

Designing Augmented Reality for the Classroom

Abstract

Augmented reality (AR) has recently received a lot of attention in education. Multiple AR systems for learning have been developed and tested through empirical studies often conducted in lab settings. While lab studies can be insightful, they leave out the complexity of a classroom environment. We developed three AR learning environments that have been used in genuine classroom contexts, some of them being now part of classroom regular practices. These systems and the learning activities they provide have been co-designed with teachers, for their own classrooms, through multiple cycles of prototyping and testing. We present here the features that emerged from these co-design cycles and abstract them into design principles.

Keywords: tangible interfaces; augmented reality; classroom orchestration; orchestration load; collaborative learning; improving classroom teaching; interactive learning environments; teaching strategies

1. Introduction

Augmented reality refers to technologies that project digital materials onto real world objects. This definition suits a large spectrum of technologies that range from a pure virtual environment to the real environment, as summarized by Milgram & Kishino (1994). In this article, we focus on augmented reality in the context of tangible user interfaces (TUIs) for tabletop environments.

The use of AR and TUIs can be an asset for learning. AR systems allow the learner to interact with the real world in ways that were not possible before. They create new situations that would be impossible to create in just the real world or just the digital environment. As pointed out by Billinghurst (2002), the use of tangibles as a metaphor for object manipulation and the ability to smoothly transition between reality and virtuality can create new educational experiences. As for TUIs, they also have the potential to improve the learning experience, although questions have been raised about their real impact on learning (Marshall, Rogers & Hornecker, 2007). Many AR and TUIs systems have been developed for education but few of them have been deployed into real classroom environments.

Following the work of Kerawalla, Luckin, Seljeflot & Woolard (2006), we report the lessons learned from bringing tangible AR systems into classrooms. While AR systems are often associated with informal learning, we explored their impact in formal education, both in primary and secondary schools. This immersion into school ecosystems produced design considerations that go beyond pedagogical criteria – whether this activity will trigger learning outcomes – and take into account the diverse constraints of classrooms, such as time, space, discipline, or curriculum. We present three different learning environments refined through many field studies and several years of participatory design involving teachers. Based on our experience designing and evaluating those systems as well as on the literature about classroom orchestration, we propose five principles for “designing a learning environment for a classroom”: *integration, awareness, empowerment, flexibility, and minimalism*. The design principles are illustrated with specific features from the three learning systems.

2. Related Work

A full history of AR is beyond the scope of this work. We focus our review of the previous work on TUIs, a branch of AR. Although TUIs became popular with Ishii and Ullmer's work (1997), tangible interfaces had been used before (Wellner, 1991; Fitzmaurice, Ishii & Buxton, 1995). Wellner (1991) created the DigitalDesk, a system meant to bridge the divide between digital and paper documents. Fitzmaurice et al. (1995) introduced the Graspable User Interface which is an interface made of small physical artifacts that allow direct control of virtual objects. In an evaluation of the graspable interfaces, Fitzmaurice & Buxton (1997) hinted that the mouse is a general all-purpose weak device that can be advantageously replaced by strong specific devices for a specific and limited task. Indeed, attaching digital contents to physical tokens allows for rich, two-handed and parallel interaction. Since then, the work on TUIs has been abundant and many TUI applications have been developed in several domains: optics (Underkoffler & Ishii, 1998), urban planning (Underkoffler & Ishii, 1999; Arias, Eden, Fischer, Gorman, & Scharff, 2000), chemistry and system dynamics (Patten, Ishii, Hines & Pangaro, 2001), or topology (Piper, Ratti & Ishii, 2002; Patten & Ishii, 2007).

TUIs involve physical mechanisms that play a role in learning, such as gesturing, physical movement and embodiment (Cuendet, Jermann, Dillenbourg, 2012; Goldin-Meadow 2003; O'Malley & Fraser, 2004; Roth, 2000). Gestures can provide external representations of a problem or an object, which play a role in problem solving and learning (Ainsworth, 1999; Larkin & Simon, 1987), by helping the learner to make inferences or by

freeing up cognitive load. The "sense of presence" described by Mantovani & Castelnuovo (2003) allows experiential, situated learning: connections made to previous knowledge is stronger because AR provides a balance of immersion and real interaction. Another important dimension for learning is the coupling between cognition and physical experience (O'Malley & Fraser, 2004).

Hornecker & Buur (2006) explained that tangible interaction goes beyond the coupling of virtual and physical objects (a *data-centered* view). For example, the *expressive-movement-centered* view focuses more on the interaction itself than the physical-digital mapping. In a pedagogical concept, it corresponds to the possible enactments of concepts by students. Furthermore, the *space-centered* view of tangible interaction focuses on the position of the user in space. This deals for example with aspects like the division of labor, which can result naturally from the spatial disposition of tangible resources, which orient the roles of learners in a group, as shown by Jermann, Zufferey, Schneider, Lucchi, Lépine & Dillenbourg (2009).

The space in which users of a pedagogical TUI evolve, the classroom, is a complex and peculiar environment. Don Norman's assertion about everyday life applies perfectly to the classroom: "*It is complex, not because any given activity is complex, but because there are so many apparently simple activities, each with its own set of idiosyncratic requirements. Take a large number of simple actions and add them up, and the overall result can be complex and confusing: the whole is greater than the sum of its parts.*" (Norman, 2010, p.64)

The testing of learning environments in lab settings ignores this complexity of the classroom, and the results of lab studies can therefore reveal to be wrong once the learning environment is deployed in the classroom. For example, Stringer, Rode, Toye & Blackwell (2005) illustrated the richness of the classroom environment through their many iterations in the development of a TUI for learning to form arguments. They trialed tangible prototypes in classrooms through eleven iterations, which contributed to invalidate their initial assumptions regarding how TUIs might be useful. Similarly, Stanton, Bayon, Neale, Ghali, Benford, Cobb et al. (2001) stressed the importance of evaluating TUIs in the classrooms rather than in a lab, and drew several lessons from the design of their system for Storytelling: physical size and props encourage collaboration by making interaction more visible, content created using one interface might be difficult to access with another, superficial changes can have effects that are hard to predict, and low technologies are more important than over-polished products, because users are more familiar with them. All these observations in a real environment only reveal the need for a higher attention to the classroom in the design of pedagogical TUIs.

3. Usability at the classroom level

Instructional design consists in elaborating learning activities and learning technologies that take into account the targeted learning outcomes, the specificity of the contents to be learned, the peculiarities of the learners, and the principles of educational psychology, i.e. how people learn. These elements can be described as *intrinsic* to the learning process. Over the recent years, many scholars started to also pay attention to *extrinsic* constraints, i.e. constraints that do not relate to any learning theory, but nonetheless shape classroom practices, such as the time budget and time segmentation, the physical constraints of the classroom space, and the need to maintain a reasonable level of discipline, to minimize the teachers' workload, or to cope with the heterogeneity of the class (Dillenbourg & Jermann, 2010; Moraveji, Morris, Morris, Czerwinski & Riche, 2011; Roschelle, Rafanan, Estrella, Nussbaum & Claro, 2009; Prieto, Villagra-Sobrino, Jorin-Abellan, Martinez-Mones & Dimitriadis, 2011). By referring to them as the 'logistics' of classrooms, Nussbaum and Diaz (2011) stress that these constraints do not correspond to a grand learning theory but to practicalities that, if they are neglected, may spoil the most effective instructional design. Simply stated, both internal and external constraints need to be satisfied to obtain a learning solution that "works well" in the classroom.

Understanding why a tool "works well" has a name in HCI: usability. In the interaction between a user and a system, this variable is estimated by measures of performance, rate of errors, or user satisfaction. The design has to take into consideration individual constraints such as the user's previous experience and his or her cognitive load. We defined this as the first circle of usability (Dillenbourg, Zufferey, Alavi, Jermann, Do-Lenh, Bonnard et al., 2011). When two or more users interact with each other through technologies, usability is measured by factors such as the quality of conversations, the richness of deictic gestures, and the smoothness of turn taking. At this second circle of usability, design considerations are for instance the WYSIWIS principle (Stefik, Bobrow, Foster, Lanning & Tatar, 1987), and the 'least collaborative effort' (Clark and Brennan, 1991). The constraints of the classroom define the third circle of usability. In the same way Hutchins defined a cockpit as a distributed cognitive system (Hutchins, 1995), the classroom can be modeled as the user of a learning environment. We define the effort necessary for the teacher – and other actors – to conduct learning activities at this third circle as the *orchestration load*.

Classroom orchestration refers to the real-time management of multi-plane pedagogical scenario that maximize constraints satisfaction. The term 'multi-plane' refers to Vygostky's discrimination of intra-individual, inter-individual and social places. Many pedagogical approaches used in the classroom integrate individual, team and

plenary activities into a consistent scenario. While an AR system may be designed only for instance for the team plane, the integration of this system in a classroom will require embedding the teamwork within a scenario that often includes individual work (e.g. readings, summaries) and class-wide activities (e.g. introductory lectures, learners making presentations, teacher-led debriefing). The design of an AR/TUI environment must therefore anticipate how what learners did before the AR-based activity will feed this activity and how this activity will then provide substance to the following ones. In many environments, the pedagogical activities stretch beyond the individual, team, and class planes, to involve three higher planes: the classroom periphery (other classes, the parents, the school, etc.), the community (e.g. visits to a local craftsman), and the entire world (via the Internet). In other words, it is not simply the interaction of individual users or teams with the AR that needs a good usability, it is the smooth integration of the AR environment in the classroom workflow that is embedded in the notion of classroom usability.

4. Design Principles

In terms of 'AR for learning', some design principles have been proposed to optimize the first two circles of usability (e.g. Fishkin, 2004; Hornecker & Buur, 2006; Billinghamurst, Kato & Poupyrev, 2001; Fjeld, Bichsel & Rauterberg, 1999). Kerawalla et al. (2006) added design requirements that address the third circle: (1) AR systems should be flexible enough for the teacher to adapt to the needs of their students; (2) the content should be taken from the curriculum and delivered in periods as short as other lessons; and (3) the system should take into account the constraints of the context. The third requirement is obviously very broad. Our contribution aims to disentangle the many elements embedded under the word 'context'.

The general design principle could be: *"Classroom usability increases if the learning environment satisfies all classroom constraints"*. This principle would then be accompanied by a list of constraints, both intrinsic (i.e. how to make people cognitively learn) and extrinsic (how to effectively shape classroom practices with the presence of technologies). This is however not very helpful as the design difficulty with novel technologies is to anticipate how some design choices may later on violate classroom constraints, especially the extrinsic ones. We propose five principles that, retrospectively, proved to be good design choices while observing how our AR systems were used in classrooms. At this point, the definition of the design principles is made broad in order to limit their number but they will be refined through examples later in this article.

- Integration** *Orchestration load decreases if the learning environment is integrated in the workflow. As previously described, the AR tool often supports one activity among a heterogeneous set of classroom activities. Some activities are individual, in teams, at the class level or at higher levels. Some of them are computer-based, some not. The extent to which data are smoothly transferred between activities will impact the orchestration load.*
- Empowerment** *Orchestration load decreases if the learning environment allows the teacher to keep a central point in the classroom interactions when it is necessary (Dillenbourg & Jermann, 2010). Consider a classroom with 25 students using a tablet computer. If a teacher is trying to give an explanation to the entire class, he will fight several minutes against the distraction of the tablet in order to get the students' attention. While it is politically correct to define the teacher's roles as a facilitator, or a 'guide on the side', it is very difficult to manage a classroom while being on the side. This is not a plea for lecturing: a guided-discovery activity requires a smart drive of the classroom by a smart teacher assisted by technology.*
- Awareness** *Orchestration load decreases if the learning environment provides the teacher with permanent awareness of the state of all students in the class. Individual learner modeling would trigger orchestration overload, when facing 30 students. Awareness reduces orchestration load by providing – within one visualization – data aggregated on the time axis and/or on the social axis (Alcoholado, Nussbaum, Tagle, Gomez, Denardin, Susaeta et al., 2011; Moraveji, Kim, Pawar, Ge & Inkpen, 2008).*
- Flexibility** *Orchestration load decreases if the learning environment is flexible enough to adapt the activities to the evolution of the scenario (e.g. the state of students, the time remaining) and to accommodate unexpected events. Rigid pedagogical scenarios do not go well with the serendipity inherent in a classroom environment. For example, if the technology requires groups of four students, the design should anticipate how the teacher could adapt the learning activities if some students drop off.*
- Minimalism** *Orchestration load decreases if the learning environment does not provide more information (and functionalities) than what is required at a given time. As developers, we tend to over-estimate the amount of information that classroom actors – in our case teachers – may handle while managing a class of 30 or 300 students. Adding functionalities that teachers never use is not neutral; it creates indirect load, as we will see.*

These principles concern external constraints, which does not mean that intrinsic constraints – learning sciences – can be neglected. Our principles expand instructional design, which is extremely well developed in terms of intrinsic constraints, to somewhat previously neglected zone of extrinsic constraints. The principles also do not address other managerial constraints such as safety, or the relevance of the pedagogical activity to the curriculum. Instead, they focus on how to reduce the orchestration load. To be most effectively applied, the principles should be instantiated through a participatory iterative design with the teacher and iterative testing in the classroom.

5. The hardware environment

Although the three systems presented below are different from each other, they rely on the same hardware: the TinkerLamp. The TinkerLamp is a camera-projector system developed in our laboratory. It consists of a camera and a projector directed at a tabletop. There are four versions of the TinkerLamp presented in Figure 1. While they all support a camera-projector system and therefore seem quite similar, the four versions of the TinkerLamp differ in ways that have implications on classroom orchestration.

The first (a) version has no mirror, and therefore, a smaller projection area than the three other versions: 50×30 centimeters instead of 70×40 centimeters. It does not have an embedded computer, which makes it less cumbersome in itself but also requires more effort since the teacher has to connect a laptop to each lamp. The three later versions (b, c, and d) all embed a computer so that the interaction with the hardware is minimal for the end user: it requires only to be plugged on an electric outlet and there is a single on/off switch.

The (a) version existed in four colors. This simple feature seems neutral regarding how people learn, but it facilitated classroom orchestration: in a classroom where four teams were working on four lamps, the teacher could simply address teams by the color. This is a detail, but indicates the very pragmatic approach that we develop here, as well as the fact that design choices that are not usually considered as educational decisions may actually have an educational importance.




Figure 1. The four versions of the Tinkerlamp.

The (b) version is the first one with an embedded computer and a mirror. It has two major flaws. First, it is hard to move around because it is heavy and because there is no natural handle on it. It is also not “see through”, which means that the body of the lamp can block the sight of the teacher who wants to see students working on the lamp. This physical feature of the hardware goes against the ‘awareness’ principle. This is acceptable in particular settings, such as the one of Swiss primary school classrooms, which are typically equipped with one or two computers located in a corner. While the teacher is working with the whole class or managing individual work, he will often ask a small group to work on the computers. The (b) design is then fine since the group is somewhat isolated from the rest of the classroom and may focus on its work. However, in the case of an entire classroom working with several lamps at the same time, especially in the example of teenagers classes described below, the teacher has to keep a visual control over the groups at any time. This is the main reason why the (c) version was adopted for classroom usage over the (b) version. Note that the (c) version is also foldaway, which makes it easier to store in a cupboard when not used.

The most recent version of the lamp (d) goes even further in integrating in the classroom environment: it is compact, easy to carry, and the lightest of all versions. This minimizes the time necessary to set up the classroom before a lesson, and – as we stressed before – time is a major constraint. This lamp has two mirrors that allow it to be used in two configurations: the standard one with the projection on the table, and a second one in which the projection can be shown on a wall. This second configuration is set up by simply removing one of the mirrors. This allows the teacher to orchestrate the transition between activities located at different social

planes, the team plane (when the lamp projects on the team table) and the classroom plane (when the lamp projects on the wall). The experiments we report hereafter made use of the (a) and (c) versions of the TinkerLamp.

The three learning environments use fiducial markers similar to ARTag markers (Fiala, 2005) to precisely track the various elements placed on the surface below the lamp. Since the projection comes from above the tabletop, the system can project information on top of the objects placed on the surface of the interface. Augmenting objects in this way is not possible with a standard multi-touch table, where projection comes from below the surface. On the other hand, the “from above” systems with a single camera are not able to track fingers accurately and finger input is thus not possible. The tracking of objects with the camera on top is also more sensitive to the light conditions. At the beginning of the development of the Tinkerlamp, changes of light (e.g. a cloud reducing sun light) generated tracking problems and made the teachers unsatisfied. Now, the algorithms are robust enough and light problems are minimal. Tracking is of course sensitive to the obstruction of markers by the users hands or other objects. However, since the markers are very visible, the students understand quickly when some information disappears that they are hiding a marker. At the opposite, we tried invisible markers (infrared ink), which led to more occlusions by hands.

6. TinkerLamp: an AR system to teach logistics

We now present the first learning environment using the TinkerLamp: an environment developed to train vocational apprentices in the domain of logistics. The prevailing organization of the vocational education in Switzerland follows a dual approach: the apprentices (16-20 years old) work four days per week in a company and spend one day per week in a vocational school. At school, logistics apprentices are required to learn concepts about the storage and transportation of goods, the design of warehouses and transportation routes, as well as the management of inventories and information. However, at the workplace, apprentices do not have opportunities to reflect on and practice these concepts. The types of tasks they perform at the workplace are relatively basic. In some warehouses, the activities assigned to apprentices are limited to moving boxes with a forklift between two locations chosen by a computer. Apprentices are usually not involved in higher-level managerial duties, which are handled either by the director or by more experienced co-workers (Jermann, Zufferey & Dillenbourg, 2008).

The objective of the TinkerLamp system is to create a stronger link between the theory taught at school and the experience acquired at the workplace. More specifically, it aims at helping logistics apprentices understand theoretical concepts presented at schools by letting them experiment these concepts on an augmented small-scale model of a warehouse (Figure 2). Apprentices interact with the warehouse model using miniature plastic shelves, docks, and offices. Each element of this small-scale warehouse is tagged with a fiducial marker that enables automatic camera object recognition. The model is augmented with visual feedback and information through a projector in the lamp head.



Figure 2. Apprentices experimenting theoretical concepts using an augmented small-scale model of a warehouse. A TinkerSheet is placed on the left of the model.

The concrete representation offered by the small-scale model is complemented with less embodied, more abstract representations given on TinkerSheets. A TinkerSheet is a piece of paper automatically tracked in real-time by fiducial markers that allows users to control the system through the use of a small black token, to set parameters for the simulation, change the size of the forklift, etc. It also serves as a visual feedback space on

which textual or graphical summary information from the simulation is projected (i.e. the warehouse statistics such as surface areas, and the degree of utilization of the warehouse).

The 3-year long curriculum of these apprentices is now implemented in the TinkerLamp. Among many others, a typical learning activity aims to guide apprentices to discover principles to organize storage in a warehouse. It involves solving the trade-off between the efficiency of the warehouse (how fast goods can be transported in and out of the warehouse) and the storage capacity (how much goods can be stored inside the warehouse). For example, a warehouse with a high storage capacity is less efficient than a warehouse with a low storage capacity, because forklifts do not have as much free alley space to maneuver and drive fast. These trade-offs are tested by running a simulation on the warehouse models. The simulation augments information directly on top of the model in real-time, e.g. how forklifts approach the shelves, the statistics about the warehouse inventory, or storage management strategies.

Figure 3 (left) shows an example of these augmentations: the drawing of navigation nodes around and on top of each shelf. A pallet is represented by a square, which is surrounded by two triangles. Each triangle is either green or red, indicating whether a forklift can access the pallet from the corresponding side or not. When two shelves are placed too close together, the shelf nodes turn red, indicating that there is not enough space for a forklift to work in the alley between these shelves.



Figure 3. Shelf and navigation node augmentation on top of the small-scale model: a red square indicates that a pallet cannot be accessed by a forklift due to the lack of space (left). Simulation of a warehouse model: three forklifts are working in the warehouse (right).

The apprentices have to learn the different ways to count the surface of a warehouse such as the raw surface, the raw storage surface, and the net storage surface. The raw surface is simply the surface of the whole warehouse. The raw storage surface is the raw surface minus annexes (offices, technical rooms, docks, etc.). By building and exploring physical models of warehouse, apprentices are expected to discover how different types of surfaces impact on work efficiency. The apprentices can choose among three types of forklifts that differ in terms of size and maximum driving speed using the TinkerSheet interface. This choice of forklifts allows apprentices to experiment with the trade-offs between the forklift types and the warehouse storage capacity. For example, they can see that forklift types have an influence on both the work efficiency (a faster forklift moves more pallets in a given time) and the storage capacity (faster forklifts are bigger, need larger alleys and thus reduce the capacity of the warehouse) (Figure 3, right).

The latest version of the TinkerLamp for logistics training includes tools called TinkerKeys and TinkerBoard that have been added to trigger reflection activities for apprentices and to facilitate classroom orchestration. TinkerKeys are small paper cards that teachers carry with them easily when touring the classroom. When a teacher sees a need to intervene with a group (for example, he finds that a group is doing too many simulations and does not reflect on the statistics) or with the class (when he wants to quickly get attention from the whole class), he places a TinkerKey on the group table (Figure 4, left). Each TinkerKey triggers a different functionality in the TinkerLamp, either changing a state, or performing an action that facilitates the teacher's orchestration work.

TinkerBoard is a public display in the classroom (Figure 4, right). It makes the whole class history visible on a big projection board. This display supports the teacher during his class debriefing, showing the layouts and their statistics that each group has built during the learning activity. It also includes an event bar showing what activity each group is doing and how busy the apprentices are with physical manipulations. TinkerBoard is an instance of the digital 'dashboards' or 'cockpits' developed in learning technologies. The specificity is that this tool is domain-specific and used in real time (not for post-hoc analysis).



Figure 4. TinkerKeys (left and middle) and TinkerBoard (right).

The TinkerLamp system has been deployed and used in several vocational schools for three years. We conducted lab experiments, many field experiments as well as participant observations with more than 350 students and 8 teachers at various occasions. These field studies provided us with important insights as to how to improve the system and the realization of the five design principles.

Illustrating the principle of Integration

The integration is pushed very far in this system. The exercise sheets used with the system (TinkerSheets) are placed in the regular A4 folder that students have together with traditional exercise sheets, the only difference being that they have fiducial markers (Figure 5, right). The left part of the sheet, placed outside the working space, contains instructions for the apprentices. The right part of the sheet includes all input and output fields required by this activity. In other words, the curriculum relevance here is tangible: it corresponds to taking a page in the curriculum and placing it under the lamp.



Figure 5. The folders regularly used by apprentices (left). The TinkerSheets are placed inside the regular folder, achieving a seamless curriculum integration (right).

The information printed or projected on the TinkerSheets was also designed in the same way it is presented in the apprentices' textbooks. It is important for apprentices to play with multiple representations of the same phenomenon. The teacher may decide to use multiple representations, but this has not been imposed to teachers by inventing problem representations that are not the standard ones. Figure 6 shows the multiple representations of the stock management concept used in the textbooks, on the blackboard and on TinkerSheets.



Figure 6: Integration of representations: Multiple representations of the stock management concept used in a textbook (left), on the blackboard (center), and on TinkerSheets (right).

Paper sheets are information container. The information projected is transient, but apprentices often write/draw above what is projected to make this information persistent. In some activities, the teacher asked a member of each group in the classroom to bring their sheets to the blackboard and copy their data. Once all data were collected, the teacher asked students to compare the four warehouse layouts in order to understand concepts such as the trade-off between storage space and alley width. A paper workflow is slower than a digital one, but it is visible by all actors (who is going where with a paper sheet) and it is more flexible than a virtual workflow (see next principle).

This paper-based workflow extends outside the classroom. The teacher may for instance annotate a sheet with information useful for next year or make a copy of an annotated sheet to a colleague. In the last experiments, the apprentices were able to save simulated warehouse layouts on the TinkerSheets. The teacher could then select and print an activity sheet (fieldwork) with instructions from the best or the most interesting warehouse layouts. The apprentices were asked to bring this sheet to their workplace and to compare the simulated warehouses to the one where they actually worked. The integration of classroom technologies with homework is something that has been neglected in our field. While teachers warned us that these apprentices never completed their homework, it occurred that 90% of the apprentices did the fieldwork, and 82% of them did discuss it with their supervisor at work. Using paper sheets allowed us to minimize the time investment: they could discuss the comparison in a few minutes, over coffee, whereas going to a computer, connecting to a server and doing the same activity would have taken much more time and would have required a more formal appointment.

Illustrating the principle of Flexibility

TinkerKeys allow for a flexible management of the classroom. The teacher walks in the classroom with TinkerKeys in his hand or pocket. He can use them in any order and in innovative ways. For example, a teacher in our field trial said that he could leave a TinkerKey with the high-performing group, allowing them to run simulation when they wanted, but still keep the blocking and control with the low-performing groups.

Because TinkerBoard is a public display, the teacher could have spontaneous debriefing discussion with the class at any time, without having to plan the collection of data from each team. We observed teachers looking at the TinkerBoard on a regular basis and then illustrating the problem-solving strategies, benefits, and drawbacks of the solutions to the students on the TinkerBoard just by walking up to it and pointing to it. Our field trials confirmed that the teachers having the TinkerBoard in their class often had 4 or 5 spontaneous debriefing sessions during the activity – more than when they did not have it.

Illustrating the principle of Empowerment

Our experiments revealed that the simulation was actually too engaging: apprentices often ran the simulation many times without much reflection. The role of the teacher is to trigger reflection in teams but this is hard to do when the apprentices are engaged with a playful material such as the TinkerLamp. The TinkerKeys were therefore developed to bring the teacher back in power.

The “Allow Simulation” TinkerKey is used by the teacher to control the students’ capability to run simulation. Our scenario is that the groups were not authorized to run a simulation without the teacher’s permission. In order to run one, they have to call the teacher, who asks them to predict whether the performance of the current layout will be more or less than the previous simulation and to justify their prediction.

The “Pause Class” TinkerKey is used at the class level, to blank out all of the projected feedback from the TinkerLamps in the whole class. The teacher can easily and quickly get full attention from the students in order to give an instruction or move to a class-wide activity such as class debriefing.

During our field studies, rather than just walking from one table to another and discussing with the apprentices, the teachers confirmed that they were now able to create a situation for a better intervention with them and ensure the productive discussions for learning to occur. They used each of the TinkerKeys for a minimum of 4 times (to pause the class) and maximum of 27 times (to block and allow simulation) during a class of 1.5 hour.

Empowering the teacher includes support for the complementarity of class lecturing and debriefing. These class-wide activities are important for summarizing the main findings of the class and triggering reflection from the students. The TinkerBoard supports this type of activities by providing the teacher with a display consisting of the main learning resources in the class, including the design of warehouse layouts and their statistics. The teacher can smoothly and quickly organize a common debriefing at the end of the class to compare the different results from each group and help the students come up with important conclusions.

Illustrating the principle of Awareness

The TinkerLamps provide teachers with two level of awareness. The first level is due to the use of tangible objects: an (unexpected) benefit of tangibles, compared for instance to multi-touch tables, is that a teacher located a few meters away from the lamp can see what is being done on the table. This means that, while glancing across the classroom, the teacher sees which teams have placed no shelf, a few shelves or many shelves. This is the first level of monitoring: a simple awareness of the degree of activity.

We provided a second interpretative level of awareness, inspired by our field experiments. We noticed that teams performing too many shelf manipulations generally did not engage in reflection. To alleviate this problem, the TinkerBoard displays an event bar per team (Figure 4, right) in which the color scale, from yellow to red, indicates the number of shelf manipulations per time unit in real-time. This condensed summary also shows the occurrence of other reflective-related activities such as the running of warehouse simulations. A sequence of simulations, without much time in between, indicates a discouraged trial and error approach by the group (building and running without thinking). This information could encourage the teacher to intervene and ask the students to do more reflection. These awareness tools are designed to reduce the orchestration load for teachers working with four TinkerLamps at the same time and have proved to be effective (Do-Lenh, 2012). Let us stress that awareness is multi-directional: even if we consider it from the teacher viewpoint, the information is also available for the other students in the classroom.

Illustrating the principle of Minimalism

There is no login in TinkerLamp. The login is usually justified by the need to keep individual traces that would hypothetically be analyzed by teachers. These extra functions come at a price: it takes usually a few minutes to have all students login (forgotten passwords, keyboard locked on capital letters, etc.). Using the TinkerSheets to login and to directly start the appropriate activity may spare 3 to 5 minutes, i.e. up to 10% of a lesson time. The principle of minimalism addresses the indirect costs of functionalities that are not strictly necessary.

The same applies to TinkerKeys: the number of TinkerKeys was designed to be minimal to accommodate the limited time the teacher had when interacting with multiple groups. Five TinkerKeys were implemented in the current version of the system. Observations from our field studies validated this design choice. We initially invented many more TinkerKeys, but the time necessary for the teacher to retrieve them from a deck of cards reduced the advantage of paper cards, namely their ease of manipulation. We also increased the graphical difference between cards to speed up the card selection process by the teacher.

The TinkerBoard also reflects minimalism. It simply shows a list of warehouse layouts built by each group, together with an event bar whose color ranges from yellow to red. It is minimalistic because it does not try to provide other sophisticated visualizations, such as the identity of which student is doing what, or automatic detection of special events. It does not provide a diagnosis for the teacher but more modestly, only some awareness of team activities. The rationale was that more complex and precise display would require the students and the teacher to spend more time and cognitive effort analyzing and interpreting what the system is showing, and would increase the risk of distracting the users from their own learning and orchestration task.

7. Tapacarp: an AR system to train carpenters

The second system illustrating our design principles is Tapacarp. Tapacarp is a learning technology developed for carpenter apprentices. Like logistics apprentices, carpenter apprentices are enrolled in a 3-year dual apprenticeship. They attend a vocational school one day a week and work in a company for the other 4 days of the week. Most of them are typically aged between 15 and 20 years old.

There are three main subjects that carpenter apprentices must master: the law of statics for buildings, three dimensional (3D) visualization, and the knowledge of construction materials and their usage. Tapacarp focuses on the subject of 3D visualization, which carpenters develop mainly through drawing classes. The 3D visualization is omnipresent in the work of carpenters. They must be able to imagine 3D objects (typically a roof structure) from a 2D drawing on paper (a plan) in order to visualize how the beams need to be carved. Going from 2D to 3D is part of the descriptive geometry field and it requires extensive training. One of the specific tasks linked to the 2D-3D transfer is to find the true size of an object based on its orthographic projections. There are various techniques to achieve this, but carpenters mainly use the rabattement technique introduced by (Monge, 1798) and graphically explained in Figure 7.

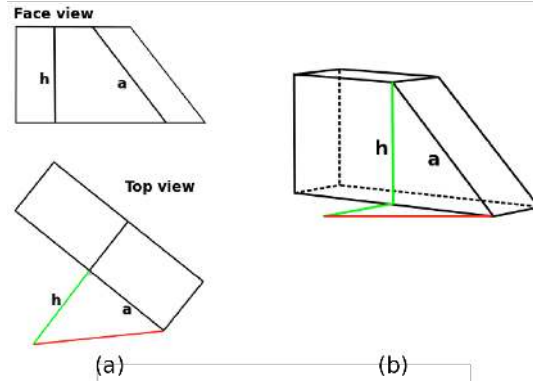


Figure 7. (a) The rabattement technique with the projections: to find the true size of the edge a by rabattement, one typically takes its height h on the face view and reports it perpendicularly on the top view; the true size is then the red line. (b) A perspective representation of a rabattement.

The goal of the activity presented here is to teach apprentices the principle of finding the true size of an edge or a face by rabattement, which is described in Figure 7. The lesson is composed of 11 activities, one per page, presented in an increasing level of difficulty. Each activity page is equipped with a fiducial marker and all activity pages are stapled together to form an activity booklet. Each activity page also holds a 3D representation of the block to use and has a six centimeters horizontal white space where dynamical instructions are projected by the system, as can be seen in Figure 8(c).

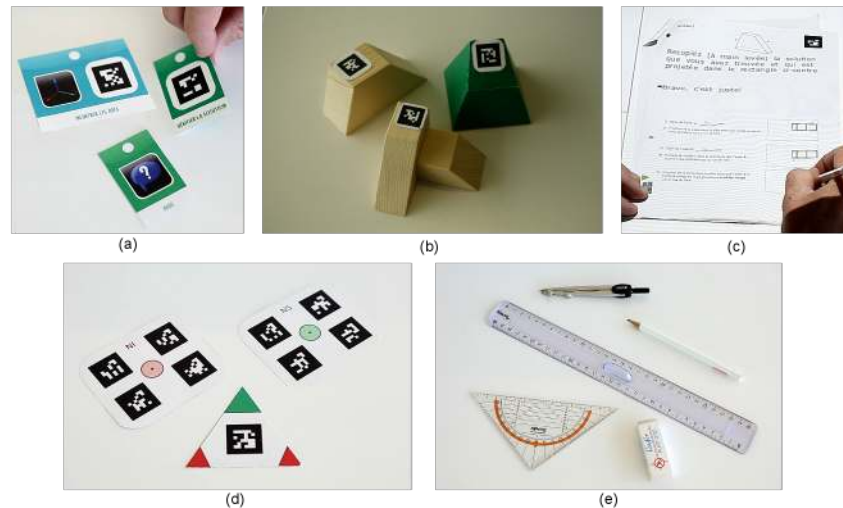


Figure 8. The modalities of the interface: cards, blocks, an activity booklet, paper tools, and drawing tools.

The system is composed of five modalities: blocks, cards, a paper activity booklet, drawing tools, and paper tools. The activities are a mix of manipulating the blocks, taking measures on them, performing virtual actions with paper tools, and drawing on paper. The cards are used to trigger actions such as checking the solution, playing an animation, and changing parameters of the display. A demonstration of the usage of the five modalities is shown in Video 1.

We outline how the interface is used through the example of the activity shown in Figure 8(c). There are two steps in this activity, both printed out on the activity sheet. The first set of instructions asks the apprentices to take the block shown at the top of the page and to draw one of its face (face A) in its true size; the face must be drawn on the top view of the block that is printed on the page. If the apprentices know how to proceed, they can

take the block and measure – with their ruler – the correct height of face *A*, and report it on the top view, completing the drawing as asked. If, on the contrary, they do not know how to proceed, they can put down the “Help” card. The system then projects instructions in the top white space and guides the apprentices as they complete the activity in a step-by-step fashion, using the paper tools. The second set of instructions simply asks the apprentices to write down their method and to call their teacher. The teacher comes, checks the answer, and if the latter is correct, gives them an additional card. When shown to the system, the card allows the apprentices to see a step-by-step 3D animation of the process they just completed. They can use it in further activities. For now, they turn the page to go to the next activity.

To validate the design of Tapacarp, we ran a user study in a classroom. Over three sessions, a total of 24 apprentices grouped in pairs used Tapacarp during a 1.5 hour lesson. The blocks, the cards, the paper tools, and the activity booklet were provided to the apprentices, and they were also asked to use their own regular tools: pencils, a ruler, a protractor, and a compass.

Illustrating the principle of Integration

The integration of paper and drawing tools resulted from a curriculum and assessment constraint: carpenter apprentices need to master their drawing tools, and the teacher needs to be able to evaluate how they use them. The usage of paper was decisive in the adoption of the technology by both the apprentices and the teacher. Apprentices simply solved the exercises on paper as usual, taking the augmentation as a nice add-on. The teacher saw that it did not change much in terms of preparation of exercises or explanations. Tapacarp also fit with the extra-classroom organization of the school that requires students who have not completed their drawing to finish them at home. Because they use the same means (paper) and tools as usual, they can finish their drawing at home.

Tapacarp made use of wooden and cardboard blocks, which were the same blocks as they used in the traditional class. The identification to wood and paper is strong among the carpenter community, and using those materials allowed Tapacarp to fit with the existing cultural environment.

Illustrating the principle of Flexibility

The affordance of paper lends great flexibility to quickly change activities, providing a better balance between routine and improvisation. The easy navigation between the activities (by turning the page) is an example of flexibility, making it easy for students to refer to a previous activity, or for the teacher to quickly access the history of what students have done.

The distribution of cards and blocks opens up new possibilities of scaffolding. Indeed, one can imagine that students start with a minimum amount of cards and blocks and gain more of them as they progress in the activities.

The activities worked well for students of various levels: the weaker students only completed the minimum drawing for each activity, while some stronger students drew the rabattement for all possible edges on the figure. The help and the animation cards also helped the weaker students without preventing stronger students from going faster.

The multi-modal interface of Tapacarp fostered the emergence of role distribution between the two members of a group. For example, one student took control of the drawing tools, and the other one controlled the block. This has the benefit of engaging both pair members and has been proven useful in terms of learning scenario (see e.g. Burton, Brna, and Treasure-Jones, 1997).

Illustrating the principle of Empowerment

Giving the teacher special cards was done so as to empower him: on completion of two of the eleven activities, the apprentices had to ask the teacher to check their answer. If the answer was correct, the teacher gave them a card. During the field study, we observed a tension between teacher empowerment and the flexibility principle. Indeed, some groups did not call the teacher when they should have and just went on with the next activity. They were able to do it because calling the teacher was not enforced for the sake of flexibility, but in this case, the higher flexibility removed power from the teacher.

Teacher empowerment is also achieved in Tapacarp through the multi-modal interface. The five modalities lower the barrier for teacher interventions, since the teacher does not have to take hold of the mouse or the keyboard and the large display makes it such that both apprentices and the teacher can interact with the system simultaneously.

Illustrating the principle of Awareness

The awareness with Tapacarp is not optimal. During the design, the thought was that thanks to the tangible interface, the teacher could easily monitor the progress of each group. However, it turned out to be difficult for

the teacher to estimate the progress of the groups from more than two meters away, mainly because of the small size of the blocks and the nature of the activity (a drawing is hard to observe from a distance). On the positive side, the teacher could nevertheless see which student was using what part of the interface and could therefore have an idea of who was playing what role in the group.

Illustrating the principle of Minimalism

The interface of Tapacarp is easily reduced to the minimal number of elements. In the field study, the apprentices often started the activity by pushing away what they did not need (for example the extra blocks and cards). The tangibility indeed made it easy to sort out what was needed from what was not.

The scaffolding of the cards is another illustration of the minimalism of Tapacarp: at the beginning the apprentices received only the strict minimum of cards needed to perform the first activities. As the activities went on they gained more cards that gave them access to additional functionalities.

The features of the system were kept to the minimum as well. For example, there is no artificial intelligence to correct the students work or give them suggestions on what they are doing wrong. There is no requirement to log into the system, and the only saving and history functions are the ones provided by the paper.

8. Kaleidoscope: an augmented paper interface for exploring symmetries

With Kaleidoscope, we explored the design space offered by paper as an interface. Paper has two aspects: it can be a document, and it has tangible properties. In its document aspect, it formats and makes a printed, written, or drawn content persistent. In its tangible aspect, it can be grasped, manipulated, and transformed easily. In both aspects, practices have been associated with paper since a time long enough to consider them natural. The document side of paper is already omnipresent at schools, in the form of textbooks, or exercise sheets, and allows a quiet deployment and a usage that minimizes cognitive overhead (Oviatt, Arthur & Cohen, 2006).

Unlike Tapacarp and the Tinker environment for logistics, Kaleidoscope is not directed at secondary school apprentices, but at primary school pupils learning geometry. The tangible aspect of paper is especially useful for geometry: rotation/translation correspond to geometrical transformations, folds correspond to symmetries, area calculations can be introduced by cutting out pieces of paper, etc. Practices based on paper are important too: geometry makes an extensive use of tools, such as compass, rulers, and protractors. Their usage is part of the expected learning outcome in primary school.

We developed several pedagogical scenarios based on paper interfaces, which we deployed in three primary schools and with more than 20 classes (Bonnard, Verma, Kaplan & Dillenbourg, 2012). The activities involve various geometry topics: classifying quadrilaterals, measuring and describing angles, finding the formulas for areas, etc. The pedagogical goal of the activity presented hereafter, Kaleidoscope, is the exploration of symmetries. The research goal is to investigate how paper interfaces can be used to foster the expertise of teachers in terms of pedagogical design, and more particularly for tailoring an activity to be used in their class. To do so, we created a generic component, the Kaleidoscope sheets, that separates the technical design from the pedagogical design.

A Kaleidoscope sheet reflects its content according to one or more symmetry axes. The symmetry axes split a predefined input area into a smaller input area and its reflection on the other side of each axis. For example, with a vertical axis, the left side of the original area will be shown reflected on the right. More axes further split the input area and show more reflections, e.g. a vertical and a horizontal axis will cause the reflection of the top-left corner onto the three other corners (see Figure 9).

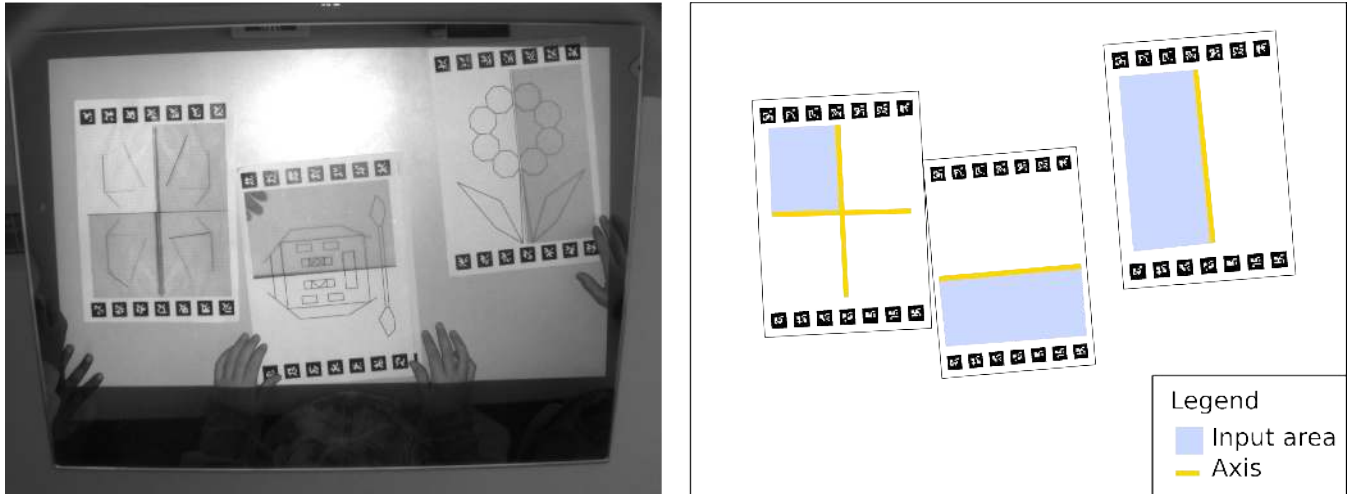


Figure 9. A view from the top of the Kaleidoscope sheets, with a schema indicating the axes reflecting the input area. The grayer sides of the sheets are actually the projection of the other side of the axes. Anything in the input area is re-projected, such as fingers, not only the ink on the paper.

We proposed to teachers to design their own activity and use it in their class. Four teachers and their classes of about twenty pupils of third, fourth and fifth graders from three different schools were involved. We prepared a sample activity with six exercises to illustrate how the sheets could be used, but very early in the demonstration, each teacher had a precise idea of what they wanted to do with them.

For each set of axes, a teacher can be given a blank page containing only fiducial markers, the axes, and the limits of the input area. A teacher can then customize this blank Kaleidoscope sheet without any help from the developer. It is enough to add free text (instructions), a drawing, or a grid to the paper. The teacher does not even need a computer: everything can be written and drawn by hand and with regular tools (ruler, transparent papers, etc.) and placed in a copy machine. One teacher requested the digital versions of Kaleidoscope sheets to prepare the activity completely by herself, and it took less than one hour for the other teachers to describe exactly what they wanted printed on the sheets (instructions, axes, or figures).

Each teacher ran the designed activities by themselves in their classes, with one TinkerLamp to be shared among all the pupils. We only observed how the teachers managed the activity and did not need to intervene. We collected many lessons on technology-enhanced orchestration in such authentic settings, because the teacher did both the design and the use of the activity.

Illustrating the principle of Integration

The main strength of a paper interface for orchestration is the integration. Paper is the main component of the workflow of the classroom: pupils read from textbooks, do their exercises in notebooks, which they bring to their parents, etc. As a consequence, several pedagogical objectives defined by the curriculum are linked to paper itself, such as using tools. In the case of geometry, the teachers expected the pupils to use the ruler to draw the figures, and be accurate. The use of sheets as an interface allowed this. Furthermore, one of the teachers included exercises defined in official pedagogical resources to design the patterns used for symmetry, because the official resources are also defined with paper in mind.

Paper interfaces also integrate the practicalities of the classroom workflow: not every pupil can use all the resources at the same time. In our case, we shared one TinkerLamp among the class, because it is easiest to fit in the classroom. With two of the teachers, the pupils worked at their desk, using the TinkerLamp only to check their solutions. This allowed for a shared, parallel interaction with the system. The distribution was naturally fair: in one of the studies, we set an expiration time after which the Kaleidoscope sheet would not display the symmetric image of the input area. This measure was not necessary, as no pupil used up all the allocated time. In other words, the orchestration load to balance the use of the TinkerLamp in the workflow is minimal.

Illustrating the principle of Empowerment

The various pedagogical scenarios designed by the teachers illustrated the empowerment allowed by paper interfaces. The teachers could be as central as they wanted. Two teachers simply explained the activity to all the pupils and roamed the room to help when it was needed. Two other teachers designed an activity where they controlled precisely the work of pupils at the TinkerLamp, and let the rest of the class work autonomously on another exercise.

The TinkerLamp does not make the teacher obsolete. It shows to pupils the expected symmetry, so that they can correct by themselves, but the teacher is also able to check the exercise: she just needs to have a look at the Kaleidoscope sheet. The TinkerLamp supported the teacher: if the explanation required it, the teacher could bring a pupil and the sheet under the lamp to illustrate in another way the explanations.

Illustrating the principle of Flexibility

During the activity itself, the paper interface provided flexibility at two levels. First, it allowed for various ways of checking and explaining mistakes. We already mentioned that the TinkerLamp and the teacher could check the work of pupils, but the pupils could also use the mirrors that were used the week before to introduce symmetries. The mirror reflected the grid to give a hint at how the reflection should look like. This complemented the augmentations, which were more precise, but more cumbersome to use, as it involved a move from the desk to the TinkerLamp.

Second, one of the teachers also focused on flexibility in the design of the activity. She did not write any instruction in order to be able to verbally explain the activity as she needed to. She also designed two variants of each of the exercises, one being easier than the other. This allowed the pupils to choose the level of challenge that they wanted, and the teacher to adjust it by simply taking away one of the sheets.

Illustrating the principle of Awareness

The TinkerLamp being shared among the whole class resulted in an important gain in awareness. One or two pupils were checking their answer at the same time in the interaction surface, and others waited in line. The teachers could easily measure the size of the line and act accordingly. It informed them on the difficulty that the pupils were having (harder exercises required more use of the lamp). In the case where an alternative activity kept the pupils busy, the teachers could easily remind pupils of the limit size of the queue, to force the pupils to work on the alternative activity.

Interestingly, the awareness resulting from the use of the TinkerLamp involved the pupils in the reduction of the orchestration load. The pupils waiting in the queue next to the TinkerLamp naturally started giving feedback on the work of pupils checking their answers. This lightened the amount of explanations required from the teacher, and stimulated peer corrections, which is generally a productive learning mechanism.

Illustrating the principle of Minimalism

Kaleidoscope sheets are very minimal: whatever is projected on one part of the sheet is reflected on the rest of the sheet. This does not put any limitation on what is reflected, be it printed ink, writing, drawing, fingers, cutout shapes, or even food. More importantly for the teachers, a Kaleidoscope sheet could be used for an axis that was not intended in the design. For example, if the teacher wanted to show that a cut-out shape has one symmetry axis but not two, she could rotate the shape on the intersection of the two axes to display the double symmetry, but then move the object along only one of the axis, so that a single reflection appeared. This is more convenient than forcing the teacher to retrieve the exact sheet corresponding to the expected reflection.

9. Summary and conclusions

We introduced five design principles for learning systems that are meant to be used in a classroom: integration, awareness, empowerment, flexibility, and minimalism. These principles stem from our research experience: a system that is pedagogically efficient is not enough to be used in a classroom. This is confirmed by the recent focus on classroom orchestration within the CSCL community. The principles focus on what makes an AR learning system work in a classroom. They describe the differences between a system that would produce learning outcomes in a lab study and a system that “works well” in the everyday life of a classroom. In order to make those principles more concrete, we illustrated them through three AR learning systems that were deployed in classrooms of children and teenagers.

In the three systems, *integration* was implemented to blend with the existing practices of the classroom and allow the system to be used together with the existing material. This was done mainly thanks to the use of paper as a part of the interface. *Flexibility* was about allowing the teacher to cope with varying levels among the

students, be it within a group or between the groups. Another facet of flexibility was to give freedom to the teacher to access each part of the systems at any time. The cards were the main tool used to *empower* the teacher, and TinkerBoard showed how a centralized tool could put the teacher back to the center of the class. The tangibility of the interface and TinkerBoard made the teacher *aware* of the progress of the groups and the dynamics within the groups. Finally, *minimalism* in both the visualization features of the interface and the number of available functionalities was consistently observed in all three systems.

The scientific status of these principles is however questionable. They rely on a large but uncontrolled empirical basis. If we cumulate the life of these three systems, we have more than a decade of experience in using AR into classrooms. However, most experiments were done in the – partly – uncontrolled context of classrooms. If we add a feature X to an AR system, any improvement of classroom orchestration between the consecutive versions of the system may be due to many other factors than X. Moreover, working with teachers who had many ideas and great expectations from us, we often changed features X, Y, and Z between two versions. We therefore have to be modest regarding the epistemological status of the five principles.

Our main contribution is probably to illustrate by these principles the big gap that exists between a system that supports learning and a system that works well in a classroom. By emphasizing and theorizing this gap, we would like to attract the attention of our community: what many would call “implementation details” actually requires careful thinking and deep knowledge. And what may sound as practical issues actually deserves a theory, yet to be built.

10. References

- Ainsworth, S. (1999). The functions of multiple representations. *Computers & Education*.
- Alcoholado, C., Nussbaum, M., Tagle, A., Gomez, F., Denardin, F., Susaeta, H., Villalta, M. and Toyama, K. (2011). One Mouse per Child: interpersonal computer for individual arithmetic practice. *Journal of Computer Assisted Learning*.
- Arias, E., Eden, H., Fischer, G., Gorman, A., and Scharff, E. (2000). Transcending the individual human mind - creating shared understanding through collaborative design. *ACM Trans. Comput.-Hum. Interact.*, 7(1), 84–113.
- Billinghurst, M., Kato, H., and Poupyrev, I. (2001). Collaboration with tangible augmented reality interfaces. *HCI International*, (1), 5-10.
- Billinghurst M., (2002). Augmented reality in education. *New Horizons for Learning* 12.
- Bonnard, Q., Verma, H., Kaplan, F. and Dillenbourg, P. (2012). Paper Interfaces for Learning Geometry. In the proceedings of the 7th European Conference on Technology Enhanced Learning.
- Burton, M., Brna, P., and Treasure-Jones, T. (1997). Splitting the collaborative atom: How to support learning about collaboration. *AI in Education*.
- Cuendet, S., Jermann, P., and Dillenbourg, P. (2012). Tangible interfaces: when physical-virtual coupling may be detrimental to learning. Proceedings of the 2012 British Computer Society Conference on Human-Computer Interaction.
- Clark, H.H. & Brennan S.E. (1991) Grounding in Communication. In L. Resnick, J. Levine & S. Teasley (Eds.), *Perspectives on Socially Shared Cognition* (127-149). Hyattsville, MD: American Psychological Association.
- Do-Lenh, S. (2012). Supporting Reflection and Classroom Orchestration with Tangible Tabletops. *PhD Thesis*.
- Dillenbourg, P., and Jermann, P. (2010). Technology for classroom orchestration. *New Science of Learning, Springer Science+Business Media*.
- Dillenbourg, P., Zufferey, G., Alavi, H., Jermann, P., Do-Lenh, S., Bonnard, Q., Cuendet, S., and Kaplan, F. (2011). Classroom orchestration : The third circle of usability, *Proceedings of CSCL 2010*.
- Fiala, M. (2005). ARTag, a fiducial marker system using digital techniques. *Computer Vision and Pattern Recognition*, vol.2, no., 590-596.
- Fishkin, K. P. (2004). A taxonomy for and analysis of tangible interfaces. *Personal Ubiquitous Comput.*, 8(5), 347–358.
- Fitzmaurice, G. W., Ishii, H., and Buxton, W. A. S. (1995). Bricks: laying the foundations for graspable user interfaces. *CHI'95: Proceedings of the SIGCHI conference on Human factors in computing systems*, 442–449.

- Fitzmaurice, G. W. and Buxton, W. (1997). An empirical evaluation of graspable user interfaces: towards specialized, space multiplexed input. *CHI '97: Proceedings of the SIGCHI conference on Human factors in computing systems*, 43–50.
- Fjeld, M., Bichsel, M., and Rauterberg, M. (1999). BUILD-IT: a brick-based tool for direct interaction. *Engineering, Psychology and Ergonomics*, (4), 205–212.
- Goldin-Meadow, S. (2003). Hearing gesture: How our hands help us think. *Harvard University Press*.
- Hutchins, E. (1995). How a cockpit remembers its speeds. *Cognitive Science*, 19, 265-288.
- Hornecker, E. and Buur, J. (2006). Getting a grip on tangible interaction: a framework on physical space and social interaction. *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, 437–446.
- Ishii, H. and Ullmer, B. (1997). Tangible bits: Towards seamless interfaces between people, bits and atoms. *CHI '97*, 234–241.
- Jermann, P., Zufferey, G., and Dillenbourg, P. (2008). Tinkering or Sketching: Apprentices' Use of Tangibles and Drawings to Solve Design Problems. *Lecture Notes in Computer Science, Times of Convergence. Technologies Across Learning Contexts*, 5192, 167-178.
- Jermann, P., Zufferey, G., Schneider, B., Lucci, A., Lépine, S., and Dillenbourg, P. (2009). Physical space and division of labor around a tabletop tangible simulation. In *CSCL '09*, 1, 345-349.
- Kerawalla, L., Luckin, R., Seljeflot, S., and Woolard, A. (2006). "Making it real": exploring the potential of augmented reality for teaching primary school science. *Virtual Reality*, 10(3-4), 163-174.
- Larkin J. H., and Simon H. A. (1987). Why a diagram is (sometimes) worth ten thousand words *Cognitive Science*.
- Mantovani, F., Castelnuovo, G. (2003). Sense of presence in virtual training: enhancing skills acquisition and transfer of knowledge through learning experience in virtual environments.
- Marshall, P., Rogers, Y. and Hornecker E. (2007). Are tangible interfaces really any better than other kinds of interfaces? *CHI'07 workshop on Tangible User Interfaces in Context & Theory*.
- Milgram, P. and Kishino, F. (1994). A taxonomy of mixed reality visual displays. *IEICE Transactions on Information Systems*, 77(12).
- Monge, G. (1798). *Géométrie Descriptive*. J. Klostermann fils, Paris.
- Moraveji, N., Kim, T., Pawar, U., Ge, J., Inkpen, K. (2008). Mischief: Supporting Remote Teaching in Developing Regions. *CHI'98: Proceedings of the SIGCHI conference on Human factors in computing systems*, 353-362.
- Moraveji, N., Ringel Morris, M., Morris, D., Mary, C., and Riche, N. (2011). Classearch: Facilitating the development of web search skills through social learning. *CHI'11: Proceedings of the SIGCHI conference on Human factors in computing systems*.
- Norman, D. A. (2010). *Living with Complexity*. MIT Press.
- Nussbaum, M. & Diaz, A. (To appear). Classroom Logistics: Integrating Digital and Non-digital Resources. *Journal of Computers in Education*.
- O'Malley, C., and Stanton Fraser, D. 2004. Literature Review in Learning With Tangible Technologies. Discussion Paper. FutureLab.
- Oviatt, S., Arthur, A., and Cohen, J. (2006). Quiet interfaces that help students think. *Proceedings of the 19th annual ACM symposium on User interface software and technology*, 191–200.
- Patten J., Ishii H., Hines J., and Pangaro G. (2001). Sensetable: a wireless object tracking platform for tangible user interfaces. *CHI'01: Proceedings of the SIGCHI conference on Human factors in computing systems*, 253-260.
- Patten, J. and Ishii, H. (2007). Mechanical constraints as computational constraints in tabletop tangible interfaces. *Proceedings of the SIGCHI conference on Human factors in computing systems*, 809–818.
- Piper, B., Ratti, C., and Ishii, H. (2002). Illuminating clay: a 3-D tangible interface for landscape analysis. *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, 355–362.

- Prieto, L. P., Villagra-Sobrino, S., Jorin-Abellan, I. M., Martinez-Mones, A., and Dimitriadis, Y. (2011). Recurrent routines: Analyzing and supporting orchestration in technology-enhanced primary classrooms. *Computers & Education*.
- Roschelle, J., Rafanan, K., Estrella, G., Nussbaum, M., and Claro, S. (2009). From handheld collaborative tool to effective classroom module: embedding CSCL in a broader design framework. *Proc. of CSCL*.
- Roth, W. (2000). From gesture to scientific language. *Journal of Pragmatics*, 32(11), 1683–1714.
- Stringer, M., Rode, J., Toye, E., and Blackwell, A. (2005). The Webkit Tangible User Interface: A Case Study of Iterative Prototyping. *Journal of Pervasive Computing*, 4(4), 35–41.
- Stanton, D., Bayon, V., Neale, H., Ghali, A., Benford, S., Cobb, S., Ingram, R., O'Malley, C., Wilson, J., Pridmore, T. (2001). Classroom collaboration in the design of tangible interfaces for storytelling. *CHI'01: Proceedings of the SIGCHI conference on Human factors in computing systems*.
- Stefik, M., Bobrow, D., Foster, G., Lanning, S., and Tatar, D. (1987). WYSIWIS revised: early experiences with multiuser interfaces. *ACM Transactions on Information Systems*, 147-167.
- Underkoffler, J. and Ishii, H. (1998). Illuminating light: an optical design tool with a luminous-tangible interface. *CHI'98: Proceedings of the SIGCHI conference on Human factors in computing systems*, 542–549.
- Underkoffler, J. and Ishii, H. (1999). Urp: a luminous-tangible workbench for urban planning and design. *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems*, 386–393.
- Wellner, P. (1991). The DigitalDesk calculator: tangible manipulation on a desk top display. *Proceedings of the 4th annual ACM symposium on User interface software and technology*, 27–33.