

PART 4

Chapter 31

Learner Modelled Environments

Kaśka Porayska-Pomsta and Sara Bernardini

Introduction

Learner modelled environments (LMEs) are digital environments that are capable of automatically detecting learner's behaviours in relation to a specific knowledge domain, to reason about those behaviours in order to assess learner's performance, skills, socio-emotional and cognitive needs, and to act accordingly in a pedagogically appropriate manner. Digital environments that possess such capabilities are typically referred to as *Intelligent Learning Environments*, or more traditionally – as *Intelligent Tutoring Systems* (ITSs).

LMEs research is motivated by the need to understand how meaningful interactions between a teacher and a learner develop, what contributes to successful learning and how the pedagogy can and ought to be adapted to the individual learners. Adapting the instruction to the needs of the individuals constitutes an important way in which to support them in their transformation from novices to experts in a specific subject domain.

Theories from diverse disciplines serve as the basis for tackling the above questions. Increasingly, the existing theories can be combined, tested and improved through computational means – more specifically through computational models of learning

and communication processes and through dynamically generated and updated models of individual learners themselves. The advantages that such computational models bring to the learning sciences are their systematic, objective and inspectable nature, and the fact that, in contrast with models created by and held in human teachers' heads, they lend themselves to being manipulated, changed and repeatedly tested. Importantly, they can be used to make predictions about the relationships between the specific elements included in the models, the pedagogy selected on the basis of such models and the learning outcomes that may ensue as a consequence of learners being exposed to the specific pedagogy. Such computational models can therefore serve to both support learning of a specific subject domain and to progress research about learning and teaching. In other words, LMEs research is of relevance to anyone interested in understanding how learners learn and especially of how successful learning and teaching can be supported through technology.

Modelling the learner

A core element of a LME is its *learner model*, traditionally also known as a *student model*. It is a core element, because it is responsible for interpreting the learner's observable behaviours against the background knowledge available, and through this, for providing a system with information necessary for making pedagogical decisions (e.g. whether or not the student is ready to progress to the next problem), as well as for choosing the appropriate feedback (e.g. whether at any specific point the learner would benefit from, say, a hint, positive feedback, or a prompt).

Learner modelling (LM) is concerned with *real-time* diagnosis of the learner's knowledge and with diagnosis of their cognitive and affective states during learning. LM is important within computer-based systems intended to promote learning, because it guides the *just-in-time* adaptation of the interaction to the individual learner (Dillenbourg and Self, 1992). Many educators consider individually adapted instruction as ideal means for facilitating effective learning, with the *2-sigma effect* (Bloom, 1986) being often cited as the main evidence of the increased learning outcomes that ensue from such instruction, and as one factor that motivates many an investment into technology-enhanced solutions for education (Koedinger and Corbett, 2006). Specifically, in comparing the effect of one-on-one tutoring with the effect of a number of other variables and treatments, Bloom found that the average student who had been tutored in one-to-one context performed better than 98% of the students who had been taught in the classroom (and that none of the other treatments produced an effect as large as 2 sigma).

The ability to model others is inherent in all human social behaviour, of which educational interactions are one instance. In general, humans engage implicitly in modelling others on a moment-by-moment basis (e.g. Sperber and Wilson, 1995). Observing and reasoning about others' intentions, beliefs and goals is also crucial to facilitating relevant and effective learning, i.e. learning that results in deep understanding of concepts and skills, that is long-term and transferrable to new situations. Teachers continually assess learners' needs, abilities and progress, and adapt the difficulty of the problem accordingly. For example, they modify the explanation, give hints, alter the form and strength of praises, and remediate

feedback according to how they interpret the learner's actions in context. Teachers are *trained* to pursue the explicit goal of helping to advance the individual learner's knowledge and skills.

For example, consider a simple problem of double-column subtraction, where the number 65 needs to be subtracted from 500. Let us suppose that the learner's answer is 565, which is incorrect: the learner's mistake is to add 65 to 500 instead of subtracting from it. In order to provide the learner with appropriate support, apart from spotting the error, the teacher needs to diagnose the root of the error: is it due to a simple slip and inattention of the learner or is it due to some deeper underlying misconception? In order to be accurate, this diagnosis needs to be based on the teachers' knowledge of the learner, the immediate circumstances in which the error occurred, the history of the previous problems given to the learner, the way the learner solved them (correctly or incorrectly), the instruction that was available to the learner prior and during to problem solving and so on. Such information enables the teacher to decide *what* to say to the learner. On the other hand, the information that the teacher infers about the learners' emotional and motivational states contributes to their decisions about *how* to communicate the support to the learner – a learner lacking confidence may need a hint and encouragement, whereas a confident, but bored learner may benefit from feedback that identifies the mistake and instructs.

Thus, the main purpose of learner modelling is to enable a LME to choose the most appropriate pedagogical actions for the moment, i.e. actions that have the best chances to lead to the desired short- and long-term learning outcomes for specific individuals.

LM also constitutes a paradigm within which to study learning and provides tools for learning about learning. In other words, the *computational* learner modelling tools embedded in LMEs provide means for representing the covert mental processes (both cognitive and affective) that relate to learning, for specifying, testing and improving theories about cognition and for making predictions about learning outcomes in the domains of interest.

Learner modelling as knowledge representation

Learner modelling is a computational approach to representing knowledge about the learner during a learning task. The result (the output) is a *learner model* – a qualitative, approximate and possibly partial representation of knowledge about the learner, and about the learner's knowledge and learning. The construction of a learner model depends on the information about the learners' specific behaviours (the input) in relation to the domain knowledge being learned.

Knowledge representation is key to enabling the construction of learner models. It is concerned with two primary questions: (i) *what* knowledge to represent and (ii) *how* to represent it. These two seemingly simple questions lead to further, fundamental questions in education and cognitive science research, including: what is knowledge?, what constitutes and influences learning?, and how do we know that someone has reached the level of mastery in a domain?

To support learner modelling in a LME, at least four types of knowledge need to be represented: (1) domain knowledge, (2) knowledge about the learner (3) pedagogic knowledge, and (4) communication knowledge. These types of knowledge are

interdependent in that domain knowledge representation is essential to enabling the modelling of the learner's progress and it determines the type of pedagogic knowledge that is required to facilitate the learner's progress. Representation of the knowledge about the learner is essential for enabling decisions about the appropriate pedagogy and the communication strategies that will best realise the pedagogy. Whilst all four types of knowledge are crucial to a full-fledged LME, with domain knowledge representation of particular importance to enabling the assessment of learner's progress and to choosing optimal support (Woolf, 2009; Sison and Masamichi, 1998), this chapter focuses specifically on learner modelling.

Knowledge about the learner

Knowledge about the learner is a representation of what the learner knows including specific skills and misconceptions at any given time during and after a tutoring interaction. It is crucial to enabling a teacher (whether human or computer tutor) to provide appropriate feedback to the learner, i.e. feedback that will promote long-term learning. We refer to this type of information as the *learner's knowledge*.

Moreover, it has been long recognised that learning does not happen purely in the learner's head, but also depends on the environment in which it takes place (Sawyer, 2006). While prior knowledge along with the learner's misconceptions provide the necessary information for assessing the progress of the learner in some domain knowledge, it does not always suffice for choosing feedback that supports learner's motivation. Learning is a process that involves the learner as a person, with emotional predispositions and transient affective states such as anxiety, boredom, frustration or joy, that impact their motivation and ultimately both their learning

experiences and their learning outcomes (Lepper et al., 1993). We refer to the information about the learner's psycho-emotional states as *learner's affect*.

Learner's knowledge

The term *learner's knowledge* encompasses both the factual knowledge, stored in a person's declarative (factual) memory, and the relevant skills needed to use, apply and update the facts (procedural knowledge). More specifically *learner's knowledge* typically refers to *what* facts the student knows, the *extent* to which the learner is believed to know them, the learner's *skills* along with the extent to which they are believed (by the tutor) to be mastered, what knowledge is erroneous (*misconceptions*), and what the learner does not know (*missing knowledge*).

There are a variety of different ways in which a model of learner's knowledge and skills can be conceptualised, represented and inferred. The type of representation and inference used often depends on the theoretical and/or computational perspectives taken by the individual researchers and the research questions they seek. Put simplistically, cognitive scientists are typically concerned with *cognitive fidelity* of the models that they create and they may use learner modelling techniques and tools as a way of verifying the validity of those models (e.g. Anderson et al., 1990). On the other hand, ITS researchers are typically concerned with supporting the learner in achieving mastery of a subject matter or a subset thereof, with learner modelling tools providing a way of diagnosing learner's knowledge to enable optimal feedback by the system in *real time*. Thus, although cognitive fidelity of a learner modelling tool is highly desirable to the ITS researchers, *computational efficiency* is of essence (Dillenbourg and Self, 1992).

It is important to bear in mind that the theoretical and pedagogical question of how knowledge is, or should, or can be represented and acquired by people, is somewhat, but not altogether, separate from the issue of how the learner's knowledge acquisition can be modelled computationally during a learning task, i.e. the specific computational techniques that can be used to infer the model of the learner. This muddled relationship (and the tension) between a theory preferred and a computational approach favoured is a known facet of the research related to building and deploying LMEs. This is because such environments are inextricably cognitive as well as computational (Ohlsson and Mitrovic, 2007). The difficulty of disentangling the cognitive and the computational considerations is also largely responsible for the apparent lack of an integrated, overarching definition of the learner modelling field as a whole and the difficulties with presenting it as a coherent collection of findings, principles and guidelines for the novice researchers who wish to learn about it and to explore it.

Representing learner's knowledge

Assuming the existence of a domain knowledge representation, building a learner model requires one to choose the appropriate way in which to represent the existing and evolving knowledge of the learner. Although there are a number of differing approaches to representing learner knowledge, many ITSs contain a representation of an *ideal* learner model, which, typically, is a representation of expert knowledge in the subject domain. Specifically, learner's knowledge can be represented as (i) overlaying the ideal learner model, i.e. as a subset of what the expert knows (fig 1a), (ii) as a differential model, whereby learner's knowledge is represented in terms of

two categories: knowledge that has been already presented to the student vs. knowledge that has never been introduced to him (fig 1b), or (iii) a perturbation model (fig. 1c), which combines the standard overlay model with a representation of learners' misconceptions.

[fig 1: 1a, 1b, 1c]

Perturbation Model and Bug Libraries (BLs), used most notably in the BUGGY systems (Brown and Burton, 1978; Burton and Brown, 1982), are the most successful of the three types of representations and are still in use (in improved forms) in some LMEs such as the Cognitive Tutors (Koedinger et al., 1997).

BLs specify common misconceptions and errors in terms of lists of *mal-rules*, reflecting the *typical* erroneous knowledge in a domain. They allow a computer tutor to identify incorrect knowledge through accounting for *errors of commission*, i.e. errors that can be observed overtly through the actions of learners, for example a step in a solution or an answer to a question. In general, BLs are useful for (i) diagnosing step-by-step the state of the learner's knowledge and possible misconceptions, (ii) predicting the reasoning path taken by the learner in a specific problem (iii) predicting the learner's future errors and answers and (iv) adapting the feedback accordingly. However, they are not practical for modelling learners' knowledge in ill-defined domains such as architecture, where the value of a solution path may not be known until the solution is reached. They are even less suitable for domains such as social interaction skills, where appropriateness of learner's actions is relative to the context of a specific situation, because the specification of mal-rules is absolute and discrete rather than, as such domains demand, relative and probabilistic.

Nevertheless, BLs serve to demonstrate the role of learner modelling and the difficulties involved in representing and inferring learner's knowledge. They highlight the relationship between domain knowledge and different types of errors that one may make in relation to it and raise questions as to what are misconceptions, what constitutes a typical misconception in a domain and what learners' behaviours count as manifesting mastery and lack thereof. As such BLs provide an invaluable tool for thinking about learning, the best ways to support learning, showing that even the simplest and best defined problems such a double-column subtraction can result in potentially an intractable number of possible errors that originate from different sources.

Cognitive Tutors and Cognitive modelling

Cognitive Tutors (CTs) are an example of an improved and hugely successful application of the BLs. CTs, originally developed as a test-bed for the ACT* and then ACT-R theory of cognition (Anderson, 1983), combine a specification of a domain knowledge, using ACT-R, with the cognitive diagnosis tools – specifically *model tracing* and *knowledge tracing*. ACT-R aims to describe human cognitive processes involved in problem solving in procedural domains such as mathematics, in order to inform our understanding of the structure of knowledge, of how knowledge is stored and how it is acquired and enhanced by humans over time when exposed to appropriate triggers. ACT-R makes a distinction between declarative knowledge (e.g. “*When both sides of the equation are divided by the same value, the resulting quantities are equal*”), and procedural knowledge (e.g. “*IF the goal is to solve an equation for variable X and the equation is of the form $aX = b$, THEN divide both*

sides of the equation by a to isolate X.”, (see Corbett et al. (1997)), and it envisions that the factual and procedural knowledge components are fundamentally distinct from one another.

Cognitive fidelity is paramount in ACT-R, because the purpose of the modelling in which it engages – the cognitive modelling – is to enhance our understanding of how humans really learn and how the cognitive processes are reflected in the neural activities of the human brain. Thus, it is important to appreciate that the purpose of cognitive modelling is not necessarily to enhance education as such, but rather it is to explain how humans acquire and use knowledge; at best it is to mimic the real processes. Ensuring cognitive fidelity of a model is therefore first on the agenda in cognitive modelling, and, i.e. learner modelling within CTs, provides the means for verifying the models that ensue.

Cognitive diagnosis consists in keeping track of (i) the learner’s progress through a solution and (ii) the growth of learner’s knowledge over time, respectively. The diagnosis is based on the specification of production rules that apply, given a particular stage in the problem-solving episode. The production rules are annotated for correctness and specificity of solutions that they offer. During problem solving a CT keeps track of the solution steps committed by the learner and identifies the production rules in its database that correspond to learner’s solution steps. The annotations associated with each production rule in the database provide the basis for the assessment of the correctness of the learner’s step and for decisions about appropriate feedback.

The system employs *model tracing* to infer a solution path for individual learners given a specific goal. CTs assume that a cognitive model represents the *ideal* learner model. Given a description of declarative and procedural knowledge for a specific domain and given the actions of the learner at the interface level (*learner's behaviours*), the system can generate (*trace*) a sequence of possible steps through a problem space alongside the steps that the learner is producing. Since different elements of the interface in a cognitive tutor correspond to a representation of the problem-solving goal and to the relevant information needed to achieve it, the tutor is able to use the information about the learner's behaviours at the interface level to infer the learner's current goal. Furthermore, by applying its production rules to the learner's goal, the tutor is able to generate a set of possible rules that could satisfy that goal (Corbett et al., 1997).

While model tracing is concerned with generating a generic learner model step-by-step, the purpose of *knowledge tracing* – the second diagnosis mechanism employed in CTs – is to assess the growth of learner's knowledge over time. Knowledge tracing provides the basis for selecting the appropriate problem given the changing mastery level of the learner. Similar to model tracing, it consists of the tutor generating a set of production rules at each level of the curriculum taught and in assessing the extent to which the learner is believed to have mastered those rules. This assessment is done through assignment of probabilities to each of the productions generated by the tutor based on the learner's behaviours whenever an opportunity arises for a specific rule to be applied.

Cognitive modelling has been very successful in supporting learning of certain aspects of well-defined domains, e.g. mathematics. It is important to bear in mind, however, that their success in learner modelled environments came out of the benefit of hindsight, in particular from the early research in intelligent tutoring systems, which highlighted the importance of aiming for cognitive fidelity in the representation of expert knowledge, because the less cognitively truthful the model the greater the chance that the pedagogic support offered might be inadequate. Cognitive fidelity seems especially important in relation to (i) a systems' ability to reason about the problem in the same way as humans do and (ii) a system's ability to apply the same mechanisms as humans do when searching for a solution. Early tutoring systems that fail with respect to the first requirement, e.g. SOPHIE I (Brown et al., 1973), tend to lack the ability to support learner's reflection: while they are able to recommend to the learner actions that are optimal in each problem solving context, they fail to describe the context and the rationale for the solution, leaving the learner to construct these for themselves. Tutoring systems that fail with respect to the second requirement, e.g. GUIDON (Clancey, 1983), also fail to advise the learner on what to do next, because they have no idea about what learning trajectory the learner follows (Corbett et al., 1997).

Constraint-Based Modelling

An approach, developed in parallel with the CTs, is constraint-based modelling (CBM) (Ohlsson, 1992). CBM emerged as a response to the issue of the *over-specificity* of the learner knowledge descriptions that were characteristic of the early LMEs such as DEBUGGY and later CTs (e.g. Ohlsson and Mitrovic (2007)). Over-

specificity of learner models refers to the observation that, in its general form, learner modelling is intractable (Holt et al., 1994; Self, 1990). Independently of the technique that is chosen to represent learner's knowledge, any mechanism that aims to infer a complete model of a learner is bound to rely on the specification of hundreds, sometimes thousands, of individual knowledge chunks. In practice, this level of specificity cannot be handled effectively, even when sophisticated authoring tools are employed to construct the appropriate knowledge representations. Indeed, if human tutors are taken as the gold standard for building digital learning environments, then the question arises of whether it is sensible to expect a fully specified learner model at all. Empirical studies show that human tutors do not tend to rely on complete and fully coherent models of learners when they teach, and yet can be very effective (Holt et al., 1994; Leinhardt and Ohlsson, 1990). Ohlsson's CBM approach builds on this observation by introducing a technique that requires only a partial, but nevertheless effective learner model.

In the CBM approach, the problem of over-specificity is addressed through abstraction, whereby domain knowledge is represented through a set of constraints corresponding to correct solutions for that domain. Each constraint is an ordered pair $\langle C_r, C_s \rangle$, where C_r is the relevance condition that identifies the class of problem states for which the constraint is relevant, and C_s is the satisfaction condition that identifies the class of states in which the constraint is satisfied. Each member of the pair can be thought of as a set of properties of a problem state, with the constraint representing a statement such that *if the properties C_r hold, then the properties C_s have to hold also*. If the properties C_s do not hold, this means that an error has occurred. Therefore,

indirectly each constraint also represents a set of erroneous solutions, i.e. the solutions that violate the constraint. A constraint-based learner model consists of the set of constraints that the learner does and does not violate. A violation of a constraint on the part of the learner signifies a missing or incomplete knowledge on their part.

Formulating a set of constraints for a domain that are also valid pedagogically is a difficult task, but by addressing the over-specificity problem this task is eased greatly, making CBM a computationally efficient solution to learner modelling. As Ohlsson argues, a set of constraints does not necessarily have to be complete in order to be useful. An incomplete set of constraints may lead to rare errors not being caught, but as long as the most common mistakes are addressed, the system is still able to provide learners with valuable feedback. Importantly, formulating the knowledge in terms of constraints (i.e. only correct knowledge requires to be fully specified) enables reusability of such representations across different populations of learners.

While ensuring computational efficiency, CBM also claims to be a cognitive truthful approach: it is based on a psychological theory of learning that asserts that learning occurs primarily when students catch themselves making mistakes. This theory makes a distinction between *generative* and *evaluative* knowledge (Ohlsson, 1996).

Generative knowledge produces problem-solving actions and can be expressed by means of a set of *rules*. Evaluative knowledge evaluates the actions' outcomes in terms of their desirability with respect to a specific goal, which can be expressed by a set of *constraints*. Ohlsson conjectures that the acquisition of a new cognitive skill is (at least, partially) based on the transfer of knowledge from the evaluative to the

generative component (Ohlsson 1996). Hence, in the context of this theory of learning, a constraint-based modelled environment constitutes an amplified evaluative knowledge base that contains the constraints representing the evaluative expert knowledge. By giving the learner access to the evaluative knowledge, the system can speed up and augment the transfer of information from the evaluative to the generative component, and so ultimately it can support the learning process.

Learner's affect

None of the approaches to learner modelling reviewed so far incorporate a mechanism that accounts in any way for the influence that emotions and context of educational interactions have on the experience and effectiveness of learning. Yet, the importance of learners' affect to learning has been long recognised in education. As early as 1908, Yerkes and Dodson discusses the relationship between learner's emotional arousal and learning performance, observing that moderate arousal is always preferable to the extremes of low or high arousal, which can be detrimental to learning (Yerkes and Dodson, 1908). This observation still holds with a number of recent studies in neuroscience and psychology providing further evidence that cognition and emotions are profoundly intertwined (Damasio, 1994). Experiencing emotions helps humans make decisions given the demands of their environment, with the three main cognitive processes involved in learning: attention, memorisation and reasoning, being strongly influenced by the specific emotions experienced. Positive emotions enhance attention, facilitate memorisation and improve efficiency and efficacy in problem solving and decision-making (Isen 2000). Negative emotions, on the other hand, can be detrimental to concentration, may disturb the retrieval of information, make

memorisation less effective and can induce convergent and sequential reasoning (Lissetti and Schiano, 2000).

An increasing number of LMEs incorporate modules that recognise, interpret and respond to learners' affective states. The goal of such modules is to generate real-time data that informs our understanding of what constitutes the experience of emotion(s) during learning and of the relationship between such experiences and learning outcomes. In such LMEs, the learner model incorporates information about learners' affective states and updates it regularly based on the interaction between the learner and the system.

Modelling learner's affect automatically requires a designer to: (i) specify the emotions of interest in a given domain, (ii) define the emotions in terms of observable behaviours, and (iii) specify a mechanism for detecting and interpreting the behaviours in terms of the emotions identified. These three requirements apply regardless of the theoretical or technical standpoint of the individual designers. Real emotional experiences are far messier than contemporary digital implementations allow for: emotions are transient, co-occurring and overlapping; they are short lived and often triggered by multiple events (D'Mello and Graesser, 2011), and they depend on medium and long-term affective predispositions, goals and beliefs of the people who experience them (Gebhard 2005). Educational contexts bring about very specific types of affective experiences, bearing heavily and explicitly on the types of goals, attitudes and beliefs that learners may harbour in a given context about a specific domain, the support they are receiving, and themselves as learners.

What affect to model?

The past hundred years of research on emotions produced many diverse theories that could potentially serve as the basis for developing systems that detect and interpret learners' affect. A recent excellent review by Calvo and D'Mello (2010) provides a taxonomy of the most prominent theories of emotion used specifically to inform automatic detection of emotions of a generic user. The theories considered include those that are used to inform the design of *generative* models, i.e. models of how emotions should be produced by artificial agents in real time (either through physical expressions of emotions (Darwin, 2009; Ekman, 1971) or a combination of physical and physiological expressions (James, 1884)). Alternatively, the existing theories are used to *predict* (i.e. model) user emotions based on specific triggers in the environment such as objects, events and other people (e.g. *cognitive appraisal* theories – Scherer (2005), or the OCC theory (Ortony et al., 1988)). Many of the theories, save for the ones developed in the *social constructivist* perspective (Averill, 1980; Stets and Turner, 2008), consider emotions outside of the environment in which they come about or the task at hand. Those that do consider the context as integral to the experience and understanding of emotions are however far less amenable to being automated than those that consider emotions in isolation.

Although most designers of learner modelled environments remain agnostic to the theories available, the main considerations taken into account in choosing a particular theory in this context are: (a) how well a theory lends itself to being implemented in a computer system, and (b) how well it is able to support the selection of appropriate feedback for a given learner. Similar to modelling learner's knowledge,

considerations of computational efficiency and psychological fidelity play a role in what approaches are adopted. An additional consideration is whether a theory is pedagogically viable: how well a theory lends itself to being implemented may conflict with its applicability to a wide range of LMEs, e.g. arguably the nature, diversity and duration of emotional experiences within an educational game environment are different from the affect experienced in exploratory environments, because the nature of tasks and rewards for achieving the tasks therein are different. The focus on educational viability also highlights the question of how well a given theory of affect can account for the affective states that users experience during learning, i.e. whether the affective states modelled are relevant to learning.

All of the six approaches reviewed by Calvo and D'Mello, individually and sometimes in combination have been used to inform various implementations of the automatic affect detection systems. Some of them such as the one adopted by D'Mello et al. (2007) combine physiological and behavioural signal detection, such as skin conductance, eye gaze tracking and posture detection, with linguistic cues analysis, in order to infer motivational states such as boredom, flow, etc. In doing so such approaches combine different theoretical perspectives – in D'Mello et al.'s case the theories of emotions as expressions, embodiments as well as theories of motivation that tend to be socio-constructivist by nature. In the rest of this chapter we focus on two relatively coherent trends that have been substantially used in the context of modelling learner's affect: the cognitive appraisal and the social constructivism.

Cognitive appraisal vs. social constructivism

The OCC theory (Ortony et al., 1988) is one of the most popular theories of emotions in the affective computing field, not least because it is computationally viable. It represents twenty-two emotions categorised into a hierarchy of six groups such as well-being: joy and distress or attraction: love and hate. The categories are based on valenced (i.e. positive/negative) responses of a person to current states of the world and to their causes. OCC accounts for how emotions arise from people's appraisal of the current situation, where any situation consists of objects, events and agents. OCC allows one to specify the intensity of emotions based on the likelihood of an event taking place or of a person's familiarity with an object. The model is very detailed, but its limitation is that it accounts only for prime emotions, leading some researchers to question its relevance to learning (Calvo and D'Mello, 2012).

In contrast with the cognitive appraisal theories, in the social constructivist perspective emotions are seen as socio-cultural constructs that are co-constructed with others through social interactions and (linguistic) communication (Averill, 1980; Glenberg et al., 2005). While, most approaches within affective computing rely on other theories available, arguably, affective computing research that focuses on education is by necessity socio-constructivist. This is because in education, the learning processes, including the affective states that impact a person's learning, are inextricably perceived as outcomes of the interaction between the learner, the teacher and the learner's performance in relation to a learning task at hand. Specifically, there is substantial and maturing research pointing to learners' *motivation* as a form of affective experience of particular relevance to learning (Craig et al., 2004).

Motivational states are sometimes understood as “the determinants of thought and action” (Weiner 1992, p. 17) and, unlike the short-term, prime emotions such as joy, distress, love or hate, etc., motivation is a composite of emotions, long- and short-term goals of a person and their beliefs, as well as their emotions experienced in specific contexts.

The two theoretical perspectives are not incompatible, even if they approach the affect modelling from different directions: the cognitive appraisal implementations use prime emotions to describe learner’s affect, whereas the social constructivists focus broadly on motivation, differences in learners’ behaviours and the influence of motivation regulation on learner’s affect and performance (see figure 2). As Conati and Maclaren (2009) point out, the ultimate goal of learner modelling research is to devise methods and approaches, which combine the low-level emotion detection and interpretation with higher-level motivational state modelling.

We use Keller’s (1983) model (figure 2) to illustrate both the relevance and compatibility of the different approaches used to date. The figure represents Keller’s understanding of motivation and motivation regulation as a loop between effort, performance and consequences, whereby the level of effort that the learner dedicates to a task depends on their values and goals. The level of effort dedicated impacts on the quality of the learner’s performance and results in the intrinsic and extrinsic outcomes (in the figure: ‘consequences’), including learner’s emotional responses, and social and material rewards. The consequences are appraised by the learner (through cognitive evaluation) and these appraisals feed back into ‘motives’ either

reinforcing or changing the values of the learner.

Keller's model shows explicitly how motivation depends on the input from the environment and especially how appropriate pedagogical design and management fits in with learner's motivation and emotional regulation (environmental inputs).

Therefore, the model is explicitly socio-constructivist in nature in Averill's sense.

Second, the model explicitly accounts for learner's cognitive appraisal of the outcomes of their actions, which is where the approaches concerned with the prime emotions, such as the OCC model, fit – these are highlighted by us in the top right part of the figure. While, none of the existing implementations include all of the aspects of this (enhanced) model, it can be thought of as representing, possibly with some modifications, the ideal to which the affective learner modelling field aspires.

In the following sections we describe two approaches to affective learner modelling as exemplars of the existing implementations that adopt different theoretical entry points.

[figure 2]

Modelling the learner's affect based on a cognitive appraisal theory

Many LMEs rely solely on the OCC model, although seldom do the learner models generated include all twenty-two emotions specified therein. Instead, the designers tend to select the most relevant emotions that fit the domain or the mode of interaction afforded by their environment. One example is the Prime Climb, two-player game developed to support children in learning number factorisation (Conati and Maclaren, 2009). Each player controls a different character, and the two players need to cooperate to allow the respective virtual characters to climb a mountain and arrive to the top. The mountain is divided into numbered sectors. Each character can only

move to a numbered sector that does not share any factors with the sector occupied by the other character. If a player makes a wrong move, then her/his corresponding character falls down the mountain.

The learner model in Prime Climb includes only the six most relevant emotions to the target domain: *joy/distress*, *admiration/reproach*, and *pride/shame*. Prime Climb uses video game paradigm to exploit the intrinsic motivation that such games have been found to afford (Johnson and Rickel, 2001). Apart from modelling the learner's progress through the number factorisation exercises, the environment is also enabled with a diagnostic model responsible for detecting and interpreting children's emotions as they progress through the game. The model also assesses how multiple goals, which the students typically have while playing the game, influence learners' appraisal of the game interaction, and consequently their emotions. Five high level student goals have been considered in Prime Climb: *Have Fun*, *Avoid Falling*, *Learn Math*, *Beat Partner*, and *Succeed By Myself*.

Because the relations between student personality, goals, game states and emotions are not deterministic, the student model in Prime Climb is implemented through Dynamic Bayesian Networks to manage uncertainty in real-time. The Prime Climb affective model represents the first known probabilistic representation of the OCC theory. The purpose of this model is to facilitate a nuanced selection of feedback and ultimately to enhance both the experience and learning outcomes. In order to maintain learner engagement while providing instructional learning support, Prime Climb incorporates an animated pedagogical agent, which responds to learners' requests for

hints or offers hints when it decides that the student needs help. The agent relies on a probabilistic model of the student's knowledge about factorisation to decide when to intervene and what hints to provide. The evaluation studies of Prime Climb with users in primary schools offer promising evidence that providing personalised instruction, in this case by means of an animated pedagogical agent, can improve learning over traditional methods (Conati and Zhao, 2004).

Modelling learner's affect based on interaction analysis (socio-constructivism)

In Porayska-Pomsta et al. (2008) the specific context of a situation along with the interaction between a learner and a teacher are integral to both regulating learners' emotions and to being able to recognise and act on them in pedagogically viable ways. Our starting point is the natural language dialogue that takes place between the teacher and the learner. In line with the socio-constructivist theories of emotions, our premise is that modelling learner's emotions and the delivery of appropriate support requires an understanding of the interaction context in which the emotional states are experienced. We define the educational interaction loop (also known as the *affective loop* – see (D'Mello, et al., 2007)) as involving a teacher: (i) observing learner's behaviours in context, (ii) modelling the learner's affective states based on such observations, (iii) deciding how to support the learner based on the model, and in turn (iv) influencing learner's affective states by acting on his/her decisions of how to provide the support.

We used the Language Enhanced, User Adaptive, Interactive e-Learning for Mathematics (LeAM) system as a context for this research. LeAM consists of a learner model, a tutorial component, an exercise repository, a domain reasoner and

natural language dialogue capabilities. We conducted a number of studies, over two years, in order to inform the design of the learner and the natural language dialogue models. In order to ensure ecological validity of the data generated, we restricted the student-teacher communication to a chat interface (figure 3), with no visual or audio channels, to resemble the interface of the final learning environment. Five experienced teachers and twenty-eight learners were asked to adapt their responses to each other by accommodating to this limited channel of communication. The data collected include natural language dialogue logs, verbal protocols and semi-structured interviews data from teachers and learners and, crucially, teachers' real-time annotations of the situations in the context of which they were providing feedback to the learners, which included affective states such as learner's confidence, interest and effort, as well as other contextual information such as the teachers' perception of the universal and relative difficulty of the task given to the learner, the correctness of learner's answers, amount of session time left, learner's aptitude, knowledge and goals. The data collected provide us with rich information not only about the context that teachers take into account when diagnosing learners' affect, but also with the specific examples of the feedback that the teachers deem appropriate for the individual learners in the *heat-of-the-moment*, given the specific tasks and given their own teaching and communication styles.

The findings from these studies informed the design of the LeAM's learner model.

The model, implemented as a Bayesian Network, represents causal relationships between a specific combination of contextual factors and enables to link between the learners' affective states that the teachers infer based on observable behaviours of

learners, and the ways in which the teachers act on such inferences (Porayska-Pomsta et al., 2008).

[figure 3]

Discussion and Conclusions

This chapter introduced the concept of LMEs as a method through which learning can be both supported and studied. In particular, learner modelling was presented as the core, defining characteristic of any learner modelled environment and as a necessary pre-requisite for a technology enhanced educational tool capable of adapting the interaction and the pedagogical support to the individual learners' needs in real-time.

Automatic learner modelling, although still an emerging discipline, is not brand new and much research is available that illustrates different approaches to it. Early learner modelling tools focused solely on modelling learners' knowledge. As the field matured and new technologies, such as physiological sensors, became available and affordable, the focus extended to learners' affect. However, few researchers in the field approach the task of learner modelling and of building LMEs from the same theoretical, technological or pedagogic perspectives. This makes the field lack uniformity and, consequently, hinders comparisons between the different implementations and a coherent account of both the trends and achievements therein.

The aim of this chapter was to bring the many approaches and points of view, as much as it is currently possible, into a coherent description of the LM field, in order to enable a novice reader to understand its main tenets and perspectives. We identified

two broad types of uses of learner models: (i) to implement a theory about human learning processes and (ii) to support human-computer interaction in educational contexts. Although these research aims are not incompatible, the choice of one over another typically requires the research focus being placed on achieving in such models either cognitive and/or psychological fidelity, or computational efficiency. Throughout the chapter we provided exemplar implementations reflecting these two foci, as well as examples of approaches which, like the constraint-based approach to learner's knowledge modelling or Conati's and Maclaren's approach to affect modelling, attempt to achieve a bit of both. We selected the approaches to learner modelling in both the cognitive and the affective branches that represent the predominant trends in a still maturing field and which allow us to highlight the main questions related not only to the design of LMEs, but also to education research in general. There are many more notable instances of LMEs than described here. There are examples of LMEs that illustrate the state-of-the-art achievements in learner modelling and many other of research that informs the design of such environments. Using the references we provide to support our conclusions throughout the chapter, we hope that the interested reader will be able to trace and pursue those approaches.

In concluding this chapter, we would like to reflect on some of the criticisms that the field sometimes receives and which is encapsulated in the question of whether learner modelling is at all a feasible endeavour. There are some who point to the intractability of the learner modelling question and to the fact, that thus far, the question has been addressed with little success: the models created are limited in scope and are difficult to port to other domains and different types of learners; they

focus only on the cognitive, but not affective, or they model learners with respect to the states that are only tangentially relevant to learning. Whilst these criticisms are valid and they point to what is the most likely reason for the lack of complete coherence in the field, we believe that the question of feasibility is misguided. Models are by definition a reduced representation of the real phenomena and this is precisely why they are so useful for research. They lend themselves to being tested, manipulated and changed and through this they inform our understanding of the phenomena studied. In relation to learner modelling, it is easy to forget that as much as they are tools through which believable and educationally beneficial interactions can be facilitated, such models are first and foremost research tools through which we can explore what it means to learn.

Bibliography

- Anderson, J. R. (1983) *The Architecture of Cognition*. Boston, MA: Harvard University Press.
- Anderson, J. R., C.F. Boyle, A.T. Corbett, and M.W Lewis. (1990) "Cognitive modelling and intelligent tutoring." *Artificial Intelligence* 42: 7-49.
- Averill, J.R. (1980) "A Constructivist View of Emotion." In *Emotion Theory, Research and Experience*, by R. Plutchik and H. Kellerman, 305-339. New York: Academic Press.
- Bloom, B. (1986) "The 2 Sigma Problem: The Search for Methods of Group Instruction as Effective as One-to-One Tutoring." *Educational Researcher* 13, no. 6: 4-16.

- Brown, J. S., and R. R. Burton (1978) "Diagnostic models for procedural bugs in basic mathematical skills." *Cognitive Science* 2: 155-192.
- Brown, J.S., Burton, R., and Zdybel, F. (1973) "A model-driven question-answering system for mixed initiative computer-assisted instruction." *IEEE Transactions on Systems, Man and Cybernetics*: 248-257.
- Burton, R.R., and Brown, J.S. (1982) "An investigation of computer coaching for informal learning activities." *Intelligent Tutoring Systems*: 157-183.
- Calvo, R. A., and S. D'Mello (2010). "Affect Detection: An Interdisciplinary Review of Models, Methods, and their Approaches." *IEEE Transactions on Affective Computing* (Published by the IEEE Computer Society) 1, no. 1: 18-37.
- Clancey, W. J (1983) "GUIDON." *Journal of Computer-Based Instruction* 10, no. 1: 8-14.
- Conati, C., and H. Maclaren (2009) "Empirically Building and Evaluating a Probabilistic Model of User Affect." *User Modeling and User-Adapted Interaction* 19, no. 3: 267-303.
- Conati, C., and X. Zhao (2004) "Building and Evaluating an Intelligent Pedagogical Agent to Improve the Effectiveness of an Educational Game." *Proceedings of IUI '04, International Conference on Intelligent User Interfaces*. Island of Madeira, Portugal, pp. 6-13.
- Corbett, A. T., Koedinger, K. R. and Anderson, J. R. (1997) "Intelligent tutoring systems." In *Handbook of Human-Computer Interaction*, by M. G. Helander, T. K. Landauer and P. Prabhu. Amsterdam, The Netherlands: Elsevier Science.

Craig, S. D., A. C. Graesser, J. Sullins, and B. Gholson (2004) "Affect and Learning: An Exploratory Look into the Role of Affect in Learning with AutoTutor." *Journal of Educational Media* 29: 241-250.

Damasio, A. R. (1994). *Descartes' error: emotion, reason, and the human brain*. New York: G. P. Putnam.

Darwin, C (2009) *The Expression of the Emotions in Man and Animals*. Penguin Classics.

Dillenbourg, P., and Self, J. (1992) "A framework for learner modelling." *Interactive Learning Environments* 2, no. 2: 111-137.

D'Mello, S. K., Picard, R. W. and Graesser, A. C. (2007). "Towards an Affect-Sensitive AutoTutor." *Special issue on Intelligent Educational Systems – IEEE Intelligent Systems* 22, no. 4: 53-61.

D'Mello, S., and Graesser, A. C. (2011) "The half-life of cognitive-affective states during complex learning." *Cognition and Emotion* 25, no. 7: 1299-1308.

Ekman, P (1971) *Universals and Cultural Differences in Facial Expressions of Emotion*. University of Nebraska Press.

Gebhard, P (2005) "ALMA - A Layered Model of Affect." *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, Utrecht, pp. 29-36.

Glenberg, A., Havas, D., Becker, R., and Rinck, M. (2005). "Grounding Language in Bodily States: The Case for Emotion." In *The Grounding of Cognition: The Role of Perception and Action in Memory, Language and Thinking*, by R. Zwaan and D. Pecher, 115-128. Cambridge: CUP.

Holt, P., Dubs, S., Jones, M., and Greer, J.E. (1994). "Student modeling: The key to individualized knowledge-based instruction." In *The state of student modelling*, by J. Greer and G. McCalla, 3-35. Berlin: Springer-Verlag.

Isen, A.M. (2000) "Positive Affect and Decision Making." In *Handbook of Emotions*, by M. Lewis and J. Haviland-Jones, 417-435. New York, Guilford.

James, W. (1884) "What is an Emotion?" *Mind*, pp. 188-205.

Johnson, W.L., and Rickel, J. (2001). "Research in Animated. Pedagogical Agents: Progress and Prospects for Training." *International Conference in Intelligent User Interfaces*. Santa Fe.

Keller, J. M. (1983) "Motivational design of instruction." In *Instructional-design Theories and Models: an Overview of their Current Status*, by C. M. Reigeluth, 383-434. Hillsdale, NJ: Lawrence Erlbaum Associates.

Koedinger, K. R., and Corbett, A. T. (2006). "Cognitive Tutors: Technology bringing learning science to the classroom." In *The Cambridge Handbook of the Learning Sciences*, by K. Sawyer, 61-78. New York: CUP.

Koedinger, K. R., Anderson, J. R., Hadley, W. H., and Mark, M. A. (1997) "Cognitive tutors: Lessons learned." *Journal of the Learning Sciences* 4, no. 2: 30-43.

Leinhardt, G., and Ohlsson, S. (1990). "Tutorials on the structure of tutoring from teachers." *Journal of Artificial Intelligence in Education* 2: 21-46.

Lepper, M., Woolverton, M., Mumme, D., and Gurtner, J. (1993). "Motivational techniques of expert human tutors: Lessons for the design of computer-based tutors." In *Computers as Cognitive Tools*, by S. Lajoie and Derry S., 75-105. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

- Lisetti, C., and Schiano, D. (2000). "Automatic Facial Expression Interpretation: Where Human- Computer Interaction, Artificial Intelligence and Cognitive Science Intersect." *Pragmatics and Cognition* 81, no. 1: 223–239.
- Lynch, C., Ashley, K., Alevan, V., and Pinkwart, N. (2006). "Defining ill-defined domains; a literature survey." *Proceedings of the Workshop on Intelligent Tutoring Systems for Ill-Defined Domains at the 8th International Conference on Intelligent Tutoring Systems*. Jhongli, pp. 1-10.
- Mitrovic, A., and Ohlsson, S. (1999). "Evaluation of a constraint-based tutor for a database language." *International Journal of Artificial Intelligence in Education* 10: 238-256.
- Mitrovic, A., Koedinger, K.R., and Martin, B. (2003) "Comparative Analysis of Cognitive Tutoring and Constraint-Based Modeling." *9th International Conference on User Modelling (UM 2003)*. Lecture Notes in Computer Science, pp. 313-322.
- Newell, A. (1990) *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.
- Newell, A., and Simon, H. A. (1995). "Gps, a program that simulates human thought." In *Computation and Intelligence, Collected Readings*, by G. F. Luger, 415–428. AAI, and MIT Press.
- Ohlsson, S. (1992) "Information-processing explanations of insight and related phenomena." In *Advances in the psychology of thinking*, by M. T. Keane and K. J. Gilhooly, 1-44. Harvester Wheatsheaf.
- Ohlsson, S. (1996) "Learning from performance errors." *Psychological Review* 103: 241- 262.

Ohlsson, S., and Mitrovic, A. (2007). "Fidelity and efficiency of knowledge representations for intelligent tutoring systems." *Technology, Instruction, Cognition and Learning* 5, no. 2: 101-132.

Ortony, A., Clore, G.L., and Collins, A. (1988) *The cognitive Structure of Emotions*. Cambridge: Cambridge University Press.

Porayska-Pomsta, K., Mavrikis, M., and Pain, H.. (2008). "Diagnosing and acting on student affect: the tutor's perspective." *User Modeling and User-Adapted Interaction*: 125-173.

Sawyer, R. K. (2006). *The Cambridge Handbook of the Learning Sciences*. Cambridge: CUP.

Scherer, K. R. (2005). "What are emotions? And how can they be measured?" *Social Science Information* 44, no. 4: 695-729.

Self, J. A. (1992). "Bypassing the intractable problem of student modeling." In *Intelligent tutoring systems: At the crossroads of artificial intelligence and education*, by Weiner, B. *Human Motivation: Metaphors, Theories, and Research*. London: Sage Publications Inc.

Sison, R., and S. Masamichi. "Student Modelling and Machine Learning." *International Journal of Artificial Intelligence in Education* 9 (1998): 128-158.

Sperber D., and Wilson D. (1995) "Relevance. Communication and Cognition", Blackwell publishing.

Woolf, B. (2009). *Building intelligent interactive tutors: Student-centered strategies for revolutionizing e-learning*. Burlington, MA: Morgan Kaufmann.

Yerkes, R.M and Dodson, J.D. (1908). "The relation of strength of stimulus to rapidity of habit formation" *Journal of Comparative Neurology and Psychology* 18: 459-482.