

The Evolution of Research on Digital Education

Pierre Dillenbourg¹

Published online: 16 February 2016

© International Artificial Intelligence in Education Society 2016

Abstract How does AI&EdAIED today compare to 25 years ago? This paper addresses this evolution by identifying six trends. The trends are ongoing and will influence learning technologies going forward. First, the physicality of interactions and the physical space of the learner became genuine components of digital education. The frontier between the digital and the physical has faded out. Similarly, the opposition between individual and social views on cognition has been subsumed by integrated learning scenarios, which means that AIED pays more attention today to social interactions than it did at its outset. Another trend is the processing of learners' behavioural particles, which do not carry very many semantics when considered individually, but are predictive of knowledge states when large data sets are processed with machine learning methods. The development of probabilistic models and the integration of crowdsourcing methods has produced another trend: the design of learning environments has become less deterministic than before. The notion of learning environment evolved from a rather closed box to an open ecosystem in which multiple components are distributed over multiple platforms and where multiple stakeholders interact. Among these stakeholders, it is important to notice that teachers play a more important role than before: they interact not only at the design phase (authoring) but also in the runtime phase (orchestration). These trends are not specific to AIED; they depict the evolution of learning technologies as a whole.

Keywords MOOCs · Digital Education · Learning Analytics

Introduction

The only way I have found to anticipate the future of AIED has been to look at the deeper trends in last 25 years and extrapolate from there. Therefore, this paper describes the evolution of AIED more than it's future. However, I will not focus only on the

✉ Pierre Dillenbourg
pierre.dillenbourg@epfl.ch

¹ Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

AI&Ed community but I will take a broader perspective. In the vicinity of AIED, several communities have emerged such as Intelligent Tutoring Systems, Computer-Supported Collaborative Learning (CSCL), Learning Sciences, and, more recently, Learning Analytics, Education Data Mining and several MOOC-related communities. My analysis cannot be restricted to research contributions that have been presented in venues with a proper AIED label. Even if one did narrow down the scope of AIED to educational applications of AI techniques, the boundaries would be unclear since AI per se has radically evolved over 25 years. For instance, several projects presented hereafter include machine learning or computer vision methods, which originally belonged to AI, but are now published in other venues. In summary, the many communities comprising learning technologies overlap too much to draw boundaries. Therefore, this paper considers the evolution of these communities in a broad sense, inside and around the AIED community.

I will refer to ‘25 years’ as the arbitrary time span defined by the anniversary of this journal, even if the foundations of AIED were laid down many years before the launch of IJAIED. For the sake of clarity, I deconstruct the evolution of learning technologies into six fundamental trends. These trends are artificially segmented from each other; in reality, they form a beam of evolution vectors.

Trend 1: More Physical

Over the years, the evolution of computer science has blurred the border between what is digital and what is physical. This evolution was especially clear in the field of Human-Computer Interaction (HCI) with concepts such as augmented reality, ‘tangible bits’ (Ishii and Ullmer 1997), tangible user interfaces, etc. Twenty-five years ago, the interactions between a learner and an AIED environment were restricted to a display, a keyboard and a mouse. The computer and the human brain were viewed as two symbolic machines interacting by exchanging symbols. Currently, a learning environment can be any device such as a treadmill that reads body signals for adapting the slope, a digital driving coach who also considers traffic events, a robot that helps with writing, etc. This trend is bidirectional: physical objects or events enter the digital realm and, conversely, digital objects populate the physical environment. Let me now consider each of these two directions, starting with the physical becoming digital. I illustrate this with two technological developments, physical interaction and multimodal learning analytics.

Physical interactions. Tangible interfaces are handmade objects that are physically manipulated by the user in order to impact their digital counterparts. Cameras, sensors or RFID readers identify the actions performed on objects. With such a definition, a keyboard and a mouse could be described as tangible interfaces, but the terms ‘*tangible interface*’ usually refer to a set of objects, moved on a surface (often only in 2D) by one or a few users placed around this surface. While learning activities in AIED were initially restricted to symbolic manipulations (except for a few projects), novel learning environments expanded learning activities to physical manipulations. The germs of this evolution can be found in our community, for instance with Papert’s advocacy for the physical implementation of constructionist principles (e.g. LegoLogo). Twenty years later, this vision of education happily married with digital fabrication technologies (e.g. 3D printers) to produce the makers/fablab movement (Blikstein and Krannich 2013), a

community that is distinct from AIED, but inherited part of its genome. Outside our community, the interest in physical actions was boosted by Varela's extension of cognitive science towards embodied cognition (Rosch et al. 1991) and, overall, by the development of a tangible interface in HCI (Ullmer and Ishii 2000). The debate regarding the intrinsic educational value of physical manipulations (versus digital manipulations) goes as far as Froebel or Montessori (Zuckerman et al. 2005). But, beyond these cognitive considerations, tangible interfaces offer overall a great level of *usability* compared to mouse-based interfaces or even multi-touch interfaces, especially when several learners have to manipulate multiple objects (Jermann et al. 2009; Lucchi et al. 2010). This evolution towards physical manipulation includes haptic devices, with force feedback (i.e. both input and output is physical) which open digital education to very accurate gestures, namely for training surgeons (Toussaint et al. 2010) or learners with physical disabilities. Finally, this evolution includes robotics since they combine the digital and physical world. Robotics appeared from the beginning of AIED, e.g. the work of Martial Vivet (Leroux et al. 1996) and the numerous initiatives that use robots for teaching programming or problem solving. Somehow, the use of robots as pedagogical agents that may physically interact with learners, for learning geometry (Walker and Burleson 2012) or writing (Hood et al. 2015) has so far been a minor stream in AIED, but I expect it will expand soon given the availability of relatively cheap humanoids and robot-independent programming environments. Another example of this evolution towards more physicality is the development of tabletops. The point is not that multi-touch interactions are more physical than mouse-based interactions, even if some multi-touch gestures could be considered as such, but rather to realise how much interactions are modified simply by moving from a vertical display to a horizontal surface. For instance, Selker et al. (2002) placed sensors that capture the height of tables and their relative positions. It is important to understand the way learners occupy the classroom space, whether they sit or stand, where they are placed with respect to the team workspace, to the other teams or to the blackboard, how easily they move from private to public space, etc. These physical issues, not addressed at the outset of learning technologies, are indeed key concerns while orchestrating classroom activities.

Multimodal analytics expanded the learner modelling process beyond explicit user answers by measuring various non-verbal signals as well as environmental information. Compared to the infancy of AIED, a huge variety of sensors currently exist, many of them being cheap, small, not intrusive and rather accurate. This hardware evolution was paralleled by the development of algorithms able to process massive amounts of data. This led researchers to pay attention to signals that did not receive much attention in the past (see Trend number 2 'Less Semantic'). There is a natural evolution from AIED to multimodal learning analytics (Blikstein 2013; Ochoa et al. 2013), for instance related to the development of collaborative tabletop environments that include multi-touch interfaces or tangible interfaces. In these environments, learner modelling not only requires tracing object manipulations but also capturing verbal and non-verbal interactions among learners around the table (Martínez et al. 2011). The input became much broader than the traditional learner-system interaction channel.

Now considering the second facet of this evolution, *the digital becomes physical*. This refers to the output of a digital environment's impacts on physical space. Of course, any output affects the physical space, which is the very notion of output, but it is not the standard output on a computer display or on a printer, but the various

innovations developed over the last decade in HCI: ambient computing, the disappearing computer, peripheral displays, interactive furniture, wearables and, more globally, the internet of things. Let us consider ambient displays. What constitutes an innovation is not the mere fact of embedding a display in a wall or a table; the salt is inventing communication channels that do not require full attention, exploring new territories opened by peripheral perception. Hence, the key question is: what is the difference between the information on a central display versus information transmitted via devices that do not look like computers? In many cases, this low resolution information (e.g. the number of connected people rather than the list of their names) provides learners (and teachers) with some *awareness* of the state of some persons, some groups, some objects, etc. The notion of awareness reflects the influence of another close community, computer-supported cooperative work (CSCW). Initially, awareness tools were displayed on the computer screen, but, since computer displays are cluttered with information, awareness tools allow designers to escape from this interaction bottleneck and to exploit peripheral perception: peripheral displays, background sounds, vibrations of the furniture, etc.

One example of furniture-based output is the Reflect table (Bachour et al. 2010). It has been designed to help teams balance participation: LEDS embedded in the table in front of team members display the amount of speech produced by each of them. Actually, what this table displays is simply a level of noise (see Trend 2 ‘Less Semantic’), without any intelligence. A priori, this lack of intelligence can hardly be considered as aligned with the goals of AIED, but it illustrates some kind of minimalism that we called ‘*modest computing*’: displayed information is not the focus of attention, the resolution of information is low (we display some coloured surface below a blurring glass instead of displaying the number of spoken seconds, which would actually have been easier), no interpretation is proposed, etc. (Dillenbourg et al. 2011). The same minimalism applies to the Lantern (Fig. 1). This device was designed after discovering the low accuracy of orchestration during recitation sections, namely the inappropriate distribution of the teaching assistant time over waiting teams (Alavi and Dillenbourg 2012). Lantern is a distributed awareness device that displays which team is working on which exercise, for how long and how long they have been waiting for the teaching assistant. Empirical studies have shown that Lantern significantly increases the efficiency of these sections, namely because students continue working while waiting for the assistant.

Awareness tools support two regulatory functions illustrated by Reflect and Lantern. On the one hand, Reflect table provides self-awareness, which triggers the learner’s reflection and self-regulation at the team level (or ‘*group mirrors*’, Jermann and Dillenbourg 2008). On the other hand, Lantern provides awareness of others, which aims to facilitate classroom orchestration (Dillenbourg 2013; see Trend 6 ‘More Teachers’). The teacher dashboards developed in the learning analytics community belong to this second category, the difference with Lantern being whether the information is displayed in a central way (in a dashboard) versus in a distributed way (in Lanterns). In both cases, the regulatory effect comes from a single principle: *to make things visible that would otherwise be invisible*, such as team balance, waiting times etc. Actually, this concept appeared in the early days of AIED with the notion of ‘reification’ (Reusser 1993): representing graphically (on the display) some abstract features



Fig. 1 The classroom is the system. Above, the classroom is the output (Lantern, Alavi and Dillenbourg 2012); below, the classroom is the input (Raca et al. 2014)

of problem solving that would otherwise remain invisible, such as the problem solving path of learners.

In summary, the digital-physical blur is bidirectional, the digital becoming more physical and the physical becoming more digital. Even if it is bidirectional, I entitled this trend ‘more physical’ because AIED started on digital grounds, and hence today has become more physical than it was in 1990. But, what is probably more important is the underlying epistemological shift that unifies both sub-trends: the digital-physical blur.

In other words, the object of (some) AIED researchers has changed the two basic science missions, discovery and invention: the focus of discovery and the goal of invention shifted from a computer-based system to a learning environment, not in the sense of a ‘virtual learning environment’ but in the sense of a physical classroom (or any equivalent space in informal education), being treated as part of the digital space.

I expect the digital-physical issue to become obsolete in the future. Instead, I expect researchers to have a global view that could be called ‘*the classroom¹ is the system*’, as illustrated by the juxtaposition of the two pictures in Fig. 1. *The input is the classroom*: the input is not restricted to the learners’ behaviour but includes any classroom signal that can be captured, e.g. measuring classroom level of attention (see Trend 2). *The output is the classroom*: the output is no longer a simple display but a *range of ways to change the physical learning space*. This is illustrated by projects such as WallCology (Malcolm et al. 2008) and smart classrooms (Slota et al. 2013). In other words, while AIED initially aimed at modelling the contents and the learner, a challenge for **the future of AIED** is to *model educational spaces*, i.e. the physical space and the diverse actors who inhabit this space, in order to make education more effective. We conceptualised this evolution by defining a *third circle of usability* (Dillenbourg et al. 2011), where the user is not an individual (first circle) or a team (second circle) but the entire classroom is viewed as a physical and sociological system.

¹ Classroom is used here as a metaphor exportable to informal education where the classroom for a sailing coach would be the sea and for a cooking club would be a kitchen.

Trend 2: Less Semantic

The hypothesis that underlies the birth of AIED could be stated as follows: if a learning environment includes an articulated model of the task the learner has to do (the domain model), then it would enable a deeper understanding of the learner's mistakes (student model) which would hence enable better pedagogical decisions (tutor model). The domain model added some 'semantics' to any action in the environment. These three models used to be familiarly referred to as the 'trilogy' of AIED. The development of the domain model relied on AI modelling languages; the tutoring model did not receive as much attention as other components and hence the main effort has been placed on learner modelling, i.e. the interpretation of learners' behaviours (solving actions, answers, utterances). One intuition was that, since a single answer can often be interpreted in multiple ways, considering a sequence of answers (or a sequence of solution steps) would enable a more reliable cognitive diagnosis. This aggregation thereby further increased the semantic value of the learner's behaviour. The growth of qualitative methods in neighbouring scientific communities, such as CSCL and Learning Science, also witnessed a thirst for semantics. At the same time, this thirst contrasted with the loss of semantics within other disciplines that formerly belonged to AI, such as computer vision or language translation. It also contrasted with the success of 'semanticless' approaches in new disciplines such as social signal processing (Vinciarelli et al. 2009). If the learner behaviour was the atom of learner modelling, new approaches dwell inside the atom, capturing behavioural fragments or behavioural particles that, considered individually, do not convey very many semantics: gaze positions, body posture, keyboard typing patterns, speech prosody, brain waves, etc. Of course, learning scientists always had interest in non-verbal behaviour, but it expanded into multiple types of *behavioural dust* (Dillenbourg 2015) due to the evolution stressed in the first trend: the availability of multiple sensors that can easily be interfaced with computers and the improvement of machine learning methods. Capturing anything that can be captured and running black box predictive algorithms may seem a caricature of science and hardly deserve the 'intelligent' qualification expected in AIED. Actually, this machine learning does not work well without the careful selection of features, i.e. abstractions of raw data, and these features do actually convey some semantics. For instance, Raca et al. (2014) estimated the global level of learners' attention in a lecture room by placing two cameras in front of students. The ratio of co-movements (students rotating their head in a near-synchronous way) predicted self-estimated attention levels. While it is hard to semantically relate head movements to attention, the explanation is that co-movements are generated by the movements of an external source of attention, namely the movements of the teacher, and therefore predict attention to this source. One could argue that these head movements have a semantic value, i.e. that any data pattern that is relevant in educational interventions inevitably encodes some semantics, that may have been defined explicitly a priori by some specific ontology (Jovanović et al. 2006) or elaborated a posteriori by a human interpretation. This is not the right place to re-open the Chinese room debate that animated AI scholars for many years (Harnad 1990) or to enter the current debate in machine learning regarding black box models. I simply point out that the shift from interpreting a correct/wrong answer to interpreting behavioural dust (e.g. head rotation) does indeed constitute a fundamental evolution of our field. This evolution is certainly

obvious for scholars engaged in the ‘learning analytics’ and ‘educational data mining’ communities, but it constitutes a deep change in the way we think about learning environments.

The described drift away from semantics has been reinforced by the power of probabilistic approaches. The use of Bayesian networks for learner modelling appeared in the early nineties, as pointed out by Self (1995) and continued to grow (Zapata-Rivera and Greer 2000; Conati et al. 2002; Desmarais and Pu 2005). The general idea is to take into account the uncertainty of cognitive diagnosis, namely the probability that the learner is in a certain state x_i given the fact that he or she produced the behaviour b_i , denoted $p(x_i|b_i)$. The same evolution also concerns the tutoring model (Mayo and Mitrovic 2001). It was initially designed as a container of the pedagogical knowledge necessary to make decisions at run time, namely individual adaptation of instruction. This knowledge was initially stored in the declarative formats that dominated AI in the eighties, such as production rules, but it evolved towards probabilistic approaches. Already in the eighties, John Self (1990) proposed the slogan ‘do not diagnose what you cannot treat’, which brought pragmatism to research on learner modelling and the pedagogical model. Translated into computational terms (Murray et al. 2004), this means that if the system has the choice between two learning activities, a_n and a_m , all it needs to know is whether the learner in state x_i has more chances to reach a state x_j with a_n than a_m : $p(x_j|x_i, a_n) > p(x_j|x_i, a_m)$. A set of numerical values, probabilities or state transition matrices (Dillenbourg 2015) indeed encompass the pedagogical knowledge that would have been stored in the tutoring model in declarative or procedural formats. These approaches are especially taking advantage of the number of participants in Massive Open Online Courses (MOOCs). Is this (probabilistic) empirically acquired pedagogical expertise more effective than the (symbolic) elicited pedagogical knowledge such as aptitude-treatment laws? One drawback of these new approaches is the ‘cold start’ problem (Eagle and Barnes 2015), i.e. the fact that these algorithms only become reliable when a sufficient number of learners have used the system. One advantage is that the acquired model is specific to the learning tasks and audiences; it avoids trying to instantiate general rules. Another advantage is that some of these methods, namely Markov models, capture the dynamics of a learning process, which was something quite difficult with declarative models. This trend is not specific to digital education but it is the side effect of a more general evolution of computer science. We return to this point in Trend 4.

I expect that these probabilistic approaches will continue to develop in education in the coming years. There is no reason to exclude semantic information: in principle any information that improves performance should be integrated. Hence, a challenge for **the future of AIED** is the integration of multiple levels of abstraction, from raw signals to domain ontologies, into multi-layered algorithms.

Trend 3: More Social

AIED was born when learning was mostly perceived as an individual activity and when technologies were designed as individual workspaces. Cognitive science played a double role, inspiring both a vision of learning as the construction of (individual) knowledge structures and the design of AI algorithms, somehow ‘naturally’ bringing

together these two facets of the design of learning environments. In this context, it was natural that the educational value of computers was conceived as the possibility to individualise instruction. So-called ‘personal computers’ emerged as the right tools for implementing state-of-the-art of educational research: given the discovery of ‘aptitude-treatment interaction’ laws, the effectiveness of a pedagogical approach varied according to the individual characteristics of the learner and computers were the appropriate machines to deliver adaptive instruction. Now, as there were more children than computers, researchers often had to put several children per computer. In this case, the individualisation could be expected to break down and the learning gains should hence have been lower. The opposite has been observed (Dickson and Vereen 1983): the lower finesse of adaptation was compensated by the higher richness of interactions among learners.

These findings came out more or less synchronously with the resurgence of social cognition theories. These theories existed long before AIED but did not become influential in our communities until the late eighties. As with individual cognition twenty years before, there was a convergence between a vision of human learning as being social and a set of technologies that foster social interaction. It is interesting to notice that, in Europe, most of the CSCL pioneers came from the AIED community, (Baker 1993; Hoppe 1995; O'Malley 1995). A few decades later the influence of Vygotsky on the development of digital education is probably greater than the influence of Piaget, whose constructivism is often reduced to active learning. This evolution has also been visible in AIED with the research on intelligent collaborative systems (Soller 2001), group student models (Hoppe 1995) and online communities (Cheng and Vassileva 2005).

A decade later, the *socio-individual pendulum* somehow bounced back to the middle as research revealed that, to be effective, collaborative learning requires some level of structuring, i.e. the definition of individual roles within team activities. Instead of engaging in free collaboration, CSCL designers ask teams to follow what is referred to as a collaboration script (Weinberger et al. 2002; Dillenbourg and Hong 2008). Learners have a personal brain, a personal motivation, some ‘agency’ that drives collaboration, as stated by Schwartz (1995).

A similar pendulum move is happening with MOOCs. At the outset, the so called ‘cMOOCs’ – in which ‘c’ refers to Siemens’ connectivism theory (Siemens 2013) – pushed the pendulum very far towards the social pole. A few years later, the large scale xMOOCs emerged, in which the discourse emphasised the individual pace of instruction. Nowadays, xMOOC practices have revealed the key role of social interaction even in xMOOCs: the main predictor of drop out is the lack of engagement, as estimated for instance by forum participation (Wen et al. 2014; Sinha et al. 2014), participants spontaneously meeting physically or online, large sets of participants being segmented into cohorts to maximise interactions with homogenous sub-populations, the crowdsourcing methods used for peer grading and other collective interactions that are ultimately social, etc. In other words, MOOCs are indeed both very individual and very social.

At the centre of this social-individual continuum, *integrated learning scenarios* combine individual, team and class-wide activities (Dillenbourg and Jermann 2007). No, instructional designers do not have to stay within the walls of one learning theory. Yes, it is perfectly fine to develop scenarios that combine lectures and teamwork and

individual activities. The term ‘integration’ is two-fold. First, these activities need to constitute a consistent whole, for instance, contrasted cases discussed in teams followed by a lecture (Schwartz and Bransford 1998). Second, these activities can be digitally integrated: the output of an activity feed into the next activity (e.g. the individual opinions collected by a questionnaire are used for group formation for the next activity). A sequence of activities connected by pedagogical relationships and dataflow operators constitute what I have formalised through orchestration graphs (Dillenbourg 2015).

This third evolution trend is similar to the first one: the opposition between the individual and social views of cognition has somehow been blurred, in the same way that the distinction between digital and virtual has been blurred. However, it is fair to say that, currently, learning technologies rely much more on social cognition than they did when AIED emerged. Functionalities that support individual learning and those enhancing social learning processes both have to be integrated into a learning environment. Even if a learner works alone, his or her reasoning is shaped by language. Even if he or she works in a team, the team processes partly depend upon his or her personal behaviour, knowledge and motivation. Therefore, a challenge for **the future of AIED** is the development of computational models in which the individual and social aspects of human cognition are taken into consideration with each learning activity. *Even if the brain is an individual piece of hardware, it operates social software. Even if we are social animals, we still have a personal brain.*

Trend 4: Less Design

Another deep evolution trend concerns the way we design learning environments. Computer-assisted instruction (CAI) originated from a branch of educational research known as ‘mastery learning’ (Bloom and Carroll 1971), which developed evidence-based principles and several methods for *instructional design*. This design is deterministic: even if a learning environment includes multiple paths, the designer needs to anticipate every question, every feedback and every branching decision. At the opposite, in a microworld such as LOGO, there was no didactic plan and no predefined path; learners may freely explore the environment. Of course, there is no ‘zero design’: even in a microworld, the cognitive activities of learners are indirectly shaped by the set of objects they manipulate. These two examples, CAI and microworlds, illustrate two different forms of design, the former being closer to train scheduling, the second being closer to architectural design (e.g. designing building entrances that afford social interactions). The field of AIED has explored these two ends, sometimes using ‘high’ design as in the cognitive tutors (Aleven et al. 2009) or softer design as with constraint-based approaches (Mitrovic and Ohlsson 1999). The very idea of AIED actually lies in the middle of this continuum: if the system has a map of the city to cross (the domain model) and an idea of where the learner stands and wants to go (the learner model), instead of planning 5 paths (hard design), one may let the student explore freely and still provide feedback and hints. The field of CSCL was initially closer to the ‘low design’ end: the word ‘supported’ referred to the fact that the collaboration tools actually mediated interactions among learners and hence their cognition. For instance, Roschelle (1992) showed that the role of the symbolic representations was

less to convey proper semantic values than to shape verbal interactions among learners. Later on, the development of macro-scripts brought CSCL closer to instructional design, since these scripts specify a sequence of activities, but also include some ‘indirect’ design, e.g. when selecting two learners with opposite opinions to trigger socio-cognitive conflicts.

Today, this pendulum continues bending towards the ‘less design’ end, which I illustrate now with 4 examples. The first example is what G. Fischer (2000) called *meta-design*, letting learners be the designers of their own environment, a concept that was later on reflected in the cMOOCs developments. Tchounikine (2008) has applied this approach to CSCL by letting learners design their own macro-scripts. This approach is boosted today by the (great) loss of control that we have over learners’ environments: even when they use a learning environment we designed, learners may open any number of other windows with a Wiki, a chat, a MOOC, etc. The second example of ‘less design’ is the use of *crowdsourcing* methods for acquiring teaching content: instead of designing 20 carefully selected positive and negative instances of the concept to be taught, as a designer would have done 30 years ago, he or she can now ask his or her 2000 online students to each upload 3 examples and thereby obtain a noisy but very rich set of examples. The third example is the use of recommender systems in education (Tang and McCalla 2003) since the selection of learning contents or resources is not based on decision rules defined by the designer but on the experience (ratings, performance) of other learners. The last example is related to the probabilistic approach presented in Trend 2: instead of defining a rule that decides if the learner who completes activity 4 should move to activity 5 or jump to activity 8, the designer may implement both activities, try them and let the system collect empirical evidence of when one activity is more suited than the other. Of course, this approach has some limits: one cannot leave all decisions open, but some key decisions can be left to the probabilistic methods described in the second trend. This idea actually emerged at the very beginning of AIED, then called ‘*self-improving tutoring systems*’ (O’Shea 1979; Kimball 1982; Dillenbourg 1989) but faded out since then, probably because machine learning methods were quite ineffective at that time and because the size of data sets were then nothing compared to those we can collect today.

A challenge for **the future of AIED** is to understand the complementarity between the craft of designing (manually) learning activities and the power of stochastic methods for optimising the initial design.

Trend 5: More Open

Learning technologies are more open today than when AIED was born, in any meaning of the word ‘open’: in the sense of free access (as in MOOCs), in the sense of open source, in the sense of being open to participants’ contributions (not all content is provided by the designer as used to be the case) and in the sense of open architecture (e.g. RSS feeds bringing information from multiple sources into learning environments). This evolution is rather obvious, I do not need to develop these points here, but it is still worth mentioning them because the current situation is radically different than 25 years ago. I will only develop three points: the openness of systems, the openness to the learner’s world and the openness of the community.

The first point concerns the software architecture of learning environments. In the early days, the AIED landscape was structured around a few exemplar intelligent systems, which constituted the landmarks of our scientific debate. At that time, a ‘system’ was indeed a set of files stored on a computer, written more or less from scratch and often incompatible with any other learning environment. This system was designed to be used by a single learner, sitting in front of a personal computer, without any other actor (teacher or peers). This was before the development of web technologies and the emergence of distributed software architecture. Currently, a learning environment is made of multiple software layers and multiple services running somewhere on multiple devices. Despite the fact that this software evolution affected all domains in computer science, many learning technology projects are still about developing ‘a system’ rather than adding, for instance, a component to an existing ecosystem. Nonetheless, recent projects have aimed at providing services to other learning environments, for instance group formation services (Isotani et al. 2009) or MOOC-add-ons for discussion such as Prollster.² Distributing systems come with some drawbacks, namely in terms of complexity. For several years, the interoperability standards were limited to metadata, but the emergence of simple standards such as LTI now paves the road to rich digital educational ecosystems. We somehow broke away from the romantic myth in which the goal of a PhD would be to develop ‘my system’ but rather to contribute to these ecosystems. This evolution implies some drawbacks in terms of complexity but offers unprecedented possibilities in terms of what can be collectively built.

The second type of openness is that learning environments are more open to the external world than in the past, since students more easily bring their world into the digital environment, for instance by uploading stories, pictures, movies or any digital object. This evolution reflects the increasing influence of social media on learning technologies. This is illustrated for instance by cMOOCs or any MOOC that relies on participants’ contributions. Another example comes from vocational education and training (VET), namely from dual systems (apprentices work for instance 4 days per week in a company and one day at school). In a VET context, opening the learning environment towards the world is critical, the challenge being to maintain a strong relationship between the school and the workplace (Schwendimann et al. 2015). Their environment illustrates this openness: apprentices are invited to take pictures or videos of specific workplace objects or events (defined by the teacher), and this material is automatically collected into a classroom ‘flow’ where they can be annotated, commented, liked, classified and compared in order to feed classroom activities.

The third example of increased openness concerns the scholars coming from different fields that enter into digital education. AIED has always been an interdisciplinary field, but with a limited set of disciplines, mostly computer science, education, psychology. Some of the recent success stories, from Khan Academy to MOOC platforms, have been initiated by people who, despite having no expertise in the field, or because of it, happened to be more creative or to dare more than we did. Digital education does not belong to anyone, of course. This opening of our field has drawbacks, such as a regular wave of over-expectations with regard to the effects of

² <https://prollster.com/en/>

learning technologies, but these are minor issues compared to the new possibilities it generates.

I believe the segmentation of our communities into small societies, which allowed us to grow for a few decades, has now become an obstacle to our societal impact. Hence, a challenge for **the future of AIED** is to *converge with the archipelago of small sibling communities* that have been mentioned in order to reach the critical mass, the voice, the visibility that we need to bring significant contributions to formal and informal education. We need stronger communities to develop worldwide norms that enable the creation of rich digital ecosystems for education.

Trend 6: More Teachers

Despite a few discording voices ('neo-Illich' gurus), educational technology researchers have never believed that their technology would suppress the need for teachers in formal education. However, it is not unfair to say that teachers have received minor attention within the research we conducted as well as a minor role in the environments we designed. For instance, many scholars used the slogan 'from a sage on the stage, to a guide on the side'. This slogan sounds politically correct from a constructivist viewpoint, but neglects how unpleasant the idea can be for professionals to be moved to 'the side'. The role of teachers became more visible during the efforts to scale up AIED environments (Koedinger et al. 1997).

Teachers may intervene at three points in time within the life cycle of a learning environment, namely before, during and after runtime. In the design phase, the role of teachers has received attention as co-designers during participatory processes (Underwood and Luckin 2011), as autonomous designers using the authoring tool (Murray 1999; Bourdeau and Mizoguchi 2002) or as content providers (Brusilovsky et al. 2006). The role of teachers after runtime has been mostly investigated in learning analytics. The role of teacher during runtime did not receive much attention for two decades, but this changed a lot over the last decade, with the growing interest for the *orchestration* of computer-enhanced learning activities (Dillenbourg 2013). The notion of classroom orchestration existed before in educational research but this term emerged in the CSCL community in the last decade. Orchestration is the real time management of multi-plane activities (individual, team and class-wide activities) as described in Trend 3: some of the activities can be implemented in an autonomous environment, but the whole graph of individual, team and classroom activities distributed over the classroom ecosystem needs to be orchestrated. Beyond what is traditionally referred to as classroom management, orchestration emphasises the many practical constraints of classroom life such as time, space, visibility, safety... something that Miguel Nussbaum appropriately called the 'logistics' of education (Nussbaum and Diaz 2013). The need for orchestration is a consequence of the previous trend: as long as AIED systems were rather closed, computationally speaking, there was little need for external intervention, but, as they became more open, they created greater need and more opportunities for teacher interventions.

My claim is not that the role of the teacher matters (most scholars would agree with this claim), but that the design of learning technologies needs to take, more than before, teachers' practical constraints into consideration. Initially, AIED was about modelling how people learn, modelling the contents to be learned and modelling how the tutoring could be facilitated, the famous 'trilogy'. Currently, the design of usable learning technologies requires taking into consideration the whole context: does the environment require login (useless waste of time) or to darken the room (some digital tabletops did), does it generate homework, does it produce grades, etc.? For instance, we found that the value of paper components in digital learning environments was not to enable learning activities but to facilitate classroom orchestration by making the workflow visible (e.g. collecting or distributing sheets). Paper is especially well suited to the classroom ecosystem (Dillenbourg 2013).

An interesting point regarding teachers that has only received attention in our communities as a methodological bias, is known as the Rosenthal effect: the subjects in an experiment tend to behave in a way that corresponds to the expectations of the experimenter. Outside experiments, this effect should not be considered as a bias but is indeed a great pedagogical mechanism: since students are sensitive to the expectations of their teachers, a good teacher will make his or her expectations visible. My point is that, even if we implement the most effective learning activities, even if students bring their own self-expectations, humans will certainly remain less sensitive to the expectations of an artificial agent than to the expectations of a human teacher.

Another trend in AIED and other communities has been to dissociate the teacher's role (McCalla 2000), from the teacher's status. Someone who has the status 'teacher' is or should be accountable for someone else's learning, which implies many constraints such as respecting a time budget, managing safety and certifying acquired skills. The teacher's responsibility is a distinctive feature of formal education. At the opposite, the role of teaching someone can be played by many stakeholders. Some pedagogical scenarios ask learners to play the role of teachers in order to benefit from the 'learning by teaching' effects: in this approach, the human learner is invited to teach either another human learner or an AI agent who simulates a learner (Biswas et al. 2005). In informal education, the teaching role may emerge naturally when peers help each other (Greer et al. 1998) or simply in an online community platform that matches those who can teach a topic to those who want to learn it.

This brings me to the debate around formal and informal education. Informal/lifelong learning was almost absent from our research at the outset, but it received a growing interest in the field of learning technologies over the last 15 years. Today I feel concerned that it attracts too much attention in developed countries, maybe because it escapes the many constraints of research in formal education. **A challenge for AIED** is to keep fighting for the improvement of formal education, still often quite ineffective, a world where our children spend about 15 years of their life but remains ineffective. Teachers are central to this goal. In parallel, in some countries, formal education cannot grow at the same speed as the population that needs to be educated. Therefore, **another challenge** for the future of AIED is to develop technologies that both empower teachers in countries that have enough teachers but can also be implemented in countries where students do not have access to (enough) teachers. In this context, a **third challenge** is to expand AIED towards vocational education, training skilled workers around the world.

Conclusions

I could describe several other trends: more video (video has now become a basic digital substance; it was not so easy to handle 25 years ago), more mobile, more students, etc. I think these trends do not require explanations. I would rather point out two meta-trends.

First, most of the trends described in this paper are not specific to the field of learning technologies; they depict the evolution of computer science as a whole and our evolution towards a digital society. Does it mean that the evolution of AIED is only driven by technological evolution? Certainly not ‘only’: the third trend (‘more social’) results from the evolution of learning theories and the last trend (‘more teachers’) mostly comes from empirical studies, namely the observation that the talent of teachers remains a condition for obtaining sustainable learning effects. However, even if this statement may shock some educational researchers (I am one of them), this paper illustrates that the evolution of learning technologies has been mostly driven by the evolution of technologies. Computer science opens new spaces and our communities are eager to explore them, to find new opportunities to support education. This technological drive has received criticisms, but it is a reality: computer science is evolving faster than education. However, our work has been inspired by the evolution of technologies, but not only by this evolution. I have illustrated the fact that the pivotal points in the history of AIED occurred precisely when a new technology matches a new vision of learning and vice-versa. Understanding these moments of convergence is the essence of our scientific community.

A second meta-trend is that most trends reported here followed the same two phase pattern. First, new approaches emerged because they were clearly different from existing approaches: physical interfaces were novel compared digital ones; the success of ‘semanticless’ approaches has been surprising compared to former AI approaches; collaborative activities were drifting away from the myth of individualisation; informal learning was a reaction against the dominance of formal education; etc. Then, in a second phase, this duality faded away; researchers bypassed the sterile opposition between ideas and the resulting litany of overstatements. Different approaches are first perceived as mutually exclusive but then integrated for their complementarity. Wisdom comes with time. Such a two-phase cycle takes between 5 and 10 years, so we should witness of few more of them in the coming decade.

Acknowledgments Thanks to Gordon McCalla and Chad Lane who commented on an earlier version of this paper.

References

- Alavi, H. S., & Dillenbourg, P. (2012). An ambient awareness tool for supporting supervised collaborative problem solving. *IEEE Transactions on, Learning Technologies*, 5(3), 264–274.
- Aleven, V., McLaren, B. M., Sewall, J., & Koedinger, K. R. (2009). A new paradigm for intelligent tutoring systems: Example-tracing tutors. *International Journal of Artificial Intelligence in Education*, 19(2), 105–154.
- Bachour, K., Kaplan, F., & Dillenbourg, P. (2010). An interactive table for supporting participation balance in face-to-face collaborative learning. *IEEE Transactions on Learning Technologies*, 3(3), 203–213.

- Baker, M. J. (1993). *Dialogic learning: Negotiation and argumentation as mediating mechanisms, Proceedings of the World Conference on Artificial Intelligence in Education*, (pp. 4–11). UK: Edinburgh.
- Biswas, G., Leelawong, K., Schwartz, D., Vye, N., & The Teachable Agents Group at Vanderbilt (2005). Learning by teaching: A new agent paradigm for educational software. *Applied Artificial Intelligence*, 19(3–4), 363–392.
- Blikstein, P. (2013). Multimodal learning analytics. In *Proceedings of the third international conference on learning analytics and knowledge* (pp. 102–106). New York: ACM.
- Blikstein, P., & Krannich, D. (2013). The makers' movement and FabLabs in education: experiences, technologies, and research. In *Proceedings of the 12th international conference on interaction design and children* (pp. 613–616). New York: ACM.
- Bloom, B. S., & Carroll, J. B. (1971). *Mastery learning: Theory and practice*. J. H. Block (Ed.). New York: Holt, Rinehart and Winston.
- Bourdeau, J., & Mizoguchi, R. (2002). Collaborative ontological engineering of instructional design knowledge for an ITS authoring environment. In *Intelligent Tutoring Systems* (pp. 399–409). Springer:Berlin Heidelberg.
- Brusilovsky, P., Knapp, J., & Gamper, J. (2006). Supporting teachers as content authors in intelligent educational systems. *International Journal of Knowledge and Learning*, 2(3–4), 191–215.
- Cheng, R., & Vassileva, J. (2005). Adaptive reward mechanism for sustainable online learning community. In *Proceedings of the 2005 conference on Artificial Intelligence in Education: Supporting Learning through Intelligent and Socially Informed Technology* (pp. 152–159). Amsterdam: IOS Press.
- Conati, C., Gertner, A., & VanLehn, K. (2002). Using Bayesian networks to manage uncertainty in student modeling. *User Modeling and User-Adapted Interaction*, 12(4), 371–417.
- Desmarais, M. C., & Pu, X. (2005). A Bayesian student model without hidden nodes and its comparison with item response theory. *International Journal of Artificial Intelligence in Education (IJAIED)*, 15, 291–323.
- Dickson, W. P. & Vereen, M. A. (1983). Two students at one microcomputer. *Theory into Practice*, 22(4), 296–300. Special Issue: *Microcomputers: A Revolution in Learning*
- Dillenbourg, P. (1989). Designing a self-improving tutor: PROTO-TEG. *Instructional Science*, 18(3), 193–216.
- Dillenbourg, P. (2013). Design for classroom orchestration. *Computers & Education*, 69, 485–492.
- Dillenbourg, P. (2015). *Orchestration graphs: Modeling scalable education*. Lausanne: EPFL Press.
- Dillenbourg, P., & Hong, F. (2008). The mechanics of CSCL macro scripts. *International Journal of Computer-Supported Collaborative Learning*, 3(1), 5–23.
- Dillenbourg, P., & Jermann, P. (2007). Designing integrative scripts. In F. Fischer, H. Mandl, J. Haake, & I. Kollar (Eds.), *Scripting Computer-Supported Collaborative Learning – Cognitive, Computational, and Educational Perspectives* (pp. 275–301). Computer-Supported Collaborative Learning Series, New York: Springer.
- Dillenbourg, P., Zufferey, G., Alavi, H., Jermann, P., Do-Lenh, S., Bonnard, Q., & Kaplan, F. (2011). Classroom orchestration: The third circle of usability. *CSCL2011 Proceedings*, 1, 510–517.
- Eagle, M., & Barnes, T. (2015). Exploring Missing Behaviors with Region-Level Interaction Network Coverage. In *Proceedings of the Artificial Intelligence in Education: 17th International Conference, AIED 2015, Madrid, Spain, June 22–26, 2015*. Volume 9112 of the series Lecture Notes in Computer Science (pp 831–835). Springer International Publishing.
- Fischer, G. (2000). Symmetry of ignorance, social creativity, and meta-design. *Knowledge-Based Systems*, 13(7), 527–537.
- Greer, J. E., McCalla, G., Collins, J. A., Kumar, V. S., Meagher, P., & Vassileva, J. (1998). Supporting peer help and collaboration in distributed workplace environments. *International Journal of Artificial Intelligence in Education (IJAIED)*, 9, 159–177.
- Hamad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1), 335–346.
- Hood, D., Lemaignan, S., & Dillenbourg, P. (2015). The CoWriter Project: Teaching a Robot how to Write. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts* (pp. 269–269). New York: ACM.
- Hoppe, U. (1995). The use of multiple student modeling to parameterize group learning. In J. Greer (Ed.), *Proceedings of AI-ED'95, 7th World Conference on Artificial Intelligence in Education*, Washington, DC; August 16–19. Waynesville (NC): AACE.
- Ishii, H., & Ullmer, B. (1997). Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems* (pp. 234–241). New York: ACM.

- Isotani, S., Inaba, A., Ikeda, M., & Mizoguchi, R. (2009). An ontology engineering approach to the realization of theory-driven group formation. *International Journal of Computer-Supported Collaborative Learning*, 4(4), 445–478.
- Jermann, P., & Dillenbourg, P. (2008). Group mirrors to support interaction regulation in collaborative problem solving. *Computers & Education*, 51(1), 279–296.
- Jermann, P., Zufferey, G., Schneider, B., Lucci, A., Lépine, S., & Dillenbourg, P. (2009). Physical space and division of labor around a tabletop tangible simulation. In *Proceedings of the 9th international conference on Computer supported collaborative learning-Volume 1* (pp. 345–349). International Society of the Learning Sciences.
- Jovanović, J., Gašević, D., Brooks, C., Knight, C., Richards, G., & McCalla, G. (2006). Ontologies to support learning design context (pp. 615–620). Springer: Berlin Heidelberg.
- Kimball, R. (1982). A self-improving tutor for symbolic integration. In *Intelligent tutoring systems* (Vol. 1). Academic Press: London.
- Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education*, 8, 30–43.
- Leroux, P., Vivet, M., & Brézillon, P. (1996). Cooperation between humans and a pedagogical assistant in a learning environment. In *Proceedings of European Conference in AI in Education (EuroAIED)*, Lisbon, Portugal (pp. 379–385).
- Lucchi, A., Jermann, P., Zufferey, G., & Dillenbourg, P. (2010). An empirical evaluation of touch and tangible interfaces for tabletop displays. In *Proceedings of the fourth international conference on Tangible, embedded, and embodied interaction* (pp. 177–184). New York: ACM.
- Malcolm, P., Moher, T., Bhatt, D., Uphoff, B., & López-Silva, B. (2008). Embodying scientific concepts in the physical space of the classroom. In *Proceedings of the 7th international conference on Interaction design and children* (pp. 234–241). New York: ACM.
- Martínez, R., Collins, A., Kay, J., & Yacef, K. (2011). Who did what? Who said that?: Collaid: An environment for capturing traces of collaborative learning at the tabletop. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces* (pp. 172–181). New York: ACM.
- Mayo, M., & Mitrovic, A. (2001). Optimising ITS behaviour with Bayesian networks and decision theory. *International Journal of Artificial Intelligence in Education*, 12, 124–153.
- McCalla, G. (2000). The fragmentation of culture, learning, teaching and technology: implications for the artificial intelligence in education research agenda in 2010. *International Journal of Artificial Intelligence in Education*, 11(2), 177–196.
- Mitrovic, A., & Ohlsson, S. (1999). Evaluation of a constraint-based tutor for a database language. *International Journal of Artificial Intelligence in Education*, 10, 238–256.
- Murray, T. (1999). Authoring intelligent tutoring systems: An analysis of the state of the art. *International Journal of Artificial Intelligence in Education (IJAIED)*, 10, 98–129.
- Murray, R. C., Vanlehn, K., & Mostow, J. (2004). Looking ahead to select tutorial actions: A decision-theoretic approach. *International Journal of Artificial Intelligence in Education (IJAIED)*, 14, 235–278.
- Nussbaum, M., & Diaz, A. (2013). Classroom logistics: Integrating digital and non-digital resources. *Computers & Education*, 69, 493–495.
- Ochoa, X., Chiluita, K., Méndez, G., Luzardo, G., Guamán, B., & Castells, J. (2013). Expertise estimation based on simple multimodal features. In *Proceedings of the 15th ACM on International conference on multimodal interaction* (pp. 583–590). New York: ACM.
- O'Malley, C. (Ed.). (1995). *Computer supported collaborative learning*. NATO ASI Series (Vol. 128). Springer: Berlin Heidelberg.
- O'Shea, T. (1979). A self-improving quadratic tutor. *International Journal of Man-Machine Studies*, 11(1), 97–124.
- Raca, M., Tormey, R., & Dillenbourg, P. (2014). Sleepers' lag-study on motion and attention. In *Proceedings of the Fourth International Conference on Learning Analytics and Knowledge* (pp. 36–43). New York: ACM.
- Reusser, K. (1993). Tutoring systems and pedagogical theory: Representational tools for understanding, planning, and reflection in problem solving. *Computers as Cognitive Tools*, 1, 143–177.
- Rosch, E., Thompson, E., & Varela, J. (1991). *The embodied mind: Cognitive science and human experience*. Cambridge: MIT Press.
- Roschelle, J. (1992). Learning by collaborating: Convergent conceptual change. *Journal of the Learning Sciences*, 2, 235–276.
- Schwartz, D. L. (1995). The productive agency that drives collaborative learning. In P. Dillenbourg (Ed.), *Collaborative learning: Cognitive and computational approaches* (pp. 197–218). NY: Elsevier Science/Pergamon.

- Schwartz, D. L., & Bransford, J. D. (1998). A time for telling. *Cognition and Instruction*, 16(4), 475–5223.
- Schwendimann, B. A., Cattaneo, A. A., Dehler Zufferey, J., Gurtner, J. L., Bétrancourt, M., & Dillenbourg, P. (2015). The ‘Erfahrungsraum’: A pedagogical model for designing educational technologies in dual vocational systems. *Journal of Vocational Education & Training*, 67(3), 367–396.
- Self, J. A. (1990). Bypassing the intractable problem of student modelling. In C. Frasson & G. Gauthoer (Eds.), *Proceedings “Intelligent tutoring systems: At the crossroads of artificial intelligence and education,”* (pp. 107–123). New Jersey: Ablex.
- Self, J. A. (1995). *Computational Mathematics. Towards a Science of Learning Systems Design*. Lancaster: Drakkar Press.
- Selker, T., Arroyo, E., & Burleson, W. (2002). Chameleon tables: Using context information in everyday objects. In *CHI02 Extended Abstracts on Human Factors in Computing Systems* (pp. 580–581). ACM.
- Siemens, G. (2013). Massive open online courses: Innovation in education. In R. McGreal, W. Kinuthia & S. Marshall (Eds.), *Open Educational Resources: Innovation, Research and Practice* (pp. 5–15). Canada: Commonwealth of Learning, Athabasca University.
- Sinha, T., Jermann, P., Li, N., & Dillenbourg, P. (2014). Your click decides your fate: Inferring information processing and attrition behavior from MOOC video clickstream interactions. In *2014 Empirical Methods in Natural Language Processing Workshop on Modeling Large Scale Social Interaction in Massively Open Online Courses*.
- Slotta, J. D., Tissenbaum, M., & Lui, M. (2013). Orchestrating of complex inquiry: three roles for learning analytics in a smart classroom infrastructure. In *Proceedings of the Third International Conference on Learning Analytics and Knowledge* (pp. 270–274). New York: ACM.
- Soller, A. (2001). Supporting social interaction in an intelligent collaborative learning system. *International Journal of Artificial Intelligence in Education (IJAIED)*, 12, 40–62.
- Tang, T. Y., & McCalla, G. (2003). Smart recommendation for an evolving e-learning system. In *Workshop on Technologies for Electronic Documents for Supporting Learning, International Conference on Artificial Intelligence in Education (AIED 2003)*. Sydney: International Conference on AI In Education.
- Tchounikine, P. (2008). Operationalizing macro-scripts in CSCL technological settings. *International Journal of Computer-Supported Collaborative Learning*, 3(2), 193–233.
- Toussaint, B. M., Luengo, V., Jambon, F., & Tonetti, J. (2010). Modeling perceptual-gestural knowledge for intelligent tutoring systems dedicated to ill-defined domains. *International Journal of Artificial Intelligence in Education (IJAIED)*, 20(3). doi:10.13140/RG.2.1.2362.9924.
- Ullmer, B., & Ishii, H. (2000). Emerging frameworks for tangible user interfaces. *IBM systems journal*, 39(3.4), 915–931.
- Underwood, J., & Luckin, R. (2011). Themes and trends in AIED research, 2000 to 2010. *A report for the UK's TLRP Technology Enhanced Learning–AIED Theme*. A report for the UK's TLRP Technology Enhanced Learning - AIED Theme. May 2011. The London Knowledge Lab. http://www.tel.ac.uk/wp-content/uploads/2011/06/telaied_part2researchthemes.pdf
- Vinciarelli, A., Pantic, M., & Bourlard, H. (2009). Social signal processing: Survey of an emerging domain. *Image and Vision Computing*, 27(12), 1743–1759.
- Walker, E., & Burleson, W. (2012). User-centered design of a teachable robot. In *Intelligent Tutoring Systems* (pp. 243–249). Springer: Berlin Heidelberg.
- Weinberger, A., Fischer, F., & Mandl, H. (2002). Fostering computer supported collaborative learning with cooperation scripts and scaffolds. In *Proceedings of the Conference on Computer Support for Collaborative Learning: Foundations for a CSCL Community* (pp. 573–574). ACM Digital Library.
- Wen, M., Yang, D., & Rose, C. P. (2014). Linguistic reflections of student engagement in massive open online courses. In *Proceedings of the International Conference on Weblogs and Social Media* (pp. 525–534). AAAI Digital Library.
- Zapata-Rivera, J. D., & Greer, J. E. (2000). Inspecting and visualizing distributed Bayesian student models. In *Intelligent Tutoring Systems* (pp. 544–553). Berlin Heidelberg: Springer.
- Zuckerman, O., Arida, S., & Resnick, M. (2005) Extending tangible interfaces for education: Digital Montessori inspired manipulatives. In *Proceedings of CHI '05* (pp. 859–868). New York: ACM Press.