

Cellulo: Versatile Handheld Robots for Education

Ayberk Özgür¹, Séverin Lemaignan^{1,2,3}, Wafa Johal^{1,2}, Maria Beltran²,
Manon Briod², Léa Pereyre², Francesco Mondada², Pierre Dillenbourg¹

¹CHILI, ²LSRO, EPFL, Switzerland
{firstname.lastname}@epfl.ch

³Centre for Robotics and Neural Systems, Plymouth University, U.K.
severin.lemaignan@plymouth.ac.uk

ABSTRACT

In this article, we present *Cellulo*, a novel robotic platform that investigates the intersection of three ideas for robotics in education: designing the robots to be versatile and generic tools; blending robots into the classroom by designing them to be pervasive objects and by creating tight interactions with (already pervasive) paper; and finally considering the practical constraints of real classrooms at every stage of the design. Our platform results from these considerations and builds on a unique combination of technologies: groups of handheld haptic-enabled robots, tablets and activity sheets printed on regular paper. The robots feature holonomic motion, haptic feedback capability and high accuracy localization through a microdot pattern overlaid on top of the activity sheets, while remaining affordable (robots cost about €125 at the prototype stage) and classroom-friendly. We present the platform and report on our first interaction studies, involving about 230 children.

Keywords

Human-Robot Interaction; Robots for Learning; Haptic Interfaces; Tangible Robots

1. INTRODUCTION

The adoption of educational robots in schools is a potential key contribution to education, but is limited by several factors: Price, integration within the existing curriculum and disciplines, actual contribution to the learning of the material, teachers' ethical concerns and fears of additional workload, the complexity of the system and difficult logistics in the classroom, among others. As a result, research in robotics for education has mainly focused on certain subjects in specific teaching situations: custom robotic platforms to teach programming and other closely related Science, Technology, Engineering and Mathematics (STEM) concepts; and studies that create rich social interactions for tutoring or peer-learning of specific topics such as language and writing. To overcome these limitations and open new perspectives,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HRI '17, March 06 - 09, 2017, Vienna, Austria

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4336-7/17/03...\$15.00

DOI: <http://dx.doi.org/10.1145/2909824.3020247>

we investigate a novel approach where the robotic education platform is designed to be *ubiquitous*, *versatile* and *practical*.

The design of a *ubiquitous* platform follows the analogy of “pen and paper”: a set of pervasive yet unremarkable tools that blend into the daily learning routine. This is achieved by designing devices to be reliable (they must be trusted to work at any time), readily replaceable, and intuitive to use. These design goals translate into desirable qualities for an ubiquitous educational robot: a small size; a plain and sturdy appearance; a low price point; a limited set of simple affordances. Besides, to be readily replaceable, the device should not elicit strong affective bonding i.e. it should not be uniquely associated to one child, it should not call for personalization, and it should not elicit projected agency. This contrasts with certain branches of research in educational robotics where children's engagement is often sought through social bonding, and therefore, social robots.

A *versatile* system should be designed so as to be applicable to a broad range of learning scenarios, instead of as a specialized tool bound to the teaching of a particular subject. It should be easy and natural for the teachers to imagine the role and added value of the robots when creating new teaching activities in different disciplines. This implies that the robots' hardware, appearance and interaction modalities must not imply or be constrained to specific use cases.

Finally, a *practical* teaching tool must be flexible yet reliable as they are meant to be used intensively by pupils; the fragile hardware and complex software typically found in robotics do not effectively support uninterrupted learning and teaching in such scenarios. In order to gain field acceptance in the classrooms, educational robots must critically represent a net educational gain and must not incur higher workload for the teachers. Time-consuming initialization, configuration or calibration of the robots, as well as any software manipulation that would prevent the teacher to focus on the classroom for a non-trivial amount of time, should be avoided. As such, designing a robot to be practical in the classroom environment requires carefully designed ergonomics to achieve a simple and robust user experience for both the pupils and the teachers.

The Cellulo project explores the intersection of these three, partially conflicting, design goals. Our approach builds on *groups of small, low-cost, haptic-enabled tangible robots*, as seen in the concepts pictured in Figure 1. Each robot has a plain and minimalist appearance, and the robots collectively operate on large sheets of printed paper representing a learning activity. The robots can autonomously move and interact with the activity sheet; interact with each other; be



Figure 1: Concept Cellulo activities: *Left:* Robots simulate planets orbiting the sun. The learners can move the planets to alter their orbits or even attempt to remove the sun. *Right:* Robots represent atoms with Brownian motion. The learners can directly interact with atoms to change their temperature or add/remove atoms to experiment with conditions under which molecules form and break down.

held by the learner or teacher to interact with the activity; and provide planar haptic feedback. In addition, all robots are wirelessly connected to tablets (or smartphones) that coordinate the activity and the robots, and optionally provide additional dynamic content or input to the activity.

In this article, we present the design of the Cellulo platform and two preliminary field studies that do not yet aim to show clear learning gains but rather to demonstrate the Cellulo’s potential as an educative tool by illustrating its capabilities. After discussing the related work in Section 2, we describe the developed technologies that allow to fulfill the aforementioned requirements in Section 3. Then, we present the field studies co-developed with designers in Section 4. Finally, we give our outlook and future plans for Cellulo in Section 5.

2. RELATED WORK

Cellulo draws inspiration from – and in return, contributes to – several fields: educational robotics, clearly, but also *manipulatives* seen in Tangible User Interfaces (TUI).

Educational robotics.

The application domain for educative robots is split between robots used in programming, robotics, or other closely related STEM subjects where robots occur as a natural and suitable tool; or robots used in other disciplines, not necessarily technical or even scientific.

Many studies have been conducted studying robots in STEM, like [22, 8, 1, 13, 21]. [2, 11] provide reviews of both the studies and the devices used for this purpose. In terms of robots, individual programmable mobile robots (often with a differential drive for locomotion) and programmable robotic construction kits are the two major choices for teaching robotics, programming and related STEM subjects.

In the case of robots used for non-programming related curricula, the learning scenarios reported in the literature usually involve robots that possess social qualities (*i.e.* social robots) and often adopt the role of a tutor or peer; humanoid robots are unsurprisingly often found in such scenarios where social aspects of interaction are exploited for learning. [10, 7, 15, 3, 23, 12] are examples where the teaching of such topics

is explored; [14] is a rare example that uses a non-social robot to teach an artificial vocabulary (where the goal is ultimately to interact with robots).

In this regard, the Cellulo project investigates the design and use of robots in education where it is not to be perceived as a robot (*i.e.* not programmed or constructed) or a social tutor/peer; but as an “on-demand” support tool for teachers the be used in a multitude of teaching domains.

Manipulatives.

Manipulatives are physical objects specifically designed to foster learning. Cellulo, being based on robotic devices, is an extension of digital manipulatives (introduced in [20]). [24] proposes a classification of manipulatives into *Froebel Manipulatives (FiMs)*, manipulatives to model the world) and *Montessori Manipulatives (MiMs)*, manipulatives to model abstract structures). We believe that Cellulo, whose design does not enforce any application case, can be equally used as *FiMs* (*e.g.* concepts in Figure 1) or as *MiMs* (*e.g.* points on a cartesian plane in a geometry activity).

Manipulatives that are accurately 2D-localized in (x,y,θ) have been studied before; our design, which includes an open-source implementation of a localization technique developed in the context of handwriting recognition [19], achieves similar performances as [17] for a fraction of the cost. Besides, as each robot localizes itself independently, it is scalable and does not require the deployment of an external apparatus.

Being based on paper, our approach is especially well-suited for the classroom compared to techniques involving special hardware (typically, a table-top AR environment like [17]). Beyond its evident advantages in terms of costs and scalability, paper integrates smoothly into classroom ecosystems. It matches several principles for minimizing classroom orchestration load, such as flexibility and minimalism ([4]).

Using tangible interfaces to teach has been extensively explored and numerous applications have been developed covering various domains (see [4] for a review). Our approach is novel in that our manipulatives are robots and can move by themselves. As highlighted in the concepts (Figure 1) and further discussed hereafter, we believe this opens a new venue of research by bridging Educational HRI and TUIs.

Table 1: Cost of the Cellulo robot prototype.

Component	Cost (€)
Processor (<i>PIC32MZ1024ECG064</i>)	8.51
Localization subsystem [9]	17.24
Locomotion subsystem [16]	49.77
Communication (<i>RN42</i>)	12.12
UI subsystem (<i>TLC5947, AT42QT1070</i>)	7.36
Battery & charging & protection	8.86
Other (terminal, resistor, capacitor <i>etc.</i>)	8.77
Housing (53.3g PLA) & fastening	1.91
PCB manufacturing	11.39
Total	125.93

3. PLATFORM DESIGN

3.1 Overview

Our platform, pictured in Figure 1 in two concept activities, is composed of low-cost mobile robots that operate on printed sheets of paper, overlaid with a microdot pattern that enables self-localization. These sheets are printable in regular home/office printers or in printing houses and their length is not bounded (in our experiments, we used paper playgrounds up to 1m×2.4m). The robots are designed to be small, sturdy, low-cost and simple to operate; all robots are connected wirelessly to a mobile device (a tablet or smartphone) that runs the activity and orchestration logic.

The current Cellulo robot is the result of two design iterations (Figure 2), and it includes: self-localization on the activity sheet; holonomic motion; 6 capacitive touch buttons (independently back-illuminated in full RGB) and wireless Bluetooth communication. All the components found in the robot can be seen in Table 1 along with their costs at the prototype stage. Below, we detail the design of the various subsystems involving the robot and the paper, and how they come together to form the Cellulo platform.

3.2 Localization & Role of the Paper

In order to achieve complex, swarm-like robot motions in the classroom environment, we require fast and accurate localization of many robots in a way that is straightforward to deploy and store away, and is robust against kidnapping and occlusions due to childrens’ manipulations. Our method, described in detail in [9], uses a dense, deterministic and well-defined optical microdot pattern printable on regular office printers, with enough x, y space to cover over 170 million km^2 with unique patterns. The robots, equipped with a downward facing camera and a low-cost microcontroller, decode a $\sim 1\text{cm}^2$ region on the printed activity sheet to obtain their own 3DOF pose (x, y, θ) at about 93Hz. This method offers global and absolute localization with $\sim 0.27\text{mm}$ and $\sim 1.5^\circ$ accuracy without the need for any calibration. Moreover, there is no limitation on the number of robots that can be localized since the entire algorithm runs onboard the robot. The physical placement of the camera beneath the robot and the hardware design (global shutter sensor, high framerate) allows any manipulation to be performed on the robot without any adverse effect to localization performance, instant recovery from kidnapping, and perfect robustness against external illumination conditions.

The interactions between the robots, the learners and the

activities rely partly on this paper-supported localization mechanism. From the larger perspective, the activity itself is printed on large paper sheets that can feature any desired graphical elements. This capacity allows the definition of “active zones”: arbitrary-defined areas of the activity sheet that are associated with specific robot behaviors (*e.g.* green zones that deactivate the robots in Figure 1, left; red and blue reactor zones that dictate different pressure/temperature conditions in Figure 1, right). The raw robot positions are also used, *e.g.* for onboard closed-loop motion control (including haptic feedback) and, on the external controller, for multi-robot formation control.

From another perspective, absolute global localization-enabled paper opens up further possibilities, similar to the ones discussed in [6], in addressing usability concerns in the classroom. Smaller paper sheets (possibly down to A4 or A5) can be overlaid with unique patterns for each learner and the teacher; other functional sheets can also be designed to affect the flow of the activities, behaviors of the robots or host subtasks within activities. These can be freely passed around and used on any flat surface without the requirement of any other device than the robot, potentially “making the educational workflow tangible”.

3.3 Locomotion & Haptics


The Cellulo robot is capable of holonomic motion by virtue of the permanent-magnet assisted omnidirectional ball drive, previously presented in [16] (Figure 2b). Holonomic motion is a significant improvement over differential drive motion found in nearly every educational mobile robot, as it allows the robot to start moving in any direction; change the direction of motion near-instantaneously; preserve a high level of motion freedom in cluttered environments like multi-robot formations; and finally be moved in any direction without damage when forcibly manipulated by the learner. The drive also offers mechanical robustness against intensive user manipulation and potential mishandling, while remaining compact enough to fit inside a handheld device with low cost and (almost exclusively) off-the-shelf components.

We have as well to address the integration of a haptics controller to a traditional motion controller that allows the robot to *move* and *be moved at the same time*. This requires a hybrid motion controller that appropriately combines different sources of commanded motor outputs depending on the activity logic, where these sources might be: (i) drive velocities for the robot to move independently from one pose to another, (ii) tactile or kinesthetic haptic feedback in the form of force/torque output to the learner’s hand, (iii) drive velocities to assist the learner to move the robot from one pose to another in order to overcome the natural impedance of the robot on the activity sheets due to the natural friction of the blocked wheels. The development and evaluation of such a controller is currently in progress.

3.4 Software Design & Orchestration

The Cellulo platform is designed to be activity-driven rather than driven by the capabilities of the platform itself. An *activity* is the combination of the paper elements, the robots with particular interaction modalities and the tablet(s) that run(s) the activity-specific software. As such, the role of the robots and paper depends on the design of each particular learning activity.

The most basic element of an activity is composed of the



(a) 1st revision: localization and illuminated touch buttons implemented. On the left: upside down robot with exposed optical system. (b) 2nd revision: locomotion added, ergonomics and appearance refined. On the left: upside down robot with open housing exposing locomotion components and optical system in the center. Two ball wheels out of sockets for visibility.

Figure 2: Two major revisions of the Cellulo robot design. 1-to-1 scale preserved between the images.

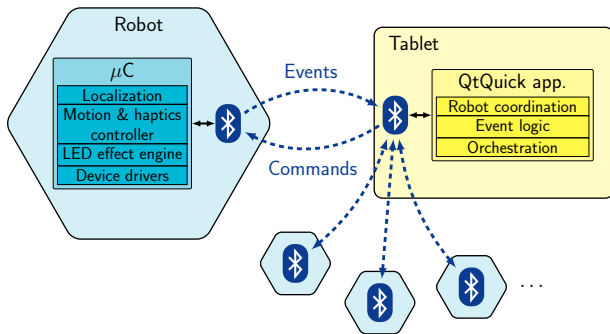


Figure 3: Software architecture of the most basic element of a Cellulo activity. Multiple robots are connected through Bluetooth to a mobile tablet where the activity logic and orchestration software resides.

activity sheet, hereafter called the *playground*, the robots that operate on the playground and the tablet. The robots run latency or bandwidth-sensitive software components onboard within the single-threaded bare-metal firmware written in C, such as the motion/haptics controller, image processing for localization and the LED effect engine. Each robot is connected through a Bluetooth 2.1 Serial Port Profile (SPP) channel to the tablet that runs a *QtQuick* application to coordinate the activity, as well as host graphical input/output elements such as buttons and often illustrations required by the activity. The architecture of these core elements can be seen in Figure 3.

From a software development and deployment point of view, QtQuick allows rapid development of the activity software using Qt Modeling Language (QML, JSON-like declarative language) with Javascript and supports deployment to many platforms and devices. In addition, any performance-sensitive components can be developed in C++ as native back-end plugins to be loaded and interfaced from within QML/Javascript. This allows us to deploy our applications to consumer mobile devices (running iOS or Android) that we develop and test on desktop computers (running Linux, Windows or Mac OS) while seamlessly using hardware resources (such as Bluetooth SPP sockets).

The paper elements, including the playground, are exported and transported as PDF files at the end of their graphics design phase. With external tools that we developed, we overlay the microdot pattern with any desired origin coordinate on these sheets. We can define the active zones with a simple custom authoring tool that operates on these PDFs and exports the zones in JSON format which are then loaded into our mobile applications. The dotted PDFs are easily printed on any platform or sent to printing houses with poster printing facilities if they are in larger sizes. It is useful to note that this workflow design, while alleviating development by providing correct tools, does not facilitate educational content creation and an education professional must be present to design the didactic sequence.

One tablet, as depicted in Figure 3, can orchestrate an activity with a maximum of 7 robots (due to Bluetooth limitations). Classroom-wide activities might require more robots. In such cases, more tablets (or other Bluetooth controllers, like smartphones) are required. For such scenarios, we envision the possibility for the teacher to have a “master” tablet connected to the pupils’ tablets over WiFi, and allowing him/her to monitor, start, pause, control, record activities at the whole classroom level. This teacher-centric approach would support the classroom orchestration by the teacher, and address some of the issues raised in [5]. This is left as future work.

4. EXPLORATORY STUDIES

4.1 Interaction Space

The platform design opens up numerous possibilities in terms of educative interactions with the activities. We investigate this interaction design space along three dimensions:

- *Passive – Autonomous robots*: at one end, the robots are passive and manually moved by the children (pure *tangible* devices); at the other end, they are entirely autonomous and the children do not manipulate them.
- *Independent robots – Multi-robot environment*. “Independent robots” act independently and unknown to other robots (it may as well be an interaction with a single robot); “Multi-robot environment” refers to

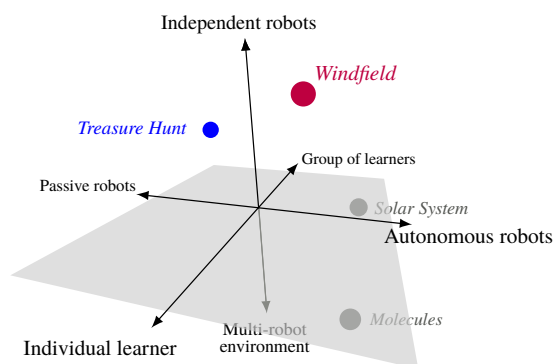


Figure 4: Three dimensions of Cellulo interaction envisioned so far. The two colored dots represent the two reported studies. The two gray dots represent the two concept activities depicted in Figure 1.

activities where the interactions between robots are numerous and essential.

- *Individual learner – Group of learners.* This dimension represents the learning process from the students’ perspective: individual or collective.

Importantly, these dimensions are not binary: learning activities might be located anywhere in this 3 dimensional interaction space, as depicted in Figure 4. We present in the following sections two exploratory studies that probe two points in this space: the “Treasure Hunt” (Section 4.2) game, and the “Windfield” activity (Section 4.3).

These two studies played a key role in the platform design process as real-world testbeds for our prototypes. The technical performance of the prototypes as well as their acceptance by the children during these studies is reported along with our observations regarding the emergent interactions with the Cellulo platform.

4.2 Treasure Hunt study

4.2.1 Activity Design

In this study, the robots are used as tangible items (*i.e. passive robots*) by a group of children. Each child is given their own robot to interact with the activity (Figure 5). We used the robots resulting from the first design iteration (Figure 2a). This instance used robots with the localization component and illuminated touch buttons as well as a 1m×2.4m playground and the mobile application with playful graphics and instructions running on the tablet. The children completed three consecutive tasks, some of which had collaborative aspects, in order to eventually let their pirate characters reach a treasure chest.

The activity was designed to explore various interactions (seen in Figure 6) enabled by the Cellulo platform in a playful situation where specific learning outcomes were not targeted. The goals of this study were first to assess the usability and reliability of the pattern-based localization system when manipulated by children “in the wild”; second, to evaluate the combination of robots, paper and tablet over a range of interaction modalities in terms of legibility and intuitiveness.

The activity was run with a total of 85 children (11 to 14 years old) with no prior experience with Cellulo. The children were split into 14 groups of 5 (each child has one



Figure 5: Treasure Hunt activity, exploring *passive* and mostly *independent* robots in a *group* activity.

robot and the tablet is shared) or 6 (5 children have one robot each and one child operates the tablet and reads the instructions). After a brief introduction, the children were instructed to follow the indications on the tablet and no further instructions were given by the experimenters. The game lasted around 12min (M=11min 47s, SD=1min 47s, min=9min 19s, max=15min 32s), and the children were invited to replay it if they wished (replay data not included in discussion).

Each child was first assigned one pirate character, represented by his/her robot and one character card (Figure 6b). Four tasks were then completed sequentially:

Task 1: Reproduce a sequence of 6 flashing LEDs by tapping the corresponding buttons (performed by all 5 children in parallel, each had random sequence); interaction pictured in Figure 6c;

Task 2: Scan the playground for the hidden key to the treasure chest, avoiding false positives (performed by all 5 children in parallel, children would spontaneously coordinate to explore different regions of the playground); interaction in Figure 6d;


Task 3: Spin the robot to pull up the treasure chest from the well (performed by the first child in each group), with the interaction in Figure 6e, similar to [18];

Task 4: Move the robot along the paths on the planks and rocks on the playground (4 paths performed sequentially by the remaining 4 children, the correct paths to follow were drawn on the back of the character cards); interaction in Figure 6f.

Before each task, the robots had to be placed on specific areas on the playground marked with “A”, “B” and “C” (as in Figure 6a). The activity on the tablet would wait until the five robots are correctly placed to start the next task. This simple multi-robot behavior enabled effective synchronization points amongst the group of children.

4.2.2 Results & Discussion

The localization subsystem was observed to have mostly satisfactory performance in terms of accuracy and responsiveness. No performance decrease could be observed after




(a) *Positioning*: Placing the robots in active zones to trigger events.




(b) *Role Assignment*: Placing the robots on character cards; correct assignment lights the robot in green.




(c) *Touch*: Pressing the illuminated capacitive touch keys on the robot, e.g. in a pre-defined sequence.



(d) *Scan*: Moving the robot on the playground to find a hidden target; as the robot gets closer to the target, the color progressively changes from the “coldest” to the “warmest”.



(e) *Spin*: Turning the robot in place to change a continuous value; an animation on the tablet tied to this value accompanies the interaction.



(f) *Follow the path*: Moving the robot through the path illustrated on an auxiliary sheet without leaving the path.

Figure 6: Interaction methods used in Treasure Hunt.

Table 2: Treasure Hunt task completion times, $N = 14$ groups, all values in seconds.

	M (SD)	Min	Max
<i>Task 1: Touch sequence</i>	131.1 (35.4)	85	199
<i>Task 2: Scan</i>	125.2 (51.6)	47	222
<i>Task 3: Spin</i>	13.6 (4.3)	5	21
<i>Task 4: Follow path</i>	162.9 (37.4)	114	224

slight progressive wear on the playground print during one day of intensive use. Three activities were run in parallel in the same room, with 3 groups on 3 tablets and 15 robots; none of the Bluetooth connections dropped at any time but slight delays were observed, which may have been caused by wireless interference from other Bluetooth devices or from the many Wi-Fi routers present on the experiment site.

The children interacting for the first time with the robot (even though some were at first reluctant to interact with the unknown technology) were observed to quickly understand that the robots “knew where they are” and the proposed interactions were easily picked up. Importantly, the constantly changing role of the robot (from a pirate character, to a pad to enter a code – Figure 6c, to a metal detector – Figure 6d, to a well handle – Figure 6e) was well accepted by the children, the tasks were all completed by all groups without any assistance (Table 2 provides the mean completion times for each of the tasks). This suggests that our goal of designing a versatile device was successfully attained in this scenario.

During the touch key sequence task (which is essentially a memory task), children who completed their sequence were

generally observed to aid those who were still struggling with theirs, resulting in a cooperative aspect that naturally emerged. Likewise, the scan task naturally fostered collaboration amongst the children to effectively “scan” the various parts of the large playground.

The spin task was explained only by a depiction of the robot with two circular arrows (similar to Figure 6e) on the back of the character card. The interaction mechanism was picked up surprisingly fast (14 seconds on average to complete the task, as fast as 5 seconds for one group).

The “follow the path” task was completed by each child in about 40 seconds on average, and the completion was observed to be delayed by both children’s mistakes caused by misperceiving the paths on the character cards and by occasional localization errors due to excessive contrast on some playground graphics.

At a higher level, the main results of the first study are: (i) Children easily engage with tangible activities with printed graphics on paper; (ii) A variety of roles can be ascribed to small robots with an unremarkable, anonymous design; children easily accept and engage in such role assignment.

The next study explores how using an active, mobile robot instead of a passive one impacts the interaction design.

4.3 Windfield study

4.3.1 Activity Design

The second activity is a learning game, and explores a different point in the interaction design space: single robots are used as haptic devices by individual children (Figure 8a). The robots use the locomotion module developed during



Figure 7: Screen capture of the Windfield activity application running on the tablet. Position of the balloon is synchronized with the position of the robot on the real map. Low and high pressure areas are represented in blue and red respectively.

the second design iteration (Figure 2b), and a preliminary haptic feedback controller to convey the force generated by the winds resulting from atmospheric pressure differences.

The activity is designed to take place on the geographical map of Europe (printed on a $0.66\text{m} \times 1.7\text{m}$ playground) where multiple atmospheric pressure points (high and low) are placed. Following their real physical counterparts, these points affect the pressure over the entire activity area and their effects diminish with distance. Given all pressure points and their positions, atmospheric pressure values are calculated over a grid spanning the entire playground (resulting in a simplified Finite Element Analysis-like method). These values are then used to calculate the pressure gradient on a given position, which dictates the intensity and direction of the wind on that position (additional factors that affect the wind such as the Coriolis effect are not considered).

The robots play the role of “hot air balloons” that can be placed on any desired position on the playground and convey the force of the wind with haptic feedback while being grasped, as illustrated in Figure 8b. This is achieved by mapping the wind velocity in the global playground frame to the motor outputs in the local robot frame through inverse kinematics. This velocity output then results in a force output when the robot is grasped and blocked in place.

Two high and two low pressure points are randomly placed on the playground, and must be discovered by the children using haptic feedback. The children explore the playground by moving the robot on various locations and “feel” how the wind blows there. By following the winds, they can locate the pressure points. They indicate their guesses by dragging and dropping the corresponding icons on the tablet (where the playground and the location of their hot air balloon is displayed), as seen in Figure 7. Prior to the activity start, the participants are told what the activity is about and what the robots are standing for. The atmospheric pressure points and how they “blow wind outwards” and “suck winds inwards” are also introduced. After playing the game, the learners are expected to gain the knowledge that: (i) the wind blows from high to low pressure; (ii) the strength of the wind diminishes with increasing distance from pressure points.

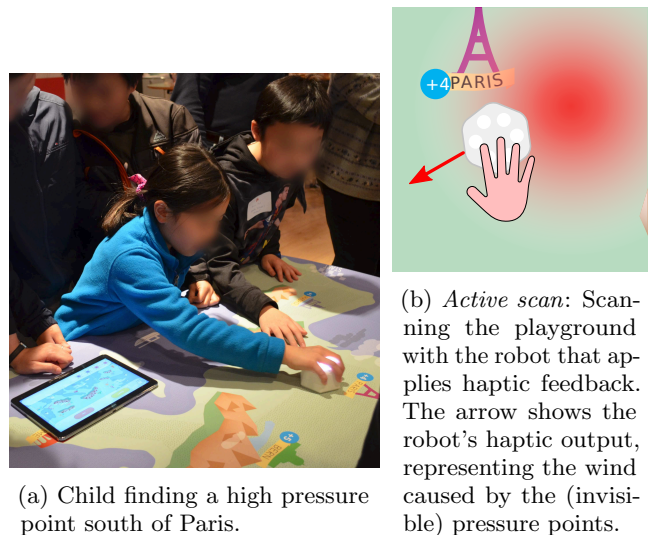


Figure 8: Haptic interaction in Windfield.

4.3.2 Results & Discussion

We ran the study during two open-door events on two consecutive days for more than 6 hours in total. The two events attracted about fifteen thousand attendees, most of which were children accompanied by their parents. The attendees came to our booths at their discretion and stayed as much as they wanted. Under these circumstances, with 3 activities running in parallel on 3 separate playgrounds, we estimate that about 150 children ranging from 5 to 14 years old interacted with our activity non-trivially (*i.e.* excluding children who were just looking and were manipulating the robots/tablets aimlessly and leaving) attempting to find the hidden pressure points for durations ranging from 2min to 15min (Figure 8a). Due to uncontrolled conditions, we report hereafter qualitative observations about the interactions.

Children engaged easily with the activity: they did not show reluctance to manipulate the robots. The haptic modality of conveying a planar force was generally well-received, most children were able to tell where in which direction the robot was pushing their hand when asked. In addition, there were multiple questions like “Where can we buy these [robots]?”, “Are you giving these [robots] away?”, “Can we take this [playground sheet] with us?” that suggest that the platform was approachable and attractive. In a few instances, the parents did not let their children approach the activity due to the presence of a tablet; these children were presumably being raised in a “mobile device-free” manner until a certain age. This needs to be accounted for in the adoption process of our platform, and remediation (like a tablet-free mode of operation) may be sought after.

From a hardware design perspective, force/torque sensors on the grasped surface are in principle required to control the force feedback in a closed loop. Since these are not present on our current robots, we detect grasps by the touch buttons on the top surface and disable the haptic feedback when the robot is not grasped. This requires the learners to grasp the robot from the top, and causes side grasps and top grasps where the palm does not touch the robot surface to remain undetected. In addition, this does not guarantee the actual transmission of forces to the learner’s hand, resulting

in approximate force feedback which occasionally results in the robot’s motion due to weak grasps.

From an interaction point of view, the effectiveness of the *active scan* (Figure 8b) interaction was observed to depend on the directionality of the force: Low pressure points are almost always found before high pressure points since ‘following the wind’ naturally leads the robot to one of the sink holes, *i.e.* a low pressure point. On the contrary, high pressure points are more difficult to find as they push the robots away and require the exploration of the surrounding area to be uncovered. In that sense, the design of a haptic-enabled activity requires care.

From the learning perspective, the age group of the participants had a clear impact on their performance. Children under 11-12 years old (the majority of our participants) often could not find the high pressure points. The ones that did find them were still not able to give satisfactory answers to questions like “What do you think the wind is doing?” or “How do you think wind is connected to the pressure points?”. Youngest participants (5-6 years old) were observed to have difficulties reaching over the far side of the playground and grasping the robot; this prevented them from properly interacting with the activity and underlines that the expected manipulations of the robots and physical size of the playground must be adjusted to the target age group.

Participants over 11-12 years old were generally able to understand that pressure points acted in all directions by scanning around the points and finding them. Some children exhaustively searched and found all pressure points by spending an abundant time in the activity. Some were able to provide answers to the questions like “It seems to blow from these high points to these low points.”, often pointing at the pressure points on the playground itself. This indicates that children were able to transfer their findings from the tablet to the physical playground, and suggest – at least for older children – that the combination paper, robot, tablet can be perceived as a single, multi-modal, educational tool.

At higher-level, the ‘Windfield’ study confirms the educational potential of the complete platform: it demonstrates a set of rich interactions between paper (the printed activity playground), robots (autonomous, self-localizing mobile devices offering haptic interaction) and an active pupil. This combination enables the child to physically experience a complex and invisible phenomenon (the interplay of the atmospheric pressures) in a natural and playful manner.

5. CONCLUSION

5.1 Summary

This article’s main contributions are two-fold: First, we detailed the design and implementation of a novel tangible robotic device that takes into account the constraints of real and unsupervised use in classrooms; to the best of our knowledge, building such an educative platform that combines mobile robots with paper-based tangible and haptic interactivity has not been attempted before. Our design process was driven by the classroom environment and the teachers’ constraints (including ease of use, low disturbance, reliability), and led to the development of a small, low-cost and robust wireless robot (*i.e. practical*) that is aimed to blend in the classroom ecosystem through its rich interactions with the traditional and pervasive paper (*i.e. ubiquitous*). In that sense, it seeks to expand (rather than replace) current

teaching practices through a multitude of activities that cover a variety of disciplines (*i.e. versatile*).

Second, we presented two exploratory studies, *Treasure Hunt* and *Windfield*, ran in ecological environments that showed the feasibility to combine robots with augmented paper sheets, to easily deploy the robots and to engage the children in these sample activities. They contained very different sets of interaction affordances utilizing both active and passive robot behaviors, and were well received by the children. Although the activities were not designed to measure clear learning gains, they are highly useful in communicating the versatility, practicality and ubiquity potential of the platform to education professionals, with whom we are now beginning to actively collaborate.

The interaction design space created by the Cellulo platform is large: building upon these two studies, we have proposed an initial formalization of this design space, along three dimensions: level of autonomy of the robot; number of robots; number of interacting pupils. Within this space, we have experimented with and identified a range of interaction patterns that combine the robots and the paper playground.

5.2 From Here to Adoption - Future Work

As we conclude these initial stages of design and testing, we acknowledge that we still have many tasks to accomplish in order to prove the potential of the Cellulo platform. For one, the haptic interaction component was designed up to now in a shallow manner, and requires more work to be adequately useful. Our current work involves the proper design of this component and explicit testing of the usability of this joint control paradigm that is *moving vs. being moved* through concrete manipulation measures.

From the educational perspective, the design process of activities should naturally involve teachers (to conceive the learning scenarios), interaction designers (to propose interaction patterns), graphic designers (to produce the visuals on paper and tablet) and software engineers (to implement the applications running the activities). The studies presented in this article were focused on exploring natural interactions with Cellulo robots rather than creating actual educative content; as such, they were co-developed by an interaction designer, a graphic designer and educational researchers who developed the software. Our current and future work involves the design of educational activities with the guidance of teachers, who are in fact also meant to be one of the main users of our platform. With these, we will aim to show that Cellulo is effective for the actual learning of multiple subjects within a real curriculum.

Following this, we plan to focus our research on the integration of Cellulo in classrooms as part of standard learning curricula (as conceptualized in Figure 1), building on co-design techniques with education professionals to ensure suitability and foster acceptance. Finally, we aim to test this integration with focused deployment studies where we will measure our platform’s practicality and overall acceptance in the long term.

6. ACKNOWLEDGEMENTS

This work has been partially supported by the Swiss National Science Foundation through the National Centre of Competence in Research Robotics, Gebert R uf Stiftung and the EU H2020 Marie Sklodowska-Curie Actions project DoRoThy (grant 657227).

7. REFERENCES

- [1] T. Balch, J. Summet, D. Blank, D. Kumar, M. Guzdial, K. O'Hara, D. Walker, M. Sweat, G. Gupta, S. Tansley, J. Jackson, M. Gupta, M. N. Muhammad, S. Prashad, N. Eilbert, and A. Gavin. Designing Personal Robots for Education: Hardware, Software, and Curriculum. *IEEE Pervasive Computing*, 7(2):5–9, 2008.
- [2] F. B. V. Benitti. Exploring the educational potential of robotics in schools: A systematic review. *Computers & Education*, 58(3):978–988, 2012.
- [3] C.-W. Chang, J.-H. Lee, P.-Y. Chao, C.-Y. Wang, and G.-D. Chen. Exploring the Possibility of Using Humanoid Robots as Instructional Tools for Teaching a Second Language in Primary School. *Educational Technology & Society*, 13(2):13–24, 2010.
- [4] S. Cuendet, Q. Bonnard, S. Do-Lenh, and P. Dillenbourg. Designing augmented reality for the classroom. *Computers & Education*, 68:557–569, 2013.
- [5] P. Dillenbourg and P. Jermann. *New Science of Learning: Cognition, Computers and Collaboration in Education*, chapter Technology for Classroom Orchestration, pages 525–552. 2010.
- [6] 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. In *Connecting Computer-Supported Collaborative Learning to Policy and Practice: CSCL2011 Conference Proceedings*, volume 1, pages 510–517, 2011.
- [7] J. Han and D. Kim. r-Learning Services for Elementary School Students with a Teaching Assistant Robot. In *Proceedings of the 2009 4th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 255–256, 2009.
- [8] K. Highfield, J. Mulligan, and J. Hedberg. Early mathematics learning through exploration with programmable toys. In *Proceedings of the Joint Meeting of PME 32 and PME-NA*, pages 169–176, 2008.
- [9] L. O. Hostettler, A. Özgür, S. Lemaignan, P. Dillenbourg, and F. Mondada. Real-Time High-Accuracy 2D Localization with Structured Patterns. In *Proceedings of the 2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4536–4543, 2016.
- [10] T. Kanda, T. Hirano, D. Eaton, and H. Ishiguro. Interactive Robots as Social Partners and Peer Tutors for Children: A Field Trial. *Human-Computer Interaction*, 19(1):61–84, 2004.
- [11] M. E. Karim, S. Lemaignan, and F. Mondada. A review: Can robots reshape K-12 STEM education? In *Proceedings of the 2015 IEEE International Workshop on Advanced Robotics and its Social Impacts (ARSO)*, pages 1–8, 2015.
- [12] S. Lemaignan, A. Jacq, D. Hood, F. Garcia, A. Paiva, and P. Dillenbourg. Learning by Teaching a Robot: The Case of Handwriting. *IEEE Robotics Automation Magazine*, 23(2):56–66, 2016.
- [13] F. Mondada, M. Bonani, X. Raemy, J. Pugh, C. Cianci, A. Klaptocz, S. Magnenat, J.-C. Zufferey, D. Floreano, and A. Martinoli. The e-puck, a Robot Designed for Education in Engineering. In *Proceedings of the 9th Conference on Autonomous Robot Systems and Competitions*, volume 1, pages 59–65, 2009.
- [14] O. Mubin, C. Bartneck, L. Feijs, H. H. van Huysduynen, J. Hu, and J. Muelver. Improving Speech Recognition with the Robot Interaction Language. *Disruptive Science and Technology*, 1(2):79–88, 2012.
- [15] S. Y. Okita, V. Ng-Thow-Hing, and R. Sarvadevabhatla. Learning Together: ASIMO Developing an Interactive Learning Partnership with Children. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 1125–1130, 2009.
- [16] A. Özgür, W. Johal, and P. Dillenbourg. Permanent Magnet-Assisted Omnidirectional Ball Drive. In *Proceedings of the 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016.
- [17] J. Patten, H. Ishii, J. Hines, and G. Pangaro. Sensetable: A Wireless Object Tracking Platform for Tangible User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '01*, pages 253–260, 2001.
- [18] E. W. Pedersen and K. Hornbæk. Tangible Bots: Interaction with Active Tangibles in Tabletop Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*, pages 2975–2984, 2011.
- [19] M. P. Pettersson. Method and Device for Decoding a Position-Coding Pattern, Dec. 5 2006. US Patent 7,145,556.
- [20] M. Resnick, F. Martin, R. Berg, R. Borovoy, V. Colella, K. Kramer, and B. Silverman. Digital Manipulatives: New Toys to Think with. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '98*, pages 281–287, 1998.
- [21] F. Riedo, P. Rétronaz, L. Bergeron, N. Nyffeler, and F. Mondada. A Two Years Informal Learning Experience Using the Thymio Robot. In *Advances in Autonomous Mini Robots*, pages 37–48. Springer, 2012.
- [22] E. Schweikardt and M. D. Gross. roBlocks: A Robotic Construction Kit for Mathematics and Science Education. In *Proceedings of the 8th International Conference on Multimodal Interfaces*, pages 72–75, 2006.
- [23] F. Tanaka and S. Matsuzoe. Children Teach a Care-Receiving Robot to Promote Their Learning: Field Experiments in a Classroom for Vocabulary Learning. *Journal of Human-Robot Interaction*, 1(1), 2012.
- [24] O. Zuckerman, S. Arida, and M. Resnick. Extending Tangible Interfaces for Education: Digital Montessori-inspired Manipulatives. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '05*, pages 859–868, 2005.