and varied aspects of collaboration. For our research we chose the rating scheme for assessing collaboration quality developed by Meier et al. [2007] as presented in Section 6.5.1.

Further, at the start of this chapter we used a top-down methodology to decompose collaborator's interactions into different kinds of actions and extracted different behaviors as the attributes of social information from within these interactions. Now, I will utilize a bottom-up methodology to find relationships between social information and the visible collaborative behaviors.

6.5.1 Meier-Spada-Rummel Rating Scheme

Meier et al. [2007] developed a rating scheme for assessing the quality of computer-supported collaboration. The scheme was established in the domain of CSCL (Computer-Supported Collaborative Learning), but is easily transferrable to CSCW. They provided a methodology to evaluate five aspects of collaboration by quantitatively rating across nine dimensions, which were identified qualitatively by the authors. The five aspects of collaboration that are evaluated in this rating scheme are *communication*, *joint information processing*, *coordination*, *interpersonal relationship*, and *motivation*. These collaboration aspects are mapped into nine dimensions as shown in Table 6.3. In order to examine the relationships with social information, we coded the videotaped recordings of the experiment sessions based on these nine dimensions as shown in Table 6.3, followed by the analysis of correlations.

Collaboration Aspect	Coding Dimensions
Communication	1. Sustaining mutual understanding 2. Dialogue management
Joint information processing	3. Information Pooling 4. Reaching consensus
Coordination	5. Task division 6. Time management 7. Technical coordination
Interpersonal relationship	8. Reciprocal interaction
Motivation	9. Individual task orientation

Table 6.3 – The five qualitative aspects of collaboration that can be evaluated by quantifying along the nine dimensions. This rating scheme was developed and presented by Meier et al. [2007].

Two experimenters separately rated video recordings of the experiment sessions in five-minute time windows based on the dimensions presented in Table 6.3. The ratings were given on a 5-point Likert scale, where 1 denoted a rater's disagreement and 5 referred to strong agreement towards the presence of a specific collaboration aspect. For example, the process of establishment of common ground among the team members was rated as one of the aspects of sustaining mutual understanding (see the first dimension in Table 6.3), and a rating of 1 referred to a failure in achieving a common ground whereas 5 signified a well established grounding among the team members. Next, we describe the aspects in each of the nine dimensions, that were rated in the video recordings.

Sustaining Mutual Understanding

Establishment of a shared understanding within groups require an easier achievement of a "common ground" about the concepts under discussion and the methods employed to complete the task. In addition, people convey their understanding about their position within the discussion through acknowledgements (or feedbacks). Therefore, we separately coded the quality of *grounding* and frequency of *acknowledgements* (both verbal and non-verbal) within groups for this dimension.

Dialogue Management

The assessment of turn taking and factors which foster the coordination in interactions among the group members were suggested by Meier et al. [2007] to be part of this dimension. Further, evaluating the amount of turn taking is easier between dyads. But as in our groups we observed dynamic creation and dissolution of sub-groups, which complicated the process of evaluating turn-taking by making it a hierarchical evaluation. Therefore, we did not evaluate the turn-taking between group members. Another factor that is not present in the original rating scheme by Meier et al. [2007] but is highly relevant for our research is *content followup*; i.e. how well the collaborators follow the content or ideas brought forward to the table by the group members. We believe that content followup forms the part of dialogue management because it represents how well the group members pursue the thoughts of others by putting it under thorough discussion or disregarding it. Therefore, we rated the quality of content followup in this dimension.

Information Pooling

Collaborators can bring forward the information that is required to complete the task by explicitly asking their peers to share the information (when the collaborators specialize in specific domains; which was not the case in our study) or externalizing their own knowledge by sharing it over the workspace. The group members must also supply appropriate elaborations to support their contribution. In addition, how relevant is the pooled information for the task at hand also plays a crucial role. However, the aspect of relevance was controlled for in our study because one experimenter was always present during the sessions, and the participants were encouraged only to talk about the task. Therefore, in this dimension we rated the *quality of information pooling* and the *level of elaboration*.

Reaching Consensus

Each participant's perception of reaching a common consensus about the solution was recorded in the post-experiment questionnaire, therefore we did not rate this aspect of collaboration. Besides consensus amongst the group members, the critical evaluation of ideas was rated as level of elaboration in the previous dimension (Information pooling).

Task Division

Task division regards the dependencies and overlaps between different sub-tasks, and the coordination inherent within the division of task. In our user-study, each group was asked to discuss about how they plan to structure the task and deciding on the number of required phases before the task started. Also, task division was controlled in our experiment design as we distributed roles to the group members in one of the experiment condition. In the other condition (NOROLE), no such task division was made. Therefore we did not rate this dimension as well.

Time Management

The participating groups were provided with a time-management (or temporal awareness) tool that provided periodic visual awareness about the remaining time in a specific phase of the task. The decision about the number of phases was also made before each session by the group members. Therefore, we did not rate this dimension.

Technical Coordination

MeetHub enabled participants to coordinate between the technical resources required to complete the task, by providing each participant with an input device of their own as well as enabling them to interact concurrently with the shared content. So, we did not rate this dimension as well, because this aspect was implicit in the design of the meeting technology.

Reciprocal Interaction

This dimension takes into account collaborators' mutual respect for each other and equal treatment of the peers. The participants in our user study were all students enrolled in the same course, and therefore had the same educational background. Further, the role of leader in one of the experimental conditions (ROLE) was devised to take care of the issues arising with conflicts between the group members. So, we did not explicitly rate this dimension. However, we recorded each group members' familiarity with the others in the post-experiment questionnaire, in order to quantify the interpersonal relationship between them.

Individual Task Orientation

This dimension concerns an individual's motivation for participating in group activities, and aligning one's actions towards a positive outcome of the group work. In our user study, the motivation to participate in the study was to complete the requirements for the CSCW course that the participants were attending. Therefore, we left out this dimension during assessment.

6.5.2 Relationship with Social Information

The ratings were done separately by two experimenters and were later averaged for each group and each task for the purpose of computing correlations. The ratings by the two experimenters were also accumulated as a time-series with a window of 5 minutes so as to analyze the interdependency trends between the social information (extracted from group's interaction with the shared workspace) and the ratings. We performed the Cohen's Kappa test to measure the inter-rater agreement, and found the Kappa to be equal to 0.83, which is regarded as a high-agreement between the raters.

Correlations

We measured Pearson's correlations between the different ratings and the attributes of social information (normalized by the session length). We found a marginally-significant positive correlation in between the total transactions on artifacts performed by the group and the quality of information pooling ($r(16)=0.42$, $p=.08$). This might indicate that the quality of joint information processing by bringing forward new information to the group is increased as collaborators exchange more artifacts among themselves.

Further, a significant negative correlation was observed in between the number of epistemic actions and the acknowledgments (both implicit and explicit) to others' utterances ($r(16)=-0.49$, $p=.04$). Acknowledgements are part of the communication aspect in the Meier-Spada-Rummel rating scheme as shown in Table 6.3. This negative correlation might signify that the epistemic actions by group members were closely monitored by the group members through discussion about the action. For example a participant asking her peer to edit some text in a specific sticky-note. This facilitated grounding of the shared information and reduced the acknowledgements that are otherwise necessary for grounding.

Next, we examined the correlations between the time-series of ratings and the social information. Such a correlation in time accounts for the changes in the group's dynamics and processes due to the task demands and work preferences. For example, in a brainstorming task we can assume that information pooling happens more in the divergent phase of the task where group members frequently share ideas, and content follow-up can be a leading activity in the convergent phase of the task where each idea is discussed and judged to reach to a relevant conclusion.

Correlations between Time-Series

In order to examine the changes in the different aspects of collaboration, and their relationship with the extracted social information over time, we computed the correlations between the time-series data with a time window of 5 minutes. The transactions and self-transactions performed over artifacts were further sub-divided into *graphical* and *textual* depending on the type of the artifact (whether it was a drawing or a textual element).

The results showed significant correlations with only two dimensions: *information pooling* and *acknowledgements* (which is part of the "Sustaining mutual understanding" dimension). The information pooling was found to be weakly correlated to the number of transactions on textual artifacts ($r(140)=0.16$, $p=.05$). Acknowledgements were observed to be negatively correlated with the self-transactions ($r(140)=-0.23$, $p=.006$), and with the self-transactions on the graphical artifacts ($r(140)=-0.23$, $p=.006$). As we see that the correlation values are not strong for us to make any conclusion, we decided to fit linear-mixed effect models, which also enabled us to study the causality in between different variables.

Linear Mixed Modeling

We fitted linear mixed models on the time-series data, considering the different time windows as the fixed effect, and the groups as the random effect variables. Our results show statistically significant relationships with three dimensions of the Meier-Spada-Rummel rating scheme: *acknowledgements* (a part of "Sustaining Mutual Understanding"), *information pooling*, and *content follow-up* (a part of "Dialogue Management").

A significant negative relationship was observed between content follow-up and the number of transactions ($F(1,135)=3.84$, $p=.05$). A negative relationship of content follow-up was also observed with transactions on graphical artifacts ($F(1,135)=2.51$, $p=.1$, marginal significance), and textual artifacts ($F(1,135)=4.64$, $p=.03$). Similarly, a marginally significant negative relationship was observed in between content follow-up and self-transactions ($F(1,135)=3.01$, $p=.08$). These results suggested that transactions and self-transactions could be complementary to the process of content follow-up. In other words, an act of referencing an artifact by dragging it with a mouse, or modifying their content could signal the group about the piece of information that an individual wishes to discuss; thus a need to verbally elaborate the information and make everyone aware of the current state of discussion is suppressed. Furthermore, the persistence and a common physical viewpoint offered to the whole group by a single shared workspace might assist in this activity, since the role of the workspace is that of a working memory for the whole group where crucial information is shared so that it can be referred to at later point of time [Dillenbourg and Traum, 2006]. Both the persistence of display and the validity of information (because the current state of workspace is a result of collective actions and contributions) enables group members to quickly ground a conversation by a reference on the workspace through a move or modify action, rather than using elaborated verbal utterances.

Next, we observed a marginally-significant positive effect of transactions on textual artifacts and the quality of information pooling (F(1,135)=2.46, p=.1). This result was also observed when we computed the correlations between these two time-series, as mentioned in the previous section (refer to "Correlations between Time-Series" in Section 6.5.2). This finding signifies that exchange of textual artifacts like sticky-notes might increase the tendency of group members to bring forward new ideas, and thus increase the quality of information pooling. One possible explanation for this phenomenon could be that transactions on artifacts induce a sense of collective (or joint) ownership of content, where no single group member owns the content within an artifact. Therefore the group members are more open and comfortable to bring forward new ideas. On the other hand, in situations where groups divide the task into sub-tasks (based on domain expertise or some other criteria), and each member works on a specific part of the problem, we can hypothesize some hesitation to contribute to other's solution.

Finally, the frequency of acknowledgements during conversations was found to be negatively influenced by the number of self-transactions (F(1,135)=6.54, p=.01), and the number of self-transactions on graphical elements (F(1,135)=8.30, p=.005). Acknowledgements within conversations, provide feedback to the speakers about the state of the conversation (are the listeners properly following the conversation), and enables the speaker to anticipate a preventive repair [Dillenbourg and Traum, 2006]. In addition, acknowledgements also provide signals about a listeners' willingness to acquire the floor next by supplying a counter-argument or an alternate idea. Keeping this in mind, we could infer from these results that manipulation of an artifact (when a group member continuously edits an artifact), can serve as a form of a feedback to the group members where the editor (or modifier) either adds the information that the speaker is conveying to the shared workspace, or shares a new information in case the collaborator wishes the group to consider an alternate perspective. In either case, the need for verbal acknowledgements is reduced because it is replaced with self-transactions over the shared workspace.

We also observed significant positive correlations in between grounding and the acknowledgements (r(16)=0.39, p=.1, marginally significant), content follow-up (r(16)=0.41, p=.09), information pooling (r(16)=0.47, p=.05), and level of elaboration (r(16)=0.56, p=.02). Furthermore, even though our linear mixed modeling results show no significant relationships directly with grounding in the group discussion; other findings suggest a possible convergence with the establishment of a common ground, which is also evident in the empirical correlations between grounding and other dimensions in the rating scheme. Therefore, we can draw empirically backed conclusions that various attributes of social information can inform us about the state of grounding within groups and can be used to model this aspect quantitatively. Besides this, we observed a complementarity in visual collaborative behavior and the actions over the shared workspace.

We discuss the practical utility of these results, to predict the episodes of poor mutual understanding, in Chapter 9.

6.6 Discussion and Limitations

This chapter focused exhaustively on the extraction, interpretation, and analyses of the varied characteristics of the social information that is hidden in the group members' interaction with the shared workspace. The nature of this social information can be considered to be *latent*, because unlike face-to-face interactions, the actions over the shared workspace might go unnoticed unless this information is explicitly presented to the group as an awareness of some kind. We hypothesized that this latent social information holds special relevance for the CSCW research because based on the principles of distributed cognition [Hutchins, 1995], the ongoing face-to-face interactions between collaborators and the interactions with external media (or external memory such as shared artifacts) cannot be separated. Previous research focused on collaborators' interactions with artifacts such as the one by Snyder [2012, 2013] regarded the aspect of externalizing conversations by sharing information on a physical artifact, and the socio-linguistic cues signifying the need to externalize the otherwise ephemeral information. However, to the best of our knowledge, no study has yet examined the nature of these interactions, the relationship of these interactions on varied group processes, and what can this information say about the state of collaboration aspects.

We designed and conducted a user-study with MeetHub to record group members' interactions with the shared content across three different types of tasks. In addition, group members in half of the participating groups were assigned different roles in order to assess the differences in social information based on the presence or absence of roles. We chose a top-down approach of analyses where the interactions with the shared workspace (referred to as *actions* in the chapter) were broken down into varied kinds, followed by the utilization of contextual information in order to classify them as different attributes of the social interaction. This approach is similar to the identification of a social signal from different behavioral cues based on the contextual information in the domain of social signal processing [Vinciarelli et al., 2009a].

We identified several characteristics of social information such as transactions, self-transactions, degree of transaction, nature of action (epistemic versus cosmetic), and artifact ownership.

Our quantitative analyses emphasized on the interactions with artifacts instead of interactions between the group members. These variables also signify various aspects of collective information processing performed over a persistent medium such as the shared workspace. Further, MeetHub provides each meeting participant with an input device of their own, this facilitated us to identify one participant's interactions with artifacts from others'.

Our results showed significant differences in these variables across different task types. These results also suggested that the social information from the interactions with the shared workspace can be successfully extracted in varied task types, and the nature of the task itself defines the presence of one kind of social information more than another. For example, a task that demands further sub-division of task into component problems will fare high on the number of self-transactions and lower on transactions and the degree of transactions. These differences across tasks can be used by the groupware to automatically detect and categorize the kind of task, and thus make intelligent context-aware decisions to support the collaborating groups.

Regarding the differences across the experimental conditions (i.e. presence or absence of roles) we found significant differences in the degree of transactions. The groups with assigned roles demonstrated a higher transaction degree than the groups without any roles. This suggests that the group leader might have regulated the group behavior in order for others to contribute equally, thus increasing the number of artifacts that were exchanged between collaborators. However, no significant differences were observed in other kinds of social information across conditions.

Next, concerning the performance (or score) of the groups, we found no statistically significant effect of any social information on the score. One possible explanation for lack of any relationship could be the nature of the solutions that were desired from the groups, which is one of the limitation in our study. All the tasks in our user study were open-ended with several valid solutions, and thus difficult to grade. In order to establish a well grounded relationship between different kinds of social information and the performance, we conducted a well controlled study with a close-ended task, which we will present in the next chapter (refer to Chapter 7).

We also examined the group member's actions over the shared workspace in order to search for different observable patterns of division of labor based on the definition provided by Jermann [2004]. The groups in our study exhibited all the three kinds of division of labor (task based, role based, and concurrent editing) across the two dimensions based on the differences in actions, and asymmetry of activity amongst group members. However, unlike the categorization of division of labor within dyads, it is a complicated process within larger groups because of the dynamic formation of sub-groups. In meetings where the roles are not well-defined, and the group members alter between different roles it might be harder to accurately identify the kind of division of labor. Despite these inherent difficulties, the findings offer crucial insights into the distribution and automatic identification of roles within the group. For example, which group member is mostly performing actions required for the spatial organization of artifacts on the workspace. Furthermore, by the means of creator-editor graphs (proportion of an individual's actions that are performed on self's artifacts, and on the artifacts created by others) we demonstrated the sense of ownership of artifacts (strongly- or jointly-owned) amongst the group members. Besides the ownership of artifacts, we also see crucial group dynamics in the edges of these plots; i.e. the group members who are working together as part of a sub-group on each other's artifacts, and the group members who are not interacting at all with each other's shared content (through a lack of an edge between the group members). These plots can also be used to provide feedback to the group about their interactions as part of a group mirror which might enable the group to regulate their behavior. For example, if all the group members work individually, these plots can make this information visible to the group and encourage them to work more collectively.

As the next step of our analyses, we investigated the relationship between the social information and a well accepted assessment scheme for computer-supported collaboration scenarios designed by [Meier et al., 2007]. We rated the video recorded sessions based on the various aspects of collaboration categorized as nine dimensions in the rating scheme of Meier et al. [2007]. We found that transactions and self-transactions influenced the quality of content follow-up, frequency of acknowledgements, and information pooling amongst group members. On the one hand, the frequency of content follow-up and acknowledgements decreased as the group members interacted and manipulated with the artifacts. On the other hand, the quality of information pooling significantly increased with the increase in number of transactions. The former relationship showed a complementarity in between verbal interactions and the collaborators' actions on the shared workspace. Such a complementarity in between verbal utterances and direct manipulation interfaces was referred by Frohlich [1993] via "referential distance", where the direct manipulation interfaces facilitate referencing an information by the act of indication or manipulation; thus reducing the referential distance as compared to conversations. In other words, transactions and self-transactions serve as "*cross-modal dialogues*" on the workspace, which reduce the need for verbal referential utterances, and thus reducing the rate of content follow-up and acknowledgements. Next, concerning the increase in the quality of information pooling with the increase in transactions might indicate towards the sense of joint ownership of artifacts. The joint ownership might induce an open mindset among the group members that might lead to increased participation by bringing in new ideas and active sharing phenomenon.

Each group member, in our study, was equipped with an input device of her own, this further reduces the effort to indicate or modify an artifact, as any group member can simply use their input device to refer to the information they wish to talk about. For example, instead of asking another group member who has the input devices (in case of a single-input system) to move an information X from one location to another; in SDG any group member with this intention can easily do this. This also indicates to a distance between the *intention* and the *action*, which is reduced while using a single display groupware. This decrease also makes it easier for group members to interact with the artifacts (thus increasing the transactions and self-transactions) while simultaneously offering verbal elaborations about the information contained in the artifact. In other words, transactions and self-transactions act as *cross-modal referents* that accompany the ongoing conversations and assists in improving the mutual understanding amongst the group members, similar to looking at an object while talking about it.

Due to its exploratory nature, there were few limitations associated with our study. As mentioned previously, the open-ended tasks made it harder for us to assess the performance of groups. Also with the presence of two independent variables (various tasks type and presence of roles) it was difficult for us to control any one of the factors strictly and the interactions between the variables. In addition, a more controlled study with more groups might reveal strong statistical relationships, some of which are marginally significant in this study.

6.7 Summary

In this chapter, we presented MeetHub 2.0 that provides group members with a shared workspace as well as a mouse & keyboard and a tablet to create and share content on the workspace. We used MeetHub 2.0 as a tool to record the various actions performed by the group members while creating and sharing artifacts. Later, these actions were characterized into various attributes of the social information that can be found within groups' interaction with shared artifacts. The extracted social information was characterized into transactions, self-transactions, degree of transactions, ownership of an artifact (strongly- or jointly-owned), and the nature of action (epistemic or cosmetic).

We also presented a user study where we examined the differences in the social information across different kinds of tasks and the presence of roles. Our results demonstrated significant differences in the social information across the different kinds of tasks. However, no differences were found across the experimental conditions, which divided the groups based on presence or absence of roles. We also studied the relationship of the social information with different dimensions of the Meier-Spada-Rummel collaboration assessment scheme [Meier et al., 2007]. Our findings showed that the extracted social information (mainly transactions and self-transactions) was correlated to the quality of grounding within groups. Therefore, we can use this social information to model the grounding aspect of collaboration.

7 Latent Social Information & Task Outcome

Other's Contribution and Mine

The user study presented in this chapter was a collaborative work with another colleague, and aimed at investigating the differences in gaze patterns across MOOC learners, and collaborative concept-map activity. Therefore, I will distinguish between my peer's contribution in contrast to my contribution. In this chapter, I will not go into details about my peers research, however I will briefly mention it where necessary.

My Contribution: I analyzed the attributes of the social information that were identified in the previous study, and investigate their relationship with the group performance in the context of a collaborative concept-map activity.

Other's Contribution: My peer was investigating the differences in gaze patterns of the learners while watching MOOC video lectures, and while performing a collaborative concept-map activity. In addition, my peer was also studying the relationship between the gaze patterns and the learning gain.

The nature of the social information that is latent within the group's interaction with the shared workspace was identified in the user-study presented in Chapter 6. Further, we identified in the previous study that some attributes of the social information can be used to model and predict the state of mutual understanding within a group. However, we could not evaluate how this social information affects the performance of a group because the tasks were open-ended and could have multiple valid solutions. In addition, a lack of single task-independent measure, and the relatively small sample size prevented us from discovering a statistical relationship with the social information. Therefore, in this chapter we present a controlled study to investigate the role of social information on the performance of groups (dyads in this study) in a collaborative concept-map activity.

7.1 Research Questions and Hypotheses

Transactions on artifacts denote asynchronous and collective processing of information contained within the artifact via exchange of the artifact amongst group members. Subsequent transactions on an artifact can be considered to be synonymous to the verbal turn-taking between two collaborators, where the state of information initially shared by one group member is repeatedly changed (through spatial reorganization, or changes in the semantics) by other group member. In broader terms, we can assume transactions to be similar to discussions around a single topic, where the topic is an artifact and the verbal contributions can be considered to be either epistemic or cosmetic in nature. Another attribute of social information which is thematically close to transactions is the ownership of an artifact. A jointly-owned artifact denotes that multiple collaborators have performed transactions (or exchanged) on the artifact. A strongly-owned artifact is never exchanged between the collaborators. Furthermore, based on the four levels of grounding (*access*, *perception*, *understanding*, and *agreement*) presented by Dillenbourg and Traum [2006], we can assume that in the presence of a shared workspace similar to MeetHub, the groups have attained the first two levels of grounding; i.e. access and perception. Transactions and joint-ownership of artifacts, as they might signify dialogue about the shared content might refer to the attainment of the third and the fourth level of grounding, and the state of mutual understanding amongst the group members.

We expect that different groups might apply different strategies to solve a problem, which in turn will lead to differences in their interactions with the shared workspace. Further, we identified in Section 6.5.2 in Chapter 6 that these interactions can predict the quality of mutual understanding amongst the group members. Therefore, we can hypothesize that this difference in mutual understanding might affect the group's performance. Next, we will discuss the research hypotheses that we wish to test in this chapter.

Hypothesis 1

A group that performs more transactions on the artifacts can be assumed to have a better mutual understanding of the shared content, due to a better grounding of the shared information. In addition, a higher frequency of transaction might also signify that the group is trying hard to reach to a mutual agreement (based on the 4 levels of grounding by Dillenbourg and Traum [2006]) about the validity of the shared content. This might result into groups performing better in the task. Therefore our first hypothesis can be formulated as: *A group that performs more transactions might also perform better in the collaborative task.*

Hypothesis 2

It is simply not the number of transactions during a collaborative session that will influence the quality of mutual understanding. A group that performs transactions on relatively more artifacts might also perform better, because it signifies that group attained mutual understanding on a higher proportion of the shared content. Therefore our second hypothesis is: *A group with higher number of jointly-owned artifacts is likely to succeed in the collaborative task.*

Hypothesis 3

The nature of interaction with the shared workspace might also influence the performance of a group. For example, in a concept-map creation task, if a group performs more epistemic actions by creating more concepts and linking them, they are more likely to perform better than a group that focuses on the visual aesthetics by making the concept map look much nicer. In addition, in a time-bounded task, a higher proportion of cosmetic actions will leave less time for meaningful actions required to create a correct concept map. Also, epistemic actions are accompanied with the discussions about the linking of concepts, and might result into a better grounding between the group members. Therefore, we can phrase our third hypothesis as: *Higher proportion of epistemic actions might increase the probability for a group to succeed in the task.*

Hypothesis 4

Verbal conversations during collaboration are the popular means of establishing a common ground, and they can be used as a resource to predict the quality of mutual understanding as demonstrated by Traum [1999] and Dillenbourg and Traum [1999, 2006]. Therefore, we can formulate our forth hypothesis as: *The groups that speak more, are putting forth more effort to attain mutual understanding, and are more likely to succeed in the task.*

In addition, if we consider the social information in the interaction with the shared workspace as a proxy for verbal acknowledgements (as discussed in Section 6.6 in Chapter 6), we can expect a complementary relationship in between speech and the various attributes of social information. Therefore, we would also like to answer a question: *What is relationship in between speech and the various attributes of social information?*

7.2 IHMC CMap Tool

In order to enable group members to collaboratively create a concept-map, we used IHMC CMap tools[1], which was developed in the Institute of Human and Machine Cognition, Florida (IHMC) as a software environment that facilitates the creation, manipulation, and sharing of concept-maps by Cañas et al. [2004]. CMap Tool allows users to create concepts in the form of boxes on a canvas as shown in Figure 7.1. In addition, the users can connect these concepts in order to specify the relationship between concepts by creating a link (both directional links with arrows, and non-directional links), and the relationship can be stated explicitly by annotating the link as shown in Figure 7.1.

CMap Tool also allows for multiple collaborators to collectively construct a concept map via CMap Server. We used this feature in the user-study presented later in this chapter. In order to use this feature, the collaborators have to connect their CMap clients to a remote server, and then choose the appropriate concept map that they wish to participate on. Once connected the collaborators get a synchronized canvas where actions of one user are simultaneously

[1] IHMC CMap Tools: http://cmap.ihmc.us/ (visited on April-05-2015)

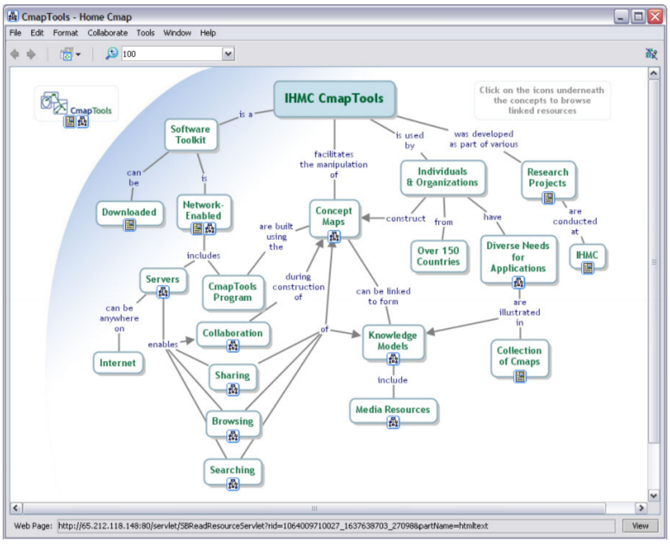

Figure 7.1 – CMap Tool: A screenshot of an example concept map presented in the article by Cañas et al. [2004]. The rectangular boxes represent the concepts, which are connected with various annotated links.

showed to others. The main difference between this setup of CMap Tool and MeetHub is that CMap Tool is not a single display groupware, and all the collaborators have their own display along with the input devices. Also, the position of each user's mouse cursor is not shown on other's display. Therefore, in order to resolve coordination breakdowns (when different users edit the same part of a concept-map) CMap Tool provides selection feedback to others. In other words, when one user selects a concept or a link to edit, other users see the object (concept box, link, or link annotation) highlighted in a different color denoting that a collaborator is currently editing this object.

Besides allowing for co-construction of concept-maps, CMap Tools also provides researchers with log files of every action performed on the canvas along with the kind of action (create, move, modify, delete, etc.) and which collaborator performed this action. These log files can later be parsed to extract suitable information about the collaborative activity.

7.3 Study III: Social Information and Group Performance

In order to test our hypotheses, we designed a user study which is presented in the next sections.

7.3.1 Participants

We recruited 98 master students from our university to participate in the user-study. Amongst these 98 participants, 20 were females and the rest 78 were males. While recruiting participants we made sure that the students were not studying biology or life-sciences, because the experimental task was a learning task related to the domain of neuroscience (resting membrane potential). Therefore, the participants belonged to various engineering domains such as computer science, electrical engineering, material sciences, and so on. The participants were paired into dyads who completed two tasks during the course of the experiment session. At the end of the experiment, each participant was compensated with 30 CHF for their participation in our study.

7.3.2 Experimental Task

Each dyad was required to complete two experimental tasks during the course of the session. The first task was accomplished individually by each participant, and the second task was collaborative. In the first phase, each participant was asked to watch a video lecture on “Resting Membrane Potential”, a topic in neuroscience concerning transmission of signals in neurons. The video lecture was taken from a famous online learning platform called Khan Academy[2]. The original video lectures on Resting Membrane Potential are split into two videos. However, for the purpose of our experiment, we combined these two videos into a single video of 17 minutes and 5 seconds length. The participants were asked to watch the video lecture in order to learn about the concept of Resting Membrane Potential, which was a pre-requisite for the subsequent task. During the video watching task, the participants were free to browse within the video contents by jumping back and re-watching certain segments of the video in order to gain a better understanding of the topic.

List of concepts	
Neuron	Na-K Pump
Axon	K^+ Ions
Electrical Gradient	K Channel
Permeability	Cl^- Ions
Diffusion	Cl Channel
Concentration Gradient	Na^+ Ions
Resting Membrane Potential	Na Channel

Table 7.1 – The initial list of concepts that were provided to the dyads to construct the concept-map.

The second task required pairs to collaboratively construct a concept-map while using CMap Tools (refer to Section 7.2). The participants were also encouraged to talk and discuss while creating the concept-map. Each pair was provided with 14 pre-created concepts on the canvas as shown in Table 7.1, and the pairs were asked to connect these 14 concepts and create links between them in order to construct a concept-map of their understanding of the topic. This was done to facilitate the evaluation of each dyads performance in the concept-map task by

[2]Khan Academy URL: https://www.khanacademy.org/ (visited on April-06-2015).

examining the number of connected concepts and the correct relationships after comparing it with a concept-map created by an expert in neuroscience.

The video-lecture watching task was not strictly time-bounded in a sense that the individuals could take a bit more time than the length of the video so as to gain a better understanding of the concept being taught in the lecture. Each group was required to finish the concept-map task in 10 minutes.

7.3.3 Procedure and Design

The steps that were followed during each experiment session are presented in the following list:

1. Upon arrival in our laboratory, each pair was welcomed and and the two participants were asked to choose their preferred seat at the opposite ends of a table. Each participant sat in front of a computer screen with an embedded eye-tracker, and also facing another participant. However, due to the presence of a screen in front of the participants, they could not see each others' faces, but they were in close proximity so that they can talk with each other, and they sat at the ends of the same table.

2. Further, the participants were asked to sign a consent form, and subsequently were given with the pre-test (refer to Appendix D.1) designed to test their background in the topic which they were supposed to learn next (by watching the video lecture).

3. After completing pre-test, the participants were asked to watch the video lectures individually.

4. Once both the participants had finished watching the video lectures, they were asked to collaboratively construct a concept map. During this phase, the collaborators were encouraged by the experimenters to talk and discuss as the task was collaborative, but were not obliged to do so. In addition, the workspaces of both the participants were synchronized in the concept-map task, and each action made by one collaborator was instantly displayed on the other's workspace.

5. After the concept-map activity, each individual was asked to complete a post-experiment questionnaire (refer to Appendix D.2) consisting of multiple-choice questions.

This user-study was a collaborative effort with another colleague, therefore examination of the role of individual's gaze during video watching on learning gain is not part of the analyses in this chapter. For the analyses of social information and its effects on the group performance, we will only regard the concept-map activity as it was the only collaborative exercise during the sessions.

7.3.4 Data Collection

The gaze of each participant was recorded by the eye-tracker during the video-lecture watching and the concept-map task. In addition to the gaze, the actions of each individual while navigating through the video lecture (such as pause/play, seek forward and backward) were also recorded. In the collaborative concept-map activity, all the interactions of the pairs on the CMap Tool workspace (such as creation of a new link, addition or modification of an annotation, changing the direction of the relationship between concepts, etc.) were recorded in the log files, along with the participant identifier signifying the user who performed the action. Finally, all the concept-map sessions were audio recorded because the group members were allowed to discuss during this task.

The concept-map logs were later parsed to extract the social information attributes (such as transactions, self-transactions, epistemic- and cosmetic actions, and ownership) for each pair. Besides these attributes, we also parsed the log-files to gather fine-grained information about actions that were performed on both the concepts and the links (such as create, modify text, move, delete, etc.). As each group was supplied with 14 concepts at the start of the activity to build their concept-map, we defined an *acquire* action because the concepts pre-existed and could not be created. When one of the participants clicked on any of the pre-existing concepts (or moved the concept on the canvas) to create a link between two concepts, the concept was considered to be acquired by this participant, and this acquisition action was regarded equivalent to the creation of an artifact (the participant was assigned to the creator of the artifact).

In addition to the recording of gaze and actions on the workspace, the pre-test and post-test (refer to Appendix D) scores were used to compute the learning gain of each individual. However, as our research focuses only on the collaborative concept-map activity, and the learning gain was computed for the whole experiment, we will not regard this measure in our analyses. Instead, the score of each dyad in the concept-map activity was computed by comparing it with a concept map created by an expert in neuroscience.

Due to some technical issues with the eye-tracker and a few errors with the recording of concept-map logs, we had to remove 12 pairs from the initial total of 49. Therefore, our analyses will be based on a sample size of 37 pairs.

7.4 Results and Analyses

We investigated linear relationship between different process variables (such as the various attributes of social information) and dependent variables (such as score, speech, gaze similarity, etc.).

Effect on Score	F-statistic F(1, 35)	p-value	Type of Correlation
Transactions	4.07	.05	Positive
Self-Transactions	0.17	>.1	-
Links Created	18.48	.0001	Positive
Concepts Acquired	4.18	.05	Positive
Concepts and Links Used	3.96	.03	Positive
Epistemic Actions	8.14	.007	Positive
Cosmetic Actions	0.005	>.1	-
Strong Ownership of Concepts	16.22	<.001	Negative
Joint Ownership of Concepts	4.43	.04	Positive
Strong Ownership of Links	5.36	.03	Positive
Joint Ownership of Links	8.61	.006	Positive

Table 7.2 – The effects of various variables on the Score of the concept-map activity. We used linear models to examine the relationships.

7.4.1 Effects on Score

In Table 7.2, we summarized the influences of different process variables on the score of dyads in the concept-map task. Figure 7.2 also summarizes the various effects on the score. The results of the linear modeling showed that the number of transactions within groups significantly affected the score of the dyads in a positive manner ($F(1,35)=4.07$, $p=.05$). Next, we divided the groups into two categories by splitting the data at the median value of the transaction counts for each group. The dyads with a transaction count over the median were categorized as the ones with higher transaction-level and vice versa. We observed that groups with higher transactions scored significantly more as compared to the groups with lower transactions ($F(1,34)=6.78$, $p=.01$, ANOVA). These findings statistically validate our first hypothesis (refer to **Hypothesis 1** in Section 7.1). However, we did not find any significant relationship between the score of a dyad and the number of self-transactions during the session.

Next, we examined the relationships between different kinds of actions on the workspace, and how they affect the score. We observed that the groups that created more links, scored significantly better in the task ($F(1,35)=18.48$, $p<.001$). Besides this, the dyads that utilized more concepts from the pre-defined set, were also found to score higher ($F(1,35)=4.18$, $p=.05$). These findings are quite apparent as the groups are highly likely to succeed if they connect more concepts with the assigned relationships among them. Furthermore, the number of epistemic actions such as adding and modifying annotations, adding or deleting links, and so on, had a significant positive correlation with the score of the group ($F(1,35)=8.14$, $p=.007$). However, no effect was found for the cosmetic actions such as spatial organization of the concept-map, changing the font size and color, and so on. These findings also validate our third hypothesis concerning the role of epistemic actions on the score (refer to **Hypothesis 3**

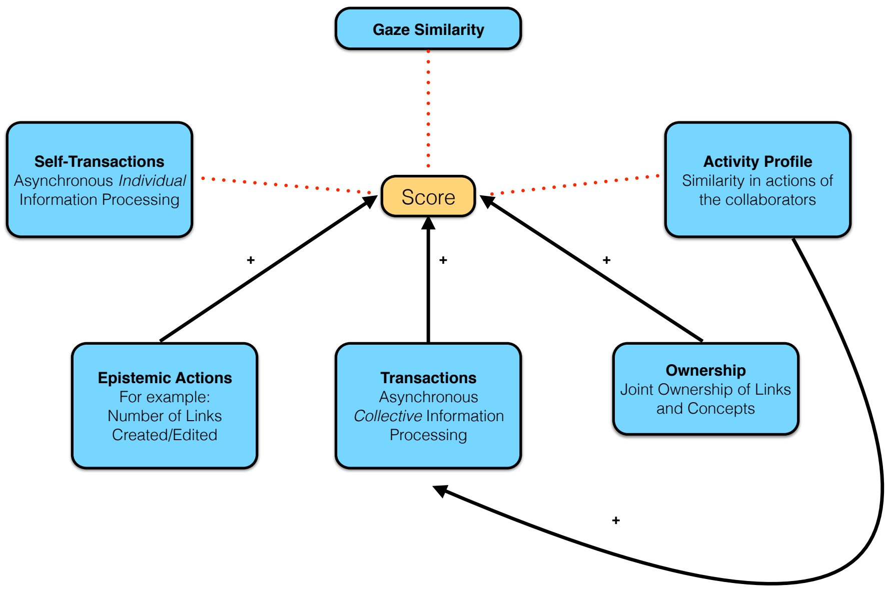

Figure 7.2 – Relationships between different kinds of social information as well as gaze similarity of participants and the score in the concept-map task. The black arrows represent significant correlations and the +/- sign denote the direction of the correlations. The dotted red lines represent no statistical relationships between the two variables.

in Section 7.1).

In order to study the effects of ownership on the score, we examined the ownership of concepts and links separately. Firstly, regarding the ownership of concepts, we observed that strong ownership of concepts was negatively correlated to the score ($F(1,35)=16.22$, $p<.001$). On the other hand, a joint ownership of concepts was found to be positively correlated to the score ($F(1,35)=4.43$, $p=.04$). These results suggested that dyads which divided the concept-map activity by working on different set of concepts (similar to disjoint sets in mathematics), scored poorly during the activity. Secondly, looking into the effects of ownership of the links, we observed positive effects of both the strong- ($F(1,35)=5.36$, $p=.03$) and joint-ownership ($F(1,35)=8.61$, $p=.006$) of links with the score. One possible explanation for this could be that the number of links that were created during the session was also found to be strongly correlated to the score. These relationships between different ownership and score are shown in Figure 7.3. Next, we measured the effects of the proportion of jointly-owned objects (concepts and links), and we found no significant effect of the proportion on the score. These results partially validate our second hypothesis (refer to **Hypothesis 2** in Section 7.1).

The gaze of both the participants on the workspace was recorded during the concept-map activity. We used the gaze data for both the participants and measured the *gaze similarity* between the two collaborators. This measure refers to the probability of looking at the same

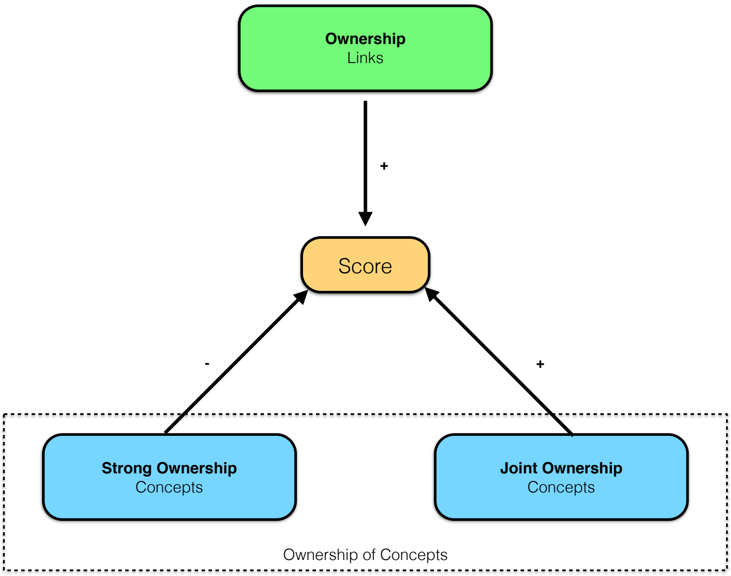

Figure 7.3 – Relationships between ownerships of concepts and links and score in the concept-map activity. The correlations are represented as arrows, and the + or - sign denotes the direction. A + sign denotes a positive correlation, whereas a - sign shows a negative correlation.

area of interest (concept, link, or annotation) at the same time. Next, we examined the effect of gaze similarity on the score of the dyads. However, we observed no relationship in between the temporal cohesion in between the two users' gaze and its influence on the score.

7.4.2 Analysis of Group Speech

The collaboration during the concept-map activity was audio-recorded as the participants were allowed to discuss. Subsequently, we coded the speech duration (pertaining to verbal conversations) of both the participants for all the dyads, in order to investigate the influence of speech on the dependent variable. The speech duration was also normalized to compute the proportion of speaking time during the session. We hypothesized that groups that speak more are exerting more effort to establish a common ground, and therefore they are more likely to succeed (refer to **Hypothesis 4** in Section 7.1). However, we did not find any significant correlation between score and the proportion of speech during the sessions. Also, no statistically significant relationships were observed between proportion of speech and the different attributes of social information.

We further coded 12 dyads' speech based on two self-defined categories of *epistemic* and *cosmetic* speech. In order to understand better the contrast in the speaking patters of dyads, we chose 6 dyads that attained the highest score during the activity, and the other 6 dyads belonged to the lowest scoring category. The epistemic speech was further sub-divided

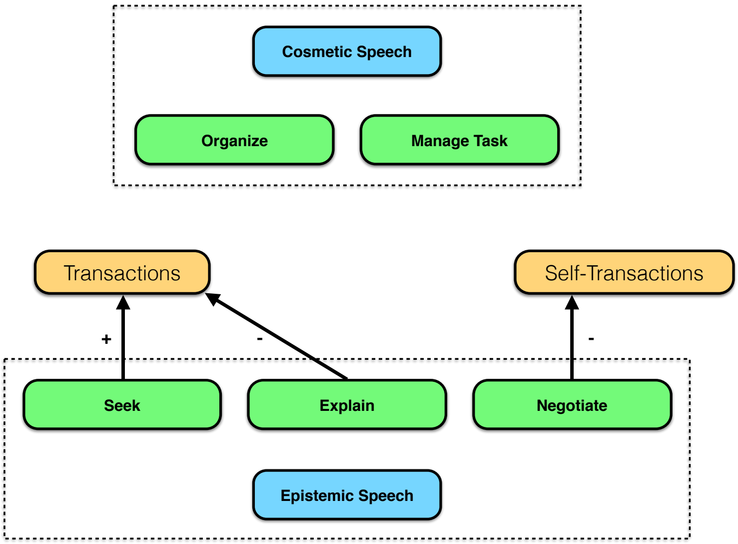

Figure 7.4 – Relationships between different kinds of coded speech and attributes of social information. The correlations are represented as arrows, and the + or - sign denotes the direction. A + sign denotes a positive correlation, whereas a - sign shows an negative correlation.

into three categories based on the nature of utterances: *information seeking*, *explanations*, and *negotiations*. Similarly, the cosmetic speech was further sub-divided into utterances concerning the *organization of concept-map*, and *task management*. The speech related to the organization of concept-map referred to the utterances that were made specifically about the visual aesthetics of the concept-map. On the other hand task management speech referred to the utterances about the strategies of accomplishing the activity and division of labor, such as asking the partner explicitly to make a link between two concepts. The coded speech was later normalized by the total duration of speech during a session in order to compute the proportion of utterances of different kind within the speech.

Figure 7.4 shows the relationships between different sub-categories of utterances and two attributes of social information - transactions and self-transactions. No correlations were found between the two sub-categories of cosmetic speech and other social information attributes. However, we found correlations between epistemic speech and social information as shown in Figure 7.4. A marginally significant positive correlation was observed between information seeking utterances and transactions ($r(10)=0.54$, $p=.07$). This might suggest that while questioning the partner about a specific idea or concept, or about the partner's specific action on the workspace (questions such as "Why did you make this link?"), the questioner grabbed the concerned object with her mouse cursor and moved it on the workspace so that the partner can see the subject of inquiry. This object was most probably created or modified before, by the partner who is questioned. This observation again reflects the process of grounding because the participants cannot see each other's face or gestures (due to the experiment setup). Therefore, an action on the workspace was used as a *proxy gesture* to reduce the chances of misunderstanding.

Explanation utterances were found to be marginally negatively-correlated with the transactions (r(10)=-0.52, p=.08). This finding might imply that as explanations were offered for questions about already grounded artifacts, so the need to exchange the object of discussion will be diminished. At this point we might expect an increase in self-transactions as participants could simultaneously modify the object under discussion while offering verbal explanations. However, we did not observe any significant relationship between explanation utterances and self-transactions.

The negotiation utterances were found to be negatively correlated to the self-transactions (r(10)=-0.61, p=.04). Negotiations might follow an explanation, or could in general signify situations where participants are trying to achieve a mutual understanding about certain concepts related to the task. This finding again suggests that as the subject of negotiations (the concept or link) was already grounded, the need to perform actions on the workspace was reduced. Each negotiation utterance by a collaborator acts as both an acknowledgement and presentation to previous negotiation utterance (refer to the "Contribution Model" of grounding by Clark and Schaefer [1989]). Further, any action on the workspace succeeding the negotiations can be assumed to be a product of the mutual agreement of the negotiation process. Hence less actions are simultaneously performed on the workspace during the negotiation process. However, no significant effect of negotiation utterances was observed on transactions.

These findings concerning the correlations between different utterance-types and social information support our complementarity hypothesis between speech and interactions on the workspace (refer to **Hypothesis 4** in Section 7.1). In other words, interactions over the shared workspace were observed to complement the need for verbal discussions and vice-versa. These findings also indicate that collaborators might choose the medium that is the cheapest in terms of the effort required to ground a certain piece of information.

The participants were asked to speak in english, and for many of our participants english was not the native language. Therefore, we observed some cases of communication breakdowns where the participants were not able to effectively communicate their thoughts in english. This could be a limitation for our study and might also introduce a bias because pairs who can effectively communicate could be more likely to perform better. On the other hand the pairs without fluent english language skills might not be able to express themselves effectively and thus their chances to perform well are inhibited.

7.4.3 Activity Profiling

Previously in Chapter 6 in Section 6.4.3, we observed the different types of division of labor among groups during our user-study with MeetHub 2.0. However, we did not examine the concrete relationships between different practices of division of labor on the performance of the group in the task. In this section, we will regard the balance of participation as an aspect of division of labor, and its influence on the dependent variables such as score as well as the relationship with other kinds of social information.

Participation balance (or imbalance) is a measure that encompasses the difference in the *quantity* of actions performed by each collaborator, as well as the *asymmetry* in the kinds of actions that are performed by different individuals. At the lower-level of granularity the kind of actions performed on the workspace might say a lot about the roles played by collaborators, and the way in which division of labor is achieved (refer to the three kinds of division of labor by Jermann [2004]: role based, task based, and concurrent editing), which in turn highlights the group dynamics. During the concept-map activity, we recorded actions of each collaborator such as create, modify, move, and delete on the workspace. Out of these actions, modify, move, and delete actions were further classified as performed on an individual's own objects or on other's objects. Based on these actions we can identify the different roles played by the different collaborators, or the way in which division of labor was achieved. For example, if a participant mostly performs move action on the artifacts, it might signify that this individual has acquired a role to improve the visual aesthetics of the concept-map. On the other hand, if the participants only perform actions on their own respective objects, it might denote that there is a division of task in a way that each collaborator works on her own sub-problem. Hence, in order to quantify these phenomenons we defined two measures, which were based on the action vectors (a list of frequency of different kinds of actions on the workspace) of each individual, as described next.

Activity Balance

Activity balance can be defined as the degree of similarity in actions of the two collaborators in terms of the different kinds of actions. We measured the activity balance for each dyad by computing the cosine similarity between the normalized action vectors of both the participants. Assuming that different kinds of actions represent a dimension each in a multi-dimensional space (similar to 3D cartesian system), an individual's frequency (or probability) of performing different kinds of actions can be represented as a vector in this multi-dimensional space. A cosine between the two vectors corresponding to the two collaborators, will inform about how different these two vectors are from each other (represented by the angle between these two vectors). An angle of 0° (shown by the cosine value of 1.0, irrespective of the magnitudes of the vectors) will show that the two vectors (or the participants) are *similar* to each other in the kind of actions they perform, and therefore their activity on the workspace is *balanced*. In other words we can say that when cosine value is 1, both the collaborators mirror each other's actions by performing the same kind of actions with different scaling factor. On the other hand, an angle of 90° (cosine value of 0) will represent that the two participants are completely orthogonal in their actions in a sense that if one performs one kind of action, the other performs a totally different kind of action; and thus the activity of both the collaborators can be considered as *unbalanced*. Such an orthogonality might also indicate towards a task-based division of labor as defined by Jermann [2004], and a cosine value of 1 denotes a role-based or concurrent-editing type of division of labor. In this measure, the scale or frequency of an individual's actions is not regarded.

Activity Intensity

Activity balance defines the similarity in activity of the two collaborators, but does not inform us about the magnitude (or scale) of differences in actions. Therefore, we defined the intensity of each collaborator's activity by computing the sum of absolute pairwise differences between the various actions in the normalized action vectors. An intensity value of 0 will denote that both the participants performed the same number of actions (same number of modifications, move, and so on), and hence we can say that the intensity of both participants' activity is *equal*. A zero-value of intensity might also refer to a concurrent-editing kind of division of labor as defined by Jermann [2004]. On the other hand a larger positive value of intensity will denote an asymmetry in the varied actions performed by the collaborators, and hence we can consider the intensity of both participants' activity to be *unequal*. For example, one collaborator might perform more modification actions as compared to the other collaborator. A higher value of activity intensity might indicate towards the presence of a role-based or task-based division of labor as defined by Jermann [2004]. Unlike activity balance, in this measure the scale is an important factor.

Next, we examined the effects of these two measures on the quality of group's performance as well as the relationship with other attributes of social information. Our results show that both the activity balance and intensity did not effect the score of a group during the activity. This finding might suggest that there is not one kind of division of labor which is better than another in terms of succeeding in the task. Subsequently, we investigated if difference in the balance and intensity of group activity reflects the social information.

Firstly, we investigated linear effects of activity balance on the attributes of social information. We observed that activity balance significantly positively influenced transactions ($F(1,35)=10.26$, $p=.003$), self-transactions ($F(1,35)=7.07$, $p=.01$), joint-ownership of concepts ($F(1,35)=16.97$, $p<.001$), and joint-ownership of links ($F(1,35)=4.11$, $p=.05$). These results highlight that dyads with concurrent-editing or role-based division of labor might have performed more transactions and self-transactions, as well as owned a higher proportion of jointly-owned artifacts.

Secondly, we looked at the impact of activity intensity on various kinds of social information. We found that activity intensity was marginally positively correlated to only strong-ownership of concepts ($F(1,35)=3.28$, $p=.08$), and strong-ownership of links ($F(1,35)=3.58$, $p=.07$). These findings might indicate that dyads with higher activity intensity (or unequal actions) might have had higher proportion of strongly-owned objects (both concepts and links). In other words, dyads that demonstrated a task-based division of labor were more likely to have a higher proportion of strongly-owned objects, and might have showed reluctance to edit or even touch the objects that were created by others.

The above findings emphasize on a very interesting transition in group dynamics based on the division of labor strategy that group members choose to adopt. On the one hand, where a choice of a role-based or concurrent editing type of division of labor might result in a higher proportion of jointly-owned objects and transitions; thus referring towards a tightly-coupled group in terms of the co-construction of content. On the other hand, a task-based strategy of division of labor refers to a loosely-coupled group preferring individual work. Furthermore, both the activity balance and intensity can be considered as attributes of the latent social information because while interacting with the workspace this information is not explicitly available to the collaborators and cannot be easily inferred from their actions over time. Since these two measures can also be computed in real time by the groupware via the varied kinds of actions performed by different group members, they provide us with a very powerful tool to identify the kind of group dynamics adopted by the group and use this information to assist group to regulate their behavior in order to assist them to perform better. This objective can be achieved through an awareness tool or a group mirror.

7.5 Discussion

The user-study presented in this chapter was designed to investigate the influence of latent social information on the performance of the group in the task. Unlike the tasks that were used in the previous study (refer to Section 6.3 in Chapter 6), it is easier to assess the quality of group's performance by comparing their solution with an expert's concept-map. In other words, the concept-map task was not as open-ended as compared to the previous collaborative tasks, as the performance can easily be assessed by evaluating the relationships between concepts. Furthermore, the previous study (refer to Section 6.3 in Chapter 6) was exploratory in nature and allowed us to understand and identify the various characteristics of the social information in group's interaction with a shared workspace. However, the exploratory nature of the previous study and the influence of two independent factors made it difficult to concretely validate our findings. Therefore, we conducted this user-study that was well controlled and enabled us to validate the effects of social information on group's performance as well its influence on group dynamics.

Our findings demonstrate that transactions, epistemic action type, and joint-ownership of artifacts positively reflects the performance of the group in the task. A higher frequency of epistemic actions (such as the number of links created and annotated during the activity) makes it more likely that the dyad might retain a relatively higher proportion of correct links, and thus score better. Further, transactions and the joint-ownership of artifacts might indicate the quality of mutual understanding concerning the shared content between the group members. In addition, we showed in Section 6.5 in Chapter 6 that some attributes of social information such as transactions and self-transactions on the shared workspace act as cross-modal acknowledgements; thus facilitating the grounding of shared content and reducing the need for verbal and gestural acknowledgements.

Transactions and joint-ownership of artifacts are closely related concepts because an artifact cannot be jointly-owned if it has not been exchanged between different collaborators. However, there is a subtle difference in how these two measures can be interpreted. The proportion of jointly-owned artifacts refer to the variance in the degree of exchanges; i.e. if only one artifact was exchanged many times or many artifacts were exchanged over time between the collaborators. The transactions represent the intensity of artifact exchanges during the collaboration. Further, a transaction on an artifact (a piece of information) might indicate that there is a dialogue concerning this piece of information, and an inherent need to reach consensus about the validity of this information. Further, the frequency of transactions can represent the effort that the group members are putting forward to reach an agreement regarding the validity of this information. The higher the frequency of transactions as well as the proportion of jointly-owned artifacts, the higher is the probability that the group has a mutual understanding regarding a significant part of the solution, and thus the group is more probable to score better in the task. Therefore, we can infer that transactions and joint-ownership refer to a process of attaining the higher level of grounding, which pertains to the third and the fourth level of grounding presented by Dillenbourg and Traum [2006].

Furthermore, the relationships between different kinds of coded epistemic utterances (information seeking, explanations, and negotiations) and social information (transactions and self-transactions) also support our inference of cross-modal acknowledgements. An increase in explanation and negotiation kind of utterances was found to be negatively correlated to transactions and self-transactions. Relating our inference to the "Contribution Model" of grounding by Clark and Schaefer [1989], we can say that each explanation and negotiation utterance acts as an acknowledgement to the previous utterance (the question in case of explanation, and previous negotiation utterance for negotiations) as well as a presentation of a different viewpoint. Therefore, if acknowledgements are verbally exchanged between the group members, a need to interact with an artifact on the workspace as a form of acknowledgement is reduced. In addition, during each explanation and negotiation the subject of the discussion (an information contained within few artifacts) was already well grounded through an action over the workspace. For example, when a collaborator asks a question and referred to a modification performed by another individual, the questioner interacted with the workspace by clicking on or moving the artifact she was referring to. However, once the context of discussion is grounded the subsequent explanation or negotiation does not require further actions on artifacts until the context of discussions is changed by referring to a different piece of information. This also explains why the information seeking utterances were found to be positively correlated to transactions.

Finally, we observed a unique and interesting relationship between the balance and intensity of group member's activity and other attributes of social information. We presented two measures - activity balance and intensity that were used to quantify the difference in the activity of the collaborators on the workspace. We observed that when both the collaborators' activities were similar (or balanced) in the kinds of actions that were performed over the workspace, the group performed more transactions and self-transactions, as well as had a higher proportion of jointly-owned artifacts. On the contrary, a higher proportion of strongly owned artifacts was observed to be related to the asymmetry in the actions performed by the collaborators. These two measures of activity balance and intensity are analogous to the different kinds of division of labor defined by Jermann [2004]. A balance in activity represents the role-type or concurrent-editing kind of division of labor, whereas the activity intensity refers to the task-based division of labor. Our findings demonstrated that depending on the choice of division of labor made by the group, we can expect differences in the social information that can be prominently collected via the groupware. This has implications for the design of awareness tools as these differences in social information have been shown to influence the performance of the group in the task in our previous findings.

7.6 Conclusion

In this chapter, we presented a user-study to investigate the relationship between various kinds of social information and their effects on the performance of the groups in the concept-map activity. We observed that transactions, and the joint-ownership of artifacts were significantly related to the task performance. In addition, the choice of division of labor was found correlated to the differences in collaborator's actions over the shared workspace. Besides these findings, by the means of this user-study we have established that the latent social information embedded in the group's interaction with the shared workspace can be utilized to reflect the performance of the group as well as aspects of group dynamics such as the division of labor. Our findings also highlighted that some attributes of social information (transactions and self-transactions) can also be used to model the quality of mutual understanding or grounding amongst the group members, as social information acts as cross-modal acknowledgements during conversations.

8 Collaborative MOOC Watching and SDG

Other's Contribution and Mine

The user study presented in this chapter was a collaborative work with another colleague, and aimed to study how study-groups watch MOOC video lectures in collocated settings, and its implication on the pedagogical design of flipped classroom scenarios. Therefore, I will distinguish between my peer's contribution in contrast to my contribution. In this chapter, I will not go into details about my peers research, however I will briefly mention it in places where necessary.

My Contribution: I analyzeed the differences in video-watching behavior across different video watching configurations, where one of the configuration consisted of a single-display groupware (SDG) to interact with the video.

Other's Contribution: My colleague was investigating the role of synchronicity in video-lecture viewing (the condition where the videos of all the individual viewers are synchronized) and its effects on the discussions amongst the study-group members.

In previous user-studies, we focused extensively on the activity of co-creation of content on a shared workspace, and how group members' interactions with shared artifacts influence various group processes. In order to study these interactions with the shared workspace, we designed and developed a single display groupware (SDG) called *MeetHub*, which enables group members to create and share content simultaneously via multiple input devices - one for each user. However, in this chapter we will present a different use case of collaboration that happens within study-groups, learning via MOOC (Massive Open Online Courses) video lectures. Previous studies on the effects of a SDG on collocated collaboration (for example Inkpen et al. [1999]; Moraveji et al. [2009]), observed an increased engagement and interactivity during the activity. Therefore, we envision that the benefits of a SDG can be present in the context of video watching activity. In this chapter, we present a user-study that compared the video-watching behavior of spontaneous study-groups subjected to different video-watching configurations, with an emphasis on the interactivity with the video content and the overall learning experience. One of the video-watching configurations employed a single display groupware.

8.1 Introduction and Motivations

Learning through MOOC video lectures is predominantly an individualistic pursuit on the learners' part. However, many learners prefer studying together and therefore form study groups spontaneously. Further, collaborative learning is often effective, because arguing and explaining results in knowledge co-creation and enhances the understanding of a domain [Vygotsky, 1980]. Furthermore, Rebetez et al. [2010] have shown that groups usually benefited more than individuals while watching dynamic content presentations. Gibbons et al. [1977] observed that learning is negatively affected while watching video lectures individually. Moreover, we believe that social facilitation within the study group might render difficult classes a more pleasing experience. Although there are numerous advantages of learning in a group, the existing MOOC platforms do not focus on creating social learning experiences and are designed for individual learning. Online forums and peer-assessment try to bring in some aspects of asynchronous collaborative learning, however they cannot substitute the richness of interactions and discussions within a study group. Web based social network services, such as *Meetup*[1], enable geographically proximal learners with common learning objectives to find and meet each other. However in its current form, Meetup is used more as a communication platform to organize a meet-up, rather than as a collaborative learning environment.

Taking into account the loosely-coupled collaboration practices around MOOC based learning, and the diversity in learning strategies adopted by the population of learners, we hypothesize that leveraging the social affordances of study groups might lead to a highly engaged and effective learning experience with MOOCs. We also hypothesize that watching MOOC video lectures in groups might lead to discussions and argumentations triggered around the difficult parts of the video; which in turn might result into a shared knowledge model. In addition, collaborative watching of video lectures can be situated at the intersection of formal and informal learning, where learners can dynamically adapt their learning experience (for example, pace) depending on their learning objectives and their prior knowledge of the subject.

Study-groups might decide upon various strategies to watch video lectures. For example, a group might decide that each leaner watches the MOOC videos individually and then they collaboratively solve exercises. On the other hand, another group might prefer to reserve an empty classroom and use the beamer to project the lecture on a big screen and watch the lectures together while simultaneously discussing whenever someone fails to understand the video content properly. These varied strategies adopted by study groups motivated us to

[1] Meetup URL: http://www.meetup.com (visited on April-13-2015)

consider them in our user-study. In addition, the observed benefits of a SDG such as increased engagement and enjoyment of the activity in addition to the learning benefits, encouraged us to include an additional condition. In this condition, group members can watch lectures on a single display groupware, where everyone has equal rights to access and control the videos. We expect that examining learners' behavior across different video-watching strategies might provide us with insights and implications regarding the most effective and acceptable video-watching strategy.

In this chapter, we will report on an exploratory and longitudinal user study that we conducted to investigate the social aspects of watching video lectures within a collocated study group. As our research focuses on the collocated collaboration within study-groups, it might have implications for the blended classroom scenarios [Tucker, 2012]. In addition, this work might inform the design guidelines for the development of tools supporting collaboration in MOOC platforms. The implications are especially relevant for collocated university students following a MOOC course in flipped format, as well as distance learning programs for developing countries where there is a lack of educational infrastructure.

8.2 Collaborative Video Watching

Previous research on *Tutored Video Instruction* (TVI), where group of learners watch video lectures together, has been identified to be advantageous over classroom lecturing as established by Sipusic et al. [1999] and Smith et al. [1999]. The authors found that TVI learners outperformed the students who attended classroom lectures as well as students who watched the video lectures individually. Cadiz et al. [2000] developed and tested a system enabling Distributed TVI (DTVI), where students could collaboratively watch video lectures while being geographically distributed. In addition, the authors also investigated the effect of communication channel richness (text chat vs. audio conferencing vs. video conferencing vs. face-to-face) on learning and interaction behavior. Their results exhibited that groups in the collocated video watching condition had longer discussions. In addition, these groups had higher average discussions per pause as compared to distributed (and collaborative) video watching. In their empirical study, Weisz et al. [2007] reported on the usage of a textual chat feature while watching online videos. They found that chatting while watching videos was perceived to be an engaging and enriching social experience by the participants and was not considered to be distracting. Furthermore, considering the formation of study groups, Tang [1993] studied the spontaneous formation of study groups, which were initiated by the students themselves. The results of the study showed that students participating in a study group exhibited higher engagement level incorporating deeper cognitive strategies as compared to individual learners.

8.3 Research Questions and Hypotheses

The research questions that we wish to answer by means of this user study are concerned with the learning experience of a study-group with MOOC. In addition, we wish to study the differences in video-watching behavior across different configurations, and to identify the configuration that significantly enhances the learning experience of learners. Therefore, we formulated the following research hypotheses.

Hypothesis 1

Watching MOOC video lectures in a study group can be assumed to be situated at the intersection of formal and informal learning, because contrary to the classroom learning, study group participants can freely discuss about the difficult or complicated parts of the lesson without hesitation or shyness. In addition, it is easier to ask questions to one's friends than to the professor in a classroom full of fellow students, where due to time constraints only few questions can be answered, and students hesitate to ask questions to prevent themselves from a situation of possible embarrassment. On the other hand, the study-group way of learning brings in some structure in the learning process as compared to individual (or self-motivated) learning, because the study group members have to decide on a specific time to meet and study rather than delaying or procrastinating. Therefore we hypothesize that *learners who watch videos in teams might have an enhanced learning experience as compared to individual learning and classroom learning scenarios.*

Hypothesis 2

Amongst different video-watching configurations, we believe that equipping each group member with a video controller (input device) of their own, might result in increased interactivity with the video lectures. Consequently, the study-group members might discuss more often as compared to conventional video watching configuration with a single video controller. Therefore we hypothesize that *a condition with a SDG might be the most interactive condition because the group members won't have to compete for resources needed to interact with the lecture.* Here, interactivity with the video lecture refers to various navigational actions such as pausing, seeking forward and backward, and browsing between different lectures.

8.4 Study IV: Comparing across Video Watching Configurations

We conducted a longitudinal and exploratory user study lasting for 5 weeks, in which groups were asked to collaboratively watch video lectures as well as solve quizzes, subjected to three video watching configurations. The study was conducted in an authentic setting with spontaneous study-groups enrolled in a bachelor level university course, which was offered in a flipped classroom format.

8.4.1 Participants

Forty (40) first-year bachelor engineering students (11 females, 29 males) were recruited from our university to participate in the study. All of the participants were enrolled in a university course called *Numerical Analysis* (in french), which was offered over MOOC for the first-half of the semester (seven weeks). Nine (9) groups of 4-5 students were spontaneously formed. All the students who participated in our user-study had no previous experience of learning with MOOCs. In addition, all except one student reported that they have participated in a study group in the past. All the participants also reported that they were well acquainted with their peers in the study-group. Finally, after the experiment the participants were compensated with 150 CHF (equivalent to 150 USD) as well as the course textbook for their participation in the study. The participants were provided with the textbook at the start of the experiment, and the financial compensation was offered at the end of 5 weeks of the study.

	Display	
Video Controller	Centralized	Distributed
Centralized	CC	-
Distributed	DC	DD

Table 8.1 – The distribution of experimental conditions by manipulating across the two dimensions of *display* and *video controller*.

8.4.2 Experiment Conditions

Based on the varied video-watching behaviors we chose to have three conditions. We used two dimensions of *display* and *video controller* to devise our conditions. These dimensions are concerned with the technological resources required to watch and interact with video lectures. The three conditions were designed after manipulating across these two (discrete) dimensions as shown in Table 8.1. The *display* dimension referred to the screen or the output channel that was used by the group members to watch the video lectures. The *video controller* dimension referred to the controller that allowed participants to navigate through the video content by pausing, seeking, or browsing between different videos. This dimension included both the hardware controller as well as the video controller supplied with the video player application. Both these dimensions were discrete in nature and their values were either *centralized* or *distributed*.

A detailed description of the three conditions is as follows:

Centralized video control and centralized display (CC)
: In this condition, the group members sat around a table and watched video lectures on a single wall-mounted projected display, connected to a single tablet computer (multi-touch enabled) as shown in Figure 8.1a. The group members could control the video (pause, seek, etc.) by using the provided video-controller in the video player application.

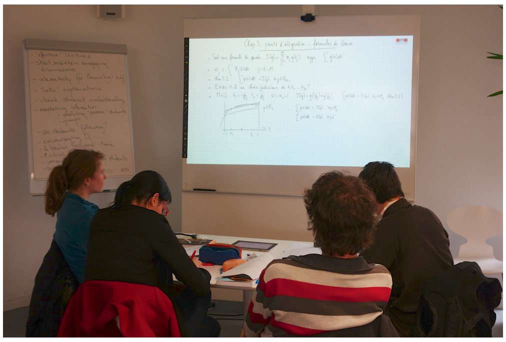

(a) CC Condition: Groups watched videos on a centralized display with a single video control.

(b) DC Condition: Groups watched videos on a centralized display, but each participant had her own video control.

(c) DD Condition: Each group member was supplied with their own tablet to watch video lectures, and control it according to their own pace.

Figure 8.1 – The three video-watching configurations that were used in the study. The conditions were designed based on manipulations across two dimensions of *display* and *video controller*.

This condition is similar to the TVI condition from the user-study of Cadiz et al. [2000].

Distributed video control and centralized display (DC)

In this condition, the group members sat around a table and watched video lectures on a single wall-mounted projected display, connected to a PC. In addition, each learner was equipped with a mouse of her own, which could be used to interact with the video player as shown in Figure 8.1b. In other words, in this condition, the study group was provided with a single display groupware (SDG) where the video-lectures were shown over the shared workspace. We developed the system for this condition as an additional component to MeetHub 2.0. Further, this condition can be considered to be analogous to a single television with multiple remote controllers, one for each viewer.

Distributed video control and distributed display (DD)

This condition was our control, where group learners sat around the table and watched the video lectures on a separate tablet computer at their own pace as shown in Figure 8.1c. Each group member could only control her own video, and no awareness of other's location within the video lecture was provided to the participants. The participants used headphones while viewing the lectures.

As shown in Table 8.1, we had no condition corresponding to centralized video controller and distributed display (CD condition) because this condition did not correspond to a realistic video watching behavior, especially in collocated setting.

8.4.3 Methodology

The study was conducted in the spring of 2013 and lasted for a period of 5 weeks. The 9 groups were divided into 3 experimental conditions as shown in Figure 8.1. The 9 study groups were asked to watch the MOOC videos from the *Numerical Analysis* (NAS) course, which was offered as a flipped course in our university (in french) and was hosted on Coursera[2]. The original duration of the NAS course was 7 weeks, during that time the students were required to watch the video lectures, and no classroom lectures were offered by the professor. However, weekly recitation (or tutorial) sessions were organized by the professor and the teaching assistants to support enrolled students with any problems related to concepts and exercises. The participants who were recruited for our user-study were enrolled in the NAS course, and this renders our study authentic because the students were also expected to take exams (including mid-term exams) at the end of semester.

For each consecutive week of the user study, each group was asked to watch the video lectures as well as solve exercises collaboratively corresponding to that specific week of the NAS course. We instructed the participants well in advance, not to watch the MOOC videos before they come for the weekly experimental session. We used one factor (video watching configuration), between-group design for our user study. The learners were encouraged to discuss anytime during the video watching session.

The experiment sessions lasted for a maximum duration of 3 hours, and the group members were asked to watch videos at their own pace, discuss about the content, and solve the quizzes corresponding to that week's lecture. Table 8.2 shows the length of video lectures (in minutes) for each week's experiment session. After each session, every participant was asked to fill in a post-experiment questionnaire meant to record the participants' perception about the video watching session (refer to Appendix E). However, as our user study was exploratory in nature and focused on the collaborative video watching of MOOC videos, we did not emphasize on the learning outcome and learning gain in collaborative settings. Therefore we also did not include any pre- and post-experiment tests measuring the learning of study group members.

[2]Analyse Numérique pour Ingénieurs: https://www.coursera.org/course/analysenumerique (visited on April-13-2015)

	Week 1	Week 2	Week 3	Week 4	Week 5	Mean (SD)
Video Length (in minutes)	39.36	56.90	59.26	48.28	59.46	52.65 (8.72)

Table 8.2 – The total video lengths (in minutes) for each weekly session of Numerical Analysis MOOC.

8.4.4 Data Collection

All the interactions of the group members with the video players (such as pause, seek, jump to a different video, etc.) were recorded in the system log files. In addition, in the DC condition, we also recorded who interacted with the video as each group learner had her own mouse device (video controller). In addition, all the participants were asked to fill in a pre-experiment questionnaire (refer to Appendix E.1), asking questions concerning basic demographics information, personality questions as well as their experience with MOOCs and study groups. After each weekly session, the participants were asked to complete a post-experiment questionnaire (refer to Appendix E.2), recording the perceptions of the group members about the video watching session, and their satisfaction with the discussions during the session. Finally, we video recorded all the experimental sessions and also conducted semi-structured interviews with the participants at the end of the user study.

8.5 Results and Analyses

8.5.1 Perceived Video Watching Experience

The post-experiment questionnaires after each weekly session as well as the semi-structured interviews with the study-group members reveal that the learners liked the study group way of learning with MOOCs. The participants reported on a 5-point Likert scale (1 denoting a poor quality, and 5 denoting a good quality of discussion) that the discussions were of high perceived quality and very beneficial as shown in Table 8.3. However, no significant statistical difference was found in the perceived discussion quality across the three conditions. Although the participants in the DD condition were watching the video lectures individually and controlling their own respective videos, they also found the discussions to be beneficial

Perceived Quality of Discussion	Mean	Standard Deviation
All Participants	4.12	0.81
CC Condition	4.17	0.79
DC Condition	4.14	0.87
DD Condition	4.03	0.76

Table 8.3 – The perceived quality of discussions during collaborative watching of MOOC video lectures.

for their learning experience. We observed that the participants in the individual watching condition waited for the other group members to finish their videos and discussed after each video. This demonstrates that being in a study group influenced the individual video viewers to synchronize their learning practice while giving importance to the discussions that followed after each video.

The participants also reported that the ability to pause the video lectures in situations where a learner required explanations about a specific concept, followed by discussions with peers was similar in experience to "*...its like pausing a professor in the classroom...*" as reported by a participant in one of the semi-structured interview. Also, the participants had positive opinions regarding the motivation for learning with MOOCs in a study group, because they can validate their understanding of the topic with their friends instantly. This was reported by one of the experiment subject as: "*...If you are alone, maybe you doubt about your results and understanding and you validate the solutions to problems two or three times, to be sure. But here [within study group] you can compare with your friends. This is much better!*". Further, the individualistic way of learning with MOOC was perceived as less interesting as compared to the study groups. One of our experiment participant reported: "*The study group is much better than studying alone. More motivating, ...asking questions and getting answers, so that we are able to understand better. Alone, it would be more difficult and less interesting*". The study group members also emphasized the fact that many MOOCs are very difficult, consume a lot of time, and study groups render this process into a pleasing learning experience. Finally, the study group members perceived an increased activeness and attentiveness while watching video lectures together; as reported by one of the participant "*...during lectures [classroom] sometimes you kind of drift away, while here [study group] because of the interaction with the group, you are obliged to become more active*".

These qualitative findings suggesting an increased engagement and enjoyment during he collaborative video-watching activity, as well as the perceived satisfaction with the learning experience and the quality of discussions were in line with our first hypothesis (refer to **Hypothesis 1** in Section 8.3).

8.5.2 Video Navigation Patterns

Examining the study groups' navigation within the video lectures might provide some useful insights into the ongoing dynamics of the study groups. For example, the time when a video was paused and the duration of the pause might indicate a discussion about the concept being taught at that time. Also, seeking back (jumping back) within a video might denote that the group failed to follow what the professor said in the video. Therefore, in order to perceptually understand the study groups' behavior while watching videos we visualized their navigational behavior as shown in Figures 8.2 and 8.3.

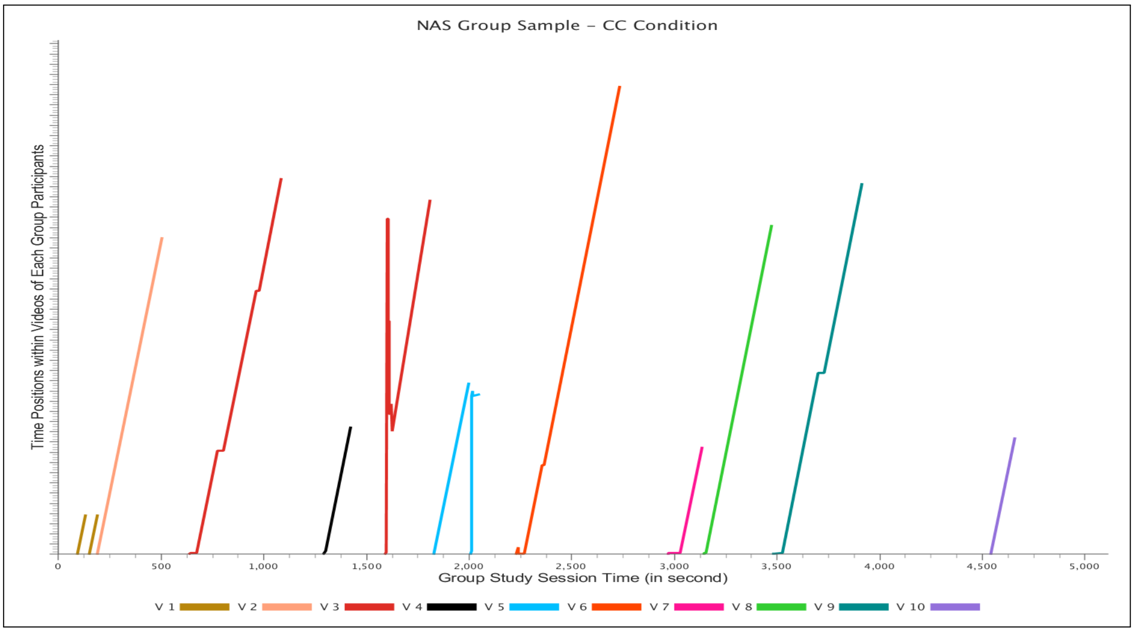

(a) A group in CC condition

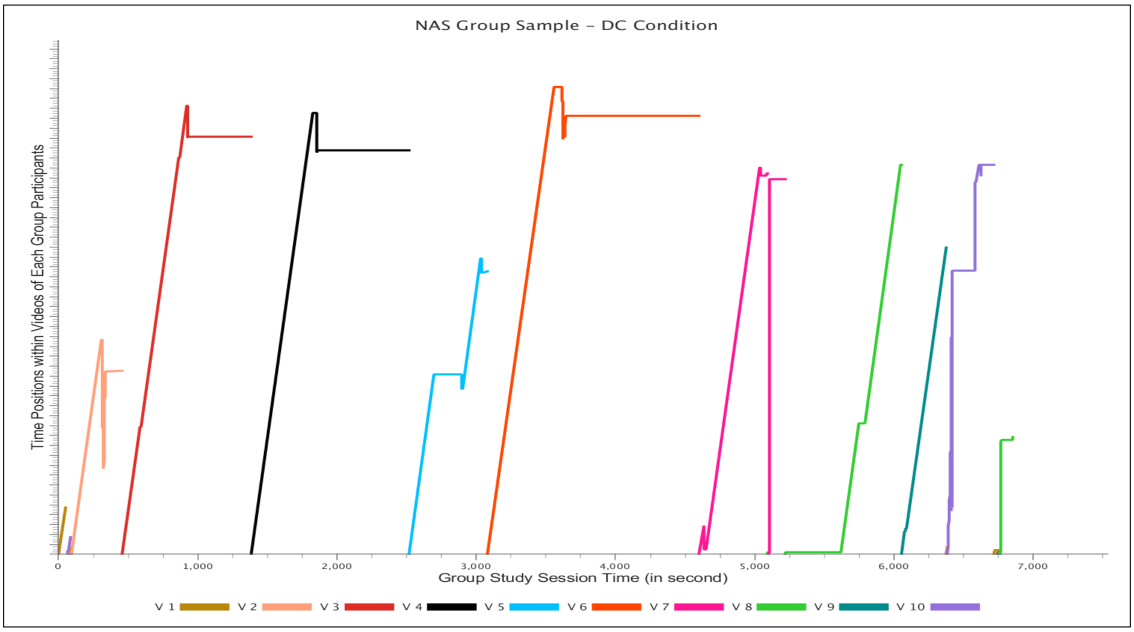

(b) A group in DC condition

Figure 8.2 – *Video navigation plots* - An example of navigation plots corresponding to one experiment session (week) for **CC** and **DC** conditions. The horizontal axis represents the session time (in seconds), and the vertical axis corresponds to the video-time position of a group (in seconds). Each video was assigned a different color for each session. It is important to note that the number of videos were different for every week of the user study.

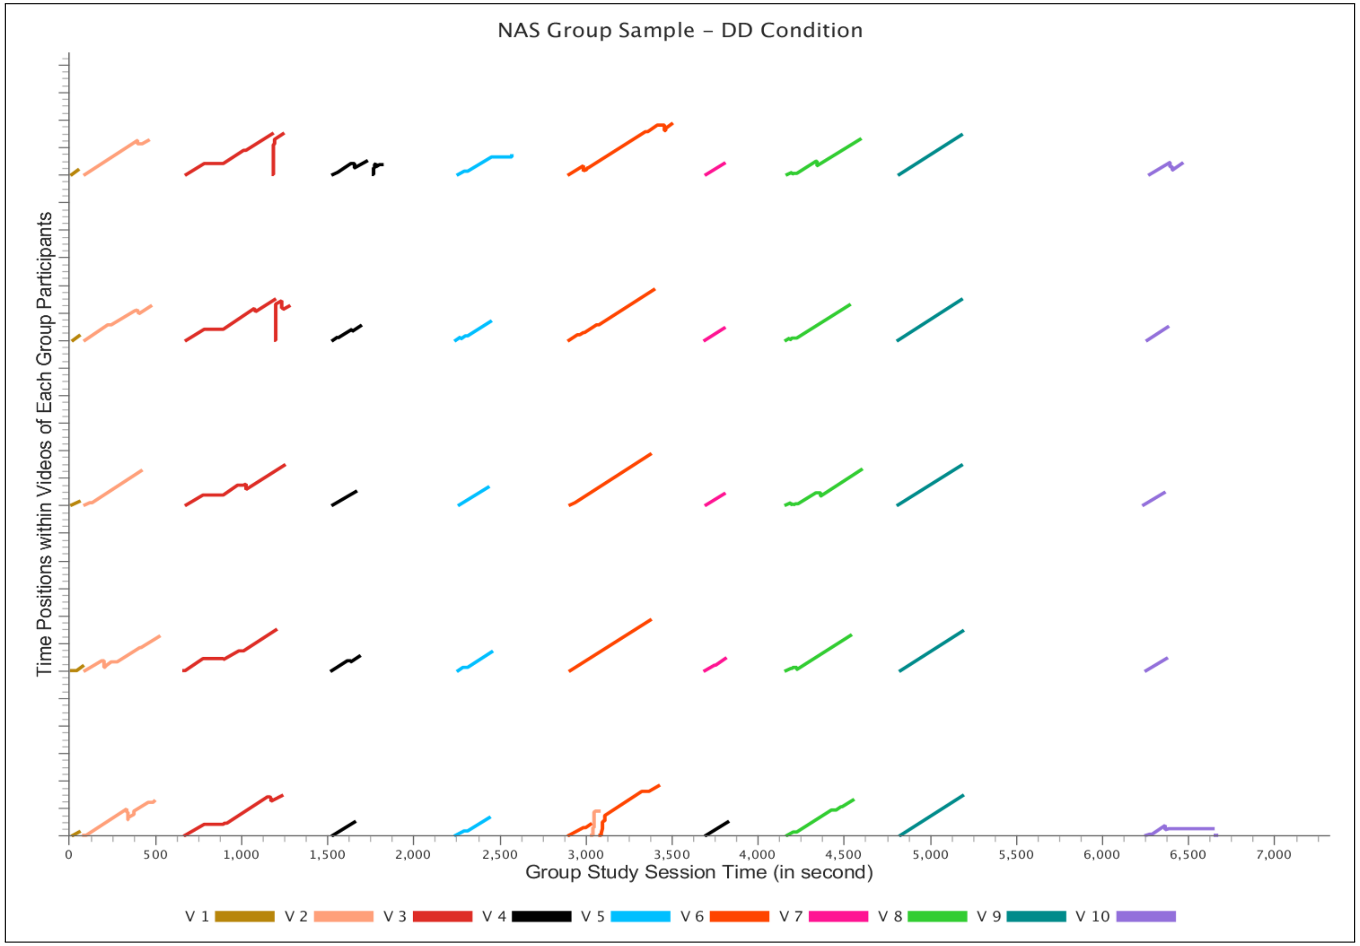

Figure 8.3 – *Video navigation plots* - An example of navigation plots corresponding to one experiment session (week) for **DD** condition. The horizontal axis represents the session time (in seconds), and the vertical axis is a compound axis consisting of video-time positions (in seconds) for each viewer. Each row of plots represent a different viewer in the study-group, and therefore there is a separate video-time (vertical) axis for each viewer. Each video was assigned a different color for each session. It is important to note that the number of videos were different for every week of the user study.

The horizontal axis in the video navigation plots corresponds to the session time, whereas the vertical axis represents the temporal position within a video. Both these axes display time in seconds. In CC and DC condition the video was synchronized for each learner, therefore the vertical axis represents the position within the video, as shown in Figure 8.2. However, in DD condition learners watched the video lectures independently, therefore the vertical axis was compounded to include the temporal video-position for each learner as shown in Figure 8.3. The video navigation plots are visually different for DD condition (see Figure 8.3) as compared to CC and DC condition (see Figures 8.2a & 8.2b), because in DD condition the navigational patterns were plotted for each group member, who were represented as a separate row in the plot. The different colored line segments in the navigational plots represent different videos that the groups watched during a single experiment session. In addition, these videos were also of different length and covered different topics.

A straight line segment with a positive slope in the video navigational plots corresponds to a video watching behavior where group members (or individuals) watch this video without interacting with it; i.e. from the beginning to its end, the video was played without any interruptions. A horizontal line segment (see Figure 8.2b) indicates a *pause*; jitters depict *fast forwards/backwards* (or seek action) within the video. The *breaks* (or gaps) between two videos (evident in Figures 8.3 and 8.2a) represent the time periods when group members were discussing or solving exercises after finishing a video and before starting another. In other words, these breaks show that during these times there was no video activity. During this study, although the session duration was quite long, the participants did not ask for any other breaks, therefore the plots show the complete picture of a video watching session.

Some groups in the DD condition showed navigational patterns which were synchronized for each group member as shown in Figure 8.3. It is interesting to see that the group members started and completed each video at approximately the same time even though they were free to watch videos at their own pace, with breaks in between two videos meant for discussions and problem solving. This indicates that belongingness to a study group encouraged the learners to synchronize their video watching behavior while favoring discussions between videos over watching lectures at one's own pace. This concept of *synchronicity* and its influence on group discussions was the research question that was studied by a colleague with whom the study was jointly conducted. Therefore, I will not discuss about synchronicity and would rather focus on the interactions with the video content across the three conditions.

Types of Video Events

In order to analyze the different kind of video events (pause within a video, break after a video was finished, seek forward or backward) that occurred across the three conditions we performed a Chi-square test. The results showed that the frequency of different video event was significantly influenced by the condition ($\chi^2(2)$=7.29, p<.01). There were significantly more than expected breaks after the video was finished in the CC and DD condition. Whereas, there were significantly more than expected pauses within the video in the DC condition and vice versa. This demonstrated that study groups where each group member had her own mouse paused and discussed whenever they encountered a difficult part in the video lecture. However, in the CC condition where there was just one video controller, the group members might have hesitated to pause during the video lecture and hence discussed after the video was over. The same was true for the DD condition where each group member watched the lectures on their respective tablets, and some groups might have synchronized their video watching behavior (an example shown in Figure 8.3) to prioritize discussions.

Considering the seek (or jumps) within the video there was no significant difference in the CC and DD condition. However, significantly less than expected fast forwards (or forward seeks) were observed in the DC condition, signifying that groups in the DC condition were watching the complete videos and were not skipping part of the video lectures. This might suggest that the groups were highly engaged with the video lectures in the DC condition.

Regarding the duration of pauses across the three conditions, we performed a linear mixed-effect analysis with repeated measures by grouping the log data by experiment session and group. Our results showed that group members in the DD condition (individual watching) paused for the least amount of time (108.17 seconds, $t(1438)=4.56$, $p<.01$). Whereas, group members in the CC condition significantly paused for the most duration (169.64 seconds, $t(29)=2.35$, $p=0.02$). The DC condition came second with the average pause duration of 142.15 seconds, however this difference was not found to be significant.

Furthermore, analyzing the frequency of pauses across the three conditions, we observed that study groups were significantly pausing twice more often in the DC (multi-mice video player) condition as compared to CC and DD condition ($\chi^2(2)=11.43$, $p<.01$) as shown in Figure 8.4b. However, no significant difference was found between the DD (individual watching) and CC (collaborative watching with one video controller) condition. This result signifies that the CC condition is similar to the DD condition. On the other hand, study groups in the DC condition could be interacting more often with the video because each group member had her own video controller, and this might have moderated the hesitation among the group members to reach out for the video controller and to interact with the video.

Time Spent On Videos

Different interactions with the video lectures contribute to the total time spent while watching videos. Therefore, we devised a measure called *time-spent-on-video index* (TSOVI) to quantify the interactivity with the video lectures. TSOVI was defined as the ratio between the total time spent on all the videos during a weekly session to the total length of video lectures that was watched. The possible values for TSOVI could be any positive real number that is greater than 1.0. As both pausing and rewinding the video content increases the overall time spent on the videos, a value of 1.0 represented that all the videos were watched exactly once without any pauses, rewinding or re-watching. A higher value would indicate that additional time was spent on the videos. The indices were computed for each group in each weekly session,

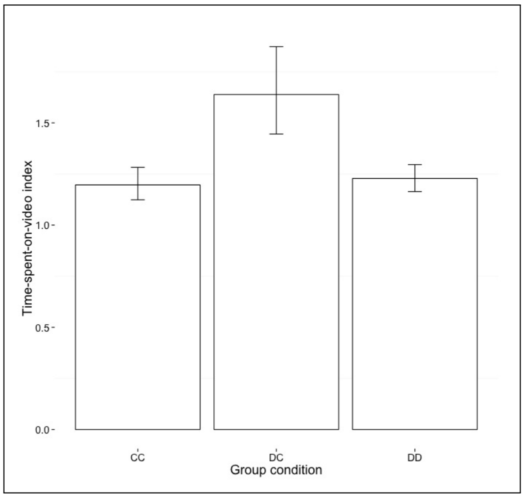

(a) Time-spent-on-video Index (TSOVI)

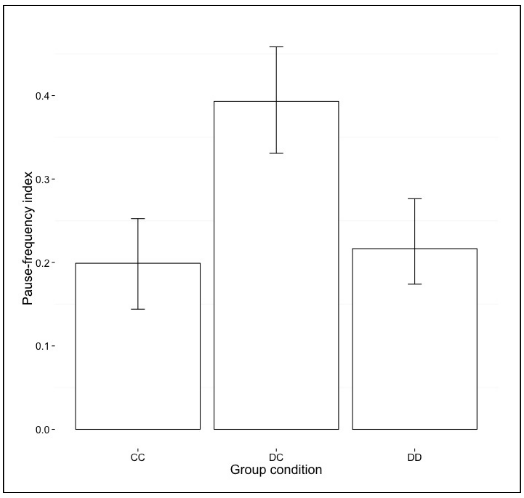

(b) Frequency of Pauses

Figure 8.4 – Bar charts with confidence intervals for time-spent-on-video index (TSOVI) and pause frequency across the three experimental conditions.

and took into account all the video lectures for that specific week during computation while excluding the unwatched content from the computations. In the DD condition, the index values for each group member were averaged to compute a single value for the whole group.

Fitting a linear mixed effect model demonstrated that on average groups spent twice more time on videos in the DC condition as compared to both the CC and DD conditions ($F(2,6)=4.53$, $p=.06$) as shown in Figure 8.4a. However, no significant differences were observed in between CC and DD conditions. One possible explanation for the significant difference observed in the DC condition could be the presence of multiple mice for group members to interact with the video content. Multiple video controllers in the DC condition enabled group members to pause, discuss, and replay the difficult parts of the video without any social pressure or hesitation, which might have been present in the CC and DC condition. Due to these in-video pauses the group members might have spent more time in the DC condition. On the other hand, groups in DD and CC conditions might have preferred to discuss only after finishing a video, this might also explain that the TSOVI was observed to be approximately similar for these conditions as shown in Figure 8.4a.

Video Play Ratio

Finally, in order to analyze the total video seconds that were watched by the study groups (without skipping) as compared to the total duration of the video lectures, we computed the play ratio. *Play ratio* was defined as the ratio of the number of video seconds watched to the total length of the video in seconds. The results showed that study groups in the DC condition (multi-mice video player) had the highest play ratio as they watched 84.67% of the total video duration (76.3% for CC condition and 76.4% for DD condition). This difference was also found to be marginally significant ($t(29)=1.88$, $p=0.07$). However, no statistical difference in play ratio was found between the CC and DD condition.

These observations concerning frequency of pauses, time spent on videos, and the total amount of video content watched without skipping, validate our second hypothesis (refer to **Hypothesis 2** in Section 8.3) that a SDG might increase the interactivity of the groups with the video content. Therefore, the DC condition can be regarded as the most interactive video-watching condition. However, we can still not assume that the DC condition was the most optimal condition for collaborative video watching because we did not measure differences in learning gains of the group members across the three conditions.

8.6 Discussion and Limitations

We conducted an exploratory longitudinal user study to investigate the collaborative watching of MOOC video lectures within a study group. The user-study participants were presented with three different collaborative video watching configurations. Our observations showed that watching MOOC lectures in a study group was widely accepted and appreciated by the learners. The study group members reported on increased levels of engagement and attentiveness. In addition, learning within a study-group rendered the learning experience with difficult courses a very pleasing experience for the learners, and assisted the group members to validate each other's understanding of the topic, thus leading to a shared knowledge model. This finding signifies the importance of leveraging the social aspects of learning with MOOCs, as it might have implications towards the design of flipped courses within the universities as well as to promote learning in under-developed and developing countries where there is scarcity of educational infrastructure.

An examination of the three video watching conditions showed that the DC condition (centralized display with multiple video controllers) was the most interactive condition. The study group members paused and discussed more often while watching the videos, as well as skipped a smaller fraction of the total video length (higher play ratio). In addition, study groups belonging to the DC condition spent twice more time watching MOOC videos as compared to the other two conditions. Higher interactivity in the DC condition as compared to the CC condition (centralized display with a single control) can be explained by the fact that each group member had her own video controller (the mouse). Therefore, the group members might hesitate less to reach out in the shared space and acquire explicit control of the video controller, as was the case with the CC condition. On the other hand the study group members in the DD condition were wearing headphones, and therefore an individual might also hesitate to grab the attention of others in order to initiate a discussion. This might suggest that the cost of initiating a discussion during the difficult parts of the video lectures might be related to the availability and access of video controllers, and the awareness of other's video watching status in case of distributed video-watching situations (DD condition). Also, the interactivity with the video lectures demonstrated that this cost of initiating discussions might be least for the DC condition as compared to the CC and DD condition.

These observations are in line with the benefits of single display groupware in peer tutoring and peer-learning as identified by Stewart et al. [1999]. In addition, the increased engagement and enjoyment of collaborators with the activity was also observed in previous SDG research such as Inkpen et al. [1999], Scott et al. [2003], and Moraveji et al. [2009]. SDG such as the one used in DC condition, also resulted into increased awareness of the learning activity amongst the group members.

On the other hand, the study groups in the DD (individual watching) condition skipped more often and discussed for the least amount of time because upon finishing with the video, an individual had to wait for other to finish, before they can start discussing. Also, the fact that a group member is waiting for others to finish the video might introduce some social pressure for the viewers to finish their video faster, and thus pressing them to skip some part of the video. Furthermore, the pressure to remain synchronized while watching videos might indicate towards fewer interactions with the video player. The observation of less interactivity with the video player as compared to DC condition makes the DD and the CC condition similar. Therefore, in terms of interactivity and group discussions the DC condition (centralized display with multiple video controllers) can be considered better than the other two conditions.

These observations also provide design guidelines for innovative tools and features to supplement the ongoing collaborative video watching within the study group, such as the ability to collaboratively annotate specific parts of the video lectures, which could later be used as group notes or exam preparation material for learners. In addition, implementation of voting tools that account for differences in opinions while solving exercises and quizzes collaboratively.

Although this user-study is one of the first to investigate the study-group behavior while watching MOOC video lectures, still there is a drawback to this study as we did not measure the learning gains of participants. The exercises that participants solved during the weekly sessions showed a ceiling effect in the scores. In addition, we did not conduct any pre- and post-tests to measure the effectiveness of various video watching configurations on learning.

8.7 Conclusion

The study presented in this chapter regarded the collaborative video-watching behavior in study groups learning via a MOOC. This longitudinal study was designed as a different use-case for single display groupware, where we examined the effects of SDG on study groups' interactivity with the video content, and compared it with two conditions: individual watching of MOOC lectures, and collaborative viewing with a centralized display and video controller.

Our findings demonstrated that groups that watched MOOC videos on a SDG interacted more with the lectures by pausing and discussing whenever the group members encountered a difficulty with the lecture content. In addition, the groups also spent significantly more time watching videos and skipped less video content with the SDG. These findings suggested on the effectiveness of SDG with respect to video interactivity. The participants also reported an engaging and satisfying experience while learning with MOOCs in a collaborative manner. Our findings hold special relevance for the design of pedagogical learning scenarios in MOOCs, that might be pertinent for developing countries where each learner cannot afford a computer and high-speed internet to learn with MOOCs.

9 Utilizing Social Information

Previously in Chapter 6, we established the relationship between different attributes of social information and the different dimensions in the Meier-Spada-Rummel rating scheme for assessing the quality of collaboration [Meier et al., 2007]. Our findings are graphically represented in Figure 9.1. We observed that transactions and self-transactions were found to be negatively correlated to two aspects of dialogue management (or common ground) - acknowledgements, and content follow-ups. In addition, transactions were also found to be correlated to the quality of joint information processing represented by information pooling (for detailed analysis refer to Section 6.5.2 in Chapter 6).

The negative correlations of transactions and self-transactions with the quantity of acknowledgements and content follow-ups, indicate their role as cross-modal feedbacks for the actual dialogues happening amongst the group members. In other words, group members' interactions with the artifacts shared over the workspace could be used as acknowledgements to ground the ongoing discussions, and thus reducing the need for verbal acknowledgements. Both the quantity of acknowledgements and the content follow-ups might also indicate the state of mutual understanding amongst the group members via grounding. Acknowledgements refer to the low-level verbal feedback to collaborators, indicating that the utterance was perceived by an individual. On the other hand, content follow-up might refer to a higher-level verbal feedback where the listener synthesizes an utterance in response to the speaker's verbal presentation. This might signify that the former utterance of the speaker was understood and a response was formulated based on the understanding of the listener. Both the quantity of acknowledgements and content follow-ups can be assumed to correspond to different levels of grounding (or mutual understanding) as presented by Dillenbourg and Traum [2006]. Therefore, we can conclude that some aspects of social information can be used to predict the episodes of poorly or well established common ground amongst the collaborators.

In this chapter, we will consider the practical utilization of the social information to model the mutual understanding of group members during collaboration. This modeling has design implications for a feedback or awareness system that notifies groups about their quality of mutual understanding.

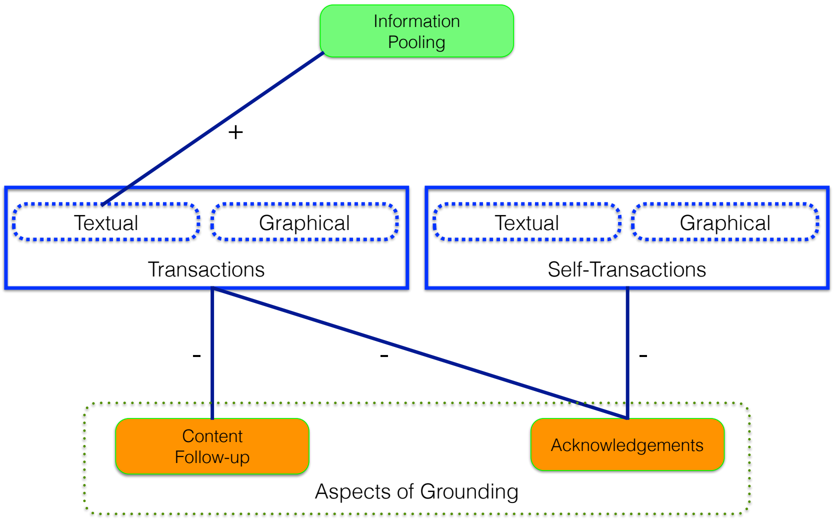

Figure 9.1 – The significant relationships between two attributes of social information - *transactions* and *self-transactions* and the three ratings based on Meier-Spada-Rummel rating scheme [Meier et al., 2007] - *information pooling*, *content follow-up*, and *acknowledgements*. The direction (correlated or negatively correlated) of the relationship is represented by the + or - sign. These findings were previously presented in Section 6.5.2 in Chapter 6.

9.1 Data Used for Modeling

In order to model mutual understanding within the group, and to design a prediction system we used the log-file data as well as the manually coded collaboration episodes, from our second user-study presented in Section 6.3 in Chapter 6. The experiment data corresponds to 6 groups of 4-5 participants each that completed three different tasks - brainstorming, decision making, and problem solving. The ratings were given to episodes of 5 minute time windows by two experimenters separately while watching the video recordings of the experiment sessions. These ratings were based on the collaboration assessment scheme proposed by Meier et al. [2007] (for a detailed description of the various dimensions of coding, refer to Section 6.5 in Chapter 6). The value of each rating belonged to a discreet range of [1, 5], where 1 indicated a poor quality, and 5 referred to a good quality of a specific collaboration aspect. The log-file data was synchronized with the video recordings of each session. These synchronized log files were then parsed to generate time-series data of the various attributes of social information in 5 minute time windows, similar to the duration of each coded episode.

We used these hand-coded ratings and the time-series of the social information for our further analyses, and to design a feedback (or alert) system for notifying groups about the poorly- or well-grounded episodes.

9.2 Collaboration Episode Quality

While coding the quality of different collaboration aspects, we assigned ratings to five dimensions concerning the communication and joint information processing aspect of collaboration. These five dimensions were - *quality of grounding*, *acknowledgements*, *content follow-up*, *level of elaboration*, and *information pooling*. Furthermore, in order to have a single succinct measure of collaboration quality for each episode, we combined these 5 ratings by summing them all up. This sum of the ratings was termed as *collaboration episode quality* (CEQ). Figure 9.2 shows the distribution of CEQ as a histogram.

Instead of attempting to predict the quality of the whole meeting session, we decided to predict the episodes of poor collaboration quality, because it is easier to predict episodes and such a feedback might be relevant for an efficient group dynamics. In addition, depending on the task requirements and the group structure, the group transitions from phases of low to high collaborative activity, and also transitions from a loosely-coupled to tightly-coupled group and vice versa [Tang et al., 2006]. Therefore, it is more reasonable to notify groups about the episodes where the collaboration quality is poor, so that the group members can respond appropriately by regulating their behavior which might be beneficial for the overall task performance.

The episodes with lower values of CEQ might refer to poor quality of mutual understanding and joint information processing. Conversely, the higher quality episodes are denoted by

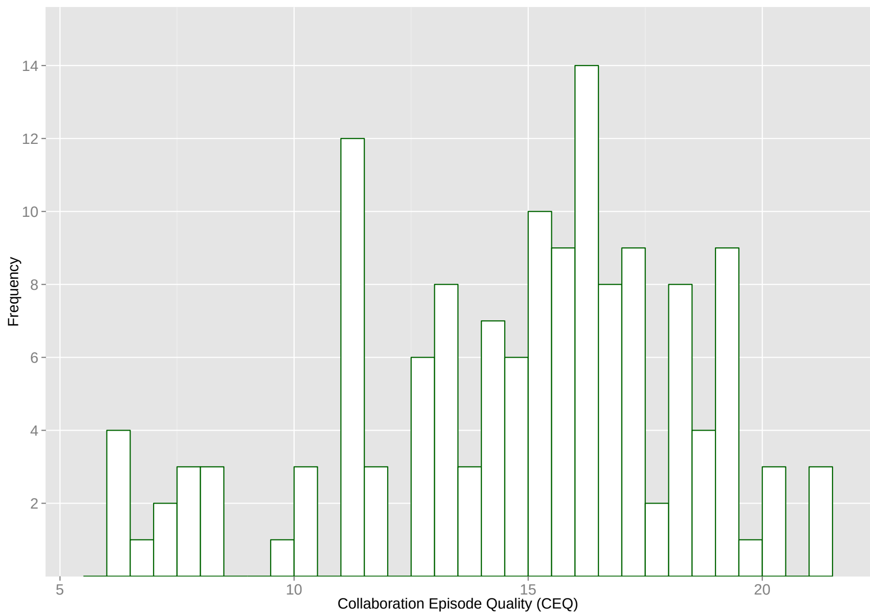

Figure 9.2 – A histogram showing the distribution of sum of rating scores. In order to assign a single concise value to each episode (time window of 5 minutes), we took the sum of all the rating dimensions.

the higher values of CEQ. Considering the distribution of CEQ in Figure 9.2, we can use the two spikes (or peaks with frequency of 12 & 14) to split the data for the purpose of prediction. The episodes with values less than the first peak (from the left) can be labeled as poor quality episodes, and the episodes with values higher than the second peak (from the left) can be labeled as good quality episodes. Therefore, we decided to take the first quartile (or 25th percentile) of the sample, and labeled these episodes as *poor quality episodes*. The rest of the data was labeled as *better quality episodes*. We made the choice of considering only the first quartile as the poor quality episodes because we believe that a prediction of these episodes is more relevant and required immediate attention of the group as compared to the better quality episodes. Besides, a higher threshold value, such as the median, might increase the prediction accuracy at the cost of higher false positives and a risk of overfitting the model.

9.2.1 Logistic Regression

The labeled data was used for modeling in order to establish the relationship of different aspects of social information (transactions, self-transactions, and epistemic & cosmetic actions) with the CEQ. We used logistic regression to model the effects of the social information and the quality of episodes because the dependent variable was binary and represented dichotomous outcomes - poor or good quality. In total, there were 142 episodes of 5 minutes duration, which corresponded to the collaboration of 6 groups completing 3 different tasks.

The results of the logistic regression are summarized in Table 9.1. Our findings suggested that only total transactions and transactions on graphical objects significantly correlated with the

Independent Variable	Coefficients	Std. Error	z-ratio	p-value	Odds Ratio
Constant	1.149	0.298	3.85	<.001	3.157
Total Transactions	0.491	0.157	3.13	.002	1.634
Graphical Transactions	-0.655	0.221	-2.96	.003	0.519
Textual Transactions	-	-	-	-	-
Total Self-Transactions	-0.009	0.012	-0.79	.428	0.990
Textual Self-Transactions	-0.008	0.008	-1.09	.277	0.991
Graphical Self-Transactions	-	-	-	-	-
Epistemic Actions	0.005	0.010	0.49	.621	1.005
Cosmetic Actions	0.011	0.011	0.99	.320	1.011
	Model χ^2 = 15.73, p = .02				
	Pseudo R^2 (Cox and Snell) = 0.156				
	n = 142				

Table 9.1 – The results of the logistic regression modeling the effects of various attributes of social information on the quality of episodes denoted by CEQ (collaboration episode quality). There were in total 142 episodes of 5 minutes duration corresponding to 6 groups completing 3 different tasks.

quality of a collaboration episode, and can be used as predictors. However, as transactions on the graphical objects form a part of the total transactions, this effect is not additive. In addition, the model was found to be statistically significant (χ^2 = 15.73, p=.02).

9.2.2 Limitations with CEQ

Basing our predictions of the quality of collaborative episode on the sum of ratings that correspond to different aspects of collaboration might have a few limitations. Therefore the aforementioned modeling might not be the optimal one. The following list summarizes the limitations with this approach.

- Computing the sum of these ratings does not account for the degree of importance of one rating over another. In other words, both acknowledgements and content follow-up refer to the quality of mutual understanding. However, they both correspond to different levels of grounding according to the levels presented by Dillenbourg and Traum [2006]. Plain acknowledgements can be considered to be belonging to the lower level of grounding signifying the perception of an utterance. On the other hand, content follow-up might represent a higher level of grounding where the listener expresses her understanding as a well formulated utterance in response to the presentation by a listener. In this approach, we did not account for this difference in the level of mutual understanding by computing a weighted sum of different ratings, it is difficult to justify the choice of weights.

- Next, Our assumption of a linear relationship between the predictors (different attributes of social information) and the response (CEQ) might not necessarily be true.

- Some aspects of collaboration that are rated across different dimensions might also be correlated, therefore summing up all the ratings is not vary informative.

9.3 Principal Component Analysis

In order to overcome the limitations that were inherent in our previous approach for predicting poor quality collaboration episodes by summing up all the rating variables, we performed Principal Component Analysis (PCA) on our collaboration ratings. PCA is a dimensionality reduction technique which maximizes the amount of non-redundant information by taking into account the correlations in between different input variables. The score of a principal component is a weighted sum of the variables, where the weights are optimally chosen to account for the explanation of the maximum variability in the data. Consequently, this score can be regarded as a more meaningful summary of the variables as compared to our previous approach.

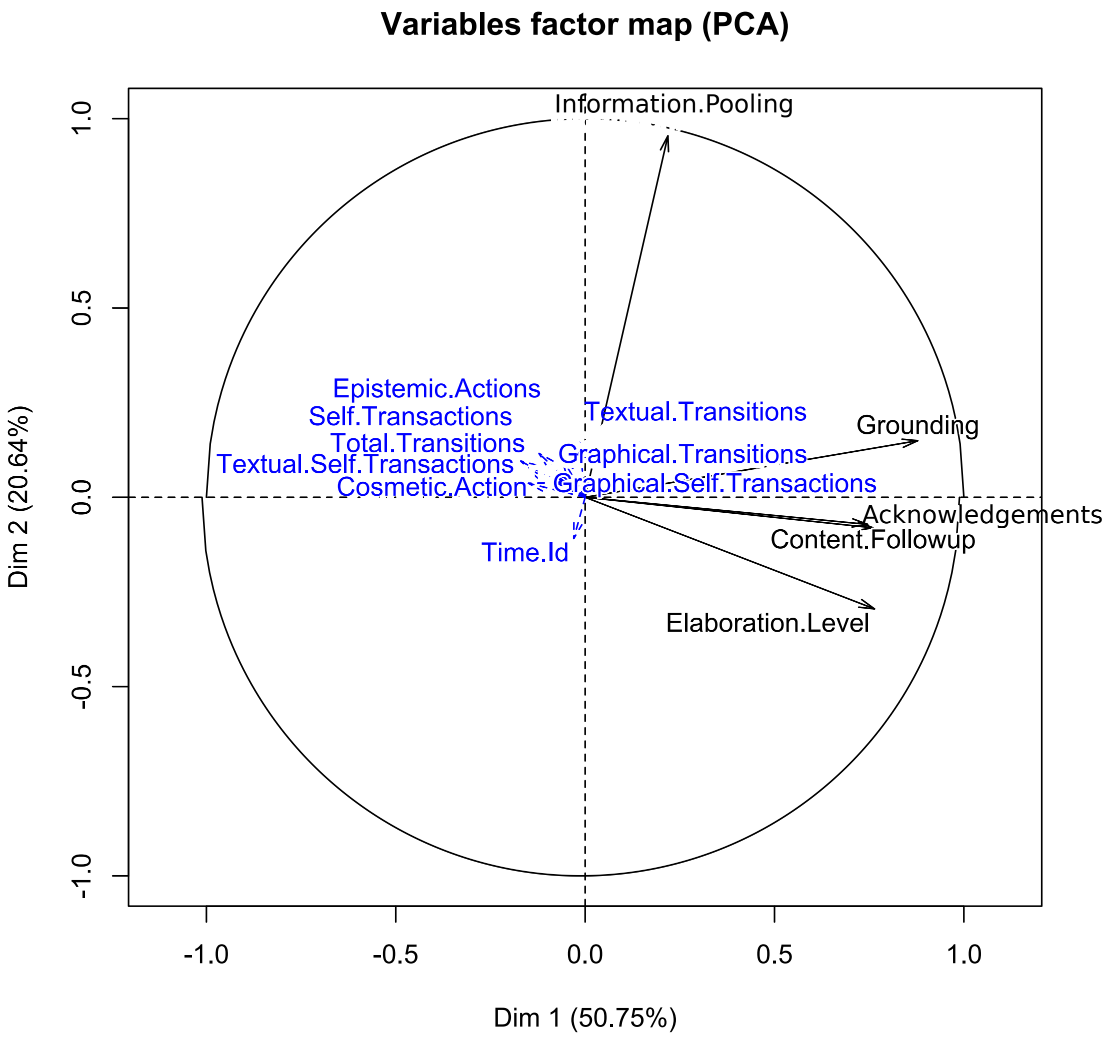

Figure 9.3 – The *Variable Scatterplot* obtained after performing the Principal Component Analysis (PCA) on all the rating variables. The plot shows the first two principal components along with the percentage variability explained by each component. The black arrows represent the different rating variables that constitute the PCA. The blue arrows are the various attributes of social information that were added as supplementary variables in the analysis; meaning that PCA was not performed with these variables but still they are plotted in order to show their relationship with the principal components.

The variable scatterplot in Figure 9.3 shows the results of the PCA where the rating variables (information pooling, grounding, acknowledgements, content follow-up, and elaboration level) are plotted along the first two principal components. Besides the ratings, we also added the different attributes of social information as supplementary variables to the analysis. The supplementary variables are not part of the PCA, but are added to the variable scatterplot in order to visualize the relationship between the supplementary variables and the principal

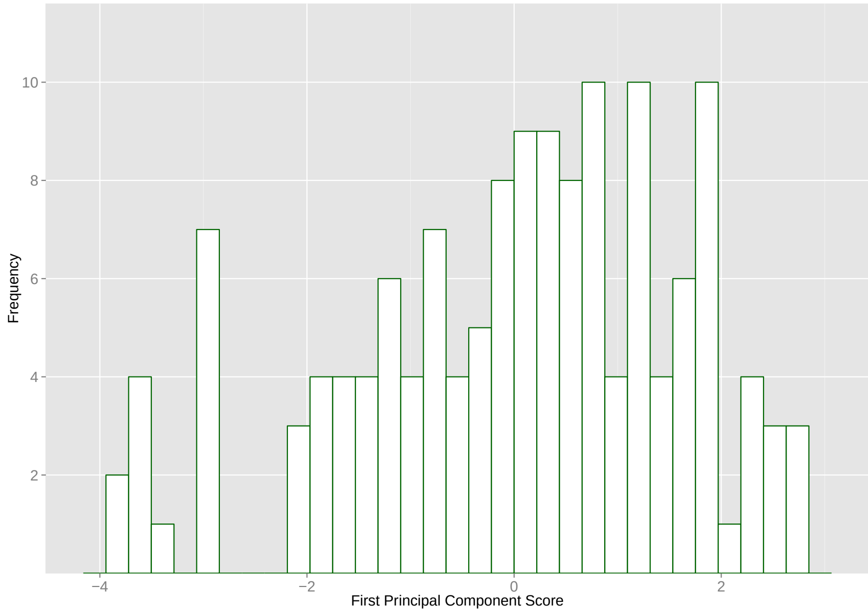

Figure 9.4 – A histogram showing the distribution of episodic scores corresponding to the first principal component, obtained after performing PCA on the rating data.

components (as shown by the blue-dashed arrows in Figure 9.3). The first two principal components jointly explain 71% of the variability in the data. The first principal component (represented as the horizontal axis in Figure 9.3) accounts for 50.75% of variability, whereas the second principal component (represented as the vertical axis in Figure 9.3) accounts for 20.64%.

Based on the projection of the rating variables along the two principal components as shown in Figure 9.3 we see that the ratings that concern the mutual understanding amongst the group members (grounding, acknowledgements, content follow-up, and elaboration level) are correlated to the first principal component. On the other hand, information pooling is strongly correlated to the second principal component. This might signify that the first principal component stands for the *mutual understanding* between the group members, and the second principal component denotes the *joint information processing*. As mutual understanding amongst collaborators is an aspect that is closely related to the quality of communication, we can use the score of each episode along the first principal component to predict the quality of communication during collaboration. We will not regard the second principal component in our prediction because only one rating variable (information pooling) was found to be strongly correlated to it, and this dimension accounts for only 21% of variability in the data. Therefore, in the next section (Section 9.4) we will design a classifier to predict the episodes with poor- or good quality of mutual understanding.

9.4 Mutual Understanding Feedback

The score for each collaboration episode of 5 minutes duration along the first principal component (refer to Section 9.3), was used to train a classifier, and predict the quality of mutual understanding. Figure 9.4 shows the distribution of the score along the first principal component as a histogram. We labeled the episodes that lie in the first quartile (25th percentile) as the sample denoting poorly-grounded episodes (or poor quality of mutual understanding). In addition, the episodes with a score higher than the median (50th percentile) were labeled as the good-quality episodes. The rest of the episodes that lie in the second quartile were ignored, and were not part of the data used for training a classifier. For the purpose of training our classifier we chose just the poor- and good-quality episodes.

We believe that the prediction of poorly grounded episodes is crucial and should be brought to the attention of the collaborators immediately. The awareness about the normal or good quality of mutual understanding is also useful for the group, but this awareness might not contribute towards a significant regulation of group's behavior, as compared to the awareness of poor quality episodes. In addition, an increase in granularity of episode quality by having more classes such as poor-, normal-, or good-quality of mutual understanding might result in a decline in the prediction accuracy. Hence, we will proceed with the prediction of poor quality episodes of mutual understanding.

We used the *random forest* machine learning method by Breiman [2001] for the purpose of predicting the poorly-grounded episodes. Random forest approach of classification constructs multiple decision trees during the training phase, and the output class is the mode of the outputs of individual decision trees within the forest. Random forest approach has a few advantages as compared to other classification methods that are especially relevant for our prediction data. Random forest is a robust approach when the number of observations in each class are significantly different. In our case, the poor-quality episodes constitute only 25 percentile of the total sample size, hence the size of this class is significantly smaller than the good-quality episodes. In addition, random forest approach trains well on a relatively smaller training set, and unlike decision trees they are less prone to overfitting.

Using our dataset of 107 labeled episodes (after removing the episodes in the second quartile) we trained our classifier, and performed cross-validation to test the prediction accuracy. As a response variable (the variable that is predicted) we used the labeled score of ratings along the first principal component. We used all the attributes of social information such as transactions, self-transactions, epistemic & cosmetic actions, task type, and the time window (if the time window belonged to the start, middle, or the end of the meeting session) as the predictors, to predict the quality of episodes. The task type and the time window were included as the predictors along with social information because these variables also introduce a degree of variability in the group dynamics and the way groups structure their activities.

We achieved a Kappa of 0.29, computed over the confusion matrix. The Kappa-value suggests that the classification accuracy is moderate. Using this classification, we can design an alarm (or notification) that notifies the groups when the quality of mutual understanding is poor. This feedback can be used by the group members to regulate their behavior in the direction of better grounding and communication. However, the prediction accuracy of the classifier could be improved over time as groups constantly use the system and provide feedback to the system about the prediction outcomes. In other words, when the group receives an alarm about the poor quality of mutual understanding from the groupware (or shared workspace), the group members can notify back to the system if the feedback was correct. In this way, the shared workspace can fine-tune itself to predict with higher accuracy if the quality of mutual understanding is poor or not, simply based on the social information contained in the group members' interactions with the shared workspace.

9.5 Design Implications & Practical Utility

In Section 9.4, we presented a classification system that is capable of predicting poor quality episodes in terms of mutual understanding and grounding. This classifier was designed based on the social information that a shared workspace can collect and interpret from the varied interactions performed by the group members while sharing and manipulating the content. Such a feedback system provides us with design implications for an alarm (or an alerting) system that notifies the group immediately about the episodes with poor quality of mutual understanding.

Next step would be to ask a question: *Who should get this feedback?* We can envision two use-cases or application scenarios for the presentation of this feedback. The first use-case is the most obvious, where the group members themselves receive the feedback from the shared workspace whenever an episode with poor quality is encountered by the system. Based on this feedback, the group members can regulate their behavior accordingly if they are trained to interpret the meaning of the alarm. Second use-case scenario could be collaborative learning or problem solving in classrooms or recitation sessions with an instructor or the teacher. In this case, the feedback can be given to the teacher instead of the students, indicating that the students are not collaborating well. Upon receiving a warning, the teacher can decide to check the status of the group and take appropriate action. In this case, the shared workspace acts as an awareness tool (for the instructors) similar to Lantern [Alavi et al., 2009], which allows teaching assistants and teachers to monitor the state of each collaborating group. Also, in certain learning scenarios where learners are performing a collaborative activity such as creating concept-maps, the shared workspace can be used by the teacher as an orchestration tool [Dillenbourg et al., 2011], to monitor the quality of each group's understanding, and subsequently assist the groups in need of help.

9.6 Conclusion

What is the practical utility of social information? We tried to answer this question in this chapter. We designed a classifier that can predict the episodes with poor quality of mutual understanding or grounding amongst the group members. We specifically regarded the prediction of poor quality collaborative episodes instead of designing a group mirror, because these episodes are of immediate importance as compared to the good quality episodes, and the group might benefit by regulating their behavior during these episodes. Also, in a previous study (refer to Chapter 7) we established a relationship between the social information and the performance of group in the concept-map task. Therefore, the significance of this feedback is increased, as a repair action on part of the group members might also influence the overall performance.

Our findings have important design implications for the design of shared workspaces and single display groupware (SDG), which can collect and use the interactions with the shared content as a diagnosis for the quality of mutual understanding between the group members. Such a coupling between a SDG and their role as awareness tools might enhance their effectiveness and influence on group dynamics.

10 General Discussion & Conclusion

In this chapter, we will summarize the research work presented in this thesis, along with its limitations and contributions to relevant domains.

10.1 Summary of the Results

We conducted several experiments to assess the quality of collaborative processes, through the analysis of group members' interactions with the shared artifacts. We developed *MeetHub* and used it as the groupware that equipped group members with a single shared workspace. Here is a list containing a brief description of our main findings:

1. **Complementarity of Input Devices:** A comparison between two single display groupware with different input configurations, multiple mice & keyboards versus multiple pen, showed no difference in the effect of input devices on the task outcome. A difference, however, was observed in the type of shared representations produced by the groups. These observed differences in the type of shared representations (sentential versus schematic) was influenced primarily by the affordance of the input device, and the usability of various input devices in collaboration was observed to be complementary.

2. **Transition in Groupware Usage:** The groups where an individual took up the role of a leader exhibited a transition in SDG usage dynamics from a multi-user (or simultaneous use) mode to a single-user mode, during the convergent phase of the task. Such a transition enabled collaborators to coordinate the task activity efficiently, especially during the decision-making phase by yielding the control of their respective input devices so that only one group member interacts with the shared workspace.

3. **Latent Social Information:** Shifting from the paradigm of focusing on the individuals and interactions between individuals to study collaborative processes, we adopted a different viewpoint by focusing on the human-artifact interaction in the domain of shared workspaces. Our studies showed that, the interactions with the shared artifacts contain

latent social information that is representative of the state of mutual understanding between the group members. We extracted and interpreted several attributes of this social information such as transactions, self-transactions, division-of-labor, artifact ownership, and the nature of interaction - epistemic or cosmetic.

4. **Cross-modal Acknowledgements:** Transactions and Self-Transactions are two attributes of social information that act as cross-modal acknowledgements for ongoing conversations between the group members. This finding signifies that interactions with artifacts that contain information related to the verbal discussions, decrease the needed effort to establish a common ground.

5. **Effects on Task Performance:** Transactions, epistemic actions, and joint ownership of artifacts were found to be the contributing factors for scoring well in the concept-mapping task. In addition, a sense of strong ownership of artifacts degraded the performance of the group during the activity.

6. **Role of Activity Balance:** The sense of joint-ownership of artifacts was observed to be related to the balanced activity of the group members. In other words, when group members mirrored each other's actions over the shared workspace (or when their individual participation was balanced), the group had a higher proportion of jointly-owned artifacts.

7. **Predicting Poorly Grounded Episodes:** The latent social information was found to be predictive of the collaboration episodes with poor quality of mutual understanding or grounding. This finding signifies the practical utility of social information to generate alerts.

8. **Effects of SDG in Collaborative MOOC Watching:** Groups that used a SDG (video player with a video controller for each group member) to watch MOOC video lectures exhibited twice more interactivity with the video content, skipped less video content, and spent more time on the video lectures and discussions. This comparison was made with the conditions where group members either watched the lectures individually, or on a single display with only a single video controller.

10.2 Reviewing the Initial Research Questions

In this section, we review our findings in regard to our general research questions presented in Chapter 4.

10.2.1 Role of different Input Modalities

Driven by affordances, varied input devices are appropriate for specific kinds of interactions, and therefore they might be effective only for a small set of task activities. This was the motivation behind our comparison of the two input configurations for collaboration with the Single Display Groupware (SDG). Our findings showed differences across the two input configurations of SDG (Mouse & Keyboard (MK) versus Pen & Paper (PP)) in the number of graphical and textual elements that were created and shared by the group members. Groups in the MK condition created more content in the sentential form. Whereas more diagrammatic forms were created in the PP configuration.

The nature of the shared representations produced across the two conditions provide the empirical evidence that an input device might be effective for a specific task demand. This is reflected by the abundance of one kind of shared representation (sentential versus diagrammatic) over another across the two conditions. As a keyboard is a faster, efficient, and robust means of entering text, groups in the MK condition shared more text as compared to the PP condition. Conversely, a pen is a more natural means of drawing and sketching, and this resulted into more graphical elements in the PP condition than the MK condition. The participants' perception of the effort required to create and manipulate artifacts (both graphical and textual) further supports the aforementioned findings. Higher effort was perceived to create text in the PP condition, while using a pen, and therefore less textual artifacts were created in this condition. Similarly, in the MK configuration, participants reported on perceiving higher effort to create graphical artifacts while using a mouse. Furthermore, the diagrammatic forms of shared representations are identified as more effective in expressing and interpreting information as compared to the sentential forms of representations [Scaife and Rogers, 1996], one can expect that groups in the PP configuration are more likely to perform better than the groups in the MK condition. However, we observed no such influence of the input devices on the group performance.

Our findings emphasized the role of input device affordances in supporting one task activity effectively over another. Additionally, during meetings, creation and sharing of artifacts is a supporting activity facilitating the communication of ideas amongst the collaborators, as well as storing relevant information for future references. Collaborators might decide to create diverse representations of their shared knowledge depending on the task requirements and the individual knowledge of the group members. Therefore, equipping collaborators with only one kind of input device might limit the type of shared representations produced by the group, and in turn decrease the usability of the SDG.

The apparent complementarity of different input modalities provided us with crucial design guidelines, that lead to the development of our second prototype of *MeetHub*, where each collaborator was equipped with a mouse & keyboard, as well as a tablet. Provision of such an ecology of input devices enables users to choose the appropriate channel of interaction with the shared workspace. This choice may be governed by the mapping made by the user opportunistically between the perceived affordances of the device and the representation demanded by the task. Further, the decision to equip collaborators with tablets was made following our observations in the PP condition, where the lack of visual feedback of others' activity on paper lead to an alteration in the natural behavior of group members while writing on paper. Instead of looking at the paper while writing, the participants were looking at the wall-mounted public display to avoid content overlaps. Therefore in MeetHub 2.0, the displays on all the tablets were synchronized with the public display in real-time, and interactions with any display were echoed on others immediately. Additional design guidelines for the development and deployment of SDG are elicited in Section 10.4.1.

10.2.2 Enhanced Interactivity via Multiple Input Devices

Previous studies by Inkpen et al. [1999], Scott et al. [2003], and Alcoholado et al. [2012] have demonstrated increased engagement with the activity, increased activeness and interactivity while using a SDG in collaboration scenarios involving children (from small groups to classrooms). In one of our user-studies comparing the different configurations for collocated watching of MOOC video lectures, we observed similar increased interactivity with the video content in the condition where study-groups were equipped with a SDG. Spontaneous study-groups participated in our longitudinal user-study consisting of three video-watching configurations: *a*) individual video watching with distributed video control and display (DD), *b*) collective video watching with centralized video control and display (CC), and *c*) collective video watching with a SDG comprising of distributed video controls and centralized display (DC).

We analyzed the groups' navigation through the lecture videos. Our findings showed that study-groups in the DC condition (with a SDG) paused the video-lectures twice more often as compared to the other two conditions. Groups with a single mouse for each member (DC condition) paused the video lectures whenever someone was faced with a difficulty, followed by discussions amongst the learners. Whereas in the other two conditions (CC and DD) group members watched the lectures until the end before proceeding with the discussions regarding the content of the video. In addition, study-groups in the DC condition spent twice more time watching lectures as compared to the CC and DD condition. Also, less video content was skipped in the condition with a SDG as compared to the other two conditions.

Our results demonstrated that the video-watching configuration with a SDG (DC condition) was the most interactive condition, where group members paused and discussed more often while watching videos. Higher interactivity in the presence of a SDG as compared to the other two video-watching configurations can be explained by the fact that each member had her own video controller (mouse). This might suggest that the cost of pausing and initiating a discussion was much less in the presence of a SDG. The easier access to the video controllers might moderate the hesitation amongst the group members to interact with the video content. In other words, the availability of an input device for each user in a SDG reduced the distance between the *intent* to ask a question and the *action* required to accomplish it (i.e. pausing), which in turn reduced the hesitation to interact with the video and discuss.

10.2.3 Role of Latent Social Information

Varied interactions with shared workspace have been shown to contain latent social information which act as a proxy for the visible collaborative processes such as dialogues and gestures. In one of our user-studies, we extracted and identified the nature of various attributes of social information. In addition, we analyzed the relationships between these attributes across different task types and presence (or absence) of roles. Our findings showed differences in these attributes across different task types, which in turn hasve implications towards the automatic detection of task type by the groupware. However, we did not pursue the goal of automatic task detection in this dissertation. Further, the groups where individuals played specific roles exhibited a higher degree of transactions as compared to groups without roles. Through social information we also showed the different instances of division-of-labor (defined by Jermann [2004]) that were adopted by the groups during meeting sessions.

Apart from these differences across tasks and roles, we demonstrated the effects of social information (transactions, epistemic actions, and joint ownership of artifacts) on group performance in closed-ended tasks, where the performance can be easily evaluated and the task does not contain multiple valid solutions, such as a learning task. This finding has implications for predicting the performance of groups based on their interactions alone with the shared content. In addition, we demonstrated that a particular choice of division-of-labor might result into differences in proportion of jointly- or strongly-owned artifacts. A role-based or concurrent-editing kind of division-of-labor resulted into higher proportion of jointly-owned artifacts, whereas a task-based division-of-labor lead to contrary results with higher proportion of strongly-owned artifacts.

Finally, we studied how the social information reflects the state of mutual understanding between the collaborators, and also explored the practical utility of social information. Analyzing the relationship between social information and hand-coded dimensions corresponding to different aspects of collaboration revealed that some attributes of social information are related to the state of mutual understanding amongst collaborators. During group discussions, Transactions and Self-Transactions served as cross-modal acknowledgements, and due to this cross-modal switching where actions on shared workspace replaced the verbal utterances, the frequency of verbal feedback and content follow-up went down. Verbal feedbacks and content follow-up denote the different levels of grounding, and therefore we concluded that we can use these attributes of social information to model the state of grounding within groups. Leveraging this relationship between social information and grounding, we developed an alert (or feedback) system that is capable of predicting the collaborative episodes with a poor quality of grounding. This feedback might enable group members to regulate their behavior upon receiving an immediate alert from the groupware. In addition, such a prediction system has potential applications in collaborative learning scenarios, where the alerting system is used by a teacher to identify and take reparative actions with the teams that are undergoing episodes of poor mutual understanding.

10.3 Contributions

In this section, we summarize the contributions of this dissertations and position them inside the relevant research domains.

10.3.1 Meeting Technology - MeetHub

Single Display Groupware (SDG) as a meeting technology has been demonstrated to have a significant influence on the quality of collaboration, but mostly involving children in educational scenarios. By means of this dissertation, we emphasized on the meeting tasks in organizational settings such as brainstorming, decision making, and problem solving. We designed and developed *MeetHub* - a meeting technology that enabled collaborators to simultaneously interact with the shared artifacts over a shared workspace. The concurrent interaction with the shared workspace was enabled via multiple input devices - one for each meeting participant.

The design and development of MeetHub underwent an iterative process where we firstly evaluated the usefulness of a SDG in decision making meetings, as well as analyzed the effects of different input devices (mice & keyboards versus pens & paper) on the collaboration in Chapter 5. Our findings suggested that driven by the affordances, the effects of various input devices are complementary with respect to the nature of artifacts that can be created, and the interactions with the artifacts that are feasible. Inspired by this finding and the associated design implications, we re-designed the final version of MeetHub that incorporated two-way interaction with the shared workspace. Firstly, each user was equipped with a mouse and a keyboard that can be used to interact with the workspace on a wall-mounted display. Secondly, each user was also equipped with a tablet and a stylus to interact with the shared artifacts. The workspace across all the displays (all tablets and the central wall-mounted display) was synchronized in a way that a change in the state of content by any user was reflected immediately on other displays. The provision of two-way interaction with the workspace extended the range of activities that can be supported by the groupware, and the group members can opportunistically map the task demands to the affordances of the input device that is most suitable for the activity.

Besides enabling group members to co-create content during collaboration, our primary use of MeetHub was to collect interactional data. It enabled us to log every fine-grained interaction with the shared artifacts during the meeting sessions, and allowed us to conduct extensive analysis on the collected interaction data.

Our focus during the design of MeetHub was not just centered on the development of the groupware application. We also developed a Software Development Kit (SDK) (software framework) that enables developers to quickly prototype SDG applications. Unlike the pre-existing SDG software frameworks, our framework is capable in identifying and parsing the input streams from a wide range of USB HID (Human Interface Device) and is not constrained to the identification of any one kind of input device. In addition, the SDG comprises of a library of simple widgets and user controls that are SDG enabled, i.e. they are responsive to the input events corresponding to the multiple input devices. One example of rapid development that we presented in this dissertation was the development of the video player application that is responsive to multiple mice, and was used as one of the condition during the user-study with collaborative MOOC watching.

10.3.2 Artifact Centered Methodology

Assessment of collaboration and its processes is a several-decade old research problem that has been addressed differently in the domain of CSCW and Social Signal Processing. The past approaches that addressed this problem have focused on the collaborators and the interactions between them.

During the course of our analyses, we adopted a different methodology to assess and examine collaborative processes. Our methodology takes into account the interactions of group members with the multiple artifacts shared over the workspace. In our methodology, instead of analyzing the direct interactions between individuals, we treat the shared workspace as an auxiliary communication channel and an external memory for group discussions, which is also coupled to the activity and interactions amongst collaborators. This new methodology entails the development of groupware such as SDG where each individual is capable of unconstrained access to the workspace, and where individuals don't have to compete to gain access to the workspace.

The socio-linguistic perspective on interactions with artifacts by Snyder [2013, 2014] is thematically close to our research methodology. However, Snyder has addressed an important preliminary question of *why do individuals create and interact with artifacts?*. Her analysis is qualitative in nature and does not consider the implications of these interactions on collaborative processes. This makes our quantitative analysis of interactions with artifacts unique in a sense that we bridge the gap between the aforementioned preliminary question and the the role and contribution of these interactions in different collaborative processes.

Unlike collaboration assessment, which is regarded as a time-consuming process (mainly due to the efforts required to code behavior and audio transcriptions), our artifact-centered analysis has implications for real-time assessment and analysis of collaborative processes, as presented in Chapter 9. Such an analysis is based on the recording and interpretation of interaction logs which can be done in real time, and innovative tools can be deployed to provide timely feedback to the group members about the group processes.

10.3.3 Dual Nature of a Shared Workspace

Our findings have reinforced the *dual nature* of groupware or shared workspaces. We demonstrated that shared workspace is not just a mediator of communication or a medium for co-constructing shared knowledge, but also it plays a crucial role of analyzing and assessing group behavior and providing process feedback to the group members at relevant times. In a way, our findings have bridged the gap between supporting collaborative activity, and tools that assist group members to regulate their behavior. This amalgamation of a shared workspace and a collaboration assessment tool might render the groupware highly efficient, and pave the way for increased penetration of groupware in organizations.

10.4 Implications for Design

10.4.1 In Organizations

The reinforced dual nature of groupware as described before in Section 10.3.3 can be particularly beneficial for organizational meetings. Designing such a groupware would require development of additional tools and functionalities that cater to the organizational work culture. For instance, presence of a private working space, where individuals can refine ideas before sharing them over the public workspace would be a necessity. Additionally, groups should be provided with the functionality to load documents (such as PDFs, presentations, etc.), annotate them, and easily transfer them back to the meeting participants after the meeting as a record of meeting proceedings and finalized products. Such tools might prove to be beneficial for decision making and product-design review meetings, where group members can simultaneously comment over the document in review. Further, simultaneous interaction and annotation would require the document viewers to be SDG enabled, which can be achieved easily through the Software Development Kit (SDK) that we developed for designing new Graphical User Interface (GUI) features that can respond to multiple input devices. Further, new conflict management strategies need to be employed in the groupware to resolve conflicts arising due to collaborators attempting to manipulate the same artifact at the same time, and other interactional conflicts arising due to the multiplicity of input devices.

How should the varied kinds of input devices be distributed amongst the collaborators? One possible answer to this question would be to address the individual preferences of the collaborators. For instance, a user might prefer to use a tablet over a mouse & keyboard. However, considering individuals' preferences might lead to scenarios where all the group members are using the same kind of input devices, and therefore the advantages of complementarity of input devices cannot be effectively leveraged. Further, mice & keyboards afford for the interactions with the wall-mounted public display, whereas tablets afford for personalized usage. Hence, there should be a balance between both kinds of input devices. Bring your own device (BYOD) is a policy supported by many organizations, where individuals can bring their own tablets (with the groupware application installed) to the meeting, and connect their respective devices with the groupware server. In collocated settings, groups can be provided with a few mice and keyboards to complement the tablet devices, and can be shared amongst the collaborators on demand. In distributed settings, the collaborators can use their respective tablets with an additional audio-video conferencing functionality.

Visualizing latent social information to inform collaborators about the quality of their mutual understanding should also be considered as an important design decision. One way to visualize this information is to implement a notification (or alarm) system, similar to the one presented in Section 9.4, Chapter 9. Another way of visualizing this information is to design a group mirror similar to the one designed by Streng et al. [2009] and the one by Bachour et al. [2010]. However, the efficacy of any visualization can only be examined by conducting user-studies and studying the effects of specific visualizations on group processes, which forms part of the future work.

10.4.2 In Educational Scenarios

Although the efficacy of a SDG has already been established in educational scenarios, especially in classroom education, by various studies conducted by Szewkis et al. [2011], Alcoholado et al. [2012], and Caballero et al. [2014]; still we demonstrated in Chapter 8 that a SDG can be beneficial in collocated watching of MOOC video lectures. Our user-study provides implications for the design of MOOC pedagogical scenarios involving collaborative learning, which might have relevance for flipped classrooms in universities as well as under-developed countries with scarce computational resources. MOOC designers can implement activities and tasks that leverage the benefits of collaborative learning, where several learners in geographical proximity can register for the course as a study-group. For example, collaborative quizzes can be designed in a way that students can solve an exercise individually, followed by discussions and voting amongst the study-group participants about the most suitable answer.

10.5 Limitations & Future Work

In this section, we discuss some of of the limitations with our meeting technology as well as the methodology used to address the relevant questions raised in this dissertation.

10.5.1 Exploratory Nature of Some Studies

Three out of four of our studies were exploratory in nature where we had limited control over experimental parameters or variables. The exploratory nature of our studies allowed us to consider the role of many factors in our experiment design, as well as the freedom to collect and analyze a lot of data. As the quantitative research analyzing interactions with shared artifacts has not been well defined in research literature and previous research work, such an approach was advantageous for us to regard the relationships between variables from different perspectives. However, this exploratory approach added some noise to the conclusions that can be drawn, and might lead to weak relationships between variables.

10.5.2 Lack of a Private Workspace

During our user-studies, MeetHub provided group members with a public workspace, which was also synchronized with all the connected tablets, so that group members had the same view to the shared artifacts. This decision was made to extensively investigate the interactions with a SDG, where group members can flexibly adopt the most suitable interaction style with the shared artifacts, which is also coherent with the task requirements. However, providing group members with a private workspace might have afforded for a specific group dynamics of task division, and does reduce the variability in group dynamics one can study.

10.5.3 Ecological Validity

All of the user-study participants were engineering students from our university. In addition, the user-studies were conducted in the lab setting with the presence of an experimenter. The different task types that the participants completed were short-duration non-recurrent tasks, and cannot be treated similar to project-management tasks, which usually have a fixed agenda and are recurrent in nature. These factors might introduce a bias in our study as well as raise questions about the ecological validity of our findings.

The shared workspace that was used in our studies enabled simultaneous interaction through multiple input devices. Therefore our findings should be carefully interpreted as they might not exactly scale for other kinds of shared workspaces such as physical whiteboards or SMART boards.

10.5.4 Future Work: Exploiting the Role of Context

During the course of our analyses we regarded the interactive actions that were performed by the group members over the shared workspace, and the social information contained within them. However, we believe that the relevance of social information would be even more pronounced if the context is taken into account. The contextual information that can supplement the social information can be related to the task as well as the semantic relationships between artifacts. We mentioned earlier that the observed differences in social information across varied task types can be utilized to automatically infer the ongoing task. This information might prove useful as groups might adopt a different dynamics for different kinds of task, and this knowledge can be used by the groupware to adjust (or fine-tune) the parameters for the accurate prediction of episode quality.

Semantic relationships between different shared artifacts such as proximity, annotations, links, etc., might indicate the importance of a certain piece of information as well as signify the role of these information pieces towards the final solution. A thorough examination of these relationships between artifacts can be used to classify the actions performed over these artifacts as more significant than other actions, thus expanding the set of attributes of social information.

10.6 Final Words

This dissertation presented the outcome of a few years of research work during which *a*) we identified the relevance of collaborators' interactions with the shared workspace in informing about the ongoing collaborative processes, *b*) developed a single display groupware, *MeetHub*, to support co-construction of knowledge within group, *c*) designed and conducted user-studies with *MeetHub* to extract the latent social information and analyze its influence on collaborative processes, and finally, *d*) used the social information to predict the quality of mutual understanding in collaborative episodes.

Finally, we would like to articulate that this dissertation has stressed on the microscopic viewpoint towards the examination of collaborative processes. Our approach has deviated from the direct observation of interactions between collaborators, and rather focused on their interactions with the shared artifacts over the workspace. Creation and sharing of artifacts has always been a part of established work culture, and we have strived to leverage this practice in an analytical way to assess the collaboration. However, the nature of these interactions has diverse and complex inter-dependencies with other factors such as the nature of the task, inter-personal relationships, roles and hierarchies within organizations, and accepted work practices. Amongst these factors, the presented research has addressed some, while others can be investigated in the future studies. In summary, we have attempted to answer the question *What do the group's interaction with a shared workspace say about the ongoing collaboration?*

A First Study: Questionnaires

The following pre- and post-experiment questionnaires were given to each participant during our first study, which compared the different input configurations of a Single Display Groupware (SDG).

A.1 Pre-Experiment Questionnaire

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Group #: ______________________

Experiment #: __________________

Task Booklet: ☐ A ☐ B ☐ C ☐ D

Pre-Experiment Questionnaire

1. Your age is: ____________________
2. Gender : ☐ Female ☐ Male
3. Do you suffer from Colorblindness: ☐ No ☐ Yes, color: ______________
4. How many hours do you use a computer per day: ☐ <1 ☐ 1-3 ☐ 3-6 ☐ >6
5. Have you ever used a digital pen before? ☐ No ☐ Yes: ______________________
   What did you use it for? ____

6. What collaboration tools have you used before?

Name of tool	Used for

7. I prefer working ☐ in groups ☐ individually
8. It is best when all members of a group participate equally in the task.
   Agree ☐—☐—☐—☐—☐ Disagree

9. I see myself as a rather discreet person when it comes to group collaboration.
   Agree ☐—☐—☐—☐—☐ Disagree

10. I usually attend ☐ 0 ☐ 1-3 ☐ 3-5 ☐ 6 or more meetings in a week
    And their average duration is usually
    ☐ Less than 30min ☐ 30min - 2 hours ☐ More than 2 hours

11. What types of meetings do you usually attend?
    (fill in using: *project* meeting / *professor or teacher or boss* / *brainstorming* / *decision making* / *social* / *presentation* / *information* / _name_a_category_that_is_not_listed_)

Please turn the page to select your position with respect to the table!

1

12. Please indicate your chair position with respect to the screen and table:

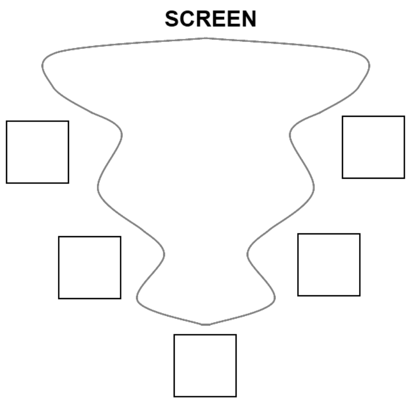

A.2 Post-Experiment Questionnaire

Post-Experiment Questionnaire

Questions regarding the given task

1. Did your team reach consensus about the murderer? ☐ Yes ☐ No

2. The time given for the task was sufficient.
   Agree ☐ ☐ ☐ ☐ ☐ Disagree

3. We all had the same opinion about the murderer.
   Agree ☐ ☐ ☐ ☐ ☐ Disagree

Questions regarding social aspects of the meeting

4. The group had a leader emerging.
   Agree ☐ ☐ ☐ ☐ ☐ Disagree

5. My contributions were taken into account by the group.
   Agree ☐ ☐ ☐ ☐ ☐ Disagree

6. I actively participated in the task.
   Agree ☐ ☐ ☐ ☐ ☐ Disagree

Questions regarding the ease of use of the tool

7. How much effort did it take to produce text with the tool?
   Very little ☐ ☐ ☐ ☐ ☐ Very much

8. How much effort did it take to produce drawings / tables with the tool?
   Very little ☐ ☐ ☐ ☐ ☐ Very much

9. How difficult was it to link information from different pages?
   Very easy ☐ ☐ ☐ ☐ ☐ Very difficult

Questions regarding coordination

10. The space available on the screen was sufficient.
    Agree ☐ ☐ ☐ ☐ ☐ Disagree

11. The participants' contents overlapped.
    Rarely ☐ ☐ ☐ ☐ ☐ Often

12. How much effort did it take to avoid overlapping of content?
    Very little ☐ ☐ ☐ ☐ ☐ Very much

Questions regarding the purpose of the tool

13. We used the tool to collect facts.

Agree ☐—☐—☐—☐—☐ Disagree

14. We used the tool to build hypotheses about suspects.

Agree ☐—☐—☐—☐—☐ Disagree

15. We used the tool to represent the story of the murder.

Agree ☐—☐—☐—☐—☐ Disagree

Open questions

16. What is your overall impression about the system? Do you find it cool / clumsy / useful / hard to manage / new / worthy / useless? Would you use such a system in meetings? In which types of meetings?

____

____

____

____

____

____

____

____

17. If you had to design such a system yourself, for a meeting, how would you design it? What features would you include? What else do you think would be appropriate to have in a meeting, before it, or after it?

____

____

____

____

____

____

____

____

B Second Study: Experimental Tasks

The following three experiment tasks were designed for our second user-study, where we extracted the different attributes of latent social information from group's interactions with the shared workspace.

B.1 Brainstorming Task

Instructions for the Thumb Problem

We don't think this is very likely to happen, but imagine for a moment what would happen if everyone born after 2013 had an extra thumb on each hand. This extra thumb will be built just as the present one is, but located on the other side of the hand. It faces inward, so that it can press against the fingers, just as the regular thumb does now. Here is a picture to help you see how it will be.

Now the question is:

What **practical benefits** and **difficulties** will arise when people start having this extra thumb? Please brainstorm in your group. You will have **TWO** phases for this task. In the first 20 minutes, you should generate as many ideas as possible **WITHOUT** judgment, criticism or evaluation, and in the last 10 minutes you will have to **JUSTIFY** your idea list and come out with a final list of the most practical ones. Free free to adjust the time allocated for each phase during your discussion with the time management tool on your iPad.

B.2 Decision Making Task

Instructions for the Energy Crisis Problem

Much of central China is enduring its worst energy crisis, with factories and residents facing power cuts as supply runs short of demand. This fast-growing country has long experienced periodic power shortages, especially in winter and summer when weather extremes boost demand for heating and cooling. Thermal power station is still the major type of power stations, and the coal fuels contribute to about three-quarters of the country's electricity generation.

China Electricity Council has published statistics about power shortfall in the past five years as illustrated in the following table:

Year	Power Shortage (beyond current installed capacity)
2007	14 million Kilowatts
2008	19 million Kilowatts
2009	22 million Kilowatts
2010	26 million Kilowatts
2011	30 million Kilowatts

Considering China's economic and population growth, it is estimated that the power shortfall will be steadily increasing for the next five years and then reach its saturation point in the year 2016. Suppose you are a group of consultants hired by China Energy Council and your task is to analyze the given statistics, estimate the power needs (only the power shortage) and design an energy plan to solve the energy crisis by the end of 2021. Please decide on the types and the numbers of power plants to be built and an estimated cost. You have to present good enough reasons both for the government and power companies. That being said, although you are not limited to a certain amount of money or a certain area in China to build power plants, you must consider environmental factors and the cost/benefit ratio.

	Installed Capacity	Construction cost	Operation cost	Selling Price
Nuclear	1-3 million Kilowatts	13000 ￥/Kilowatt	0.5 ￥/Kwh	1.3 ￥/Kwh
Wind farm	0.1-0.3 million Kilowatts	6500 ￥/Kilowatt	0.4 ￥/Kwh	1.2 ￥/Kwh
Solar Energy	0.01-0.15 million Kilowatts	9000 ￥/Kilowatt	0.6 ￥/Kwh	1.4 ￥/Kwh
Fossil fuel (coal)	0.6-4 million Kilowatts	4800 ￥/Kilowatt	0.8 ￥/Kwh	1.2 ￥/Kwh
Fossil fuel (gas)	0.5-2.3 million Kilowatts	3400 ￥/Kilowatt	0.9 ￥/Kwh	1.4 ￥/Kwh
Hydroelectric	0.5-3 million Kilowatts	3000 ￥/Kilowatt	0.7 ￥/Kwh	1.3 ￥/Kwh
Tidal	0.0005-0.004 million Kilowatts	8000 ￥/Kilowatt	0.8 ￥/Kwh	1.3 ￥/Kwh

B.3 Problem Solving Task

Instructions for the Neuroscience Task

This time you will be given 15 minutes to read a neuroscience text before starting the experiment.

Every neuron has an electrical voltage on both sides of the membrane that is called the "membrane potential". The neuron at rest (which does not transmit nerve impulses) generally has a membrane potential of about -65 mV. The membrane potential of a non-stimulated neuron is called the "resting potential". This negative potential can be explained by the fact that the interior of the neuron is negatively charged while its exterior is positively charged. Thus it is said that the neuron is polarized.

The resting potential exists only across the membrane; in other words, the liquids that are inside and outside the neuron are electrically neutral. The resting potential is generated by differences in the ionic composition of interior and exterior environments. Thus, the inside of the neuron contains a smaller concentration of sodium (Na +) and a higher concentration of potassium (K +) than the outside. In the extracellular fluid, the positive charges of sodium ions are generally balanced by chloride ions (Cl-). In the intracellular fluid, negatively charged proteins (A-) facilitate the equalization/balancing of the positive charges of potassium ions (K +).

The ionic differences arise on the one hand from the difference in ionic permeability of the membrane, and on the other hand from the operation of the sodium-potassium pump. In the resting state, the membrane is about 75 times more permeable to K + than to Na +. This resting permeability is bound to the properties of passive ion channels in the membrane.

The concentration gradients of K + and Na + ions explain their diffusion from the medium where they are most concentrated to the medium where they are least concentrated, that means towards the exterior of the neuron for the K + ions and towards the interior for the Na + ions. Furthermore, K + ions diffuse more rapidly than sodium ions. From this follows that the positive ions that diffuse outward are a little more numerous than those which diffuse inward, leaving a small surplus of negative charges inside the neuron; this phenomenon leads to an imbalance of electric charges (electrical gradient) which causes the resting potential.

As there is always a certain quantity of K + leaving the cell and a certain amount of Na + that enters it, one might think that the concentration of Na + and K + ions on both sides of the membrane will equalize, which would lead to the disappearance of their respective concentration gradients. However, this is not the case because the sodium-potassium pump exchanges (the) Na + ions from the interior with the K+ ions from the exterior of the neuron. In other words, the K + ions are pumped into the neuron at the same time as the Na + ions are released to the outside.

You have the following three tasks to complete:

1. Compare the roles of Na+ and K+ in neuro-transmission.
2. Draw a schematic neuron which illustrates the generation of resting potential.
3. Assume that you are a group of TAs for a neuron science course. Please design an assignment to check whether or not your students understand the concept illustrated in this article.

C Second Study: Questionnaires

The following experiment questionnaire was given to each participant during our second study, where we extracted the different attributes of latent social information. Each group completed three tasks (brainstorming, decision making, & problem solving), and the questionnaire was given to each participant after each task.

Post-Experiment Questionnaire

Experiment – 2

Name:
User ID:
Group:

The following questions are about content sharing in your meeting.

- Which tool did you use more for writing.
  ☐ Keyboard ☐ iPad
  ☐ Both iPad and Keyboard ☐ Neither of them (Didn't Write)

- Which tool did you use more for creating new objects, moving and deleting.
  ☐ Mouse ☐ iPad
  ☐ Both iPad and Mouse ☐ Neither of them

- Which input tool would you prefer in meetings.
  ☐ iPad ☐ Mouse and Keyboard
  ☐ Both ☐ No Preference

- The usage of Pen/Stylus was intuitive with the iPad.
  Strongly Disagree ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 Strongly Agree

- The group reached a consensus at the end of the meeting.
  Strongly Disagree ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 Strongly Agree

- I feel that my contributions were taken into account by the group.
  Strongly Disagree ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 Strongly Agree

- At which display did you look more during the experiment.
  ☐ iPad ☐ Whiteboard (Public Display)

- During the discussions, in which direction most of your gestures were made to.
  ☐ iPad ☐ Whiteboard (Public Display)
  ☐ Other participants

- I think that the meeting environment (Table, Public Display, iPads and Stylus, Mouse & Keyboards) facilitated group coordination effectively (Disagree … Agree).
  Strongly Disagree ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 Strongly Agree

The following questions are about awareness of time management in your meeting.

- In how many parts did you split your meeting ?
  ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5

- Did you observe the blinking time management notification on the public display ?
  ☐ Yes ☐ No

- What did you mostly do when time was up for one of the phases ?
  ☐ We continued to discuss regardless of bypassing the allocated time.
  ☐ We extended the discussion a little, but kept in mind that we needed to move on.
  ☐ We quickly discussed/noted down some conclusions in order to move on fast.
  ☐ We moved on immediately because there was already a consensus.
  ☐ We moved on immediately even though information or consensus was still lacking.

- For how long did you typically extend your discussion over the allocated time ?
  ☐ 0 minutes (not at all)
  ☐ At most 5 minutes
  ☐ 5 to 15 minutes
  ☐ We were not aware that we bypassed the allocated time.

- At the end of the meeting, did you feel that you managed your time proficiently?
  ☐ Yes, the pre-suggested time allocation actually corresponded to our management.
  ☐ We tended to over-discuss items but the awareness helped us keep on time track.
  ☐ We tended to finish our collaboration faster than the allocated times.
  ☐ We did not care about the time management.

- Please rank the overall utility/efficiency/ appropriateness of having a time management awareness element in your meeting.
  ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5

- Please let us know if you have any comments regarding how you addressed the time management in your meeting.

  ____

  ____

  ____

The following questions are about the searches in your meetings
The term "search suggestion" refers to the "moving rectangle" on the whiteboard which contains two image blocks and one Wikipedia block.

- The search suggestions are NOT disturbing or intrusive to the discussion (Disagree ... Agree)
  Strongly Disagree ☐ ☐ ☐ ☐ ☐ Strongly Agree
  1 2 3 4 5

- The search suggestions helped me with accomplishing the task (Disagree … Agree)
  Strongly Disagree ☐ ☐ ☐ ☐ ☐ Strongly Agree
  1 2 3 4 5

- I am satisfied with the number of suggestions for each search (Disagree … Agree)
  Strongly Disagree ☐ ☐ ☐ ☐ ☐ Strongly Agree
  1 2 3 4 5

- I think Wikipedia suggestions were more helpful than images (Disagree … Agree)
  Strongly Disagree ☐ ☐ ☐ ☐ ☐ Strongly Agree
  1 2 3 4 5

- I think the keyword suggestions extracted for websites (on the browser window) were useful (Disagree … Agree)
  Strongly Disagree ☐ ☐ ☐ ☐ ☐ Strongly Agree
  1 2 3 4 5

- From the search suggestions, list the image/Wikipedia link that you thought were useful during the discussion (if there were any), and explain why it was useful. Was the suggestion (image/wikipedia) itself useful or the web page containing the suggestion is useful?

  ____
  ____
  ____
  ____

- List keywords that you thought might be useful for making searches during the discussion (if there were any)

  ____
  ____
  ____
  ____

D Third Study: Questionnaires

The following test questionnaires were given to each participant in our third study, where we studied the effects of latent social information on group performance. Each dyad was asked to take a pre-test on the subject of Resting Membrane Potential in neuroscience. Next, the participants watched a video lecture on the same subject, followed by a collaborative concept-map activity. Finally, each participant took a post-experiment test.

D.1 Pre-Experiment Questionnaire

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Pretest

Participant ID: __________

Instructions:

1. Please answer the questions you are sure about. Please do not make random guesses.
2. Please put a cross (X) on the correct option in the answer box. If you don't mark your answer in the answer box it will be considered as an incorrect response.

1. The membrane potential of the neuron is constant.

☐ True ☐ False

2. The original cause of the resting potential is the fact that the amount of the positive ions which diffuse to the interior is slightly more than the amount of the positive ions which diffuse to the exterior.

☐ True ☐ False

3. The original cause of the resting potential is the fact that the potassium ions diffuse faster than sodium ions.

☐ True ☐ False

4. Sodium-Potassium pump brings the sodium ions in and potassium ions are expelled through the membrane.

☐ True ☐ False

5. Which of the following phenomena explains that the resting potential is negative?

a. There are more negative ions than positive ions in the liquid that is in the interior of the neuron.

☐ True ☐ False

b. The negative ions that diffuse into the interior of the neuron are more than those which diffuse outward.

☐ True ☐ False

6. What would happen if the sodium-potassium pump is artificially blocked?

a. This would lead to the disappearance of the concentration gradients of K+ and Na+ ions on either side of the membrane.

☐ True ☐ False

b. Many potassium ions would accumulate in the interior of the neuron and the neuron no longer works.

☐ True ☐ False

7. The diffusion of positive ions to the outside is faster than the diffusion of positive ions to the inside of the neuron.

☐ True ☐ False

D.2 Post-Experiment Questionnaire

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Post Test

Participant ID: ___________

Instructions:

1. Please answer the questions you are sure about. Please do not make random guesses.
2. Please put a cross (X) on the correct option in the answer box. If you don't mark your answer in the answer box it will be considered as an incorrect response.

1. The higher the concentration of Na+ ions in the interior of the neuron is, the more positive the resting potential is.

☐ True ☐ False

2. The most important cause of the resting potential is the fact that Na+ channels are highly permeable for Na+ and let many Na+ ions diffuse inside.

☐ True ☐ False

3. When the membrane is at resting potential sodium ions are attracted towards interior of the neuron due to an electrical gradient and a concentration gradient.

☐ True ☐ False

4. If the membrane was permeable only to sodium, assuming normal concentrations of ions inside and outside the membrane, the resting potential would be about +50mV.

☐ True ☐ False

5. What would happen if the sodium-potassium pump is artificially blocked?

a. The membrane would have a more positive potential than normal rest.

☐ True ☐ False

b. This would lead to a decrease in membrane potential between the inside and outside areas of the neuron.

☐ True ☐ False

6. The electric potential is equal to zero as long as the recording electrode is positioned outside of the membrane of the neuron.

☐ True ☐ False

7. At rest, the positive ions are attracted by the charges outside the membrane and the negative ions are attracted by the charges inside the membrane.

☐ True ☐ False

8. The sodium-potassium pump pumps the sodium and potassium ions in the same direction as the concentration gradient.

☐ True ☐ False

9. The higher the concentration of K+ ions outside of the neuron is, the more negative the resting potential is, all other conditions being equal.

☐ True ☐ False

E Fourth Study: Questionnaires

The following questionnaires were given to each participant in our fourth study, where we studied the collaborative MOOC video lecture watching of groups. The participants were asked to complete the pre-experiment questionnaire before the start of the experiment. However, the participants were asked to complete the post-experiment questionnaire after each weekly lecture watching session.

E.1 Pre-Experiment Questionnaire

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Color: ______________________________
Group: ______________________________
ID: ______________________________

Pre-Experiment Questionnaire

Please read through the following questions and answer them.

Section I: Personal Information

- Full Name (Capital Letters):______________________________
- Gender: Male □ Female □
- Age: ______________________________
- Study Major
  - □ Mechanical Engineering
  - □ Communication Systems
  - □ Computer Science
  - □ Physics
  - □ Management of Technology
  - □ Other (Please specify) ______________________________
- Semester: ______________________________

Secion II: Personality questions

The following questions assess your personality. For each statement, please place your opinion on the scale, ranging from "*Strongly Disagree*" to "*Strongly Agree*".

1. I consider myself as extraverted, enthusiastic.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

2. I consider myself as critical.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

3. I consider myself as dependable, self-disciplined.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

4. I consider myself as anxious, easily upset.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

5. I consider myself as open to new experiences.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

6. I consider myself as reserved, quiet.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

7. I consider myself as sympathetic, warm.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

8. I consider myself as disorganized, careless.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

9. I consider myself as calm, emotionally stable.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

10. I consider myself as conventional, uncreative.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

Section III: Familiarity with other group members

How well do you know each person in this study group? For each person, please rate the degree to which you know her/him. For your own name, please ignore the question.

- Please name your current seat
  - □ A
  - □ B
  - □ C
  - □ D
  - □ E

- I know person A very well

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- I know person B very well

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- I know person C very well

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- I know person D very well

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- I know person E very well

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

Section IV: Studying in groups

- During the semester, how often do you study as part of a group?
  - □ Never (I always study alone)
  - □ Less often than once a month
  - □ Once a month or more
  - □ Once a week or more
  - □ Once a day or more

- If you have been involved in any study groups before, please describe your typical group practices. Eg. How often you meet, for which course, with how many people, where you meet and what you usually do?

  ____
  ____
  ____
  ____

- I take a lot of notes during the lectures.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- During the semester, you prefer to take notes on which medium
  - □ Loose paper sheets
  - □ Notebooks
  - □ Textbooks
  - □ Computer/Laptop/Tablet
  - □ Others (please specify) ________________________________

- I prefer sharing my notes with my friends/colleagues.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- I prefer using notes that my friends took during semester.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- What materials do you usually use while preparing for exams (multiple choices can apply)
  - □ Lecture Slides
  - □ Notes (personal OR friend's)
  - □ Textbooks
  - □ Others (please specify) ________________________________

Section V: MOOC use and perceptions

- Before the current MOOC course, how many MOOC courses have you attempted or completed before?
  ☐ None, this is my first MOOC course
  ☐ One
  ☐ Two
  ☐ Three
  ☐ Four or more

- If this is NOT your first MOOC course, please tell us about the previous courses. For instance, which courses did you register for, did you complete them, etc.?
  ____
  ____
  ____
  ____

- I regard MOOC-based learning as being more effective and efficient than attending real classes.

	1	2	3	4	5	
Strongly disagree	☐	☐	☐	☐	☐	*Strongly agree*

Section VI: Use of iPad and other tablet devices

- Experience in using a tablet device. How long have you used a tablet device (iPad or other)? Please select the MOST suitable option from the list below.
  ☐ I have never used a tablet device.
  ☐ I have tested a tablet device a few times.
  ☐ I have used a tablet device for less than half a year.
  ☐ I have used a tablet device from half-a-year to 1 year.
  ☐ I have used a tablet device from 1 to 2 years.
  ☐ I have used a tablet device from 2 to 3 years.
  ☐ I have used a tablet device for more than 3 years.

- Access to iPad or other tablet device (Please select one or more suitable options from the list below)
  ☐ I never used to own or have access to an iPad or other tablet devices.
  ☐ I currently have and iPad in personal use.
  ☐ I currently have another kind of tablet device (eg. Android) in personal use.
  ☐ I currently have access to a shared iPad or other kind of tablet device.
  ☐ I used to have access to an iPad or other tablet device but am not currently using one.

E.2 Post-Experiment Questionnaire

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Name: ______________________________
Group: ______________________________
Session: ______________________________

Post-Experiment Questionnaire

Please read the following statements and provide us with your opinion.

I. **Video Lecture**

- I prefer watching MOOC video lectures in a group OR study group.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- I would prefer to watch the video lectures alone and discuss later (with a colleague or teacher or teaching assistant) about the parts which I did not understand.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- When I misunderstood (or did not understood) a part of the video lecture, I preferred to (multiple choices might apply):
  - □ Pause the lecture and replay the misunderstood part
  - □ Pause the lecture and ask my group members to provide an explanation
  - □ Read the concept from the textbook (or other sources on internet) later.
  - □ Ask the teacher or teaching assistants later.
  - □ Did not care (or did nothing)
  - □ Others (please specify) ______________________________

- I feel that group discussions while watching videos are beneficial for learning and facilitate better understanding of the lesson.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- When I had trouble understanding a part of the lecture, the explanations offered by other group members:
  - □ Were satisfactory and helped me understand the concept better.
  - □ Were partially satisfactory, as the group members themselves did not understand the concept better and no discussions followed.
  - □ Were partially satisfactory, but mostly resulted into discussion which helped us all to understand the concept better.
  - □ Were totally unsatisfactory and I did not agree with them.
  - □ Were confusing to me.
  - □ My group members offered NO explanations.

- I feel that the lecture (or the chapter being taught) is hard to understand.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

II. Note – Taking

- I preferred taking notes while watching MOOC session.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- During the MOOC session, I took notes because (multiple choices might apply):
  - □ I want to use these notes for my personal usage before exams.
  - □ Notes help me remember what was being said during the lecture.
  - □ Notes help me assimilate (or understand) the content better.
  - □ I took notes because the professor was writing during the MOOC session.
  - □ I took notes because my group members were taking notes.
  - □ Others (please specify) ________________________________

- If I did not take notes during the MOOC session, it was mainly because (multiple choices might apply):
  - □ I did not want to waste time taking notes as there was already little time.
  - □ My colleagues were taking notes, so I can borrow or use their notes later.
  - □ I can always re-watch the MOOC content, so I feel that there is no need to take notes.
  - □ I prefer reading the course textbook for the exam and that's why I did not take notes.
  - □ I feel that my notes are not very good as compared to my friends'.
  - □ I do not like taking notes.
  - □ Others (please specify) ________________________________

- I don't have a problem if my friends borrow my notes after the lectures.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- I feel that by sharing notes among us, while preparing for exams OR while studying, helps me understand and learn better.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

III. Input Devices

- The presence of multiple cursors on the public display was distracting or disturbing.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- When other group members were pausing (or rewinding) the video to discuss, this was frustrating for me.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- The presence of multiple cursors on the public display (video player) motivates me to participate more with regard to the interaction with the system.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- I preferred to take control of the mouse device and when my fellow group members asked to Play/Pause/Seek/interact with the video etc., I was the one doing it.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- I felt that one of the group members took responsibility to interact with the video (play, pause or replay part of the video), when someone did not understand a part of the lecture.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

- While we were discussing about the lecture or quiz, we used our mouse cursors to point (or gesture) towards a specific part of public display.

	1	2	3	4	5	
Strongly disagree	□	□	□	□	□	*Strongly agree*

Bibliography

H. S. Alavi, P. Dillenbourg, and F. Kaplan. Distributed awareness for class orchestration. In *Learning in the synergy of multiple disciplines*, pages 211–225. Springer, 2009.

C. Alcoholado, M. Nussbaum, A. Tagle, F. Gómez, F. Denardin, H. Susaeta, M. Villalta, and K. Toyama. One mouse per child: Interpersonal computer for individual arithmetic practice. *Journal of computer assisted learning*, 28(4):295–309, 2012.

T. J. Allen et al. Managing the flow of technology: Technology transfer and the dissemination of technological information with the r&d organization. *Cambridge (US)*, 1977.

R. Anson, R. Bostrom, and B. Wynne. An experiment assessing group support system and facilitator effects on meeting outcomes. *Management Science*, 41(2):189–208, Feb. 1995. ISSN 0025-1909. doi: 10.1287/mnsc.41.2.189. URL http://pubsonline.informs.org/doi/abs/10.1287/mnsc.41.2.189.

P. Antunes and C. J. Costa. Handheld cscw in the meeting environment. In *Groupware: Design, Implementation, and Use*, pages 47–60. Springer, 2002.

H. Arrow, J. E. McGrath, and J. L. Berdahl. *Small groups as complex systems: Formation, coordination, development, and adaptation.* Sage Publications, 2000.

K. Bachour. *Augmenting face-to-face collaboration with low-resolution semi-ambient feedback.* PhD thesis, ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE, 2010.

K. Bachour, F. Kaplan, and P. Dillenbourg. An interactive table for supporting participation balance in face-to-face collaborative learning. *IEEE Transactions on Learning Technologies*, 3(3):203–213, July 2010. ISSN 1939-1382. doi: 10.1109/TLT.2010.18.

R. Baeza-Yates and J. A. Pino. A first step to formally evaluate collaborative work. In *Proceedings of the international ACM SIGGROUP conference on Supporting group work: the integration challenge*, pages 56–60. ACM, 1997.

G. Bafoutsou and G. Mentzas. Review and functional classification of collaborative systems. *International Journal of Information Management*, 22(4):281–305, Aug. 2002. ISSN 0268-4012. doi: 10.1016/S0268-4012(02)00013-0. URL http://www.sciencedirect.com/science/article/pii/S0268401202000130.

B. Bederson and J. Hourcade. Architecture and implementation of a java package for multiple input devices (mid). Technical report, HCIL Technical Report, 1999.

M. M. Bekker, J. S. Olson, and G. M. Olson. Analysis of gestures in face-to-face design teams provides guidance for how to use groupware in design. In *Proceedings of the 1st Conference on Designing Interactive Systems: Processes, Practices, Methods, & Techniques*, DIS '95, pages 157–166, New York, NY, USA, 1995. ACM. ISBN 0-89791-673-5. doi: 10.1145/225434.225452. URL http://doi.acm.org/10.1145/225434.225452.

T. Bergstrom and K. Karahalios. Conversation clock: Visualizing audio patterns in co-located groups. In *40th Annual Hawaii International Conference on System Sciences, 2007. HICSS 2007*, pages 78–78, Jan. 2007a. doi: 10.1109/HICSS.2007.151.

T. Bergstrom and K. Karahalios. Conversation votes: enabling anonymous cues. In *CHI'07 Extended Abstracts on Human Factors in Computing Systems*, pages 2279–2284. ACM, 2007b.

T. Bergstrom and K. Karahalios. Seeing more: visualizing audio cues. In *Human-Computer Interaction–INTERACT 2007*, pages 29–42. Springer, 2007c.

E. A. Bier and S. Freeman. Mmm: A user interface architecture for shared editors on a single screen. In *Proceedings of the 4th annual ACM symposium on User interface software and technology*, pages 79–86. ACM, 1991.

Q. Bonnard, H. Verma, F. Kaplan, and P. Dillenbourg. Paper interfaces for learning geometry. In *21st Century Learning for 21st Century Skills*, pages 37–50. Springer, 2012.

N. S. Borenstein. Computational mail as network infrastructure for computer-supported cooperative work. In *Proceedings of the 1992 ACM Conference on Computer-supported Cooperative Work*, CSCW '92, pages 67–74, New York, NY, USA, 1992. ACM. ISBN 0-89791-542-9. doi: 10.1145/143457.143463. URL http://doi.acm.org/10.1145/143457.143463.

L. Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.

D. Caballero, S. A. Van Riesen, S. Álvarez, M. Nussbaum, T. De Jong, and C. Alario-Hoyos. The effects of whole-class interactive instruction with single display groupware for triangles. *Computers & education*, 70:203–211, 2014.

J. J. Cadiz, A. Balachandran, E. Sanocki, A. Gupta, J. Grudin, and G. Jancke. Distance learning through distributed collaborative video viewing. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work*, pages 135–144. ACM, 2000.

A. J. Cañas, G. Hill, R. Carff, N. Suri, J. Lott, T. Eskridge, G. Gomez, M. Arroyo, and R. Carvajal. Cmaptools: A knowledge modeling and sharing environment. In *Concept maps: Theory, methodology, technology. Proceedings of the first international conference on concept mapping*, volume 1, pages 125–133, 2004.

J. M. Carroll, D. C. Neale, P. L. Isenhour, M. B. Rosson, and D. S. McCrickard. Notification and awareness: synchronizing task-oriented collaborative activity. *International Journal of Human-Computer Studies*, 58(5):605–632, 2003.

M. Cherubini, M.-A. Nüssli, and P. Dillenbourg. Deixis and gaze in collaborative work at a distance (over a shared map): A computational model to detect misunderstandings. In *Proceedings of the 2008 Symposium on Eye Tracking Research & Applications*, ETRA '08, pages 173–180, New York, NY, USA, 2008. ACM. ISBN 978-1-59593-982-1. doi: 10.1145/1344471.1344515. URL http://doi.acm.org/10.1145/1344471.1344515.

P. Chiu, A. Kapuskar, S. Reitmeier, and L. Wilcox. Notelook: Taking notes in meetings with digital video and ink. In *Proceedings of the seventh ACM international conference on Multimedia (Part 1)*, pages 149–158. ACM, 1999.

A. Clark. *Being there: Putting brain, body, and world together again.* MIT press, 1997.

H. H. Clark. Language use and language users. *Handbook of social psychology*, 2:179–232, 1985.

H. H. Clark and S. E. Brennan. Grounding in communication. *Perspectives on socially shared cognition*, 13(1991):127–149, 1991.

H. H. Clark and E. F. Schaefer. Contributing to discourse. *Cognitive science*, 13(2):259–294, 1989.

S. Converse. Shared mental models in expert team decision making. *Individual and group decision making: Current issues*, 221, 1993.

J. N. Cummings and S. Kiesler. Who collaborates successfully?: Prior experience reduces collaboration barriers in distributed interdisciplinary research. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work*, CSCW '08, pages 437–446, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-007-4. doi: 10.1145/1460563.1460633. URL http://doi.acm.org/10.1145/1460563.1460633.

R. C. Davis, J. A. Landay, V. Chen, J. Huang, R. B. Lee, F. C. Li, J. Lin, C. B. Morrey III, B. Schleimer, M. N. Price, et al. Notepals: Lightweight note sharing by the group, for the group. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 338–345. ACM, 1999.

G. DeSanctis and R. B. Gallupe. A foundation for the study of group decision support systems. *Management Science*, 33(5):589–609, May 1987. ISSN 0025-1909. doi: 10.1287/mnsc.33.5.589. URL http://pubsonline.informs.org/doi/abs/10.1287/mnsc.33.5.589.

A. Dielmann and S. Renals. Automatic meeting segmentation using dynamic bayesian networks. *Multimedia, IEEE Transactions on*, 9(1):25–36, 2007.

P. Dillenbourg and D. Traum. Does a shared screen make a shared solution? In *Proceedings of the 1999 conference on Computer support for collaborative learning*, page 14. International Society of the Learning Sciences, 1999.

P. Dillenbourg and D. Traum. Sharing solutions: Persistence and grounding in multimodal collaborative problem solving. *Journal of the Learning Sciences*, 15(1):121–151, Jan. 2006. ISSN 1050-8406. doi: 10.1207/s15327809jls1501_9. URL http://dx.doi.org/10.1207/s15327809jls1501_9.

P. Dillenbourg, M. J. Baker, A. Blaye, and C. O'Malley. The evolution of research on collaborative learning. *Learning in Humans and Machine: Towards an interdisciplinary learning science.*, pages 189–211, 1995.

P. Dillenbourg, G. Zufferey, H. Alavi, P. Jermann, S. Do-Lenh, Q. Bonnard, S. Cuendet, and F. Kaplan. Classroom orchestration: The third circle of usability. *CSCL2011 Proceedings*, 1: 510–517, 2011.

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, pages 149–158, Swinton, UK, UK, 2009. British Computer Society. URL http://dl.acm.org/citation.cfm?id=1671011.1671028.

P. Dourish. Extending awareness beyond synchronous collaboration. *Position Paper for CHI*, 97:249–256, 1997.

P. Dourish and V. Bellotti. Awareness and coordination in shared workspaces. In *Proceedings of the 1992 ACM Conference on Computer-supported Cooperative Work*, CSCW '92, pages 107–114, New York, NY, USA, 1992. ACM. ISBN 0-89791-542-9. doi: 10.1145/143457.143468. URL http://doi.acm.org/10.1145/143457.143468.

A. Druin, J. Stewart, D. Proft, B. Bederson, and J. Hollan. Kidpad: a design collaboration between children, technologists, and educators. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems*, pages 463–470. ACM, 1997.

C. A. Ellis, S. J. Gibbs, and G. Rein. Groupware: Some issues and experiences. *Commun. ACM*, 34(1):39–58, Jan. 1991. ISSN 0001-0782. doi: 10.1145/99977.99987. URL http://doi.acm.org/10.1145/99977.99987.

D. Fisher and P. Dourish. Social and temporal structures in everyday collaboration. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '04, pages 551–558, New York, NY, USA, 2004. ACM. ISBN 1-58113-702-8. doi: 10.1145/985692.985762. URL http://doi.acm.org/10.1145/985692.985762.

D. M. Frohlich. The history and future of direct manipulation. *Behaviour & Information Technology*, 12(6):315–329, 1993.

S. R. Fussell and R. M. Krauss. Accuracy and bias in estimates of others' knowledge. *European Journal of Social Psychology*, 21(5):445–454, 1991.

S. R. Fussell, R. E. Kraut, and J. Siegel. Coordination of communication: Effects of shared visual context on collaborative work. In *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work*, CSCW '00, pages 21–30, New York, NY, USA, 2000. ACM. ISBN 1-58113-222-0. doi: 10.1145/358916.358947. URL http://doi.acm.org/10.1145/358916.358947.

S. R. Fussell, L. D. Setlock, J. Yang, J. Ou, E. Mauer, and A. D. I. Kramer. Gestures over video streams to support remote collaboration on physical tasks. *Hum.-Comput. Interact.*, 19 (3):273–309, Sept. 2004. ISSN 0737-0024. doi: 10.1207/s15327051hci1903_3. URL http://dx.doi.org/10.1207/s15327051hci1903_3.

N. P. Garg, S. Favre, H. Salamin, D. Hakkani Tür, and A. Vinciarelli. Role recognition for meeting participants: an approach based on lexical information and social network analysis. In *Proceedings of the 16th ACM international conference on Multimedia*, pages 693–696. ACM, 2008.

J. Gibbons, W. Kincheloe, and K. Down. Tutored videotape instruction: A new use of electronics media in education. *Science*, 195:1139–1146, 1977.

P. S. Goodman and S. Shah. Familiarity and work group outcomes. *Group process and productivity*, pages 276 – 298, 1992.

S. Greenberg, C. Gutwin, and A. Cockburn. Awareness through fisheye views in relaxed-wysiwis groupware. In *Graphics interface*, volume 96, pages 28–38, 1996.

J. Grudin. Why groupware applications fail: Problems in design and evaluation. *Office: Technology and people*, 4(3):245–264, 1989.

J. Grudin. Computer-supported cooperative work: History and focus. *Computer*, 27(5):19–26, May 1994a. ISSN 0018-9162. doi: 10.1109/2.291294. URL http://dx.doi.org/10.1109/2.291294.

J. Grudin. Groupware and social dynamics: Eight challenges for developers. *Commun. ACM*, 37(1):92–105, Jan. 1994b. ISSN 0001-0782. doi: 10.1145/175222.175230. URL http://doi.acm.org/10.1145/175222.175230.

C. Gutwin and S. Greenberg. A descriptive framework of workspace awareness for real-time groupware. *Computer Supported Cooperative Work (CSCW)*, 11(3-4):411–446, 2002.

E. T. Hall. The hidden dimension: man's use of space in public and private the bodley head, 1969.

E. T. Hall, R. L. Birdwhistell, B. Bock, P. Bohannan, A. R. Diebold Jr, M. Durbin, M. S. Edmonson, J. Fischer, D. Hymes, S. T. Kimball, et al. Proxemics [and comments and replies]. *Current anthropology*, pages 83–108, 1968.

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 SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 609–618, New York, NY, USA, 2010. ACM. ISBN 978-1-60558-929-9. doi: 10.1145/1753326.1753418. URL http://doi.acm.org/10.1145/1753326.1753418.

T. E. Hansen and J. P. Hourcade. Comparing multi-touch tabletops and multi-mouse single-display groupware setups. In *Proceedings of the 3rd Mexican Workshop on Human Computer Interaction*, pages 36–43. Universidad Politécnica de San Luis Potosí, 2010.

B. Hartmann, M. R. Morris, H. Benko, and A. D. Wilson. Augmenting interactive tables with mice & keyboards. In *Proceedings of the 22nd annual ACM symposium on User interface software and technology*, pages 149–152. ACM, 2009.

J.-B. Haué and P. Dillenbourg. Do fewer laptops make a better team? In *Interactive artifacts and furniture supporting collaborative work and learning*, pages 1–24. Springer, 2009.

K. Hawkey, M. Kellar, D. Reilly, T. Whalen, and K. M. Inkpen. The proximity factor: Impact of distance on co-located collaboration. In *Proceedings of the 2005 International ACM SIGGROUP Conference on Supporting Group Work*, GROUP '05, pages 31–40, New York, NY, USA, 2005. ACM. ISBN 1-59593-223-2. doi: 10.1145/1099203.1099209. URL http://doi.acm.org/10.1145/1099203.1099209.

C. Heath and P. Luff. Collaboration and controlCrisis management and multimedia technology in london underground line control rooms. *Computer Supported Cooperative Work (CSCW)*, 1(1-2):69–94, Mar. 1992. ISSN 0925-9724, 1573-7551. doi: 10.1007/BF00752451. URL http://link.springer.com/article/10.1007/BF00752451.

E. T. Higgins and W. S. Rholes. "saying is believing": Effects of message modification on memory and liking for the person described. *Journal of Experimental Social Psychology*, 14 (4):363–378, 1978.

T. E. Higgins. The "communication game": Implications for social cognition and persuasion. In E. T. Higgins, C. P. Herman, and M. P. Zanna, editors, *Social Cognition: The Ontario Symposium, Vol. 1*, pages 343–392. Lawrence Erlbaum Associates, Hillsdale, NJ, 1981.

G. W. Hill. Group versus individual performance: Are n+ 1 heads better than one? *Psychological Bulletin*, 91(3):517, 1982.

W. C. Hill and J. D. Hollan. History-enriched digital objects: Prototypes and policy issues. *The Information Society*, 10(2):139–145, 1994.

J. Hindmarsh and C. Heath. Embodied reference: A study of deixis in workplace interaction. *Journal of Pragmatics*, 32(12):1855–1878, Nov. 2000. ISSN 0378-2166. doi: 10.1016/S0378-2166(99)00122-8. URL http://www.sciencedirect.com/science/article/pii/S0378216699001228.

U. Hinrichs, M. Hancock, S. Carpendale, and C. Collins. Examination of text-entry methods for tabletop displays. In *Horizontal Interactive Human-Computer Systems, 2007. TABLETOP'07. Second Annual IEEE International Workshop on*, pages 105–112. IEEE, 2007.

J. Hollan, E. Hutchins, and D. Kirsh. Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 7(2):174–196, 2000.

E. Hutchins. How a cockpit remembers its speeds. *Cognitive Science*, 19(3):265–288, 1995. ISSN 1551-6709. doi: 10.1207/s15516709cog1903_1. URL http://dx.doi.org/10.1207/s15516709cog1903_1.

P. Hutterer and B. H. Thomas. Groupware support in the windowing system. In *Proceedings of the eight Australasian conference on User interface-Volume 64*, pages 39–46. Australian Computer Society, Inc., 2007.

K. M. Inkpen, W.-l. Ho-Ching, O. Kuederle, S. D. Scott, and G. B. Shoemaker. This is fun! we're all best friends and we're all playing: supporting children's synchronous collaboration. In *Proceedings of the 1999 conference on Computer support for collaborative learning*, page 31. International Society of the Learning Sciences, 1999.

H. Ishii, M. Kobayashi, and J. Grudin. Integration of interpersonal space and shared workspace: ClearBoard design and experiments. *ACM Trans. Inf. Syst.*, 11(4):349–375, Oct. 1993. ISSN 1046-8188. doi: 10.1145/159764.159762. URL http://doi.acm.org/10.1145/159764.159762.

S. Izadi, H. Brignull, T. Rodden, Y. Rogers, and M. Underwood. Dynamo: a public interactive surface supporting the cooperative sharing and exchange of media. In *Proceedings of the 16th annual ACM symposium on User interface software and technology*, pages 159–168. ACM, 2003.

D. B. Jayagopi, H. Hung, C. Yeo, and D. Gatica-Perez. Modeling dominance in group conversations using nonverbal activity cues. *Audio, Speech, and Language Processing, IEEE Transactions on*, 17(3):501–513, 2009.

P. Jermann. Computer support for interaction regulation in collaborative problem-solving. *Unpublished Ph. D. thesis, University of Geneva, Switzerland*, 2004.

P. Jermann, A. Soller, and M. Muehlenbrock. From mirroring to guiding: A review of the state of art technology for supporting collaborative learning. In . K. H. P. Dillenbourg, A. Eurelings, editor, *European Conference on Computer-Supported Collaborative Learning EuroCSCL-2001*, pages 324–331, Maastricht, Netherlands, 2001. URL https://telearn.archives-ouvertes.fr/hal-00197377.

B. Johanson, O. Fox, and T. Winograd. The interactive workspaces project: Experiences with ubiquitous computing rooms. In *IEEE Pervasive Computing*, 2002a.

B. Johanson, G. Hutchins, T. Winograd, and M. Stone. Pointright: experience with flexible input redirection in interactive workspaces. In *Proceedings of the 15th annual ACM symposium on User interface software and technology*, pages 227–234. ACM, 2002b.

G. Kahrimanis, A. Meier, I.-A. Chounta, E. Voyiatzaki, H. Spada, N. Rummel, and N. Avouris. Assessing collaboration quality in synchronous cscl problem-solving activities: Adaptation and empirical evaluation of a rating scheme. In *Learning in the Synergy of Multiple Disciplines*, pages 267–272. Springer, 2009.

F. Kaplan, S. DoLenh, K. Bachour, G. Y.-i. Kao, C. Gault, and P. Dillenbourg. Interpersonal computers for higher education. In *Interactive artifacts and furniture supporting collaborative work and learning*, pages 1–17. Springer, 2009.

S. Kiesler and J. N. Cummings. What do we know about proximity and distance in work groups? a legacy of research. *Distributed work*, 1:57, 2002.

T. Kim, A. Chang, L. Holland, and A. S. Pentland. Meeting mediator: Enhancing group collaborationusing sociometric feedback. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work*, CSCW '08, pages 457–466, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-007-4. doi: 10.1145/1460563.1460636. URL http://doi.acm.org/10.1145/1460563.1460636.

D. Kirk, T. Rodden, and D. S. Fraser. Turn it this way: Grounding collaborative action with remote gestures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, pages 1039–1048, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-593-9. doi: 10.1145/1240624.1240782. URL http://doi.acm.org/10.1145/1240624.1240782.

M. Knapp and J. Hall. *Nonverbal communication in human interaction.* Harcourt Brace College Publishers, New York, 1972.

R. M. Krauss and S. Weinheimer. Concurrent feedback, confirmation, and the encoding of referents in verbal communication. *Journal of Personality and Social Psychology*, 4(3):343, 1966.

R. E. Kraut, S. R. Fussell, S. E. Brennan, and J. Siegel. Understanding effects of proximity on collaboration: Implications for technologies to support remote collaborative work. *Distributed work*, pages 137–162, 2002.

R. Kruger, S. Carpendale, S. D. Scott, and S. Greenberg. How people use orientation on tables: Comprehension, coordination and communication. In *Proceedings of the 2003 International ACM SIGGROUP Conference on Supporting Group Work*, GROUP '03, pages 369–378, New York, NY, USA, 2003. ACM. ISBN 1-58113-693-5. doi: 10.1145/958160.958219. URL http://doi.acm.org/10.1145/958160.958219.

D. Lalanne, D. Mekhaldi, and R. Ingold. Talking about documents: revealing a missing link to multimedia meeting archives. In *Electronic Imaging 2004*, pages 82–91. International Society for Optics and Photonics, 2003.

J. H. Larkin and H. A. Simon. Why a diagram is (sometimes) worth ten thousand words. *Cognitive science*, 11(1):65–100, 1987.

K. Laskowski, M. Ostendorf, and T. Schultz. Modeling vocal interaction for text-independent participant characterization in multi-party conversation. In *Proceedings of the 9th SIGdial Workshop on Discourse and Dialogue*, pages 148–155. Association for Computational Linguistics, 2008.

B. Latane. The psychology of social impact. *American psychologist*, 36(4):343, 1981.

J. M. Levine and R. L. Moreland. Small groups. *The handbook of social psychology*, 2:415–469, 1998.

J. M. Levine, L. B. Resnick, and E. Tory. Social foundations of cognition. *Annual Review of Psychology*, 44:585–612, 1993. ISSN 0066-4308(Print). doi: 10.1146/annurev.ps.44.020193.003101.

M. Lombard and T. Ditton. At the heart of it all: The concept of presence. *Journal of Computer-Mediated Communication*, 1997.

C.-K. Looi, C.-P. Lin, and K.-P. Liu. Group scribbles to support knowledge building in jigsaw method. *Learning Technologies, IEEE Transactions on*, 1(3):157–164, 2008.

T. W. Malone and K. Crowston. The interdisciplinary study of coordination. *ACM Comput. Surv.*, 26(1):87–119, Mar. 1994. ISSN 0360-0300. doi: 10.1145/174666.174668. URL http://doi.acm.org/10.1145/174666.174668.

R. L. Mandryk, S. D. Scott, and K. M. Inkpen. Display factors influencing co-located collaboration. *Comference Supplement to ACM CSCW*, 2, 2002.

M. Mandviwalla and L. Olfman. What do groups need? a proposed set of generic groupware requirements. *ACM Trans. Comput.-Hum. Interact.*, 1(3):245–268, Sept. 1994. ISSN 1073-0516. doi: 10.1145/196699.196715. URL http://doi.acm.org/10.1145/196699.196715.

I. McCowan, S. Bengio, D. Gatica-Perez, G. Lathoud, F. Monay, D. Moore, P. Wellner, and H. Bourlard. Modeling human interaction in meetings. In *Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP'03). 2003 IEEE International Conference on*, volume 4, pages IV–748. IEEE, 2003.

J. E. McGrath and A. B. Hollingshead. *Groups interacting with technology: Ideas, evidence, issues, and an agenda*, volume ix of *Sage library of social research, 194.* Sage Publications, Inc, Thousand Oaks, CA, US, 1994. ISBN 0-8039-4897-2 (Hardcover); 0-8039-4898-0 (Paperback).

A. Meier, H. Spada, and N. Rummel. A rating scheme for assessing the quality of computer-supported collaboration processes. *International Journal of Computer-Supported Collaborative Learning*, 2(1):63–86, 2007.

D. Mekhaldi, D. Lalanne, and R. Ingold. Thematic segmentation of meetings through document/speech alignment. In *Proceedings of the 12th annual ACM international conference on Multimedia*, pages 804–811. ACM, 2004.

B. E. Mennecke. Using group support systems to discover hidden profiles: An examination of the influence of group size and meeting structures on information sharing and decision quality. *International Journal of Human-Computer Studies*, 47(3):387–405, 1997.

N. Moraveji, K. Inkpen, E. Cutrell, and R. Balakrishnan. A mischief of mice: examining children's performance in single display groupware systems with 1 to 32 mice. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 2157–2166. ACM, 2009.

R. Moreland. Transactive memory and job performance: Helping workers learn who knows what. *Shared Cognition in Organizations: The Management of Knowledge, Mahwah, NJ: Lawrence Erlbaum Associates, Inc*, pages 3–32, 1999.

M. R. Morris, K. Ryall, C. Shen, C. Forlines, and F. Vernier. Beyond social protocols: Multi-user coordination policies for co-located groupware. In *Proceedings of the 2004 ACM conference on Computer supported cooperative work*, pages 262–265. ACM, 2004.

B. Mullen. Group composition, salience, and cognitive representations: The phenomenology of being in a group. *Journal of Experimental Social Psychology*, 27(4):297–323, July 1991. ISSN 0022-1031. doi: 10.1016/0022-1031(91)90028-5. URL http://www.sciencedirect.com/science/article/pii/0022103191900285.

D. C. Neale, J. M. Carroll, and M. B. Rosson. Evaluating computer-supported cooperative work: Models and frameworks. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work*, CSCW '04, pages 112–121, New York, NY, USA, 2004. ACM. ISBN 1-58113-810-5. doi: 10.1145/1031607.1031626. URL http://doi.acm.org/10.1145/1031607.1031626.

D. A. Norman and S. W. Draper. User centered system design. *New Perspectives on Human-Computer Interaction, L. Erlbaum Associates Inc., Hillsdale, NJ*, 1986.

A. Pentland. Social signal processing. *IEEE Signal Processing Magazine*, 24(4):108, 2007.

R. M. Pilkington. *Analysing educational discourse: The DISCOUNT scheme.* University of Leeds, Computer Based Learning Unit, 1999.

C. Rebetez, M. Bétrancourt, M. Sangin, and P. Dillenbourg. Learning from animation enabled by collaboration. *Instructional science*, 38(5):471–485, 2010.

R. Rienks and D. Heylen. Dominance detection in meetings using easily obtainable features. In *Machine Learning for Multimodal Interaction*, pages 76–86. Springer, 2006.

R. Rienks, D. Zhang, D. Gatica-Perez, and W. Post. Detection and application of influence rankings in small group meetings. In *Proceedings of the 8th international conference on Multimodal interfaces*, pages 257–264. ACM, 2006.

Y. Rogers and T. Rodden. Configuring spaces and surfaces to support collaborative interactions. In *Public and situated displays*, pages 45–79. Springer, 2003.

Y. Rogers, Y.-k. Lim, W. R. Hazlewood, and P. Marshall. Equal opportunities: do shareable interfaces promote more group participation than single user displays? *Human–Computer Interaction*, 24(1-2):79–116, 2009.

R. Rommetveit. *On message structure: A framework for the study of language and communication.* John Wiley & Sons, 1974.

W. B. Rouse and N. M. Morris. On looking into the black box: Prospects and limits in the search for mental models. *Psychological bulletin*, 100(3):349, 1986.

K. Ryall, C. Forlines, C. Shen, and M. R. Morris. Exploring the effects of group size and table size on interactions with tabletop shared-display groupware. In *Proceedings of the 2004 ACM conference on Computer supported cooperative work*, pages 284–293. ACM, 2004.

H. Salamin, S. Favre, and A. Vinciarelli. Automatic role recognition in multiparty recordings: Using social affiliation networks for feature extraction. *Multimedia, IEEE Transactions on*, 11(7):1373–1380, 2009.

M. Sangin, G. Molinari, M.-A. Nüssli, and P. Dillenbourg. Facilitating peer knowledge modeling: Effects of a knowledge awareness tool on collaborative learning outcomes and processes. *Computers in Human Behavior*, 27(3):1059–1067, 2011.

M. Scaife and Y. Rogers. External cognition: how do graphical representations work? *International journal of human-computer studies*, 45(2):185–213, 1996.

E. M. Schooler. Conferencing and collaborative computing. *Multimedia Systems*, 4(5):210–225, Oct. 1996. ISSN 0942-4962, 1432-1882. doi: 10.1007/s005300050025. URL http://link.springer.com/article/10.1007/s005300050025.

S. D. Scott, R. L. Mandryk, and K. M. Inkpen. Understanding children's collaborative interactions in shared environments. *Journal of Computer Assisted Learning*, 19(2):220–228, 2003.

S. D. Scott, M. S. T. Carpendale, and K. M. Inkpen. Territoriality in collaborative tabletop workspaces. In *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work*, CSCW '04, pages 294–303, New York, NY, USA, 2004. ACM. ISBN 1-58113-810-5. doi: 10.1145/1031607.1031655. URL http://doi.acm.org/10.1145/1031607.1031655.

C. Shen, F. D. Vernier, C. Forlines, and M. Ringel. Diamondspin: an extensible toolkit for around-the-table interaction. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 167–174. ACM, 2004.

C. Shen, K. Ryall, C. Forlines, A. Esenther, F. D. Vernier, K. Everitt, M. Wu, D. Wigdor, M. R. Morris, M. HancocK, et al. Collaborative tabletop research and evaluation. In *Interactive* *artifacts and furniture supporting collaborative work and learning*, pages 1–17. Springer, 2009.

M. J. Sipusic, R. L. Pannoni, R. B. Smith, J. Dutra, J. F. Gibbons, and W. R. Sutherland. Virtual collaborative learning: a comparison between face-to-face tutored video instruction (tvi) and distributed tutored video instruction (dtvi). 1999.

R. B. Smith, M. J. Sipusic, and R. L. Pannoni. Experiments comparing face-to-face with virtual collaborative learning. In *Proceedings of the 1999 conference on Computer support for collaborative learning*, page 68. International Society of the Learning Sciences, 1999.

J. Snyder. Let me draw you a picture: Coordination in image-enabled conversation. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work Companion*, CSCW '12, pages 219–222, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1051-2. doi: 10.1145/2141512.2141582. URL http://doi.acm.org/10.1145/2141512.2141582.

J. Snyder. Drawing practices in image-enabled collaboration. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, CSCW '13, pages 741–752, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-1331-5. doi: 10.1145/2441776.2441858. URL http://doi.acm.org/10.1145/2441776.2441858.

J. Snyder. Visual representation of information as communicative practice. *Journal of the Association for Information Science and Technology*, 65(11):2233–2247, 2014.

J. Snyder, E. P. Baumer, S. Voida, P. Adams, M. Halpern, T. Choudhury, and G. Gay. Making things visible: Opportunities and tensions in visual approaches for design research and practice. *Human–Computer Interaction*, 29(5-6):451–486, 2014.

G. Stasser and D. Stewart. Discovery of hidden profiles by decision-making groups: Solving a problem versus making a judgment. *Journal of personality and social psychology*, 63(3):426, 1992.

M. Stefik, G. Foster, D. G. Bobrow, K. Kahn, S. Lanning, and L. Suchman. Beyond the chalkboard: Computer support for collaboration and problem solving in meetings. *Commun. ACM*, 30 (1):32–47, Jan. 1987. ISSN 0001-0782. doi: 10.1145/7885.7887. URL http://doi.acm.org/10.1145/7885.7887.

J. Stewart, E. M. Raybourn, B. Bederson, and A. Druin. When two hands are better than one: Enhancing collaboration using single display groupware. In *CHI 98 Cconference Summary on Human Factors in Computing Systems*, pages 287–288. ACM, 1998.

J. Stewart, B. B. Bederson, and A. Druin. Single display groupware: a model for co-present collaboration. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 286–293. ACM, 1999.

N. A. Streitz, J. Geißler, J. M. Haake, and J. Hol. Dolphin: integrated meeting support across local and remote desktop environments and liveboards. In *Proceedings of the 1994 ACM conference on Computer supported cooperative work*, pages 345–358. ACM, 1994.

N. A. Streitz, J. Geißler, T. Holmer, S. Konomi, C. Müller-Tomfelde, W. Reischl, P. Rexroth, P. Seitz, and R. Steinmetz. i-LAND: An interactive landscape for creativity and innovation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '99, pages 120–127, New York, NY, USA, 1999. ACM. ISBN 0-201-48559-1. doi: 10.1145/302979.303010. URL http://doi.acm.org/10.1145/302979.303010.

N. A. Streitz, P. Tandler, C. Müller-Tomfelde, and S. Konomi. Roomware: Towards the next generation of human-computer: Interaction based on an integrated design of real and virtual worlds. *Human-Computer Interaction in the New Millenium, Addison Wesley*, pages 551–576, 2001.

S. Streng. *The role of personal and shared displays in scripted collaborative learning.* Springer, 2009.

S. Streng, K. Stegmann, H. Hußmann, and F. Fischer. Metaphor or diagram?: comparing different representations for group mirrors. In *Proceedings of the 21st Annual Conference of the Australian Computer-Human Interaction Special Interest Group: Design: Open 24/7*, pages 249–256. ACM, 2009.

D. Suthers, L. Girardeau, and C. Hundhausen. Deictic roles of external representations in face-to-face and online collaboration. In B. Wasson, S. Ludvigsen, and U. Hoppe, editors, *Designing for Change in Networked Learning Environments*, number 2 in Computer-Supported Collaborative Learning, pages 173–182. Springer Netherlands, Jan. 2003. ISBN 978-90-481-6321-2, 978-94-017-0195-2. URL http://link.springer.com/chapter/10.1007/978-94-017-0195-2_23.

E. Szewkis, M. Nussbaum, T. Rosen, J. Abalos, F. Denardin, D. Caballero, A. Tagle, and C. Alcoholado. Collaboration within large groups in the classroom. *International Journal of Computer-Supported Collaborative Learning*, 6(4):561–575, 2011.

P. Tandler, T. Prante, C. Müller-Tomfelde, N. Streitz, and R. Steinmetz. Connectables: Dynamic coupling of displays for the flexible creation of shared workspaces. In *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology*, UIST '01, pages 11–20, New York, NY, USA, 2001. ACM. ISBN 1-58113-438-X. doi: 10.1145/502348.502351. URL http://doi.acm.org/10.1145/502348.502351.

A. Tang, M. Tory, B. Po, P. Neumann, and S. Carpendale. Collaborative coupling over tabletop displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, pages 1181–1190, New York, NY, USA, 2006. ACM. ISBN 1-59593-372-7. doi: 10.1145/1124772.1124950. URL http://doi.acm.org/10.1145/1124772.1124950.

A. Tang, J. Lanir, S. Greenberg, and S. Fels. Supporting transitions in work: informing large display application design by understanding whiteboard use. In *Proceedings of the ACM 2009 international conference on Supporting group work*, pages 149–158. ACM, 2009.

J. C. Tang. Findings from observational studies of collaborative work. *International Journal of Man-Machine Studies*, 34(2):143–160, Feb. 1991. ISSN 0020-7373. doi: 10.1016/0020-7373(91)90039-A. URL http://www.sciencedirect.com/science/article/pii/002073739190039A.

K. Tang. Spontaneous collaborative learning: A new dimension in student learning experience? *Higher Education Research and Development*, 12(2):115–130, 1993.

D. W. Taylor, P. C. Berry, and C. H. Block. Does group participation when using brainstorming facilitate or inhibit creative thinking? *Administrative Science Quarterly*, pages 23–47, 1958.

L. Terrenghi, D. Kirk, A. Sellen, and S. Izadi. Affordances for manipulation of physical versus digital media on interactive surfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 1157–1166. ACM, 2007.

L. Thompson and G. A. Fine. Socially shared cognition, affect, and behavior: A review and integration. *Personality and Social Psychology Review*, 3(4):278–302, Nov. 1999. ISSN 1088-8683, 1532-7957. doi: 10.1207/s15327957pspr0304_1. URL http://psr.sagepub.com/content/3/4/278.

R. S. Tindale, C. M. Smith, L. S. Thomas, J. Filkins, and S. Sheffey. Shared representations and asymmetric social influence processes in small groups. *Understanding group behavior, Consensual action by small groups*, 1:81–103, 1996.

H. Tischler. *Introduction to Sociology*. Harcourt Brace College Publishers, New York, 1990.

D. R. Traum. Computational models of grounding in collaborative systems. In *Psychological Models of Communication in Collaborative Systems-Papers from the AAAI Fall Symposium*, pages 124–131, 1999.

F. Tschan. Ideal cycles of communication (or cognitions) in triads, dyads, and individuals. *Small Group Research*, 33(6):615–643, 2002.

E. Tse and S. Greenberg. Rapidly prototyping single display groupware through the sdgtoolkit. In *Proceedings of the fifth conference on Australasian user interface-Volume 28*, pages 101–110. Australian Computer Society, Inc., 2004.

B. Tucker. The flipped classroom. *Education Next*, 12(1):82–83, 2012.

B. Tversky and P. U. Lee. How space structures language. In *Spatial cognition*, pages 157–175. Springer, 1998.

H. Verma, F. Roman, S. Magrelli, P. Jermann, and P. Dillenbourg. Complementarity of input devices to achieve knowledge sharing in meetings. In *Proceedings of the 2013 conference on computer supported cooperative work*, pages 701–714. ACM, 2013.

A. Vinciarelli. Speakers role recognition in multiparty audio recordings using social network analysis and duration distribution modeling. *Multimedia, IEEE Transactions on*, 9(6):1215–1226, 2007.

A. Vinciarelli, M. Pantic, and H. Bourlard. Social signal processing: Survey of an emerging domain. *Image and Vision Computing*, 27(12):1743–1759, Nov. 2009a. ISSN 0262-8856. doi: 10.1016/j.imavis.2008.11.007. URL http://www.sciencedirect.com/science/article/pii/S0262885608002485.

A. Vinciarelli, H. Salamin, and M. Pantic. Social signal processing: Understanding social interactions through nonverbal behavior analysis. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, 2009. CVPR Workshops 2009*, pages 42–49, June 2009b. doi: 10.1109/CVPRW.2009.5204290.

A. Vinciarelli, M. Pantic, D. Heylen, C. Pelachaud, I. Poggi, F. D'Errico, and M. Schroeder. Bridging the gap between social animal and unsocial machine: A survey of social signal processing. *IEEE Transactions on Affective Computing*, 3(1):69–87, Jan. 2012. ISSN 1949-3045. doi: 10.1109/T-AFFC.2011.27.

L. S. Vygotsky. *Mind in society: The development of higher psychological processes*. Harvard university press, 1980.

J. R. Wallace, S. D. Scott, T. Stutz, T. Enns, and K. Inkpen. Investigating teamwork and taskwork in single-and multi-display groupware systems. *Personal and Ubiquitous Computing*, 13(8): 569–581, 2009.

D. M. Wegner. Transactive memory: A contemporary analysis of the group mind. In *Theories of group behavior*, pages 185–208. Springer, 1987.

D. M. Wegner, T. Giuliano, and P. T. Hertel. Cognitive interdependence in close relationships. In *Compatible and incompatible relationships*, pages 253–276. Springer, 1985.

D. M. Wegner, R. Erber, and P. Raymond. Transactive memory in close relationships. *Journal of personality and social psychology*, 61(6):923, 1991.

M. Weiser. Ubiquitous computing. *Computer*, 26(10):71–72, 1993.

J. D. Weisz, S. Kiesler, H. Zhang, Y. Ren, R. E. Kraut, and J. A. Konstan. Watching together: integrating text chat with video. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 877–886. ACM, 2007.

S. Whittaker, E. Geelhoed, and E. Robinson. Shared workspaces: how do they work and when are they useful? *International Journal of Man-Machine Studies*, 39(5):813–842, 1993.

D. Wigdor, H. Jiang, C. Forlines, M. Borkin, and C. Shen. Wespace: the design development and deployment of a walk-up and share multi-surface visual collaboration system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1237–1246. ACM, 2009.

H. A. Wilke and R. W. Meertens. Group performance. 1994.

C. G. Wolf, J. R. Rhyne, and L. K. Briggs. Communication and information retrieval with a pen-based meeting support tool. In *Proceedings of the 1992 ACM conference on Computer-supported cooperative work*, pages 322–329. ACM, 1992.

S. Worchel. You can go home again returning group research to the group context with an eye on developmental issues. *Small Group Research*, 25(2):205–223, May 1994. ISSN 1046-4964, 1552-8278. doi: 10.1177/1046496494252004. URL http://sgr.sagepub.com/content/25/2/205.

R. Zajonc and P. K. Adelmann. Cognition and communication: A story of missed opportunities. *Social Science Information/sur les sciences sociales*, 1987.

A. Zanella and S. Greenberg. Reducing interference in single display groupware through transparency. In *ECSCW 2001*, pages 339–358. Springer, 2001.

D. Zhang, D. Gatica-Perez, S. Bengio, and I. McCowan. Modeling individual and group actions in meetings with layered hmms. *Multimedia, IEEE Transactions on*, 8(3):509–520, 2006.

S. Zhao. Toward a taxonomy of copresence. *Presence: Teleoperators and Virtual Environments*, 12(5):445–455, 2003.

Rue de Lausanne, 19
CH-1020 Renens
+41 78 644 38 34
himanshu.verma@epfl.ch

Himanshu Verma

Education

2010–2015 **PhD Student in Computer Science**, *École Polytechnique Fédérale de Lausanne*, Switzerland.
Title: Latent Social Information in Group's Interactions with a Shared Workspace

2008–2010 **Masters in Information Technology**, *Indian Institute of Information Technology*, Allahabad, India.
Specialization in Human-Computer Interaction

2004–2008 **Bachelor of Technology**, *Gurukula Kangri University*, Haridwar, India.
Computer Science and Engineering

Languages

Hindi	**Native speaker**	
English	**Advanced**	*Bilingual proficiency*
French	**Intermediate**	*Average reading and writing skills*
German	**Beginner**	

Work Experience

Research

2010–2015 **Research Assistant**, *École Polytechnique Fédérale de Lausanne*, Switzerland.
Computer-Supported Collaborative Work

Teaching

2011–2013 **Computer-Supported Collaborative Work**, *Teaching Assistant*, School of Computer and Communication Sciences, École Polytechnique Fédérale de Lausanne.
Course coordination, project supervision, conducting user-studies

2014 **Digital Education & Learning Analytics**, *Teaching Assistant*, School of Computer and Communication Sciences, École Polytechnique Fédérale de Lausanne.
Managing online learning platform: EdX

Computer Skills

Languages C#, Java, C++

Statistical Tools R

Others LaTeX, ELAN

Publications

Journal Articles

2015 K. Sharma, H. Verma, D. Caballero, P. Jermann, P. Dillenbourg, *"Shaping Learners' Attention in Massive Open Online Courses"*, Submitted in International Journal in Higher Education (Special Issue on MOOCs), 2015.

2014 N. Li, H. Verma, A. Skevi, G. Zufferey, J. Blom, P. Dillenbourg, *"Watching MOOCs Together: Investigating Co-located MOOCStudy Groups"*, Distance Education, 35(2), 2014.

Other Papers

2015 K. Sharma, D. Caballero, H. Verma, P. Jermann, P. Dillenbourg, *"Looking AT versus Looking THROUGH: A Dual Eye-Tracking Study in MOOC Context"*, Accepted in Proceedings of 11th International Conference of Computer Supported Collaborative Learning, Gothenburg, Sweden, CSCL, 2015.

2014 N. Li, H. Verma, A. Skevi, G. Zufferey, P. Dillenbourg, *"MOOC Learning in Spontaneous Study Groups: Does Synchronously Watching Videos Make a Difference?"*, In Proceedings of the European MOOC Stakeholder Summit, EMOOC, 2014.

2013 H. Verma, F. Roman, S. Magrelli, P. Jermann, P. Dillenbourg, *"Complementarity of Input Devices to Achieve Knowledge Sharing in Meetings"*, In Proceedings of the 2013 conference on Computer Supported Cooperative Work, San Antonio, TX, USA, CSCW, 2013.

2012 Q. Bonnard, H. Verma, F. Kaplan, P. Dillenbourg, *"Paper Interfaces for Learning Geometry"*, In Proceedings of 7th European Conference of Technology Enhanced Learning, Saarbrücken, Germany, EC-TEL, 2012.

2012 H. Verma, F. Roman, P. Jermann, P. Dillenbourg, *"Effects of Input Device Familiarity on Content Creation and Sharing in Meetings"*, In CHI'12 Extended Abstracts on Human Factors in Computing Systems, Austin, TX, USA, CHI EA, 2012.