

AI as a Methodology for Supporting Educational Praxis and Teacher Metacognition

Kaśka Porayska-Pomsta¹

© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract Evidence-based practice (EBP) is of critical importance in education where emphasis is placed on the need to equip educators with an ability to independently generate and reflect on evidence of their practices in situ – a process also known as *praxis*. This paper examines existing research related to teachers' metacognitive skills and, using two exemplar projects, it discusses the utility and relevance of AI methods of knowledge representation and knowledge elicitation as methodologies for supporting EBP. Research related to technology-enhanced communities of practice as a means for teachers to share and compare their knowledge with others is also examined. Suggestions for the key considerations in supporting teachers' metacognition in praxis are made based on the review of literature and discussion of the specific projects, with the aim to highlight potential future research directions for AIEd. A proposal is made that a crucial part of AIEd's future resides in its curating the role of AI as a methodology for supporting teacher training and continuous professional development, especially as relates to their developing metacognitive skills in relation to their practices.

Keywords AI as a methodology · Metacognition · Knowledge representation · Knowledge elicitation · Evidence-based practice in education

Introduction

Evidence-based practice (EBP) is of critical importance in education where emphasis is placed on the need to equip educators¹ with an ability to independently generate and reflect on evidence of their practices in situ (Hargreaves 1999a, b) – a process also

¹Educators are understood here as teachers, tutors, coaches and other front-line educational practitioners.

✉ Kaśka Porayska-Pomsta
K.Porayska-Pomsta@ucl.ac.uk

¹ Institute of Education, Knowledge Lab, University College London, 23-29 Emerald Street, London WC1N 3QS, UK

known as *praxis*. Although evidence is seen as key to educational praxis, and hence to educational practitioners' professional development and innovation (Hewitt et al. 2003), it is also fundamental to AIED research where it informs the design, deployment and evaluation of its technologies.

There are a number of substantial challenges related to EBP and to supporting educators in developing independent EBP skills. One such challenge derives from a lack of consensus as to what constitutes 'good' evidence in Education (e.g. Biesta 2007; 2013). This is further amplified by a growing interdisciplinary appreciation that learning and teaching processes give rise to varied forms of evidence that can be utilised at multiple levels. For example, they can be used to support collective vs. individual learning in different contexts, e.g. formal vs. informal, by different (constellations of) stakeholders and for different purposes such as policy makers and funders to set funding priorities, or researchers and practitioners to contribute to knowledge or to innovate their own practice. Disciplinary identities, professional allegiances and the perceived pragmatic utility of the available evidence are all instrumental in the value that individual stakeholders assign to different forms of evidence. For example, experimental methods such as randomised controlled trials (RCTs) are favoured by many researchers, by policy makers and educational administrators owing to their offering quantifiable insights, e.g. with respect to student attainment, that are akin to the outcomes of biological and medical sciences – presently the gold standard of scientific evidence. However, given the intersubjective nature of learning and teaching, i.e. one that requires a willing interaction and development of some common ground between a learner and a teacher, and the inextricable dependency of learning outcomes on the context in which learning and teaching takes place, there are also strong voices within the educational research community (e.g. Biesta 2007; Slavin 2004) pointing to frequent difficulties in generalising and reproducing the results emerging from RCTs alone to real learning and teaching practices, and to individual learners. While RCTs fulfil an important confirmatory function (e.g. that a particular form of intervention as applied in a specific and tightly controlled manner has/has not worked), many question the actual ability of conclusions gleaned through RCTs alone to drive real or long-lasting educational innovation *in practice* and to support front-line educational practitioners' professional development, which is at the heart of such innovation (e.g. Carr and Kemmis 2004).

To support such innovation in a balanced way, dedicated methods and technologies need to be developed, combining RCTs and qualitative (case-based) methods, to account for the multifarious, transactional and context-dependent nature of learning and teaching and the conditions under which they may be successful (Carr and Kemmis 2004; Parsons et al. 2013). Methods that enable practitioner introspection, reflection *in* and *on* action (Schön 1987) and adaptive metacognition (Lin et al. 2005) that can be shared with, inspected and reproduced or compared by other practitioners are deemed particularly important to supporting both the understanding and the establishment of best and innovative teaching practices and to offering an effective basis for educators' life-long professional development (Lin and Schwartz 2003; Hewitt et al. 2003; Lin et al. 2005; Laurillard 2012; Cohen and Manion 1980; Conlon and Pain 1996). Enabling such a sharable metacognitive exchange, however, presents a further challenge for EBP. This challenge relates to constructing an infrastructure that is both intellectual and technological, comprising of methods, tools and common knowledge representation

standards to enable educational practitioners to generate evidence of their practices independently from researchers, for different subject domains, using diverse pedagogical methods, and at different levels of granularity, i.e. from high level pedagogic designs, right down to the level of individual actions in context.

Speaking to the focus of the present special issue, in this paper I propose that a crucial part of the future of Artificial Intelligence in Education (AIEd) is in developing a techno-intellectual infrastructure to support EBP in a way that combines different forms of evidence with the relatively neglected area of educators' metacognition in praxis. The premise of this paper is that while evidence of good practice is fundamental to educational innovation, the way that such evidence is generated and represented and by whom, i.e. by researchers vs. practitioners, is key to our ability to link such knowledge to concrete situated actions and for such knowledge to be understandable, inspectable and usable by front-line educators. Methods used to elicit knowledge of teaching and learning processes and to represent it computationally, e.g. in terms of key components of an intelligent tutoring system, offer at least working prototypes of the tools that can help teachers to incrementally identify, externalise, test and share systematic and detailed evidence of their practices.

Crucial to my proposal is the transitive nature of AIEd's engagement with key issues and debates in Education. Specifically, investing in educational practitioners using AI techniques to generate and inspect teaching and learning practices is not entirely altruistic in that the specificity of the evidence gathered using those techniques creates an important opportunity for AIEd to tap into those practices in a way that supports the implementation of AIEd systems sustainably and over long term. Such investment carries a promise of creating a dynamically generated knowledge infrastructure, thus reducing the often-prohibitive cost of developing AIEd systems. Moreover, making the AI methods available to practitioners also offers a much needed opportunity to continuously re-interrogate AIEd's engagement with existing educational practices, along with its future goals and aspirations more generally. The latter includes the way in which the field may contribute more actively to educational policies and to teachers' education *of the future*.

Viewing AI as a methodology for enhancing educational practitioners' situated introspection and metacognition allows us to bank on the AIEd field's unique focus on generating, representing and utilising knowledge about the relationship between individual learning processes and adaptive pedagogical support. Crucially, it allows us to see AI not solely, albeit importantly, as the driver of back-end functionality of AIEd technologies (e.g. Bundy 1986), but equally as a front-end *technology-of-the-mind* through which educators can represent, experiment with and compare their practices at a fine-grained level of detail and engage in predictive analyses of the potential impact of their actions on individual learners.

Related Work

The proposal to create a techno-intellectual infrastructure for supporting educators' metacognition brings to sharp focus several areas of pertinence, namely, academic and industrial attempts to create communities of educational practice such as ASSISTments (Heffernan and Heffernan 2014), ALEKS (Falmagne et al. 2006), Learning Designer

(Laurillard 2012), or STAR./LBD.Legacy (Schwartz et al. 1998); work related to teachers' metacognition and its importance to continuous professional development (CPD) and pedagogical innovation (Lin et al. 2005; Lin and Schwartz 2003); and AI, especially knowledge representation and engineering (e.g. Davis et al. 1993; Conlon and Pain 1996).

Communities of Practice and Technology-Enhanced Teaching and Learning

The idea of communities of practice as a means for (situated) learning and co-creation of knowledge is not new in itself (see e.g. Lave and Wenger 1991). The advent of web technologies, social networking sites and the increase in the volume, diversity and access to data has led over the past decade to the emergence of a whole spectrum of examples of relevance to the present proposal. Some of these examples are inspired or even driven by AI, and there are also applications, whose explicit goal is to contribute to the development of intelligent systems for learning. For instance, ALEKS (Falmagne et al. 2006) is AI-driven software for assessment and learning of maths and science from primary to university level, which utilises adaptive questioning, assessment and selection of targeted instruction and feedback to help learners. Knewton (Wilson and Nichols 2015) is also an AI-enabled system (employing advanced ontologies and Bayesian reasoning) for assessing students' fine-grained conceptual progress in maths, English language art, and biology. An important feature of Knewton that is of particular relevance to my proposal is that it allows teachers and parents to reuse or to create instructional content to both tailor the support to the needs of individual learners and at the same time to deposit new instructional knowledge representations for Knewton and other users to inspect and recycle. It is notable that both ALEKS and Knewton are commercial systems with a claimed usage by thousands of learners and teachers across the world. The two examples form part of an ever increasing number of technologies – the so-called e-assessment technologies – intended mainly to support classroom teachers' differentiated understanding of individual learners' progress and learning trajectories in specific subject domains (mostly in maths and science). A key characteristic of these technologies is that they spotlight learning analytics and assessment of the student and focus much less so, if at all, on teachers' explicit and systematic introspection of their own practices.

ASSISTments (Heffernan and Heffernan 2014) is another system, which claims a high volume of usage in the US and further afield. As with the previous two examples it can be broadly categorised as an e-assessment system, which has been at least inspired by AI approaches, in particular the model- and example-tracing paradigms developed as part of the Cognitive Tutors endeavour (e.g. Alevan et al. 2009). Similarly to Knewton, ASSISTments is open to teachers and parents, it offers data analytics and a way to deposit, alter and reuse the available content. Apart from having a mixed pedigree of both academic research and real-world teaching expertise, two features make ASSISTments particularly pertinent to the proposed focus on facilitating teachers' metacognition in praxis, namely the fact that (i) it utilises authoring capabilities for teachers to tailor content, feedback and assessment for the specific groups of students and (ii) the data generated through its means is used to confirm the effectiveness of the specific content and its delivery through large RCTs, with outcomes of such RCTs aiming to inform both research and practice. The combination of having the

effectiveness data and the flexibility to act on this data through authoring the environment, supports teachers' reflection on their specific pedagogic designs and improvements attempts, all of which get subsequently recorded ready for analysis if needed. Nevertheless, as with systems such as ALEKS and Knewton, ASSISTments' main focus is on the student's learning rather than explicitly on the learning and professional development of the teacher, the latter being emphasised here as a key factor in delivering innovative and successful learning support. Although teachers' reflections are captured in ASSISTments through the changes they make to the pedagogical designs and to the types of feedback for the given problems, the system does not per se provide tools for supporting teachers' metacognition, for articulating and recording their reflective journeys, say, through examining the same situations from multiple perspectives, or for comparing those journeys with those of other teachers.

By contrast, Learning Designer (*LD* – Laurillard 2012) focuses on providing an open resource to aid teachers in (i) articulating their teaching ideas for other teachers to adopt, (ii) using 'pedagogical patterns' of good teaching presented by others and (iii) assessing pedagogical and logistical gains and demands using data visualisation tools. LD contains several tools to help teachers engage in the design of their pedagogies and to both share and improve them. It comprises a dedicated tool to allow teachers to express their pedagogical approaches and to examine them in terms of how they could be used across different topics. LD also offers a facility for teachers to browse sample patterns that can be either adopted as they are or edited as needed. As such, LD provides teachers with a way in which to import existing designs, use advice and guidance from other practitioners on what designs may work best and in what circumstances, consider alternative designs as well as adapt the designs to suit their own pedagogical contexts. An important feature of LD is that it also offers teachers an ability to analyse computationally the pragmatic consequences of their pedagogical decisions, e.g. the workload needed to implement their design ideas in practice. LD exemplifies a tool, which goes a long way towards supporting teachers' explicit reflection through comparison of learning designs by others and of contexts in which these were applied previously. However, the designs created are delivered at a high level of specification, i.e. at the level of overall lesson plans and activities therein, and do not seem to support a fine-grained reflection of the actual momentary decisions in specific contexts. Although LD has already proven a useful tool for learning design management by teachers and for explicit reflections thereupon, unlike ASSISTments it seems to lack the means to evaluate the effectiveness of the different designs, instead relying on case-by-case and anecdotal reports of what worked for individual teachers.² Nevertheless, the explicit focus of LD on teachers' practices and the ability for teachers to compare and reuse existing designs stored in LD aligns well with the spirit of the present proposal.

LBD.Legacy (Schwartz et al. 1998) provides another example of relevance here. LBD.Legacy is a version of STAR.Legacy (henceforth referred to as *SL*) for pre-service teachers. It is a software shell designed to help *teacher*-students to engage in understanding the benefits of case-based, problem-based and project-based learning and to allow them to innovate their practices through generating ideas, through reflection, and

² Note that this subjectivity may also serve to enhance a sense of community and familiarity among its members, and consequently trust and buy-in, which makes it important not to dismiss this subjective element.

through comparison of multiple perspectives on the same pedagogic challenges. Three features of SL are especially relevant to the proposal to improve evidence-based practices by catering for teachers' situated reflection. First is the importance that SL assigns to different forms of knowledge: learner-centred (knowledge, skills and attitudes that students – in this case teacher-students – bring to the situation); knowledge-centred (knowledge related to the core concepts); assessment-centred (externalising knowledge to students and teachers); and community-centred (building on local knowledge to support collaboration between members of the community). This feature is important here, because it highlights the complexity of the teaching task in which typically, educators have to manage the teaching environment and the support they deliver at several levels simultaneously. While different forms of evidence and knowledge representations may emerge at these different levels, the fact that so many different factors compete for teachers' attention at any given time means that teachers may and often do miss many important cues and links there-between (Lin et al. 2005). Thus, one of SL's crucial aspects is its emphasis on creating a legacy of explicitly recorded contextualised pedagogical decision-making related to the different types of knowledge represented in SL, which can facilitate later situated and coordinated recall and dissemination of key decisions.

The second feature (related to the first) concerns the variety of and targeted use of resources, including videos, simulations, the web, and electronic notebooks for individuals to record their in-the-moment observations. All of these resources are employed in different ways to help elicit and externalise teacher-students' reflections, i.e. they serve as triggers first for teachers realising what they think and how they think and then to allow them to critically evaluate and perfect their approaches. For example, videos of teachers tackling specific classroom challenges are used as anchors for enquiry into the types of challenges that they or others have encountered, for exposing teachers to different, from their own, forms of pedagogies, and as a trigger for their generating new ideas. The third feature of importance is the SL's structured reflection cycle which is used to support the pre-testing of teachers' existing knowledge, generation of ideas for increasingly harder challenges, comparison of different perspectives and self-assessment, all of which support purposeful scaffolding of the development of teachers' metacognitive skills (see also Lin et al. 2005).

The approaches reviewed are only but a few in a growing spectrum of technologies aimed to support educators – mainly school and university teachers – in their daily tasks, highlighting both a growing need for such technologies and their feasibility. Naturally, each approach has their strengths and weaknesses: ALEKS, Knewton and ASSISTments support teachers in gaining quick, on-demand and systematic insight into students' achievements and difficulties and as such they improve educational practice by highlighting challenge areas in the classroom and free teachers' time so that they can dedicate it to tailoring their instruction to the demands of the specific student cohorts. However, these approaches are still relatively limited in their focus, understandably concentrating mainly on well-defined domains such as maths, with notably less content being available in other areas, especially those targeting ill-defined skills, such as design, creative enquiry or social interaction skills. ASSISTments is quite unique in its explicit aim to capture (through RCTs) evidence of what may constitute good practice and in utilising formal thinking and representations derived at least in part from the cognitive tutors' paradigm. Such evidence allows for the

pedagogic support data that is captured, to be related directly to the learning outcomes and subsequently for such data to serve as the basis for further implementation of intelligent (cognitive) tutors. In the context of the present proposal, ASSISTments' chief weakness seems to lie in its limited breadth of subjects covered, and dependence on a particular approach to representing knowledge (which may be a result of ASSISTments' close relationship with the cognitive tutors and their characteristically idiosyncratic pedagogic approach and focus), and crucially, its lack of explicit support for teachers' externalising their metacognitive enquiries. By contrast LD and SL support teachers' externalisations explicitly, offering some enlightened insights into key features of an educators' community of practice, namely the need not only for a technological infrastructure to allow the representation of multiple and diverse pedagogical approaches and perspectives, but crucially a means for a structured comparison of those perspectives against several different, albeit related, types of knowledge (e.g. *as per* SL's proposal) of relevance to delivering adaptive learning support.

Educational practitioners' Metacognition and Adaptive Decision-Making

Although research on learners' metacognition and its importance to learning has a long history (e.g. Brown et al. 1983; Hacker et al. 1998) that has already motivated AIED field to invest in open and scrutable learner modelling (for an overview see e.g. Bull and Kay 2013), research related specifically to educators' self-reflective skills and their impact on their professional development is relatively sparse.³ According to Lin et al. (2005), the nature of educators' metacognitive abilities differs substantially from the nature of metacognitive skills as understood conventionally in relation to monitoring and controlling one's individual thoughts and understanding in relatively stable contexts of particular subject domains. Specifically, educators, especially those practicing in the traditional classroom contexts, but increasingly also in contexts such as MOOCs or online learning environments (e.g. MOODLE), are routinely faced with highly variable situations that change between individual students and classes (Schwartz et al. 2005). As such Lin et al. (*ibid*) refer to many teaching situations as *unstable environments*, which require agile adaptation and consequently *adaptive metacognitive skills* on the part of the teacher.

One of the primary challenges in supporting teachers' metacognition relates to helping them recognise that apparently routine situations often have a number of hidden features. Frequently, teachers' practices are entrenched in their gestalts, i.e. their perceptual abilities and habits through which they make sense of complex teaching situations and which rarely if at all involve conscious reflection or judicious application of principles of good practice (see also Hewitt et al. 2003 for a discussion of 'gestalt'). "Instead, [teachers'] decisions are often based on a split-second product of emotion, needs, values, habit and sense of the affordances of the situation and constraints of the situation." (Hewitt et al. 2003, p.2). Nevertheless, Korthagen and Kessels (1999) suggest and Hewitt et al. (2003) and Lin et al. (2005) show that teachers' habitual

³ I differentiate here between the vast literature related to action research and EBP which proclaims the need for teachers to engage in reflective practice and significantly more limited research which I review here that aims to demonstrate empirically the importance of supporting educators' metacognition in praxis and proposes concrete technological approaches to delivering such support.

practices and interpretations of the teaching situations may be changed when externalised in the form of conscious mental representations and when critically reflected upon. Lin et al. (2005) and Hewitt et al. (2003) provide some compelling evidence of the relationship between teachers' purposeful rooting for and observing the hidden features in the teaching situations they encounter and their adaptive metacognition, and ultimately adaptation skills (see also Lin and Schwartz 2003; Dweck 1999). Importantly, both sets of authors experiment with methods tailored explicitly to facilitate teachers' successful reflection in and on action. In particular, Hewitt et al. focus on use of multimedia tools such as video-based methods (e.g. Bencze et al. 2001; Marx et al. 1998) as anchors for situated externalising by pre-service teachers of their in-the-moment decisions, as a tangible basis for comparative reflection (against different perspectives) and as a basis for teachers re-inventing their pedagogical actions. In their study, Hewitt et al. enhanced videos of real science lessons by simple interface which controlled when the videos were paused to allow pre-service teachers to (i) reflect on what they think should be done next and why, and (ii) compare their decisions for a specific episode in a given lesson with the decisions of other teachers. The individual reflections were used first to allow the teachers to specify what it is that they *actually* think of the situation, as not knowing exactly what their perspective may in fact be has been identified as a surprisingly common initial hurdle for teachers to have to overcome (Abell et al. 1998), and second how they would tackle such situation if they were confronted with it. The comparison with other teachers' reflections exposed the participants to different ways of addressing the same teaching challenges, ultimately leading the majority of them (70 % to 80 %) to either modify their original responses or to re-invent them altogether.

In a study inspired by Hewitt et al.'s work, Lin et al. (2005) used a multimedia learning shell (called CEBLE) combining video vignettes of genuine classroom events with tools for structured reflection. Specifically, the shell, which was developed based on Star.Legacy, aimed to support teachers' reflection through the cycle of: (i) observe an event/a challenge (using a video vignette); (ii) generate responses to specific questions related to whether any new/unusual features in a given event can be observed, and indicate if additional information is needed to allow reflection, explaining why/why not it is needed (the last question aimed to help teachers consider any potential sources of hidden variability in the situations studied and allow them to zoom in on the different aspects of those situations); (iii) listen to multiple different perspectives of other teachers whose values, goals and backgrounds differ and who have also gone through steps (i) – (ii); (iv) act on the selected perspectives by creating possible solutions; (v) reflect on the effectiveness of the solutions generated and share the choices of solutions and reflections with other teachers. The results of a controlled study, involving 30 participants, revealed that the teachers in the study condition (in which they proceeded through all five steps described above) produced substantially more new and more specific solutions than the control group in which teachers did not zoom in on the situations and did not compare their solutions with others. Furthermore significantly more teachers in the study group asked 'Why/How' and 'If/then' questions as opposed to 'What' questions asked by the majority of the control group, suggesting that the comparative critique cycle led the study condition teachers to engage in a more detailed and more specific reflection than the vicarious reflection task of the control group.

There are a number of important conclusions that emerge from both Hewitt et al.'s and Lin et al.'s work. The most significant one relates to the relationship between teachers' reflective and analytical skills as triggered by their being confronted with unexpected situations, also identified in previous research as conducive to people's greater reflectivity (Flavell 1979). Lin et al.'s results suggest that teachers' rooting for more detailed information in the situations studied increases the specificity of their analysis of those situations, as well as it reveals to them hidden aspects of those situations. Hewitt et al. highlight that the timing of the reflections is of crucial importance, with reflections immediately following the events of interest facilitating greater and more situated recall and precision. Moreover, both studies demonstrate the importance of teachers being able to contrast, through discussion with others, different possible solutions and other people's reflections to the same critical events.⁴ Of crucial importance referred to in both studies is the social element of those comparisons, with a striking conclusion that over-reliance on *vicarious* or the so-called *detached observer* approach carries a danger of leading to less analytical and less reflective observations (Lin et al. 2005; Hewitt et al. 2003). Finally, Lin et al. also highlight the key role of digital technology in capturing and accessing both the critical episodes and in scaffolding teachers' perceptions thereof, especially in helping them to home in on the absence of important information. Concrete representation of knowledge, e.g. as elicited through answers to targeted questions such as presented in steps (i) – (v) in Hewitt et al.'s experiment, seems fundamental to facilitating educators' praxis and to their informed improving or completely reinventing their solutions to critical events.

I now turn to examining the affordances of knowledge representation (KR) as conceived in AI research and attempt to highlight parallels and a seemingly natural fit between KR as used in AI and the methods explored in relation to scaffolding teachers' metacognition in praxis.

Knowledge Representation

Different approaches to knowledge representation (KR) have been developed and applied across the AI field, but most evidently in knowledge-based systems, including Intelligent Tutoring and Expert Systems. KR is so fundamental to AI that its functions and affordances are often taken for granted by AI researchers. As Davis et al. observed back in 1993, everyone is using KR, but no one actually says what it is. Thus, in order to examine the relevance of KR to supporting educators' metacognition it is worth recapping its key features. I rely on Davis et al.'s (1993) definition of KR to provide an overview of relevance to the present proposal – the interested reader is referred to the original AI Magazine article for further details.

At the most general KR is seen in AI as a *surrogate* of the world being represented, offering us (or a computer system) a means for reasoning about the world and for determining the consequences in the world without having to take action in it. More precisely, KR acts as a substitute for abstract concepts like actions, processes, beliefs, causality, categories, etc., forcing us to make *ontological commitments* that define our

⁴ In a different, intercultural study involving Chinese and American teachers, Lin and Schwartz (2003) also examined the importance of perspective contrasting and of exposing teachers to new and unexpected situations and solutions.

point of view on the world. Such commitments allow us to depict the same phenomenon in different ways without having to fundamentally change the way we act in the real world. As was discussed earlier, being faced with multiple perspectives of the same situations is a key enabler of teachers' adaptive metacognition and informed (or intelligent) action, making such ontological commitment-making particularly relevant in the present context.

In turn, such intelligent actions, which also constitute a key AI concept, are dependent on our goals, values and beliefs, which determine our choice of theories through the prism of which we can reason about the world. Given multiple possible theories available, our choices may yield very different conclusions and hence, yet again, different views of the world. The specific role of KR is to allow us to spell out precisely not only *what* we think, but also *how* we think about the world, thus supporting the two main pre-requisites of teachers' being able to develop adaptive metacognition (see again Hewitt et al. 2003). Logic offers an example of a theory used in AI where intelligent reasoning is viewed as a form of calculation, such as deduction. By contrast, a theory derived from psychology will view intelligent reasoning as a combination of partially observable human behaviours, plausibly involving structures such as goals, plans, expectations or emotional predispositions, some or all of which may require us to cope with uncertainty. Education too offers a variety of different theories of learning, each engendering inferences that are possible and needed, with pedagogic theories such as constructivism or collaborative learning emphasising different aspects of learning as more or less important, and different types of pedagogic support as more or less necessary.

KR has two further crucial advantages. It provides a *medium for pragmatically efficient computation*, i.e. an environment in which thinking can be accomplished (and conclusions drawn), and it acts as a *medium of human expression* i.e. a language through which we convey and ground our view of the world. Ontological and inferential representations jointly contribute to the definition of an environment in which reasoning can be accomplished, although they do not in themselves guarantee full computational efficiency. Offering an environment in which thinking can be accomplished is of particular relevance to supporting teachers' metacognition. The five-step reflection support by Lin et al. (2005) as described earlier is an example of an environment in which teachers' situated and targeted reflection and re-invention of action could be achieved. I propose that educators' critical and connected reflection can be further supported through explicit framing of the knowledge that needs to be represented in relation to the core components of intelligent tutoring systems (ITS): domain knowledge, knowledge about the learner, pedagogic and communication knowledge. All of these components individually and together offer a tangible and precise way in which to examine, manipulate and observe the mutual impact of the different interpretations of the world within the individual components on each other and on the possible external environments or contexts. Finally, as a medium of human expression, KR allows us to share the different representations with other people and opens a possibility of generating rich critiques of the multiple viewpoints, also providing a trace of our own views of the world over time and a basis for reflection on how our interpretations and knowledge evolved.

The general functions that KR fulfils in AI align well with the research on teachers' metacognition and its importance to understanding and inventing best pedagogical

practice. The research related to teachers' metacognition provides some compelling motivation for why teachers must be able to identify, externalise and codify their knowledge to challenge their gestalts and to be better equipped to deal with dynamic and frequently only partly predictable teaching and learning situations. While it may be tempting to dismiss the presently quite unfashionable symbolic approaches to KR to replace them, say, with machine-learning, symbolic KR offers important advantages that are analogous to those discussed earlier in relation to randomised controlled studies vs. qualitative case-based approaches, not least because they can be inspected at a fine-grained level of detail. Furthermore, while new forms of tapping into educators' expertise emerge (e.g. as in ASSISTments) that may make the task less cumbersome and less time-consuming for the teachers, there are important reasons as reviewed in the previous subsection, for why those methods may not provide adequate service to teachers' praxis and their pedagogic innovation and why educators may need to make a conscious and sometimes even heroic effort to represent the *what*, the *why*, and the *how* of their thinking to understand and to enhance their practices. Thus, my proposal is that KR as conceptualised and applied in AI and as encapsulated in the implementations of ITSs may provide an important mechanism for scaffolding educators' praxis and their (thought) experimentations in a computationally efficient and exploitable way.

Knowledge Elicitation

Knowledge elicitation (KE) is an inseparable companion of knowledge representation in that it is through KE that we engage in reflection about the world. KE is a *process* in which we can engage alone (through self questioning) or with others, either collaboratively or as respondents to someone else's queries, as in the approaches to eliciting teachers self-understanding and critique described earlier. Traditionally, in AIED KE is used as a means of accessing, making sense of and representing teachers' and learners' tacit knowledge and experiences in computational models.

There are many different forms of knowledge elicitation instruments that have been adopted, developed and tested in the context of AIED. For example, questionnaires or interviews, have been borrowed directly from the social sciences, whereas methods such as post-hoc cognitive walkthroughs, gained in power and applicability with the advent of audio and video technologies, and further through the increased focus on and access to logs of man-machine interactions. Other methods, such as Wizard of Oz (WoZ) applications, have been devised as placeholders for yet-to-be-developed fully functional learning environments or components thereof, with the specific purpose of informing the design of technologies at a fine-grained level of detail in a situated way (e.g. see Porayska-Pomsta et al. 2013). Although KE is standardly employed in AIED to inform the design of its technologies, its role as a means of explicitly informing educational practice is less well understood and it may be even regarded as somewhat out of AIED's focus. Yet, it is precisely in examining both how real educational practices may benefit from engaging in structured knowledge representation, that the idea of AI as a methodology comes to life. Two research projects – LeActiveMaths (in short *LeAM*, e.g. Porayska-Pomsta et al. 2008) and TARDIS (e.g. Porayska-Pomsta et al. 2014) – serve to illustrate some of these points.

Example 1: LeActiveMaths Project

LeAM is a system in which learners at different stages in their education can engage with mathematical problems through natural language dialogue. LeAM consists of a learner model, a tutorial component, an exercise repository, a domain reasoner and natural language dialogue capabilities. Its design is based on the premise that the context of a situation along with the learner-teacher interaction are integral to both regulating learners' emotions and to recognising and acting on them in pedagogically viable ways.

To inform the learner and the natural language dialogue models, studies were conducted using WoZ design and a bespoke chat interface. Specifically, the student-teacher communication channel was restricted to a typed interface with no visual or audio inputs to resemble the interface of the final learning environment. Five experienced tutors participated in the studies where they had to tutor individual learners in real time, delivering natural language feedback. They were told that the final goal of the study was to inform the specific components of the LeAM tutoring system, especially the user model and the dialogue model, which provided them with an overall frame within which to examine their own and their students' decisions and actions.

The tutors were asked to talk aloud about their feedback decisions as they engaged in tutoring and to further qualify those decisions by selecting situational factors, e.g. student confidence or difficulty of material, that they considered important in those decisions. The tutors were asked to make their factor selections through a purpose-built tool every time they provided feedback. To aid them in this task some factors were predefined (based on previous research), but these were not mandatory as the tutors could add their own factors to the existing set. The tutors could access and represent the situational factors through drop-down lists, with each containing fuzzy-linguistic values such as *very high*, *high*, *medium*, etc., each value reflecting a relative degree to which they believed a factor expressed the current state of the world. For example, the factor *student confidence* could have five possible values from *very high* to *very low*, with the tutor being able to add further values if necessary. This factor-value selection was used directly to implement the Bayesian network (both its structure and the prior probabilities therein), which was responsible in the LeAM system for mimicking the fine-grained situational diagnoses performed by the human tutors.

Students' screens were captured during each session for the purpose of replay and tutors' post-task walkthroughs, following each completed interaction. In post-task walkthroughs, the recording of the student screen, the tutors' verbal protocol, and the selected situational factors-values for the given interaction were synchronised to facilitate replay. Walkthroughs allowed the tutors and the researchers to view specific interactions again, to discuss them in detail, to explain their in-the-moment choices of factors, and to change their assessment of the situations. Any changes made during walkthroughs were recorded in addition to the original factors' selections.

The data elicited provided rich information about the relationship between tutors' feedback and the specific contexts that they take into account when diagnosing learners' cognitive and affective states. It also provided a concrete basis for the implementation of the user and dialogue models in the system and the corresponding knowledge representations. However, the studies also provided important insights into the potential impact that the KE/KR process had on the participating tutors.

Specifically, and in line with the teachers' metacognition research reviewed earlier, the demand on teachers to report on the situational factors of importance to their feedback decisions brought to their attention that such factors may indeed play a role and forced them to think explicitly about them while making those decisions. Verbal protocols facilitated verbalisation of those decisions *while* they were made and later provided an important tool for facilitating recall. Although initially, all tutors had a clear understanding of and an ability to identify the factors such as the difficulty of the material or correctness of student answer, they were much less fluent in diagnosing and explaining student's affective states. However, after an initial familiarisation period, involving up to two sessions, their willingness to engage in situational analysis and the fluency of their reports increased, while the tentativeness in identifying student behaviours at a fine level of detail seems to have decreased. This was evidenced in the increased speed at which they offered feedback to students, the level of elaboration in and the targeted quality of their verbal protocols and post-hoc interviews. For example, during the initial interactions, all of the tutors had their attention fixed on the correctness of students' answers, with the selection of the next problems or sub-problems to give as feedback to the students having been their chief concern. Initially, the tutors found it difficult (and at least 3 out of 5 of them even unnecessary) to pay attention to the language used by the individual students and to having to pay such persistently detailed attention to the different factors which were seemingly unrelated to the task of supporting learners in solving differential equations. Yet, on average, by the third interaction, most tutors, save for one,⁵ began to verbalise explicitly their observations with respect to multiple situational dimensions such as content matter, students' possible cognitive states (e.g. confusion), emotional predispositions and states (e.g. confidence). The tutors were also able to identify reasons in the actual dialogue interactions for particular students' diagnoses, e.g. they reflected on some students using question marks at the end of statements as a potential sign of lack of confidence. Importantly, although the tutors were not burdened with having to understand the intricacies of the formal implementation intended within the LeAM system, i.e. the Bayesian knowledge representation and reasoning that was eventually employed to capture the dynamics of the situational diagnoses, having to represent their selections of situational factors in terms of the degree to which they believed those factor-values to be manifest in learners' behaviours highlighted to them the nuances in individual situations' and learners' idiosyncratic needs in the face of seemingly the same learning challenges and/or student misconceptions. This was reflected in the tutors' feedback to the learners, which over time became more positive, more elaborate, and more targeted to the actual factor diagnoses made than was the case initially. This was especially visible with respect to partially correct situations, in which some tutors (3 out of 5) made increasingly consistent effort first to provide praise for the correct part of the answer, e.g. *"This is so nearly right. Maybe you can spot the mistake before I send the right answer..."* or encouragement, e.g. *"You are doing well up to now..."* (for details of the analysis see Porayska-Pomsta et al. 2008).

The use of verbal protocols during the interactions, followed by semi-structured interviews, and then post-task walkthroughs provided tutors with an opportunity to first

⁵ This tutor found it difficult to accept that her reporting on what, how and why she was doing in this context was not intended to evaluate her as a tutor, but to learn more about the tutoring process more generally.

formulate and record their diagnoses of the student in context, then reflect on and finally re-examine them. The post-hoc walkthroughs allowed the tutors to assess the consistency of their in situ interpretations and further, to analyse those situations where they did not agree with their earlier interpretations. According to some tutors, this led them to deep reflection and grounding of their understanding of (a) what matters to them the most in tutoring situations and (b) the kinds of tutoring they would like to deliver *ideally*. At the end of the study some tutors expressed a need for a tutoring system for tutors, through which they could rehearse, experiment with and perfect their understanding of the different nuances of educational interactions, showing a real appreciation of the value of having to explicitly go through the effort of externalising, explaining and critiquing their practices.

Example 2: The TARDIS Project

The TARDIS project used many insights gained during LeAM by employing KE methods throughout, even though its focus was on an ill-defined domain of social interaction. TARDIS is a simulation game for coaching young adults in job interview skills through interactions with intelligent conversational agents. The game is underpinned with models and associated knowledge representations of emotions and emotional expression rendering for the agents acting as recruiters (e.g. Youssef et al. 2015), and with real-time detection and interpretation models of observable verbal (e.g. voice features) and non-verbal (e.g. gestures and body pose) behaviours (Porayska-Pomsta et al. 2014; Baur et al. 2013). A key feature of the TARDIS project was its additional aim to inform the *design of use* of the game in real contexts of youth employment associations across Europe, where it was crucial to facilitate the practitioners' independence in using this game as part of their typical practices. Thus, in TARDIS, KE and KR provided the basis for developing practitioners' self-observation and self-reporting skills. These skills were then built on in the formative evaluation studies, in which the practitioners increasingly participated as researchers, with the support by researchers being gradually removed. The whole process was divided into three stages, roughly corresponding to the three years of the project.

The first stage (*familiarisation*) involved gradual preparation and training of practitioners, through KE, in systematic and fine-grained representation of knowledge related to different aspects of job interview coaching interactions. Specifically the KR concerned the domain of interaction, the user and the virtual recruiter models. Although no formal KR was employed to achieve this, the team had a good idea as to the likely final implementations of the different models, e.g. Bayesian network formalism for the representation of the user model, or planning for the generation of virtual recruiters' behaviours. This meant that the guidance offered to the practitioners as to the type of knowledge required was biased towards generating representations that would allow the team to kick-start such implementations, even if practitioners themselves did not have to construct the actual networks or write planning operators. The representations generated at this stage included detailed specification of job interviews' structures, involving definitions of distinct interview parts and the development of turn-by-turn interview scenarios (this specification was used directly in the implementation of the

TARDIS' scenarios using the SceneMaker⁶ modelling tool – see e.g. Damian et al. 2013). These scenarios were subsequently tagged by practitioners with appropriate types of interview questions along with the expected verbal and non-verbal candidates' response behaviours, e.g. the level of elaboration needed for each type of question, topics covered and non-verbal behaviours such as eye-contact, quality of voice etc. The ideal/intended recruiter behaviours were also specified and were associated with the individual interview parts and dialogue turns therein. Thus, at this first KE/KR stage, post-hoc walkthroughs, using video replays of mock job interview sessions between learners and practitioners were used to (a) access the practitioners' expert knowledge needed to be represented in TARDIS; (b) allow the practitioners to make overt to themselves, and to the researchers, the types of knowledge and interpretation processes of particular interest in the context of job interview skills coaching and (c) allow them to reflect on their and the learners' needs, leading to the specification of the necessary and sufficient elements of the TARDIS environment. Practitioners' reflections were recorded as their own annotations of videos.⁷ The annotations focused on the identification of the key episodes in each interview and on determining a detailed list, ordered according to importance to the practitioners' decision making, of learners' observable behaviours (e.g. looking away, amplitude and duration of their laughter, fidgeting, etc.), and the interpretation thereof in terms of complex mental states such as anxiety, hesitation etc. The result was a definition of the interaction domain, which provided the basis for the specification of the knowledge in terms of the game interaction scenarios and the social cue detection as well as complex mental states model (the latter having been eventually implemented as Bayesian networks, as planned originally).

The second stage (*testing, critique and design of use*) involved a period of continuous cycles of reflection, observation, design and action, scaffolded by researchers and guided by the Persistent Collaboration Methodology (Conlon and Pain 1996). This stage was crucial not only to the TARDIS researchers who were able to implement ever more sophisticated prototypes and refine different aspects of the knowledge encapsulated in the system's models, but it was also fundamental to the practitioners' growing confidence in providing targeted critique of those prototypes, to their increased emancipation in using TARDIS, and to their ability to experiment with its different set-ups. Crucially, the knowledge self-elicitation and representation skills, developed in the first year, along with their rehearsed focus on the type and form of information needed by the researchers to create the various computational models, provided the practitioners with a structure against which to report their observations and reflections to the researchers and a common language for both that was akin to the type of structured questioning applied by Lin et al. (2005). A key outcome of this was a growing sense of co-ownership of the tools and knowledge developed, which was reflected in the independent curation of TARDIS by the participating practitioners to their colleagues, with whom they began to co-design the use of TARDIS in their everyday practices and to consider viable redesigns of their existing pedagogical methods. This has led to a specification of knowledge in terms of procedures for how, when and to whom

⁶ SceneMaker is freely available for research purposes at: <http://scenemaker.dfki.de/>

⁷ An off-the-shelf tool called *Elan* was used: <https://tla.mpi.nl/tools/tla-tools/elan/>

TARDIS could be administered and how its use may fit best the existing support given to learners.

In TARDIS' third and final stage the practitioners engaged in summative evaluation of the system with minimal support from the researchers (*independent use and research*). As well as being able to use the system independently and to explore new ways in which to utilise it within their existing practices, a key consequence was the practitioners' confidently vocal involvement in the development and testing of a schema for annotating data of learners' job interview skills. This schema encapsulates the knowledge about the relationship between (different combinations and manifestations of) learners' observable behaviours and practitioners' interpretations thereof in terms of performance along some key dimensions related to emotional traits and states (e.g. self-confidence) and quality of learners' responses (e.g. degree of answers' elaboration vis á vis the minimum expected elaboration, given a particular type of question and interview phase). This schema, which to my best knowledge is thus far unique in its aim to aid the modelling of job interview skills at the level of detail needed to implement computational models of user and of artificial agents, was used directly in the analysis of the TARDIS' evaluation data (Chryssafidou and Porayska-Pomsta, in preparation). As such this schema provides a basis for further refinement and implementation of the TARDIS user modelling tools as well as for developing standardised guidelines for youth association practitioners working with unemployed youth in different institutional and national cultures. Looking through the prism of formal KR, at the very least, the schema constitutes an initial ontological representation of the skills, along with the pre-requisite skills and related valid possible evidence of their existence, for the domain of job interview training that could find its use in many a system concerned with training social interaction skills more generally.

Throughout TARDIS, the practitioners' roles and competencies have evidently changed from those of willing informants (the beginning of the project), through advisors and co-designers of the TARDIS system (middle of the project), to lead-practitioners who initiate projects independently in a bid to make best use of the tools given to them (end of the project). At the core of this change was a gradual shift in the practitioners' way of thinking and viewing the world of their practice, aided by their engagement in knowledge elicitation and its eventual representation in terms of design recommendations and fine-grained specification of the domain and related inferences (annotation schema). The practitioners demonstrated that their approach to technology use in their practices changed from that of mere consumers to its co-creators and owners. I opine that this may be interpreted at least as a seed of an ability to think about their domain and practices in terms that are by nature both computational (low level knowledge representation) and design (design of the technology's look-and-feel, functionality and use in situ, as well as pedagogical design⁸).

The observations reported in relation to the LeAM and the TARDIS projects suggest the potential important role of KR and KE methods to enhancing teachers' critical reflection and creativity with respect to their own practices. The examination of the effectiveness of these methods was not the focus in either of these projects and hence no objective measures can be offered at this stage. Nevertheless, at the very least, they serve as compelling examples of AI as a methodology for enabling analytical

⁸ Note that some researchers in Education view teaching as a design science, e.g. Laurillard (2012).

expression and practical consequences of such expression to educators' front-line decision-making, as well as for facilitating the implementation of AI models. The utility of these methods in supporting teachers' metacognition and reflective practices has emerged first in LeAM in the form of tutors' increased verbalisation and change in the quality of their feedback (e.g. changing from purely corrective/confirmatory to pastoral). In TARDIS this was evident, among other KRs generated, in the annotation and interpretation schema developed for analysing and assessing learners' observable behaviours as well as in practitioners' readiness to re-invent their existing practices and the procedures through which they currently support learners at their centres. Together with the evidence from research reviewed in relation to teachers' metacognition, this points the way to a potentially important area for AIEd, namely one which focuses much more explicitly than is currently the case on the development of tools for supporting educators' metacognition in praxis, which bridges between different types of evidence (RCT, case-based), and which connects teachers' reflections with respect to different considerations of relevance to their adapting their pedagogy.

Discussion and Conclusions

In this paper I reviewed research related to educational practitioners' metacognition and reflective practice, from which some guidelines as to the necessary preconditions for enabling educational praxis and innovation emerge. It is important to note that the validity of those guidelines along with any additional nuance that may serve to refine them remains an open research question. As such I see those guidelines as contributing to the definition of the future research agenda for the AIEd in the next 25 years. The initial indications are as follows.

1. Defining a formal framework within which educators can engage in reflection and where their thinking can be accomplished is of paramount importance to supporting their metacognition and adaptive pedagogy. Such a framework should scaffold the practitioners in identifying first what it is that they think about the specific challenge situations before helping them to determine how they think. Examples from studies by Hewitt et al. (2003), Lin et al. (2005) as well as LeAM (Porayska-Pomsta et al. 2008), illustrate some ways in which this could be achieved. Although in principle KR, as understood in AI, aligns well with the desired nature of such a framework, further targeted research is necessary to explore what forms of existing as well as new KR formalisms can support educators' metacognition. Moreover, research is needed to investigate optimal levels, if such exist, of formal specification that may help engender educators' greatest creativity, relative to different contexts that such creativity may benefit and in which it may be evaluated, e.g. institutional, subject domains, specific student cohorts or individual learners and teachers. Both LeAM and TARDIS examples suggest that it may not be necessary for practitioners to learn the KR formalisms at their lowest level of specification. Instead, the key seems to be in our facilitating the different forms of reasoning that is normally supported by the KR formalisms. Thus, to engender computational thinking and the accomplishment of explicitly represented reflections, it may be sufficient to allow practitioners to work with

higher level specifications such as the expression of diagnoses using probabilistic language, or (partially) ordering actions given a specified set of well-defined constraints or pre-conditions.

2. Identification of missing information and whether/why it is needed to clarify challenge situations may help educators in diagnosing them in tangible terms. Such identification can be accomplished through sharing perspectives with other practitioners and, equally, it can serve as a basis for sharing different points of view. Either way, it can help practitioners make informed decisions with respect to diverse possible actions that address the particular challenges, while also allowing them to contemplate the possible consequences of those actions, either through thought experiments or through technology-enhanced simulations, if such are available. Developing the latter to better support the former offers an interesting research opportunity for AIEd community to explore and develop, and one where the community could demonstrate and further enhance its unique interdisciplinary strengths.
3. Comparison of perspectives and the social aspect of such comparisons seems key to helping practitioners to shift their gestalts, to observe and then in turn be able to generate novel solutions to the specific challenges, leading them to more in-depth questioning, while also increasing their willingness to innovate, experiment and reinvent what they do. Having other practitioners to position themselves against one's perspective may also help evaluate the potential effects of acting in particular ways. In this case, it would seem natural for the AIEd field to build on its intimate relationship with AI to invest in the development of autonomously or semi-autonomously harvested, sorted and generated case studies, to enable such comparisons on demand, anytime and anywhere. Such AI-driven tools would not only allow their users to retrieve different exemplars of practice in contexts of similar pedagogic challenges, but as already signalled under point 2, they would allow for predictive simulations of the consequences of the educators' hypothetical/desired solutions for the different cases to be performed and evaluated along multiple dimensions, e.g. cognitive, socio-affective or interaction. Much relevant work already exists both in data driven information retrieval and case-based reasoning as well as simulation environments, which could be leveraged and re-purposed to connect and scaffold interactions and collaboration between practitioners from seemingly disparate subject domains, institutional as well as national cultures or even educational specialisations.
4. Detailed analysis and representation of individual teachers' perspectives seems key to their developing an informed understanding of what it is that they believe and know and, crucially to their being able to link such understanding to concrete possible solutions with which to experiment. AI knowledge representations, including the different types of ontologies and reasoning frameworks related to the core components of intelligent tutoring systems, provide at least conceptual tools that can serve to externalise and systematise educators' knowledge in a guided way, as I tried to demonstrate through LeActiveMaths and TARDIS projects. The need for high specificity and inspectability of the representations generated motivates the use of symbolic KR, as it ensures that the inferential procedures of relevance are not obscured from view, i.e. that they are explicitly represented. However,

this does not exclude the use, nor does it negate the utility of sub-symbolic or stochastic computations to educational practices. Instead, the key is in our developing a good understanding of when and exactly how the respective methods might serve educational practitioners best (again relative to different contexts and educational specialisations), as well as in ensuring that this understanding is shared, and ideally – co-developed, with practitioners.

5. Use of multimedia (e.g. as in Hewitt et al., or TARDIS) and multimodal access to the learning situations (e.g. as in LeAM's synchronised WOZ interface) are important for anchoring practitioners' reflections, but they have to be coupled with active and targeted questioning (as per point 1 above), with comparison with the perspectives of others (point 3), and with explicit knowledge representation effort (point 4), in order to be effective. AIEd has a long history of connecting with other cognate areas such as HCI, where methods such as participatory design build heavily on the need for anchors to enable the process of co-design (see e.g. Conlon and Pain 1996; Simonsen and Roberston (2013)). HCI researchers increasingly recognise that engaging in the process of (co-)design can be as valuable, as the design outcomes (e.g. Friedland and Yamauchi 2011). TARDIS provides a good example of how using co-design methods with practitioners may lead to the emancipation of the participants and to their increased sense of ownership of the technology and of the knowledge created. To my mind, the strength of the approach adopted in TARDIS lay in the combined application of participatory and knowledge representation methods: the practitioners' implicit knowledge was represented explicitly in the form that lent itself to different forms of reasoning demanded by the TARDIS models, but the practitioners' ability to engage in such KR was an outcome of their being engaged in TARDIS' design process from the start. The quality of the knowledge generated was not accidental in that it resulted from researchers' targeted questioning and training of the practitioners that primed them to think and express themselves in computationally-ready manner, e.g. in terms of evidence observed and degree of certainty as to the possible diagnoses. In this case HCI inspired a particular application of AI methodology, whereas AI strengthened and made concrete the outcomes of the design process, which in effect became a form of knowledge co-engineering. AIEd's future investment in building new knowledge engineering methods that leverage best practices in diverse disciplines such as HCI and AI seems essential not only to AIEd producing tools that are pedagogically effective and relevant to real world users, but also to opening new perspectives on the utility and application of AIEd's way of questioning and representing the world (which as a community we take for granted) beyond its own disciplinary boundaries. I believe, that investing in participatory knowledge co-engineering, aimed to mutually challenge and cross-fertilise existing thinking paradigms within educational, engineering and social sciences practices and which fundamentally builds on the representational rigour demanded by AI, is key to creating formalisms that are natural and accessible to a wide range of practitioners and that are relevant and useful long-term.

The experience of the LeAM and the TARDIS projects suggests that with appropriate set ups and willing practitioners the traditional AI methods can be learned and

applied by those practitioners independently of researchers. Studies by Hewitt et al. and Lin et al. seem to corroborate these conclusions showing that teachers' established and long-unquestioned points of view can be changed through careful scaffolding, that pedagogic creativity can be fostered by exposing educators to different perspectives and that the two can serve to elicit indepth, explicit descriptions of what the practitioners consider best practice in specific contexts.

However, an important trade-off in undertaking projects such as reviewed is that their outcomes are confined to those projects, with the methods used also being time-consuming to learn and apply, potentially rendering them impractical for everyday use by in-service educators. Although, educators' need to share and compare their experiences with other practitioners motivates an investment in technologically enabled communities of practice, it also points to a need to distinguish between educators in their different roles, professional stages and specific specialisations (e.g. as experts, as users and informants of the systems we create, and as continuous explorers and learners of their trade) in order to enhance the real-world utility and survival of such communities. Consequently, this means that the technologies that we develop need to be able to cater for educators in those different roles, with both the exact nature of the technology and the time and place of its use being critically dependent on the distinctions we make. To date, with the notable exception of LBD.Legacy, the predominant trend is to focus on easing teachers' job, e.g. by automating its more tedious aspects such as assessment (e.g. ALEKS or ASSISTments). Such approaches and associated technologies are naturally much loved by many educators, but they assume teachers to be experts in their trade. As such they go 'with the grain' of what teachers already know and are good at. Unfortunately, in the context of educators' developing metacognitive skills going with the grain would seem to undermine the intended outcome of practitioners' critically reflecting on their practice and their being able to link those reflections directly to diverse possible solutions. Research reviewed points to the need for educators to expend considerable amount of effort to externalise their thinking and to reflect on their practices through sharing them with others in a way that is actively structured, social and not merely vicarious. AI's KR and KE methods provide some concrete instruments, which seem philosophically and practically aligned with evidence-based praxis and educators' metacognitive abilities such as discussed throughout this paper. However, to be successful in their application and uptake they may need to be used at strategic times in educators' professional development, such as at pre-service or CPD training from where the knowledge and adaptive metacognitive skills, along with the uptake of the AIEd approaches and tools, can percolate to real-world and real-time education.

Given the maturity of the AI methodologies and its community's understanding of their merits and limitations, the AIEd as a field is ready to share them and to reinvent them as tools for supporting educational practitioners' understanding and innovative practices. The field is ready to take an important step into the future where AIEd's research can co-evolve with and respond to the real-world needs of educational practice.

Acknowledgments My colleagues Manolis Mavrikis, Karen Guldborg, Sarah Parsons, Helen Pain, Mina Vasalou and Sveta Mayer have all contributed over the years in different ways and at different times to the development of the ideas presented in this paper, as have all of my students taking the Learning and Teaching

with Technologies module at the UCL Institute of Education. LeAM and TARDIS were both funded by the European Commission (FP6-IST-2003-507826 and FP7-ICT2011-7-288578 respectively).

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Abell, S., Bryan, L., & Anderson, M. (1998). Investigating preservice elementary science teacher reflective thinking using integrated media case-based instruction in elementary science teacher preparation. *Science Education*, 82, 491–509.
- Aleven, V., McLaren, B., Sewall, J. & Koedinger R K. (2009). A new paradigm for intelligent tutoring systems: example-tracing tutors. *International Journal of Artificial Intelligence in Education*, 19(2009), 105–154. IOS Press.
- Baur, T., Damian, I., Gebhard, P., Porayska-Pomsta, K., & André, E. (2013). *A job interview simulation: social cue-based interaction with a virtual character*, Proceedings of the IEEE/ASE International Conference on Social Computing (SocialCom 2013), (pp. 220–227). Washington.
- Bencze, J. L., Hewitt, J., & Pedretti, E. (2001). Multi-media cases in pre-service science education: an apprenticeship for praxis. *Research in Science Education*, 31, 191–209.
- Biesta, G. (2007). Why 'what works' won't work: evidence-based practice and the democratic deficit in educational research. *Educational Theory*, 57(1), 1–22.
- Biesta, G. (2013). *The beautiful risk of education*. Boulder and London: Paradigm Publishers.
- Brown, A. L., Bransford, J. D., Ferrera, R. A., & Campione, J. C. (1983). Learning, remembering, and understanding. In J. H. Flavell & E. M. Markman (Eds.), *Handbook of child psychology: Vol. 3, cognitive development* (4th ed., pp. 77–166). New York: Wiley.
- Bull S., & Kay, J. (2013). Open learner models as drivers for metacognitive processes. In R. Azevedo, V. Aleven (Eds.), *International Handbook on Metacognition and Learning Technologies* (2013), (pp. 349–366). Springer.
- Bundy, A., (1986). What kind of field is artificial intelligence? In: *DAI Research Paper No.305*. Department of Artificial Intelligence, University of Edinburgh.
- Carr, W., & Kemmis, S. (2004). *Becoming critical. education, knowledge and action research*. Deakin University Press.
- Cohen, L., & Manion, L. (1980). *Research methods in education*. London: Croom Helm.
- Conlon, T., & Pain, H. (1996). Persistent collaboration: methodology for applied AIES. *International Journal of Artificial Intelligence in Education*, 7, 219–252.
- Damian, I., Baur, T., Gebhard, P., Porayska-Pomsta, K., & André, E. *A software framework for social cue-based interaction with a virtual recruiter*. In Proceedings of the 13th International Conference on Intelligent Virtual Agents (IVA 2013), vol. 8108, Edinburgh.
- Davis, R., Shrobe, H., & Szolovits, P. (1993). What is knowledge representation? *AI Magazine*, 14(1), 17–33.
- Dweck, C. S. (1999). *Self-theories: their role in motivation, personality, and development*. Ann Arbor: Edwards Brothers.
- Falmagne, J.C., Cosyn, E. Doignon, J.P., & Thiery, N. (2006). *The assessment of knowledge, in theory and in practice*. In: R. Missaoui, J. Schmid (eds.): ICFCA, Vol. 3874 of Lecture Notes in Computer Science, (pp. 61–79).
- Flavell, J. H. (1979). Metacognition and cognitive monitoring: a new area of cognitive-developmental inquiry. *American Psychologist*, 34, 906–911.
- Friedland, & Yamauchi. (2011). Reflexive design thinking: putting more human in human-centred design practices. *Interactions*, 18(2), 66–71.
- Hacker, D. J., Dunlosky, J., & Graesser, A. C. (1998). *Metacognition in educational theory and practice*. Hillsdale: Lawrence Erlbaum Associates, Inc.
- Hargreaves, A. (1999a). *The realities of teachers work: never a dull moment*. London: Cassell.
- Hargreaves, A. (1999b). Revitalising educational research: lessons from the past and proposals for the future. *Cambridge Journal of Education*, 29(2), 239–249.

- Heffernan, N., & Heffernan, C. (2014). The ASSISTments ecosystem: building a platform that brings Scientists and teachers together for minimally invasive research on human learning and teaching. *International Journal of Artificial Intelligence in Education*, 24(4), 470–497.
- Hewitt, J., Pedretti, E., Bencze, L., Vaillaincourt, B. D., & Yoon, S. (2003). New applications for multimedia cases: promoting reflective practice in preservice teacher education. *Journal of Technology and Teacher Education*, 11, 483–500.
- Korthagen, F., & Kessels, J. (1999). Linking theory and practice: changing the pedagogy of teacher education. *Educational Researcher*, 28(4), 4–17.
- Laurillard, D. (2012). *Teaching as a design science: building pedagogical patterns for learning and technology*. Routledge.
- Lave, J., & Wenger, E. (1991). *Situated learning: legitimate peripheral participation*. Cambridge: Cambridge University Press.
- Lin, X., & Schwartz, D. L. (2003). Reflection at the crossroads of cultures. *Mind, Culture & Activities*, 10(1), 9–25.
- Lin, X., Schwartz, D. L., & Hatano, G. (2005). Toward teachers' adaptive metacognition. *Educational Psychologist*, 40(4), 245–255.
- Marx, R., Blumenfeld, P., Krajcik, J., & Soloway, E. (1998). New technologies for teacher professional development. *Teaching and Teacher Education*, 14(1), 33–52.
- Parsons, S., Charman, T., Faulkner, R., Ragan, J., Wallace, S., & Wittmeyer, K. (2013). Bridging the research and practice gap in autism: the importance of creating research partnerships with schools. *Autism*, 17(3), 268–280.
- Porayska-Pomsta, K., Mavrikis, M., & Pain, H. (2008). Diagnosing and acting on student affect: the tutors perspective. *User Modeling and User-Adapted Interaction*, 18(1), 125–173.
- Porayska-Pomsta, K., Mavrikis, M., D'Mello, S., Conati, C., & Baker, R. (2013). Knowledge elicitation methods for affect modelling in education. *International Journal of Artificial Intelligence in Education*, 22, 107–140.
- Porayska-Pomsta, K., Rizzo, P., Damian, I., Baur, T., Andre, E., Sabouret, N., Jones, H., Anderson, K., & Chryssafidou, E. (2014). *Who's afraid of job interviews? Definitely a question for user modelling*. In: Proc. Conference on User Modeling, Adaptation and Personalization, (pp. 411–422).
- Schön, D. (1987). *Educating the reflective practitioner*. Washington: Meeting of the American Educational Research Association.
- Schwartz, D. L., Brophy, S., Lin, X., & Bransford, J. D. (1998). Software for managing complex learning: examples from an educational psychology course. *Educational Technology Research and Development*, 47, 39–60.
- Schwartz, D. L., Bransford, J. D., & Sears, D. L. (2005). Efficiency and innovation in transfer. In J. Mestre (Ed.), *Transfer of learning from a modern multidisciplinary perspective* (pp. 1–51). CT: Information Age.
- Simonsen, & Roberston (eds.) (2013). *Routledge international handbook of participatory design*, Routledge.
- Slavin, R. (2004). Evidence-based educational policies: transforming educational practice and research. *Educational Researcher*, 31(7), 15–21.
- Wilson, K., & Nichols, Z. (2015). *Technical white paper*, <https://www.knewton.com/wp-content/uploads/knewton-technical-white-paper-201501.pdf>.
- Youssef, A.B., Chollet, M., Jones, H., Sabouret, N., Pelachaud, C., & Ochs, M., (2015), *Towards a socially adaptive virtual agent*. In Proc. 15th International Conference on Intelligent Virtual Agents (IVA) 2015, (pp. 3–16).