# Open Research Online

The Open University's repository of research publications
and other research outputs

## Opening up the interpretation process in an open learner model

## Journal Item

For guidance on citations see FAQs.

Version: Version of Record

Link(s) to article on publisher's website:
http://ijaied.org/pub/1113/

oro.open.ac.uk

# Opening up the Interpretation Process in an Open Learner Model

**Nicolas Van Labeke**, **Paul Brna**, *The SCRE Centre, University of Glasgow, St. Andrews' Building, 11 Eldon Street, Glasgow G3 6NH, Scotland*
*n.vanlabeke@scre.ac.uk, paul.brna@scre.ac.uk*

**Rafael Morales,** *Sistema de Universidad Virtual, Universidad de Guadalajara, Escuela Militar de Aviación 16, Col. Ladrón de Guevara, 44170 Guadalajara, Mexico*
*rmorales@udgvirtual.udg.mx*

**Abstract.** Opening a model of the learner is a potentially complex operation. There are many aspects of the learner that can be modelled, and many of these aspects may need to be opened in different ways. In addition, there may be complicated interactions between these aspects which raise questions both about the accuracy of the underlying model and the methods for representing a holistic view of the model. There can also be complex processes involved in inferring the learner's state, and opening up views onto these processes raises the issues connected with the main focus of this paper – namely, how can we open up the process of interpreting the learner's behaviour in such a manner that the learner can both understand the process and challenge the interpretation in a meaningful manner. The paper provides a description of the design and implementation of an open learner model (termed the xOLM) which features an approach to breaking free from the limitations of "black box" interpretation. This approach is based on a Toulmin-like argumentation structure together with a form of data fusion based on an adaptation of Dempster-Shafer. However, the approach is not without its problems. The paper ends with a discussion of the possible ways in which open learner models might open up the interpretation process even more effectively.

## INTRODUCTION

It appears to be a common belief that opening the learner model has significant educational value. The notion of an Open Learner Model (OLM) was initially developed by Cumming and Self when the learner model was recast as an inspectable resource for the learner, under the control of the learner and possibly modifiable by the learner (Cumming & Self, 1991). From this re-emphasis on the education/personal utility of opening the learner model to the learner, many different aspects of the learner model have been opened.

Bull, Brna and Pain worked on a learner model that took into account a number of factors including learning styles and misconceptions/errors that extended the idea of the learner model simply recording what the student knew about the topic being taught (Bull, Brna, & Pain, 1995; Bull, Pain, & Brna, 1996). Kay began to examine the notion of learner models as a reflection tool by looking at their scrutability (Kay, 1997), and Bull, Dimitrova and Brna sought to enhance communication with, and

about, open learner models with a view to improving reflection (Bull, Dimitrova, & Brna, 2002; Dimitrova, 2003; Dimitrova, Brna, & Self, 2002). Morales and his colleagues worked on the cognitive effects of inspecting open learner models that sought to make explicit the implicit sensorimotor knowledge of the learner (Morales, Pain, & Conlon, 2000). This work is a fairly rare example of seeking to open up a process in which the student might be engaged rather than focusing on a belief with propositional content which could be articulated reasonably clearly.

In 1999, Brna, Bull, Pain and Self suggested an approach to assessment that was termed "Negotiated Collaborative Assessment" (NCA), a process which opened up a discussion between the learner and the system about the results of the assessment through the use of an open model of the student's knowledge (Brna, Bull, Pain, & Self, 1999). In this work, it was proposed that it would be necessary to examine:

- the kinds of assessment criteria involved;

- the reasons why the criteria are selected;

- the degree to which the student can challenge the criteria;

- the further evidence that can (and should) be collected during an interaction;

- other sources of material that can be consulted;

- the ground rules for negotiation; the ways in which these rules are selected and communicated;

- the extent to which the student influences the final decision;

- the degree to which the student learns during the assessment process.

At the heart of this NCA process is the need to examine the raw data obtained directly (or indirectly) from the learner, interpret the data and then make a value judgement based on this interpretation. Typically, this process has received relatively little attention, and, indeed is often hard to represent in any way that does more than point out that some response by the learner is wrong (or right). This is perhaps fine from some perspectives – typically ones where the focus of learning is more on content that has been suitably atomised. When all we are concerned with is assessing mastery on a set of tasks that cover some curriculum/body of knowledge, then, traditionally, assessment is not concerned with the ways in which the learner might benefit from the assessment process itself. However, increasingly, there is a move which seeks to shift the balance from assessment for the benefit of some third party to assessment for learning (Birenbaum et al., 2006).

So the issue from work on assessment that this paper takes up is "how to make clear to the learner exactly how the system (teacher) might examine their work and form a value judgement about it" with potential advantages that include the prospect that the learner becomes more aware of the assessment process. This may be of benefit to some students though, perhaps, not to all. The problems to be solved are serious enough – as Tanimoto points out, there is a major problem opening out the assessment process to learners "... requires not only a means to reveal the names and values of model variables but an interpretive mechanism that translates the information from a pedagogical perspective to a learner's perspective" (Tanimoto, 2005). In the next section, we present an approach to exposing the process of making a value judgement about the learner's performance/behaviour.

The context of the work is the development of an open learner model as part of LeActiveMath, an EU funded project in the area of technology enhanced learning. LeActiveMath is a web-based learning environment designed primarily for use by students of mathematics in their last year before

university and their first year at university. The current content consists of examples, explanations, definitions and problems in differential calculus. The open learner model is responsible for receiving data from a wide range of activities and sources, performing appropriate interpretations and making the results available to other components of the system.

This paper is organised as follow. The first section describes the underlying Learner Model and establishes some of the concepts that will be – in one way or another – presented to the learner through the extended Open Learner Model (xOLM). The second section introduces the variant of the Toulmin Argumentation Pattern deployed to coordinate the interactions between the learner and the xOLM. The third section describes the interface of the xOLM, emphasising the various external representations used to convey to the learners parts of the interpretation process. In the fourth section, aspects of the verbalisation of the interpretation process are given. Some results from the evaluation of the system are presented in the fifth section. Finally, we conclude this paper by addressing some of the issues that came up during these initial stages of the project.

## THE UNDERLYING LEARNER MODEL

The learner model used is an "extended" one - that is, the kinds of content in the model are more extensive than ones containing only beliefs about the student's knowledge of the domain. The extended Learner Model (xLM) is itself an interesting piece of work, and maintains a model of the learner which holds the accumulated evidence of the learner's state in relation to knowledge of the domain being learned, the assessment of the learner's competence, motivational and affective factors as well as metacognitive information and knowledge of conceptual and performance errors (termed CAPEs). The xLM is a probabilistic model, updated using a formalism known as a Transferable Belief Model – TBM (Smets & Kennes, 1994), a variation of the Dempster-Shafer Theory (Shafer, 1976).

The TBM is designed for problems where the actual state of affairs (e.g. the learner's ability) is not exactly known, but only known to belong to some subset of possible states. It expresses in a quantified and axiomatically well founded theory the strength of an agent's opinion (commonly called the belief) about which of the possible states corresponds to the perceived state of affairs (called a belief function). The measure expressed by a probability in the TBM has the same purpose as those in classical probability approaches, but the TBM is more general and more flexible. The major characteristic of a belief function is that the belief given to the union of two disjoint events can be larger or equal to the sum of the beliefs given to each event individually. Other advantages include the possibility to represent every state of partial or total ignorance, as well as the possibility to deal with undefined, imprecise or missing values.

Probability approaches have been used for some time in learner modelling and, more relevant to our own approach, have started to be deployed in open learner models (see notably Zapata-Rivera & Greer, 2003; 2004). A comparison of a variety of mechanisms *from a learner modelling perspective* is a worthy task, but outside of the scope of this paper. Our interest – and the ultimate aim of this paper – is the effectiveness of an Open Learner Model based on these mechanisms to convey useful information to the learner.

The specifics of the inference mechanism of the xLM and a first study of its behaviour have been described in other publications (Morales, Van Labeke, & Brna, 2006a; 2006b). But a couple of issues need to be clarified here in order to make the xOLM fully understandable: the layered aspect of the Learner Model and the relations between belief, evidence and interactions.

## A Five-Layer Learner Model

The current structure of the xLM is shown in Figure 1. It is broadly hierarchical with the model of the learner's metacognitive state at the top and details of the learner's domain knowledge at the bottom.

The extended Learner Model builds updates and returns a "portrait" of the learner on a number of inter-related aspects of the learner. These aspects – described as layers of the model from now on – are represented in Figure 1. They are metacognition, motivation and affect, competency, conceptual and procedural errors (CAPEs) and the domain on which all of the previous layers rely.
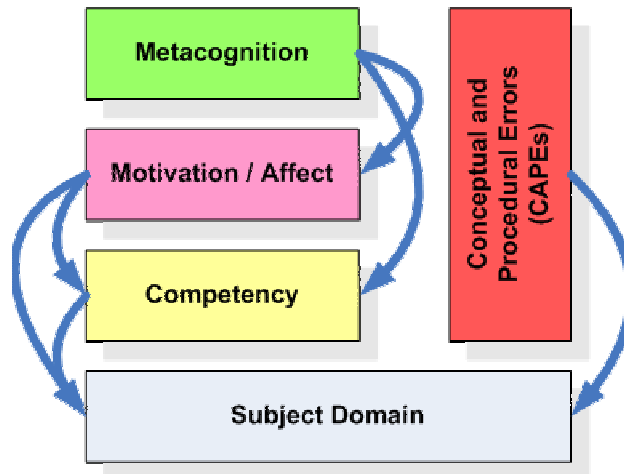


Fig. 1. The five layers of the Learner Model.

Several layers could be combined in order to extend the expressiveness of the Learner Model. Combination rules are indicated in Figure 1 by arrows, working in a waterfall fashion, with the subject domain as a mandatory base layer. The top-most layer in any combination, called hereafter the dominant layer, indicates the nature of the belief, whereas the underlying layers specify the context of the belief. For example, considering *interest* as a motivational factor, *solve problems* as a mathematical competency and *difference quotient* as a topic of the domain, the learner's *interest in his/her ability to solve problems on the difference quotient* is a motivation-related belief that the Learner Model could handle, as is the learner's *interest in the difference quotient*; the learner's *ability in solving problems on the difference quotient* is, on the other hand, a competency-related belief. Each combination of layers uniquely describes a particular belief held by the system on a particular ability of the learner; in this document; such description acts as an identifier of a belief and will be referred in this document as a *belief descriptor*.

This combination feature of the xLM is particularly useful for addressing the ambiguities that often occur when trying to determine the exact nature of a diagnosed ability of the learner. For example, assuming that some motivational factor – such as *interest* – is diagnosed in the learning environment, the xLM may not be able to detect the aspect of the activity to which the evidence is related. Solving such ambiguity in the xLM is done by considering all the possibilities of relating the evidence in the multi-layer model. Since the diagnosis is taking place within a particular learning object (e.g. associated with *difference quotient* as the domain topic and *solve problems* as a

competency), the xLM is using such evidence for both the beliefs on *interest in the learner's ability to solve problems on the difference quotient* (indicating a competency-related motivational trait) and on *interest in the difference quotient* (indicating a domain-related trait).

A major issue is how to feed the xLM with streams of information that relate to the categories of beliefs held about the learner. Currently, in the LeActiveMath project, domain knowledge (with topics related to differential calculus), misconceptions (CAPEs) and competency (adapted from the mathematical competencies described in PISA – see (OECD, 1999)) are derived from the metadata associated with exercises authored for the learning environment, affective information (such as satisfaction and liking) comes from the use of a self report tool available to the learner at the front-end of LeActiveMath, and metacognitive information (such as control and monitoring) is primarily obtained from the learner's use of the xOLM. Motivational factors (such as interest and confidence) are partly obtained by self report, and as part of the work done by the situation model to generate values for autonomy and approval (Porayska-Pomsta & Pain, 2004).

## Beliefs and Levels

The extended Learner Model represents beliefs by building a probabilistic distribution on different levels of abilities of the learner. These levels were initially extracted from the Programme for International Student Assessment – PISA (OECD, 1999), which defines mathematical skills in terms of competencies and competency levels (see Adams & Wu, 2002, p. 208 for more details) but, for consistency, this scheme has been extended to cover all layers in the xLM. We have used four levels to represent all the kinds of belief, mapping the variation between low and high ability in the relevant layer; they are abstractly termed *Level I*, *Level II*, *Level III* and *Level IV* and could be instantiated with different terminologies in the xOLM, depending on the layer they represent (see verbalisation below).

A main difference between the Transferable Belief Model (TBM) and classic probabilistic distributions is that, whereas the latter assign probabilities to each possible state of the world (e.g. the probability for the learner to be at *level II*), the TBM considers every set of possible states in the world (e.g. the learner to be somewhere between *Level I* and *Level III*). These sets – termed level sets in the extended Learner Model – represent all the possibilities to re-organise these four ability levels, while keeping their implicit order. They are represented by their width, i.e. the number of individual levels they include: the singletons (i.e. the sets only containing one of the level), the doubletons (i.e. sets containing two levels such as *Level I-II*), and 'tripletons' (the sets containing three levels such as *Level I-II-III*). The width of the sets can be interpreted as *pertinence*: the wider a set is, the less pertinent or focused the belief is.

The inference formalism of the xLM also makes use of two special sets. The *empty set* (i.e. containing no level at all) accumulates all information resulting from the combination of non-overlapping evidence (for example when the inference engine receives evidence about the learner's ability being both at *Level II* and at *Level IV*); it is interpreted as an indication of *conflict* in the belief: the higher the amount of information it contains, the less reliable the belief is. On the other hand, the full set (i.e. *Level I-II-III-IV*) represents the complete world; being a truism (the learner is *necessarily* between *level I* and *level IV*), it is therefore interpreted as *total ignorance*: the higher the amount of information it contains, the less certain the belief is.

A belief is therefore represented by a mass function, i.e. a distribution of the [0,1] probability on every possible set such that the sum of all mass (the term referring to the probability assigned to one particular set) is equal to 1. We now look at how such a mass function is obtained from the evidence.

## Beliefs and Evidence

There is a chain of subjective decisions/judgments made by the system, from a basic interaction in LeActiveMath generating an event, to its final interpretation by the xLM and presentation to the learner in the xOLM. At the various stages, different procedures operate, each of them dealing with a particular type of information (Figure 2).

Evidence for learner modelling comes into xLM in the shape of events representing what has happened in the learner's interaction with educational material and the rest of the learning environment. Several types of interaction are currently taken into account: an exercise is finished by the learner, one particular step of an exercise is performed by the learner, a self-assessment is made by the learner, etc. These events are qualitatively and/or quantitatively described by different kinds of information (termed attributes in this document) such as the exercise difficulty, associated competency and competency level, the overall performance of the learner on the exercise, the misconception – if any – associated with an incorrect answer from the learner.

Events are raw evidence for the learner's ability on particular layer that need to be interpreted in order to produce a mass function that can be incorporated in the inference engine. Two categories of events are accordingly interpreted by xLM: behavioural events (reporting what the learner has done or achieved such as performance in an exercise) and diagnostic events (reporting a judgement of learner levels produced by some diagnostic component of the learning environment, such as the affect self-report tool). The latter is much simpler to interpret since an estimation of the learner's ability is already given by the diagnostic component. Details for the interpretation mechanisms are to be found in (Morales, Van Labeke, & Brna, 2006a), but the result is a *mass function*, i.e. a distribution of the assumed learner's ability that could explain the reported performance. This mass function is now considered by the xLM as evidence for the underlying learner's behaviour.

This piece of evidence is incorporated in the appropriate belief – together with any other piece of evidence obtained by the interpretation of previous events – using the combination rule of the Transferable Belief Model. Internally, this belief is represented by a mass distribution similar to the evidence's numerical interpretation.
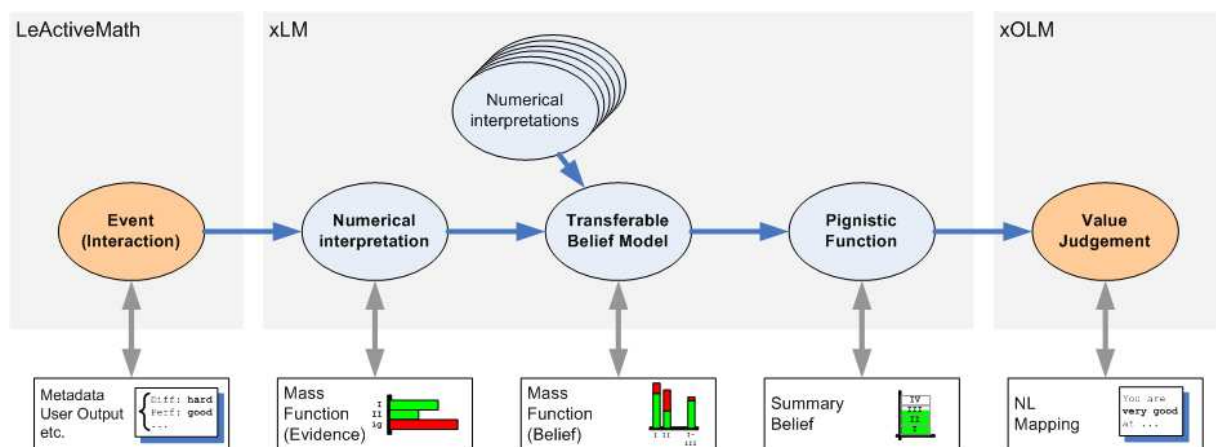


Fig 2. Overview of the belief building/updating process of the xLM.

From this broad mass function, expressing the system's belief in terms of probability on sets of learner's level of ability, the xLM can now apply a much narrower decision making process called the *pignistic function*. This pignistic function breaks down the masses associated with each set into a simple probability distribution on individual levels (by comparing sets' intersections and unions). The decision making process can be taken a step further by summarizing the four-level distribution into a single continuous value, the *summary belief*[1]. This value can be seen as the end-product of the xLM; it is available for external components as a straight-forward guess about the learner's current ability (for example the Tutorial Component of the learning environment could use it in order to adapt its delivery of the learning materials to the learner). From an Open Learner Model perspective, the summary belief is used in the xOLM – given appropriate verbalisations or externalisations – to present the learner with a readable value judgement on his/her ability.

At one basic level, it can be argued that the only important information – from the xOLM point of view – is the summary belief (i.e. the final statement made by the xLM) and the facts (i.e. the interaction between the learner and the system) used to reach that statement. But, as can be seen from the modelling process depicted in Figure 2, various decisions and interpretations are made at various stages in the process. Therefore, if the aim of the xOLM is to support the learner in understanding why and how the Learner Model reached its conclusion, then it becomes important that most – if not all – of these intermediate steps in the modelling process are presented to the learner.

However, this means that we need a mechanism to control the delivery of all this information in a way that maintains its significance. In the extended Open Learner Model, this mechanism is inspired by the Toulmin Argumentation Pattern which, besides its (superficial) simplicity, does indeed provides us with the possibility for managing both the exploration of the Learner Model and the challenge of its judgements. It also quite nicely supports a dynamic reorganisation of the evidence that helps to establish and clarify the justifications presented to the learner.


## THE TOULMIN ARGUMENTATION PATTERN

Toulmin (1959) laid out a theory and model of argumentation that went against the accepted model of Formal Logic at the time but was, in his opinion, better equipped to manage the complexities of legal argument. It promoted a more abstract approach in dealing with arguments which lends itself to the kind of argumentation that occurs in daily life[2]. We have represented Toulmin's model as the Toulmin Argumentation Pattern (TAP).

In brief (see Figure 3), *"Data"* ( or *"Grounds"*) are the facts and information that are the reason for the claim in the first place – a reasoned beginning; *"Claim"* is the position on the issue, the purpose behind the argument, the conclusion that the arguer is advocating; *"Warrant"* is the component of the argument that establishes the logical connection between the data and the claim, i.e.

---

[1] There are several ways of producing such a summary: average, highest probability, etc. There is also no reason why all methods could not be used in different contexts, for example by taking into account the current motivation of the learner. This decision-making process is basically a bet on the most 'appropriate' level of ability that characterise the learner; as such, several ways of balancing risk and gain could be envisaged. This is an issue that will be considered in future development of the project. The current implementation of the summary belief is a simple – but conservative – weighted average.

[2] Various AI-ED systems have adapted Toulmin to their purposes – for example (Suthers, Weiner, Connelly, & Paolucci, 1995).

the reasoning process used to arrive at the claim; *"Rebuttal"* is any exception to the claim presented by the arguer; *"Backing"* is any material that supports the warrant or the rebuttal in the argument; *"Qualifier"* represents the verbalisation of the relative strength of an argument, its soundness.
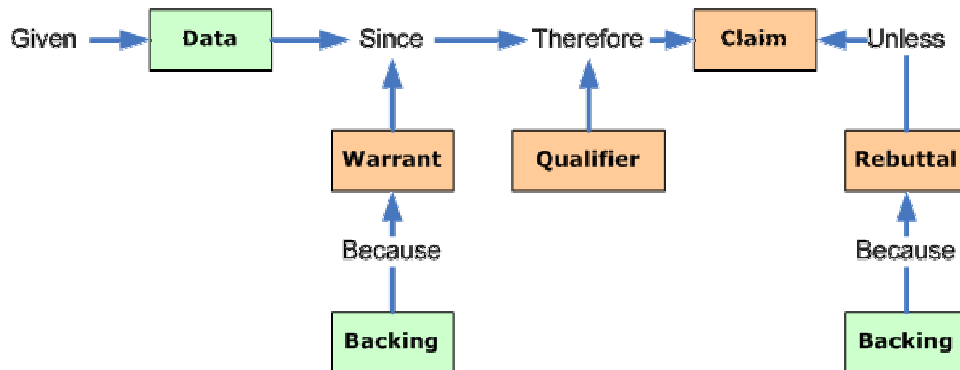


Fig. 3. The Toulmin Argumentation Pattern.

Applying the Toulmin Argumentation Pattern in the context of the extended Open Learner Model has been done by the following mapping between each element of the TAP and the extended Learner Model's internal representations (see Figure 4):

- The *Claim* is associated with the summary belief, i.e. a short, straightforward judgement about the learner's ability on a given topic (i.e. *"I think you are Level II"*).

- The *Data* is associated with the belief itself, represented both by its pignistic function (i.e. its simplest internal encoding) and its mass function (its most complex internal encoding).

- *Warrants* are associated with the evidence supporting the belief, represented by their mass function. There will be one warrant for every piece of evidence used by the Learner Model to build its current belief.

- *Backings* are associated with the attributes, both qualitative and quantitative, of the event used by the Learner Model to build the mass function of the relevant piece of evidence. Backings and Warrants come in pairs.

In the approach currently implemented in the extended Open Learner Model, we are not considering any *Rebuttal*, in the sense that all the evidence gathered by the Learner Model is evidence for the learner's abilities. Gathering and manipulating evidence against the learner abilities is an aspect worth investigating. As for the usage and definition of the *Qualifier*, it is an issue that we have decided not to consider explicitly. The Learner Model being a probabilistic model, information about certainty and plausibility are implicitly included in the belief (i.e. the data); extracting this information in a separate entity (i.e. subject to query and challenge from the learner) seems not to improve the usefulness of the approach. To summarise, an argumentation by the xOLM about some belief can be handed over to the learner according to the following general pattern:

*Given this belief since it is supported by the following list of mass functions, each of them obtained because of what you did/said during this particular activity, I therefore claim that your ability can be described by this summary belief.*
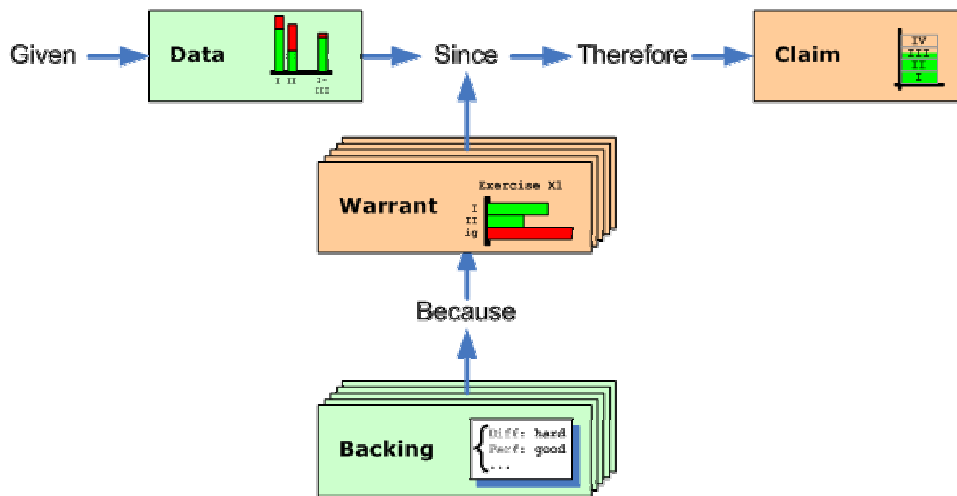
Fig. 4. Applying Toulmin Argumentation Pattern to the xOLM.

## Exploring and Challenging the Learner Model

The Toulmin Argument Pattern can be easily and helpfully used to control and contextualise the learner's interaction with the xOLM, or to be more precise, with the xOLM's judgements and justifications (i.e. the xOLM argumentation). Two aspects of it are of importance here: the exploration of the judgement and its challenge.

A mock interaction between a learner and the xOLM can be seen in Figure 5. It represents steps in the exploration of a belief – or, more accurately, of its justification. Every step of the discussion is made manifest as a result of the learner requesting explanations as to why the xOLM made its judgement. The initial stage of the exploration takes place when the learner questions the judgement made by the xOLM (*"Why do you think I'm Level III at my ability to Solve Problems on Difference Quotient?"*). The question is localised clearly on the claim, allowing the xOLM to try to justify it by "expanding" the pattern and presenting a deeper justification of the claim, i.e. the data. If the learner is still questioning the judgement, then the data can in turn be expanded to present the first element of the warrant, which can also be expanded to show the reasons (i.e. the backing) for such evidence, etc. At every stage of the exploration, both the xOLM and the learner know exactly what has been questioned, what the outcome of the question was and what remains to be justified.

The Toulmin Argumentation pattern can also help us to understand the source – and therefore the reason – of a disagreement between the xOLM and the learner. In the context of such an argumentation pattern, different interpretations of a challenge of the judgement made by the xOLM can be considered, depending on the "target"[3] of the challenge.

---

[3] Different interpretations of challenging the Data could, and should be devised. Some alternatives have already been tried but all of them relied on making explicit the processes used to build the belief, i.e. the Transferable Belief Model. This was significantly outside of the scope of the xOLM. For the near future, we decided not to let the Data be subject to challenge, despite the "inconsistency" of the approach.

| | |
|---|---|
| **Louise** | Show me what you think about my ability to **Solve Problems** on **Difference Quotient**? |



| | |
|---|---|
| **OLM** | I think you are **Level III**. |
| **Louise** | Why do you think I'm **Level III**? |



| | |
|---|---|
| **OLM** | Because that is the dominant trait in my view of your behaviour. |
| **Louise** | How did you reach this conclusion? |



| | |
|---|---|
| **OLM** | Here are the two pieces of evidence supporting my claim. |
| **Louise** | I don't understand how you obtained that particular piece of evidence. |



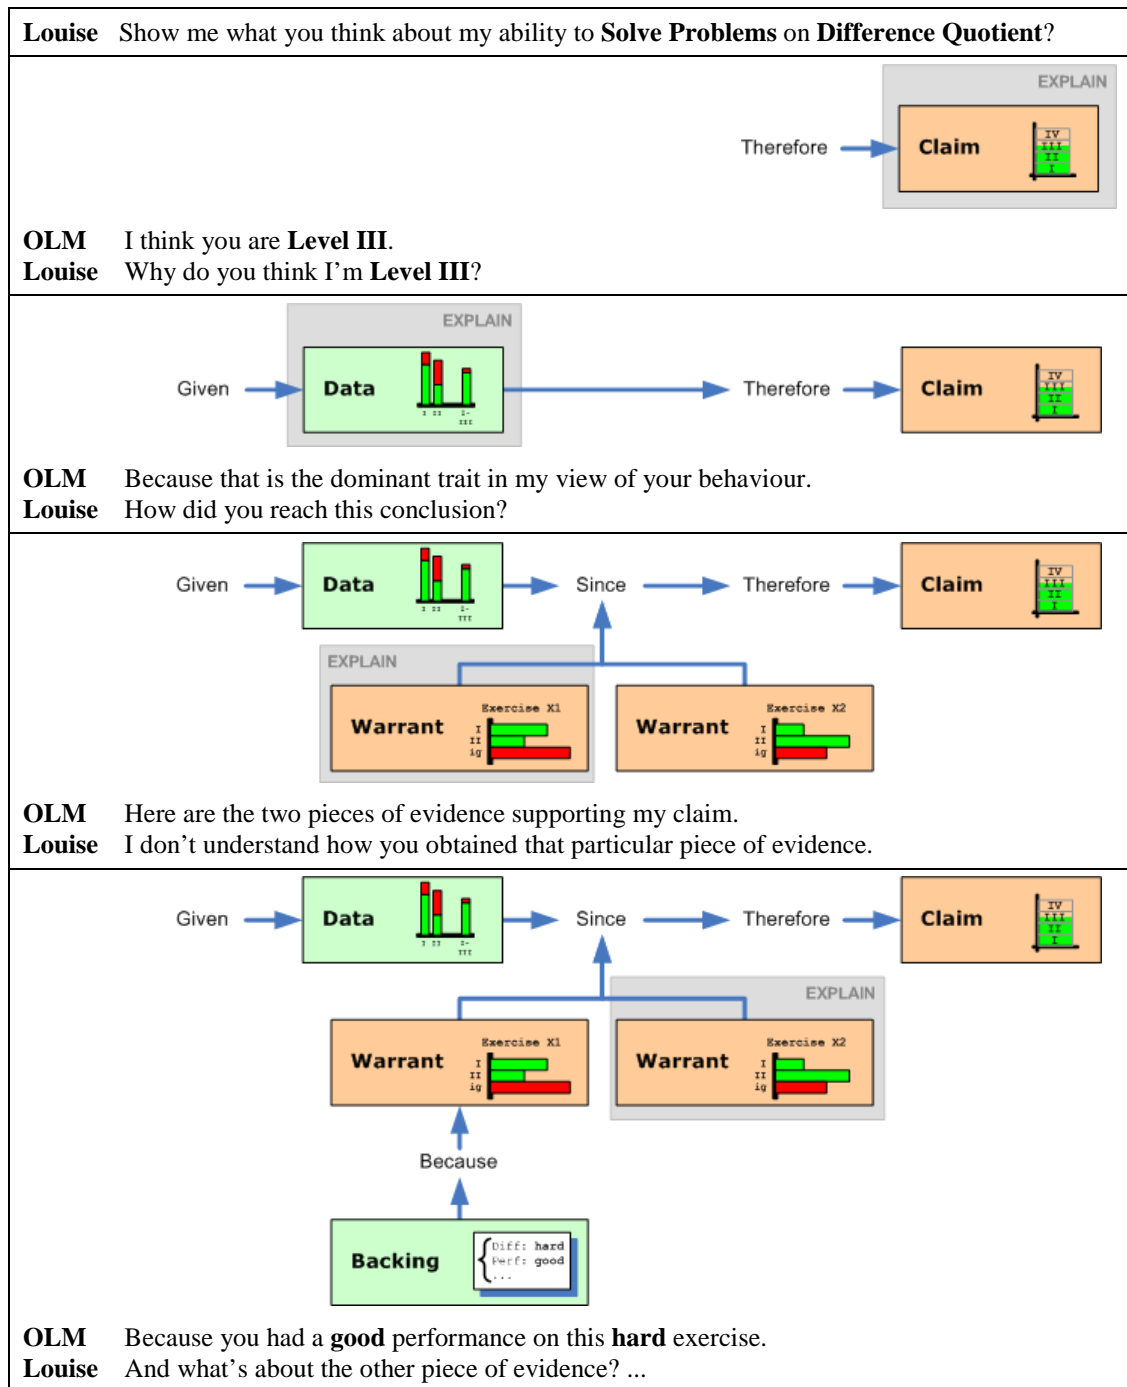| | |
|---|---|
| **OLM** | Because you had a **good** performance on this **hard** exercise. |
| **Louise** | And what's about the other piece of evidence? ... |

Fig. 5. Using the Toulmin Argumentation Pattern to explore a judgement made by the xOLM. The element of the pattern currently questioned by the learner is indicated by the 'explain' tag.

- When challenging the claim made by the xOLM, learners state their disagreement with the overall judgements. Such a challenge is not merely on the evidence used to sustain the judgement but rather on the way that evidence is combined and interpreted, making grounds for subjective factors that may not have been – or could not – considered (*"I understand your point, but I still think that you are missing something and that I am not that good on the chain rule"*).

- Challenging a warrant means that the learners disagree with this evidence being used by the xOLM to reach its conclusion. The disagreement does not apply to the justification of the evidence itself (as it will be done by challenging the backing) but the presence of the evidence itself in the justification of the belief. This possibility is offered in order to take care of situations that are usually outside of the diagnosis (*"I misunderstood the goal of the exercise; I don't think you should consider it as an indication to my ability"*).

- Challenging the backing of a claim consists in questioning the validity of one – or many – of the attributes used by the xOLM to interpret the evidence. It could be a qualitative attribute (*"I don't think that my performance was so low on this exercise"*) or a quantitative (*"I don't think this exercise was so easy"*).

## Re-clustering the Set of Evidence

A final interesting property of the Toulmin Argumentation Pattern to exploit is its ability for nesting arguments. A belief, at its basic understanding, is nothing more than the combination of all the evidence supporting it. The justification of a belief could therefore be seen as just the exploration and appropriation of every piece of evidence, one by one and in the order in which they occurred. But this approach suffers from a very obvious limitation: the sheer volume of evidence that could be potentially gathered by the extended Learner Model. Hoping that the learner will be able to access and assimilate them one by one is an assumption that will clearly not stand.

One way of addressing this problem is to re-cluster the set of evidence according to some criteria, partitioning it into several intermediary judgements, which in turn, can be represented by a Toulmin Argumentation Pattern on their own (or an abridged version). The original claim, instead of introducing a 'flat' set of warrants/backings, can now by expanded with sub-claims introducing the new sets of evidence thus reorganised (as illustrated in Figure 6).

In the current implementation of the extended Open Learner Model, two criteria are used for such re-clustering: the use of one of the attributes of an event – deemed significant for the belief's understanding – to divide the evidence space in meaningful categories and the use of the evidence's *impact factor* to discard the less important pieces of evidence supporting a belief.

For example, a belief about a competency such as the learner's ability to *solve problems* will be obtained mostly from evidence related to their performance in doing exercises. It is possible therefore to consider the justification of such a belief by presenting, on the one hand all pieces of evidence showing a low performance and, on the other hand, those showing a high performance (low and high performance being specified by a particular threshold that could be arbitrarily or dynamically determined). With this example, the *sub-claim 1* in Figure 6 is representing a "low evidence" claim, associated with the relevant pieces of evidence, whereas the *sub-claim 2* gives access to the "high performance" partition of the initial sets. Because of the heterogeneous nature of the sources of evidence used in LeActiveMath, it may often be the case that individual pieces of evidence does not fit

in one of the space created by the partition (for example evidence related to a challenge made by the learner). This is why an extra sub-claim, universally labelled "others", will be introduced by the partition to cover pieces of evidence not related to the partition attribute.
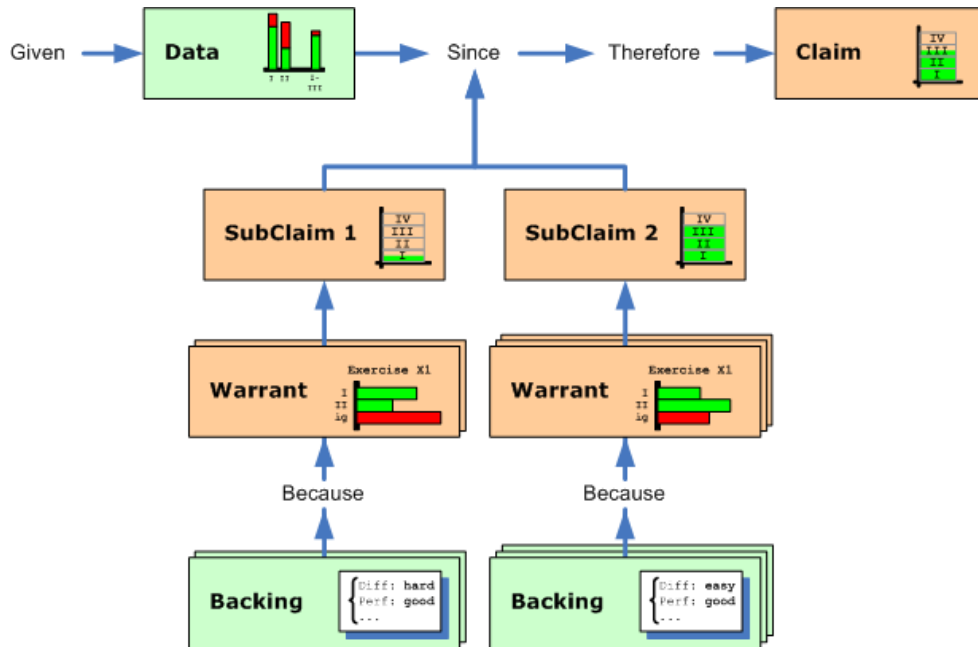


Fig. 6. Re-clustering the evidence in a Toulmin Argumentation Pattern.

The *impact factor*[4] of a piece of evidence is an indicator of the importance of that particular piece in shaping the belief it contributes to. A piece of evidence with a high level of total ignorance for example will have a very small impact on the final belief. But, since the learner modelling mechanism used in the xLM is using a gradual decay to weaken older pieces of evidence in favour of more recent ones, even a highly significant piece of evidence will become more and more unreliable over time. It means that the evidence space could be partitioned again, given a certain threshold, between *relevant* (i.e. high impact factor) and *irrelevant* (i.e. low impact factor) evidence.

It has to be noted that a re-clustering of the evidence – whatever the criteria used for the partition – means that the chronological order in which evidence were introduced in the belief is not guaranteed anymore. At this stage of this project, it is not clear if the loss of such temporal information has an impact on the learner or not. Further evaluation comparing both approaches will be needed to clarify this issue.

---

[4] The impact factor of a piece of evidence is computed by the xOLM on the basis of the *current* value of the belief and is obviously changing every time a new piece of evidence is incorporated in the belief. It is basically obtained by computing the distance between the belief with and without that particular piece of evidence.

## THE INTERFACE OF THE XOLM

The Graphical User Interface (GUI) of the xOLM is built around two distinct parts (see Figure 7[5]):

- The top half of the GUI contains the main source of interaction between the learners and the xOLM. The nature and exact content of every external representation depends on the context and stage of the dialogue between the learner and the xOLM and is described in details in the following sections. In general, this part of the interface will include, on the left-hand side, the Argument view (labelled **A** in the figure), i.e. a dynamic graphical representation of the Toulmin Argumentation Pattern, and on the right hand side, an external representation (labelled **B** in the figure) of the element of the Toulmin Argumentation Pattern selected by the learner (i.e. the current topic of discussion).

- The bottom part of the interface (called the Dialogue View, labelled **C** in the figure) contains a scrolling text pane used to provide the learners with a transcription of their dialogue with the xOLM. This pane has two purposes: first maintaining a log – in a readable format – of the history of the learners' interactions and, second, complementing – on demand – the external representations displayed in the upper-part of the xOLM with additional verbal explanations.

The usage of xOLM can be summarised into three successive tasks, representing the three stages of the dialogue between the learners and the xOLM.

- The first stage is the exploration of the content of the Learner Model, i.e. the navigation through the various belief held in the system and the selection of a particular topic of discussion.

- The second stage is the justification by the xOLM of the judgement made on the topic selected.

- The last stage is the challenge by the learners of some elements of the justification presented by the xOLM.

None of these stages are imposed on the learners – it is left to their own initiative (within the limit of the mixed-initiative supported by the xOLM). Nonetheless, they have to be operated in this particular order, i.e. a belief has to be selected for justification before the learners could challenge the xOLM's decisions.

In order to describe these stages and their dedicated interface in details, let's consider a hypothetic learner, named Louise, and observe her going through the various stages of her interaction with the xOLM. Before accessing the xOLM, Louise had a go with the learning environment by performing an exercise (a multiple-choice question) on the difference quotient:

*Compute the difference quotient of the function f(x) = 2x+5, between the points x=1 and x=3.*

She tried two wrong answers (4/3 and 4) before giving the right answer (2). At the conclusion of the exercise, Louise is prompted to have a look at her own model in order to assess her achievements with the system and therefore launches the xOLM.

---

[5] Full-screen colour version of all the screenshots in this paper can be found online at http://www.scre.ac.uk/personal/pb/ijaied07/
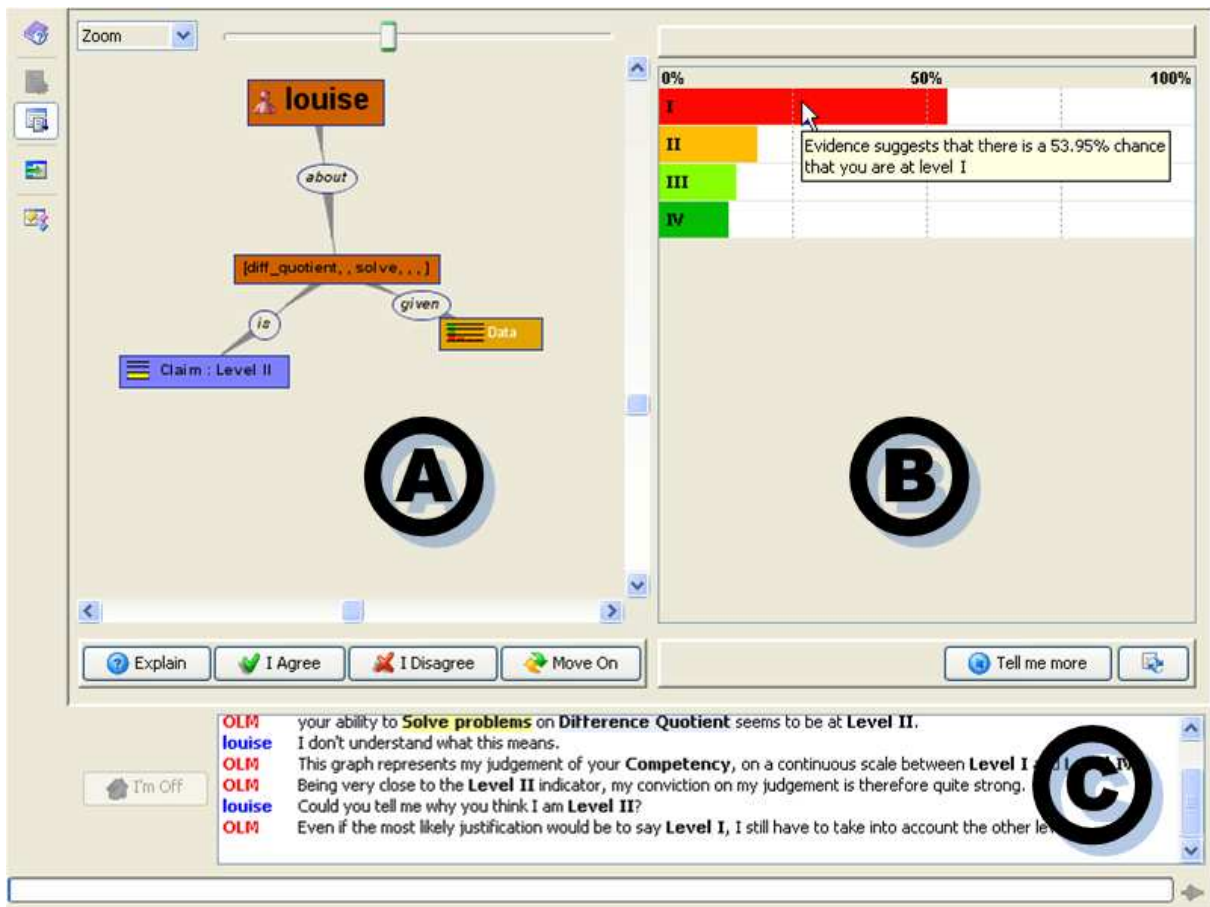
Fig. 7. A snapshot of the xOLM Graphical User Interface.

## Exploring the Content of the Learner Model

When the xOLM is launched, Louise is initially presented with the Descriptor view, whose interface she will use to specify the topic of her discussion with the xOLM. Upon the selection of an adequate topic, the xOLM will retrieve the corresponding belief from the system and present it to Louise, using both the Claim view to display the initial summary of the belief and the Argument view to materialise the justification of such a claim. These three external representations are described in turn in the following sections.

### The Descriptor View

The current implementation of the Descriptor view applies a very straightforward interface (see Figure 8), supporting Louise in explicitly constructing a belief descriptor as a way of indicating which topic of discussion she wishes to initiate.
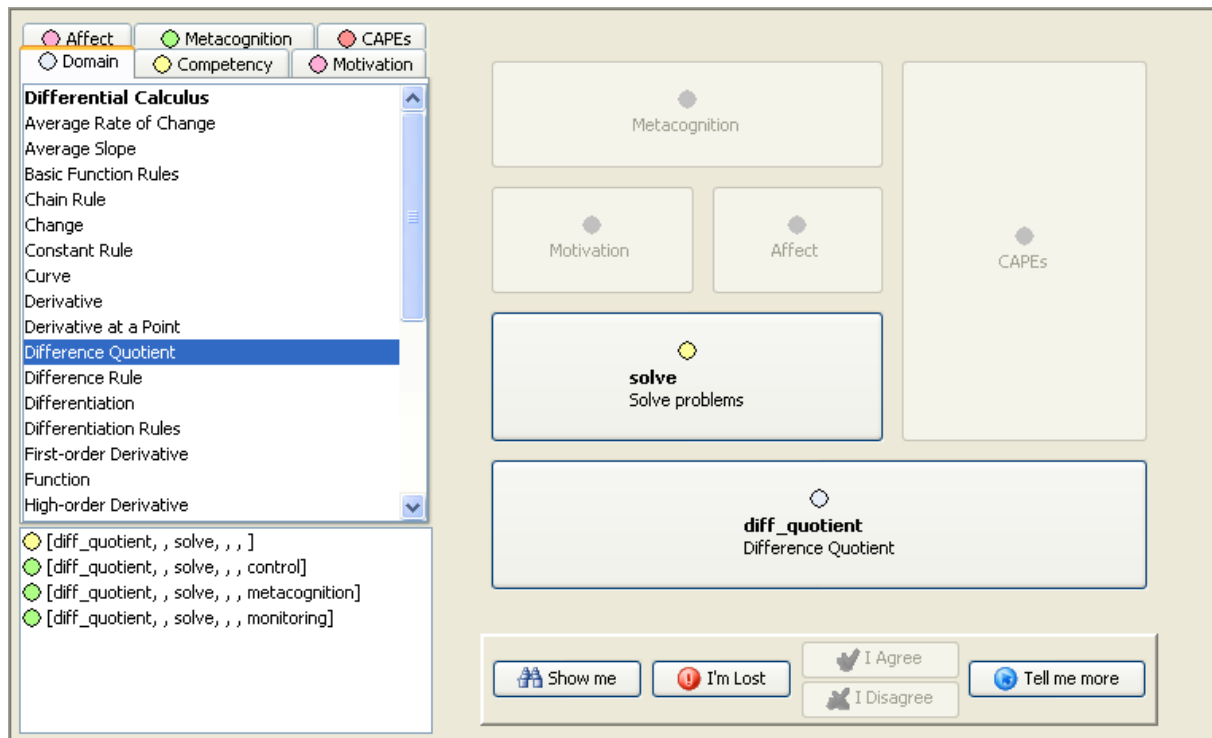
Fig. 8. The Descriptor View used to specify the belief to explore in the xOLM.

The right-hand side of the view contains a couple of placeholders arranged in a way to reflect the organisation and dependency of the layers in the Learner Model, as presented in Figure 1: the Domain topic at the bottom, supporting on the one hand the Competency topic, the Affect and Motivation topics and finally the Metacognition topic, on the other hand the CAPEs.

The intuitive idea is for Louise to pick up the topics of their interest and manually build a descriptor to submit to the xOLM. In order to do it, the left-hand side of the view presents a couple of tabbed lists, each of them containing the topics for each of the layers of the Learner Model (Domain, Competency, etc.). By simply selecting one topic from a list and dropping it into the proper placeholder, the descriptor is built step-by-step.

Once Louise is happy with the belief she want to explore (in our example, she is looking for a judgement on her ability to *solve problems on the difference quotient*, both competency and domain topics being associated with the exercise she underwent before launching the xOLM), a simple click on the *Show Me* button will validate her choice and prompt the xOLM to present its judgement on this belief, and, when pushed further on, why it is the case.

A couple of points are worth mentioning about the Descriptor view. As mentioned previously, there are constraints on how a belief descriptor is structured (no CAPEs and Competency together, a domain topic always present, etc.). It is still not clear what is the best way to make learners aware of these restrictions: either by preventing any wrong configuration before or indicating such problems after validation. In the current implementation, we opted for the second approach: in the case of a wrong configuration, the xOLM will make a statement such as *"Sorry, I don't understand your question"*, providing information about why the question was badly formulated. At every stage of

building a belief descriptor, the *Tell Me More* button provides learner with contextual help about what they are currently doing (e.g. the description of the belief currently under construction, the definition of any topic selected from one of the layers, etc.).

## *The Claim View*

The Claim view is the initial externalisation of the belief selected by Louise. Its main purpose is to provide her with a succinct and immediate overview of her overall ability, i.e. the *summary belief*. Two complementing pieces of information are simultaneously conveyed with this external representation (see Figure 9): the discrete and the continuous values of the summary belief.
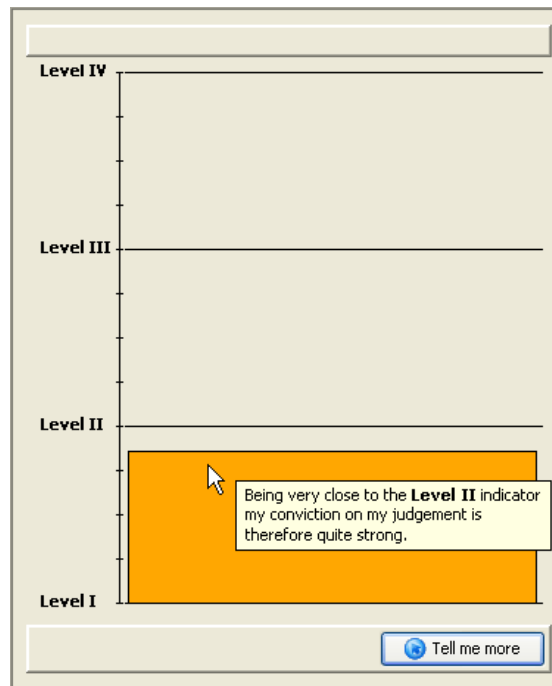


Fig. 9. The Claim view displaying the summary belief.

The discrete value corresponds to the ability level proper and is given both by the bar's proximity with the appropriate level indicator (from *Level I* to *Level IV*) and by its colour (throughout the xOLM, each of the four levels is associated with a distinct colour, using a street-light analogy: red, orange, light and dark green). The continuous value of the summary belief is given by the bar chart itself, it's size and proximity to the indicators reflecting both the dominant level (i.e. close to the *Level II* indicator), the tendency of Louise's ability (i.e. *Level II* but on the lower quadrant) and the conviction of the xOLM on its judgement (i.e. the closer the bar is to an indicator, the more certain the xOLM is about its own judgement).

A tooltip on the bar chart provides Louise with a direct and succinct description of the summary belief but, as with every external representation in the xOLM, the *Tell Me More* button allows Louise

to access a more detailed analysis (in terms of dominant level, tendency and conviction), displayed in the Dialogue view.

## The Argument View

The presentation of the *summary belief* – and its ultimate justification step-by-step – is controlled by the Argument view, displayed side-by-side with the Claim view. The purpose of the Argument view is twofold: to provide Louise both with a representation of the logic of the justification of the judgement made by the xOLM and with an interface to navigate between the various external representations associated with each component of the justification.
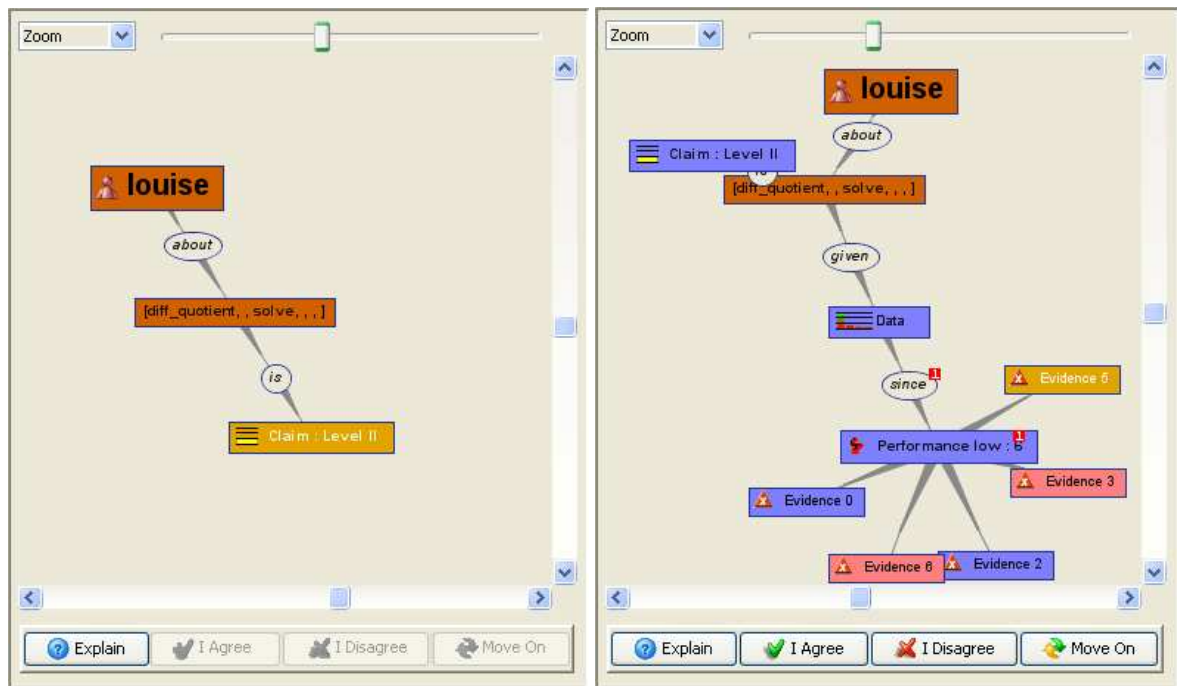


Fig. 10. Two snapshots of the Argument View, at different stages of the argumentation. On the left-hand side only the claim is initially presented to the learner; on the right-hand side, part of the pieces of evidence associated with a sub-claim (partition according to performance) are expanded.

As mentioned before, judgement about learners are organised in a Toulmin-inspired argumentation pattern, which is represented as such in this view, under the appearance of a dynamic and interactive graph. Each of the nodes of the graph is associated with one of the elements of the argumentation pattern: the claim node associated with the summary belief, the data node associated with the belief itself, the warrant and backing nodes associated with individual evidence, etc. The shapes, colours, labels and icons of the nodes are appropriately designed in order to provide a quick and unambiguous identification of the corresponding element. These nodes are also reactive to learners' interaction and act as a shortcut to access the appropriate external representation, which will be immediately displayed beside the Argument view. Thus, selecting the claim node – or a sub-claim,

if any – will open the Claim view and update its content accordingly; selecting one of the evidence nodes will open the Warrant/Backing view and update its content based on the selected evidence, etc.

Some intermediary nodes, not reactive to learner's interaction, have been added to introduce meaningful associations between the important parts of the graph (e.g. *"about"*, *"given"*, etc.); they are mostly linguistic add-ons for improving both the readability and the layout of the graph.

The content of the Argument view is expanded on demand, following Louise's request for further justification about the xOLM judgement (by using the *Explain* button, seen below the graph). Figure 10 shows two stages of the expansion of the argument pattern. On the left, the initial argument pattern is presented, corresponding to the current stage of the discussion, the target of the discussion (as selected by Louise in the Descriptor view), as well as the claim made by the xOLM on this belief, are the only visible nodes on the graph. On the right, the argument pattern has been further expanded by Louise requesting more justifications for the judgement. The data node and one sub-claim (*"Performance low"*) can be seen on the graph, as well as a partial cluster of evidence supporting it (the *"Evidence N"* nodes).

## Justifying the Content of the Learner Model

Once the claim of a belief has been displayed, one of the aims of the xOLM can now be considered achieved: to provide Louise with some information about her current estimated ability on a particular topic. Louise could feel satisfied by the answer given by the system and go on to do something else (e.g. exploring a different belief by swapping back to the descriptor view, closing down the xOLM and going back to the learning environment for more exercise, etc.).

However, she could feel puzzled by the judgement made by the xOLM and decide to explore further down its justifications for such a claim. This is basically done by selecting one of the node of the argument pattern and using the *Explain* button (see Figure 10) in the Argument view. Such a move will force the expansion of the argumentation pattern into further details: when requesting explanations about the claim node, the data node will be expanded; when requesting explanations about the data, the set of evidence will be expanded (or the sub-claims – if any – clustering the evidence); etc. By exploring in turn each of the newly expanded nodes, Louise has now access to deeper justifications for the judgement. Two elements of the argumentation pattern are particularly important: the data node (and its associated external representations the Data view) and the individual pieces of evidence (and their associated Backing/Warrant view); both views are described in this section.

### The Data View

The Data View is used by the xOLM to support the exploration and analysis of the part of the Toulmin Argument Pattern associated with the belief itself, as held by the extended Learner Model. The data view itself contains two different external representations, associated with the two internal representation of the belief (the pignistic function and the mass distribution), each of them detailing a bit further the meaning of the belief.

The Pignistic External Representation (Figure 11, left-hand side) is a natural extension of the summary belief, by displaying the normalised distribution on the four singletons (*Level I*, *Level II*, *Level III* and *Level IV*). Its purpose is to provide Louise with a first shallow justification for the summary belief stated by the xOLM. In the situation described in the figure, its message is *"I said you*

*are Level II because, although the dominant trait for my belief is level I, evidence supporting higher levels is non-negligible".*

The Mass Distribution External Representation (Figure 11, right-hand side) extends the expressiveness – and the complexity – of the pignistic function by presenting the complete distribution of information across all the level sets. The empty set (labelled 0 in the figure) and the full set (labelled I, II, III, IV) are distinctly materialised in the distribution, to put an emphasis on the conflict and total ignorance associated with the belief, two of the most important pieces of information associated with the mass distribution[6].

Even if the mass distribution is the result of the accumulation of all the evidence gathered over time, it does not explain how individual pieces of evidence had an impact on the resulting belief. To be able to access this level of description, Louise has to keep asking for further explanation about the judgement made by the xOLM, in order to reach individual pieces of evidence and their associated Warrant/Backing View.
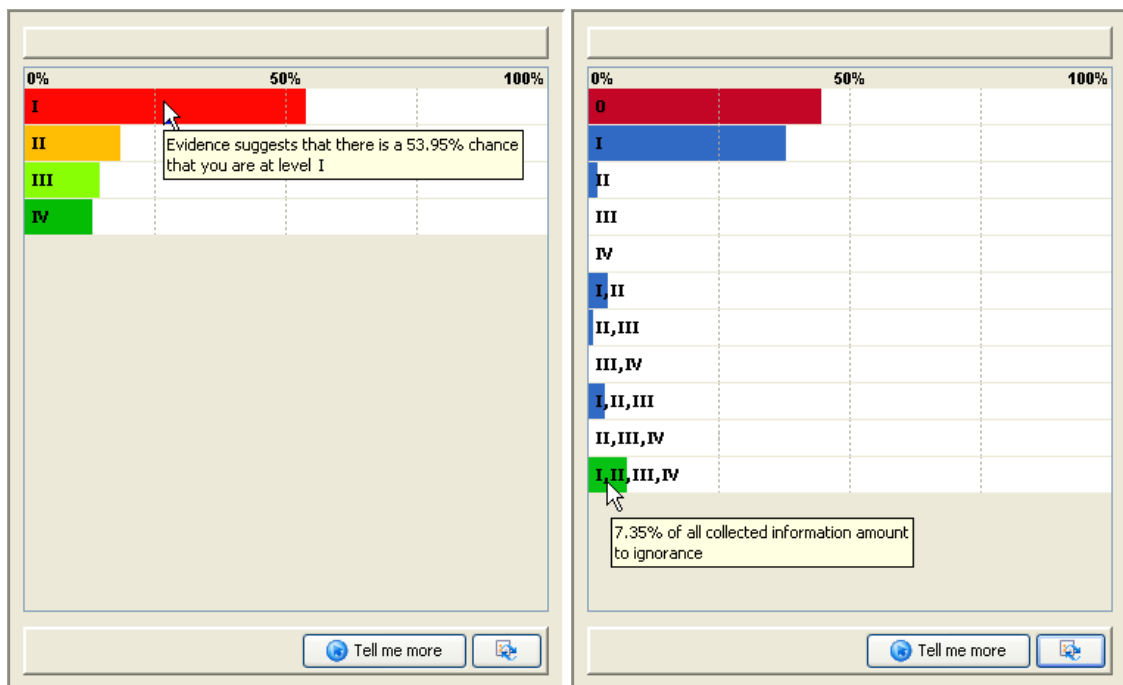


Fig. 11. The Data view and its different external representations; on the left-hand side the pignistic function, summarizing the mass function on the four levels; on the right-hand side the complete mass function, highlighting the conflict and total ignorance.

---

[6] Alternative ways of presenting the mass distribution are currently being investigated, in order to highlight other important features of the mass distribution, notably the pertinence of a set (i.e. wide versus narrow sets). Sorting the sets according to their pertinence (width and probability) is one possibility but the bar chart representation may not be optimal. Designing appropriate Law-Encoding Diagrams that will enhance learners' conceptual understanding – such as Probability Space Diagrams, see (Cheng & Pitt, 2003) – is a more ambitious but challenging goal.

### The Warrant/Backing View

The Warrant/Backing View (Figure 12) is a dual external representation intended to display information about an individual piece of evidence justifying the belief.

On the top part of the view lies the mass distribution representing the numerical interpretation of the evidence. Similar in its reading to the mass distribution view of the belief (see the Data view above), it put an emphasis on the sets that are the most likely to be the explanation of the "performance"[7] of the learner, as diagnosed during the event that generated this evidence.
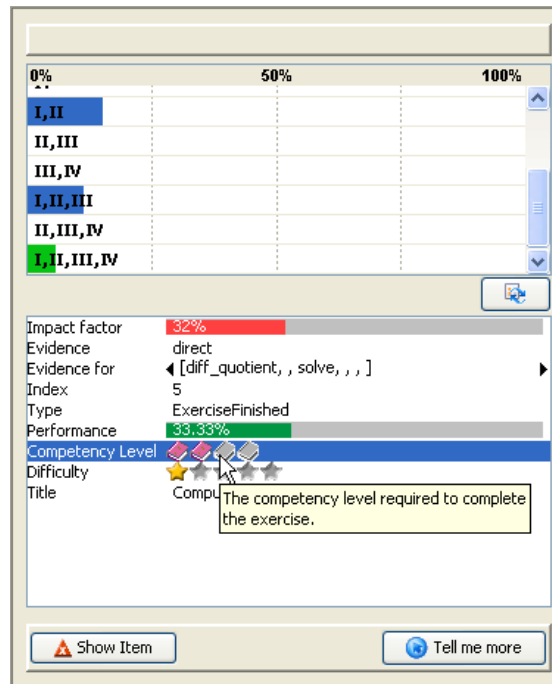


Fig. 12. The Warrant/Backing view, displaying the mass distribution generated by the corresponding event.

On the bottom part of the view, the qualitative and quantitative attributes categorising the event that generated this piece of evidence is presented to Louise. The number, nature and content of these parameters will vary from one piece of evidence to another one, depending on the nature of the interaction in the learning environment. The figure represents a piece of evidence generated by Louise finishing the exercise mentioned before. It gives an indication of her *performance* (33%), as diagnosed by the exercise component of the learning environment; of the *difficulty* of the exercise (easy, i.e. one star among five) and its expected **competency level** (level II, i.e. two books among four), both specified by the author of the exercise; of the *impact factor* (32%) of this piece of evidence, as computed by the Learner Model; etc. Tooltips are provided for both the mass distribution and the attributes list, giving a short description of the nature of the elements. As usual, the *Tell Me More*

---

[7] The term "performance" is here to be taken in a broad sense, as the underlying event could be of many different types: exercise finished by the learner, self-report about their affective state, propagation of evidence from a close topic, etc.

button can be used to access deeper and explanations about the information involved in this external representation.

Externalising the individual pieces of evidence, so that it allows Louise both to understand its content, to remember the context in which they were generated and to enable her to easily translate between evidence, attributes and belief, has proved to be a difficult task. For the current implementation of the view, we opted for a cheap and easy solution: a simple table displaying attributes and values, using – when appropriate – graphical widgets such as meters and icons to improve its readability. Recontextualising the events within the xOLM relies totally on the functionalities of the learning environment it complements. Whatever the nature of an event presented in the Warrant/Backing view, it is associated with one of the learning object of LeActiveMath (e.g. the exercise that has been achieved, the exercise at the end of which Louise expressed her satisfaction, etc.). She therefore has always the possibility to display this object again, in a dedicated LeActiveMath browser, and on her own initiative, the "Show Item" button. Whether this display is static (i.e presenting the current stage referenced by the event) or is offering sophisticated features such as a replay mechanism is a matter of the features available in the learning environment.

## Challenging the Content of the Learner Model

At any time during her exploration of a belief and its justification, Louise could decide to challenge the xOLM on its judgement. Performing this action is enabled by the three buttons (*I agree*, *I Disagree* and *Move On*) that can be seen under the Argumentation view (see Figure 10).

The notion of challenge in the xOLM is a particular one. Displaying the claim of a selected belief is a neutral operation from which the xOLM does not draw any conclusion. Louise could explore every single belief in her Learner Model without interference. But as soon as she starts questioning the xOLM for further justification, it assumes some doubt about the judgement presented and expects some form of acknowledgement of these doubts by Louise. This is a situation where user-initiative is constrained by the system: it will prevent the exploration of any other topic until the discussion around the current one has been concluded, one way or another.

Three possibilities are offered to Louise to resolve the discussion. She could finally abide to the justifications presented by the xOLM and accept its judgement (*I agree*). This decision is used by the xOLM to generate a new piece of evidence strengthening its belief. She could decide to end the discussion without committing herself to accept or reject the judgement (*Move On*). This decision is used by the xOLM to highlight this topic of discussion as unresolved; in future session, it will suggests Louise to have a look again at the same topic, in the hope that the issue could be closed. Or she could decide to reject the judgement by challenging some aspect of its justification (*I disagree*). In this case, Louise is presented to the Challenge view in order to express her claim.

### *The Challenge View*

There are three elements of the argumentation pattern that the learner could challenge: the claim (*"I am not Level II"*), one of the pieces of evidence (*"I don't accept this evidence"*) and one of the attributes of one of the events (*"I don't think that exercise was easy"*). Since they require different tasks and have different impacts in the Learner Model (i.e. on how a new evidence is being generated and incorporated in the belief), each of these challenges has its own interface for the learner to perform

it. Figure 13 shows the interface for challenging a claim, asking Louise to provide her own alternative statement on her ability.

The current interface for challenging the xOLM judgements is a short-term compromise, supporting a shallow challenge without any formal negotiation. Louise is expected to explicitly provide the xOLM with a qualification of her disagreement, stating her own confidence in her alternative claim (how strongly she believes herself to be right) and her own intransigence with respect to the conflict with the xOLM's own claim (how strongly she believes the xOLM to be right).

Ultimately, we would expect to provide a different negotiation framework, diagnosing this information by running a "real", in-depth, negotiation with the learner, using evidence, justifications and uncertainty to balance the execution of the challenge and to decide whether the xOLM should give up in front of a confident learner or to stick to its judgement when met with fluctuating confidence. A simple negotiation process has been envisaged, similar to those supported by Mr Collins (Bull & Pain, 1995) or Prolog-Tutor (Tchetagni, Nkambou, & Bourdeau, 2005), where learners are presented with an appropriate sequence of activities to perform in order to support their challenge – or to contradict themselves and realise this contradiction. Currently we ensure that, as a proof-of-concept, challenge events are properly generated and suitably processed by the extended Learner Model by introducing a new piece of evidence that can be, in turn, explored by the learner. A consequence of this design compromise means that the xOLM always gives up in favour of the learner.



Fig. 13. The Challenge view allowing the learner to dispute the claim made by the xOLM.

**The Implementation of the xOLM**

The xOLM has been implemented using JAVA, following a client-server architecture. The server-side component is closely integrated with LeActiveMath and the Learner Model proper and deals with data retrieving and forwarding on behalf of the Graphical User Interface on the client side. The GUI has been implemented as a JAVA applet, using JAVA SWING – and several open-source libraries – for the interface widgets. Communications between the two parts are supported by XML-RPC exchanges.

There were some inevitable restrictions imposed on the design and implementation of the xOLM GUI by the environment. LeActiveMath is a web-based environment with the xOLM being deployed in a separate browser window. It means that a multiple-windows GUI has been ruled out for the implementation of the xOLM in order to avoid confusion and overload of the workspace. This decision reduces the potential of the GUI significantly, in particular by preventing the use of simultaneous external representations and their supporting/complementing roles. A second self-imposed restriction concerned the optimum size of the main GUI. Particular attention has been given to the layout of the interface so that it will fit in a 'decent' minimum space (basically a 800x600 resolution screen). This decision – supported by the 'external tool' nature of the xOLM in LeActiveMath – meant in particular that some of the external representations turn out to be difficult to display completely, as can be seen for example with the Argument View (see Figure 10, right-hand side). However, the xOLM is not restricted to that minimum size and can be maximize to benefit from full-screen capacity, overcoming this difficulty to some extent.

The current layout of the various widgets and external representations in the GUI reflects that constant trade-off between readability, usability and informativity. This is why, for example, we opted for a flip-card (i.e. multi-tabs) approach for the interface, why the verbalisation pane is always visible at the interface (despite the obvious counter-argument that it does cost quite a lot of space for possibly no added value), etc. But it also means that space-consuming external representations have been difficult to implement and did need (and still need) fine-tuning for best usage in quite a confined area. This is in particular the case for the graph-based external representation used in both the alternative navigation interface and the Argumentation view.

**VERBALISATION**

To conclude the description of the interface of the xOLM, a final important aspect is worth mentioning in this paper: the verbalisation of the interaction taking place between the learner and the xOLM the lower part of the GUI, as seen in Figure 7).

The initial aim for this verbalisation in the xOLM was to complement the numerous external representations (for both the internal data of the learner model and the structure of the argumentation between the learner and the xOLM) with a verbal description, trying to overcome their intrinsic difficulty. This was seen as a necessary first step toward the promotion of dialogue-based reflective learning, encouraging the learners in contributing to the dialogue, to the diagnosis and to the justification of their decisions – see STyLE-OLM (Dimitrova, 2003) and W-ReTuDiS (Grigoriadou, Tsaganou, & Cavoura, 2004) for two examples of systems who have implemented such an approach.

The framework used to implement such a device is a simple but flexible context-dependent template-based approach. The main reason for this choice, in preference to much powerful mechanisms such as Natural Language Generation, was to guarantee the internationalisation of the

xOLM, a strong requirement for the overall project to which we were contributing. The key elements of this framework are the dialogue moves used to contextualise the discussion between the learner and the xOLM.

## Dialogue and Dialogue Moves

The learner's dialogue moves (e.g. *Explain*, *I Disagree*, *I'm Lost*, etc.) are performed by using the relevant button on the GUI; the move is immediately followed by an xOLM dialogue move (e.g. *Here Is*, *Windup*, *Unravel*, etc.) that interprets the learner's move (using the context when needed), prepares the response to the request and displays it. At that stage, the interface is again ready for the learner to use one of the valid dialogue moves. From a Finite-State Machine point of view, it roughly means that the learner's moves are transitions whereas the xOLM's moves are states.
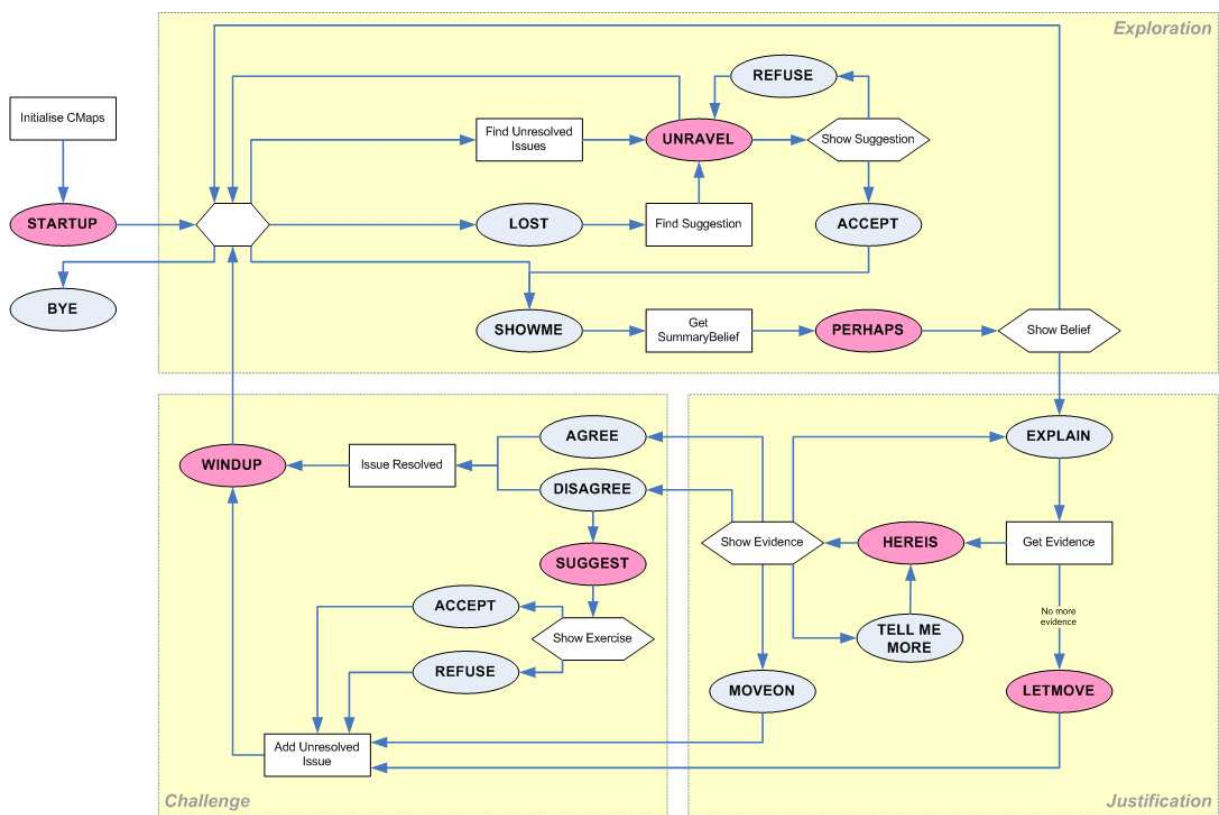


Fig. 14. Overview of the various Dialogue Moves deployed in the xOLM.

Figure 14 represents all the dialogue moves currently implemented in the xOLM, as well as their overall organisation. The dialogue moves are basically organised into three groups. The first supports the *exploration* of the xOLM and the belief it holds, by running a series of *Show Me-Perhaps* moves (i.e. by building a belief descriptor and showing the corresponding claim). The second supports the *justification* of a claim made by the xOLM, by running a series of *Explain-Here Is* moves (i.e. expanding the argumentation graph and showing the relevant external representation). The final group

implements the *challenge* proper, by organising the negotiation on the selected topic and its outcome (either by reaching an agreement between the xOLM and the learner – *I Agree* or *I Disagree* – or by one of them giving up the discussion – *Move On* or *Let Move*).

## Templates and Context

The dialogue moves framework presented above, by clearly identifying the context and stage of the discussion between the learner and the xOLM, provides us with the necessary basis to implement a template-based approach for verbalising this discussion.

In a nutshell, every single dialogue moves deployed in the xOLM is associated with one of several templates, whose selection can be conditioned by the learner having designated one of the elements of the xOLM interface (such as a node in the Toulmin Argumentation Pattern). For example, one of the templates associated with the Perhaps move – i.e. with the xOLM giving its summary judgement about the learner ability – could be *"I think you are #1 on your ability to #2"*, with *#1* and *#2* representing the placeholders to be replaced by, respectively, the verbalisation of the summary level (e.g. *Level II*) and the verbalisation of the target belief (e.g. *Solve problems on Difference Quotient*). To be operational, several sources of information used by the xOLM have to be associated with proper elements of verbalisations: each of the four ability levels for each of the dominant layers (e.g. *Level I* to *Level IV* for a belief on an affective factor could be associated with vocabularies ranging from *low* to *high*[8]); the name of every topic in every layers of the Learner Model (e.g. *Difference Quotient* or *Chain Rule* for the domain, *Solve problems* for the competencies, *Satisfaction* for the affective factors, etc.), every attribute of each event sent to an interpreted by the Learner Model, as well as their possible values (e.g. the *difficulty* of an exercise is associated with five possible values, ranging from *very easy* to *very difficult*), etc. More information about the template mechanism deployed in the xOLM can be found in (Van Labeke, 2006).

If dialogues moves are useful to disambiguate the context of the interaction between the learners and the xOLM, it does not necessarily mean that their associated templates are useful to explain this context back to the learner. The main source of problems for such verbalisation is clearly the nature of the underlying learner model, i.e. the complexity of information (probabilistic distributions) and of the processes involved (Transferable Belief Model). To illustrate the complexity of the task, let's take an example: the justification of a claim (summary belief) by highlighting the properties reflected by the pignistic distribution of the data. To support such an explanation, three elements have to be considered and verbalised:

- the summary belief, by looking at the value of the summary on a continuous scale and taking into account its proximity with the discrete thresholds (Level I, Level II, etc.), this proximity being an indication of the strength of a belief (the thresholds basically act as attractors);

- the pignistic function, which summarises the belief by a normalised distribution of the accumulated information on the four considered levels and whose relative (i.e. which level is the more probable?) and absolute (i.e. how marginal is that probability?) scores are decisive for the decision-making process;

---

[8] In the current release of the xOLM, since there is no "standard" taxonomy to which learners are familiar with, we decided to use the generic terms *Level I*, *Level II*, *Level III* and *Level IV* to verbalise all abilities. Whether this decision is accepted and understood by learners will be one of the aims of on-going evaluations.

- the decision-making process itself which extract the summary belief from the pignistic and, in the current implementation of the learner model, is roughly based on the gravity centre of the pignistic.

Combining all these sources of information into a single template is a task that proved to be extremely difficult and not necessarily helpful for learners since it often resulted in long and repetitive statements. To overcome this difficulty, a divide and conquer approach has been preferred: splitting the explanations into separate but complementary descriptions that can be obtained on the learners' own initiative. For example, two different dialogue moves are available to the learner for requesting detailed information about a given element of the Toulmin Argumentation Pattern (and its associated external representation): the *Tell Me More* move for explanations on the current instalment of the element (for example by describing the dominant feature of the pignistic function) and the *Explain* move for explanations on the transition between two elements (for example how the dominant feature of the pignistic is used to infer the summary belief).

| | |
|---|---|
| **OLM** | Welcome **Louise**. |
| **Louise** | Give me your judgement on my ability to **Solve problems** on **Difference Quotient**. |
| **OLM** | I think you are **Level II** on your ability to **Solve problems** on **Difference Quotient**. |
| **Louise** | Could you remind me what this graph is about? |
| **OLM** | This graph represents my judgement of your **Competency**, on a continuous scale between **Level I** and **Level IV**. |
| **OLM** | Being very close to the **Level II** indicator, my conviction on my judgement is therefore quite strong. |
| **Louise** | I don't understand why you think I'm **Level II**. |
| **OLM** | The dominant trait for my belief may be **Level I**, but evidence supporting higher levels is non-negligible. |
| **Louise** | Show me another view of this information. |
| **OLM** | On the right, you can see the mass distribution of my belief, showing all information from accumulated evidence. |
| **Louise** | I don't understand what this means. |
| **OLM** | This mass distribution represents my guess about your ability to **Solve problems**, in terms of the probability for you to be at the various combinations of **Competency** levels. |
| **Louise** | I don't understand how you obtained this profile. |
| **OLM** | Looking at the evidence according to your **Performance**, my judgement is mostly based on the **4** piece(s) in the **low** group. |
| **Louise** | Could you tell me more about your evidence on the **low** group? |
| **OLM** | Among the evidence showing **Performance low**, these **3** are the most important. |
| **Louise** | Show me another view of this information. |
| **OLM** | Here is the complete mass of this piece of evidence, as it was initially generated. |
| **Louise** | Show me another view of this information. |
| **OLM** | Since evidence tends to be less certain over time, here is the mass as it is now considered in the belief. |
| **Louise** | Please explain what this interaction was about. |
| **OLM** | Because you did this **very easy** exercise with a **very low** performance. |
| **Louise** | Please explain what this interaction was about. |
| **OLM** | … |

Fig. 15. Extract of the verbalisation generated from the interaction between a learner and the xOLM.

Figure 15 presents an extract of the verbalisation of Louise's interaction with the xOLM, as described in the previous sections. It highlights the flexibility of the current implementation of the mechanism but also some of its inherent limitations (such as gender/number accordance that can be dealt with – but usually at great cost and in detriment to internationalisation – and an inevitable mechanistic/repetitive style – dealt with to some extent by introducing a random selection of alternate templates for a given context).


## EVALUATION

Developing the xOLM inevitably raises the question as to how we are to know when the system is sufficient successful to merit further work and deployment in other learning environments. The first point to mention is that the purpose of the xOLM within the current context of a European project is that it should be a component in a large system while the purpose of the research reported here is to examine the claim that the xOLM can be used to open up the interpretation process. For the former purpose, an important question might be how the xOLM component contributes to learning gains. For the latter purpose, there are more immediate aims.

As usual with such systems as the xOLM, analysing learning gains and seeking to understand how the student model makes a contribution is a difficult business, e.g. (Weibelzahl, 2005). Weibelzahl points out that there are several ways of approaching the evaluation issue including "layered evaluation" (Brusilovsky, Karagiannidis, & Sampson, 2001), which attends to a number of criteria including learner satisfaction, learner motivation as well as system factors such as "the reliability and validity of the input data, the precision of the student model, or the appropriateness of the adaptation decision."

For this paper, the most appropriate and immediate questions relate to whether the xOLM really can open up the interpretation process. To determine this, we need - at the very least - to consider whether learners can understand the xOLM, whether they can trust the xOLM and whether there is any perception of the xOLM's utility.

In a preliminary study, we adopted a variation of the method of collaborative evaluation to examine the ways in which the learner understood the xOLM and its representations (Monk, Wright, Haber, & Davenport, 1993; Wright & Monk, 1991). The participant, designer and experimenter worked together in a task-oriented evaluation. To set the scene, the participant was asked to imagine that they had done a few exercises in college level mathematics and now wanted to see what the OLM can provide. From the results obtained, usability issues were identified for the main study.

The main study was undertaken by our colleague, Tim Smith, as part of the work by the team at Edinburgh University. Full details can be found in (Brna et al., 2006). Here, we concentrate on the results of the work that relates most closely to the xOLM.

A small in-depth comparative study was set up using two groups of students - those who were given an opportunity to use the xOLM embedded in the project's system for the first time, and those who had already used a previous version of the system which made no use of the xOLM but did feature a "traffic light" method of indicating estimates of knowledge. This made it possible to factor out some of the problems which arise from meeting the system for the first time (e.g. learning how to use the interface, being unfamiliar with the context in which the xOLM can be used, having too positive/negative a view based on limited experience).

Seven male and three female students (mean age 22.5 years) agreed to take part. All these participants were first-year Mathematics students at Edinburgh University having "average" confidence with calculus (average rating of 3.50 on a 5-point Likert scale). Five of these participants had taken part in previous evaluations of the learning environment without the xOLM being available (the *E group*) and five new participants from the same course (the *N group*). These participants can be regarded as highly capable computer users who should be able to use any well designed computer interface since their familiarity with the web was estimated as "very familiar" (average 4.60) and "average" familiarity with Java applets (average 2.70).

Providing a reasonably realistic context of use, the participants were asked to read a series of pages and perform a few exercises. The participants had the opportunity to explore the xOLM and discovered its working. After using the system, participants completed a questionnaire about their experience and also one about the xOLM.

Other work with the system which did not feature the xOLM had showed that participants sometimes thought that the knowledge estimates made were inaccurate. Distrust of the knowledge estimates could have resulted from a lack of transparency in the learner model. The xOLM was believed to have the potential to justify the learner model so use of the xOLM should lead to improvements in the perceived accuracy of the knowledge estimates.

Participants were asked (before using the xOLM) how accurate they thought an ITS would be and, after using it, how accurate they believed knowledge estimated by the xOLM were. Both members of the *N* and *E groups* rated ITS' accuracy as "medium" (mean = 3.4). *N group* people rated xOLM's accuracy as being higher than they had expected ("very" accurate, 4.6 out of 5); however, *E group* users were more cynical, rating it as "accurate" (3.8 of five). Considering that the experts had previously encountered the system, their rating of an ITS' accuracy could be seen as a rating for the version of the system they had previously encountered without the xOLM. The increase in rating can reasonably be attributed to the xOLM. In brief, both *N* and *E group* users believe the system's learner model to be accurate and the perceived accuracy improves after they have worked with the xOLM.

Learners did take some time to become familiar with the xOLM interface but, once this stage was past, it became evident that some aspects of the system were successful. After getting used to the Toulin view, participants found it very easy to use (rating it as 4.6 out of 5 on a scale of usefulness). Another successful aspect was the topic map which was found to be a good way to understand the links between concepts. Participants claimed that the map helped them in a way that books did not do. The Topic map was identified as useful for revision. Participants rated the Topic Map as "very useful" (4.6 out of 5) and "quite intuitive" (3.8).

An example can illustrate the potential of the xOLM. Participants began by not understanding the relationship between the concept "average slope" and "difference quotient". They were aware that there was a conceptual relationship as they observed their mastery propagating from a page on "average slope" in which they had completed an exercise to an unseen page on "difference quotients". When using the xOLM to examine what the xLM believed about their knowledge of "average slopes" they noticed that some indirect evidence was coming from an exercise they had completed about "difference quotients". This motivated them to explore this relationship using the Topic Map. This revealed a conceptual link between the two concepts which described them as being identical.

The evaluation enabled a number of usability problems to be identified, and some fixed for the version of xOLM described in this paper. These problems apart, it was found by Smith that the users enjoyed using the xOLM and were able to explore the learner model with ease (Brna et al., 2006).

## DISCUSSION

The design and implementation of a new open learner model has been outlined, and we have gone into some detail about the way in which the Toulmin Argumentation Pattern has been used to drive the argumentation about the content of the extended Learner Model. The interaction of the learner with the xOLM has been described in detail. We have described elsewhere the way in which to decouple the xLM and xOLM from LeActiveMath by parameterising the xLM, and improving its usage of Semantic Web standards and technologies (Morales, Van Labeke, & Brna, 2006c). To conclude this paper, we return to the issue of "Opening up the Interpretation Process in an Open Learner Model".

The interpretation process has been made more complex by the decision to develop an extended learner model with the structure as found in Figure 1. In particular, the adoption of the notion of competence had a profound effect on the system. The judgement of competence of any specific kind (e.g. mathematical thinking), as expounded by PISA, is not simple. The statement that a learner is at Level II for mathematical thinking requires a far more sophisticated, and less well defined, judgement than just whether or not the last question involving mathematical thinking was right or wrong[9].

We now consider three aspects of future work: to provide better representations of the dynamics of the learner model, methods of managing the amount of information, and ideas for mitigating some of the problems associated with the complexity of the approach.

### Dynamic versus Static Open Learner Model

Most of the External Representations described in this document provide the learners with an overview of the current state of the belief and not of its evolution across time. Various (inconclusive) attempts to give access to a dynamic Learner Model have been tried.

But by repeated usage and testing, it became evident that this initial attempt for representing the dynamic process raised more issues than it solved: the question of consistency across external representations (how to dynamically represent complex information such as the pignistic, the mass distribution, etc.); the question of controlling the dynamic representation (replaying the sequence forward and backward, stopping at the introduction of a given piece of evidence, etc.); the question of supporting the translation between dynamic external representation (selecting a step of the sequence to access the related evidence).

As with most learner models, the assumption that is currently driving the implementation of the xOLM is that the focus of the learners will be directed toward the actual state of the beliefs rather than toward the trajectory of their abilities; justifications will be supported by providing the learner with access to every individual evidence.

This assumption needs to be carefully challenged in the future, by introducing dynamic aspects of the extended Learner Model wherever they are likely to provide different information and support different (and improved) reflection from the learners.

### Exploration and Navigation

Exploring the content of the extended Learner Model (in other word, navigating through the Open Learner Model) is a complex but important task. Finding a proper paradigm for it and building an

---

[9] Not least, because many questions involve some elements of mathematical thinking.

appropriate interface is an issue that has still to be resolved. The current approach has been developed for its simplicity and straightforwardness but has always been intended as a compromise until a better, more suitable solution could be devised.

Two issues have to be taken into account at the same time. On the one hand, such an interface should allow learners to find and select a belief descriptor (i.e. the underlying identifier behind any information in the Learner Model) that is relevant to their current goal, task and/or desire. On the other hand, the interface has to clearly and easily relate the selected belief with the overall portrait that the system holds about the learner. If the current approach does support the first issue fairly well (by explicit listing of all the possible topics, organised among the five distinct layers of the extended Learner Model), we have to admit that it does fail on the second aspect. Shortcuts and ad-hoc scaffolds on the existing interface (such as tool-tips, access to description and definition of the terms used, etc.) have been implemented to allow the learners to build a better connection between the beliefs and the topic maps but we thought (and are still thinking) that a total redesign is needed.

A first attempt for an alternative interface has been investigated, using dynamic graphs to display and present to learners, not only all the relevant topics but also their connections and inter-dependencies. Figure 16 represents the learner as the central node, with all the topic maps used by the xOLM spreading from it: metacognition, affect and motivation maps are fully expanded; mathematical competencies are partially expanded, only the sub-competencies being hidden; the domain map (with its top-node "differential calculus") is partially visible on the top of the figure.

Using such a "interactive map" will certainly increase the usability of the interface (introducing both topics and associations to develop the narrative of the map, dynamically hide/expand/collapse sections of the graph that are not of an immediate usage by the learner, etc.. But issues such as the potential unfamiliarity of a graph approach, the difficulty of manipulation, the whole readability of the approach, etc. have also to be taken into account to make a balanced judgement.

The graph, implemented with an open-source library[10], supports very useful features such as automatic layout (spring model), expansion/collapsing of nodes and sub-graphs on-the-fly, totally configurable and re-implementable UI (both representation and manipulation of graphs). It is already available in the Open Learner Model but has limited functionalities (only belief descriptors are dynamically added on-the-fly, no filtering of graph, expansion and collapse of nodes/sub-graphs at learner's request).

Further extensive works will be necessary to investigate if the problems highlighted above could be effectively addressed by some "fine-tuning" of the library and if this does provides us with an improved alternative interface.

## A Complexity/Usefulness Trade-off

The issue of hiding part of the information that the Learner Model is holding about the learner has been mentioned as a difficult topic. The initial approach was to assume that everything should be presented to the learner.

---

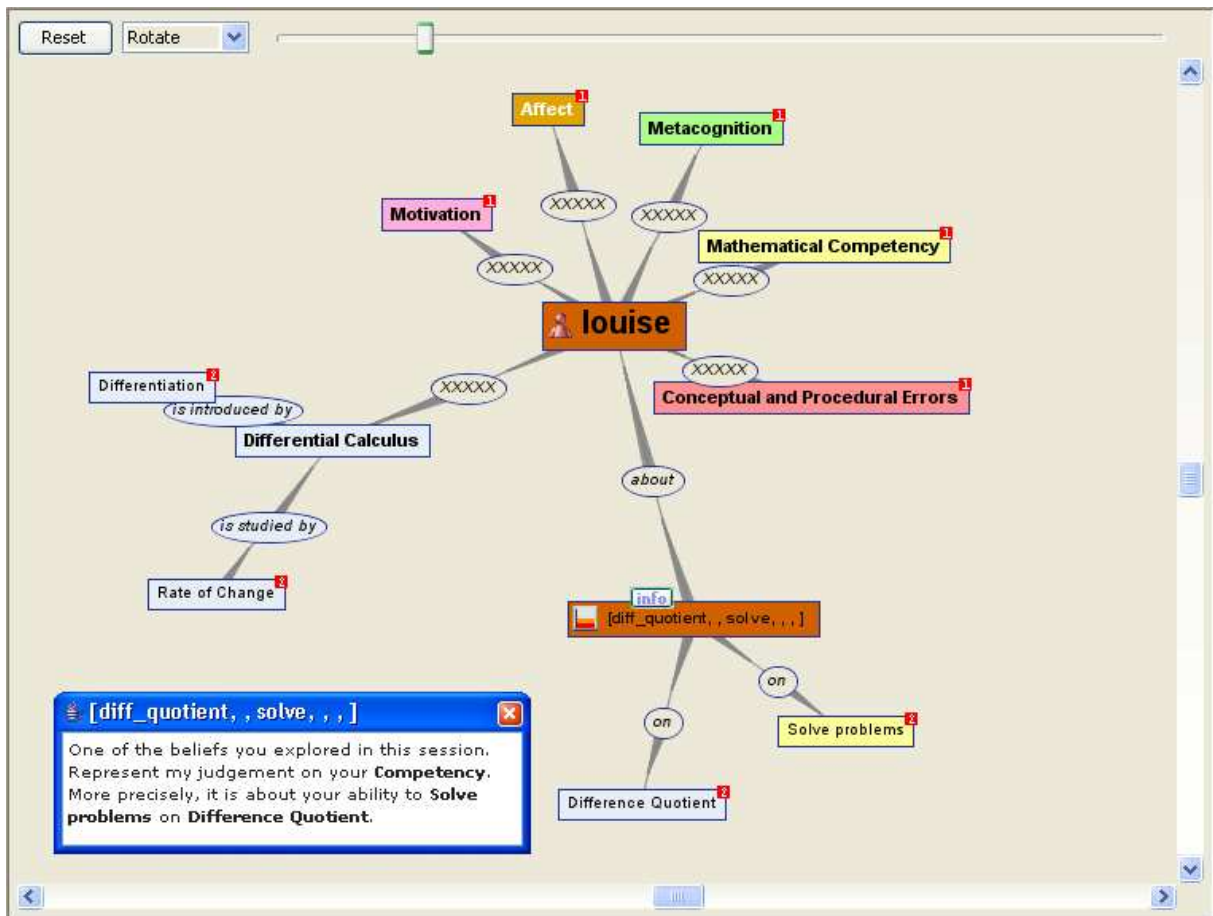[10] TouchGraph, see http://touchgraph.sourceforge.net/

Fig. 16. An alternative view for the exploration of the content of the Learner Model.

We are now able to make a couple of points about what to hide and when – though we have more questions than answers at the moment. Deploying the xOLM in a less complex manner may well make the interface more accessible to the learner. For example, the Claim view is reasonably accessible to the learner, i.e. the externalisation of the fairly straightforward and simplistic (but accurate) summary of a belief held by the Learner Model.

On the other hand, by performing such an abridgment, a significant amount of information would not be presented to the learner, therefore blurring the "logic" of the reasoning followed by the Learner Model when establishing its belief and, consequently, impeding the possibilities for the learner to efficiently challenge its decisions. To allow for such challenges, a more detailed external representation of the internal data is therefore needed – e.g. the Certainty view of the belief provides an example. By presenting all probabilities for every range of levels, the ER clearly extends the potential for justifying the model but, as a result, makes it much harder to restore some reasonable level of readability.

The current state of the extended Open Learner Model, as described in this paper, is a first sustained attempt in finding a trade-off between the two (apparently contradictory) aims of readability

and justifiability. The introduction of several complementary external representations (such as the summary belief, the pignistic function, the mass distribution, etc.), the use of a Toulmin-inspired argumentation framework for organising and controlling the usage of these representations, the supportive role that the natural-language "transcription" of the learners dialogue moves are all steps in reaching this objective.

## ACKNOWLEDGEMENTS

## REFERENCES

Adams, R., & Wu, M. (2002). *Pisa 2000 Technical Report*. Paris, France: OECD.

Birenbaum, M., Breuer, K., Cascallar, E., Dochy, F., Dori, Y., Ridgway, J., & Wiesemes, R. (2006). A Learning Integrated Assessment System. *Educational Research Review,* 1(1), 61-67.

Brna, P., Bull, S., Pain, H., & Self, J. (1999). Negotiated Collaborative Assessment through Collaborative Student Modelling. In *Proceedings of the AIED'99 - Workshop on Open Interactive and other overt Approaches to Learner Modelling* (pp. 35-42).

Brna, P., Van Labeke, N., Morales, R., Pain, H., Porayska-Pomsta, K., & Smith, T. J. (2006). *Evaluation of the Student Model* (LeActiveMath Deliverable No. D38). Technical Report, available from the authors.

Brusilovsky, P., Karagiannidis, C., & Sampson, D. (2001). The Benefits of Layered Evaluation of Adaptive Applications and Services. In S. Weibelzahl, D. N. Chin & G. Weber (Eds.) *Proceedings of the UM'01 - Workshop on Empirical Evaluations of Adaptive Systems*, pp. 1-8.

Bull, S., Brna, P., & Pain, H. (1995). Extending the Scope of Student Models. *User Modeling and User-Adapted Interaction,* 5(1), 45–65.

Bull, S., Dimitrova, V., & Brna, P. (2002). Enhancing Reflective Modelling through Communicative Interaction in Learning Environments. In P. Brna, M. Backer, K. Stenning & A. Tiberghien (Eds.) *The Role of Communication in Learning to Model*. Mahwah, NJ: Lawrence Erlbaum Associates.

Bull, S., & Pain, H. (1995). 'Did I Say What I Think I Said, and Do You Agree with Me?': Inspecting and Questioning the Student Model. In J. Greer (Ed.) *Proceedings of the AIED'95 - 7th World Conference on Artificial Intelligence in Education* (pp. 501-508). Charlottesville VA: AACE.

Bull, S., Pain, H., & Brna, P. (1996). A Student Model 'for Its Own Sake'. In P. Brna, A. Paiva & J. Self (Eds.) *Proceedings of the European Conference on Artificial Intelligence in Education*. Edicões Colibri, Lisbon.

Cheng, P. C.-H., & Pitt, N. G. (2003). Diagrams for Difficult Problems in Probability. *Mathematical Gazette,* 87(508), 86-97.

Cumming, G., & Self, J. (1991). Learner Modelling in Collaborative Intelligent Educational Systems. In P. Goodyear (Ed.) *Teaching Knowledge and Intelligent Tutoring* (pp. 85–104). Norwood, NJ: Ablex.

Dimitrova, V. (2003). Style-OLM: Interactive Open Learner Modelling. *International Journal of Artificial Intelligence in Education,* 13, 35-78.

Dimitrova, V., Brna, P., & Self, J. (2002). The Design and Implementation of a Graphical Communication Medium for Interactive Open Learner Modelling. In S. A. Cerri, G. Gouarderes & F. Paraguacu (Eds.)

*Proceedings of the ITS'02 - 6th International Conference on Intelligent Tutoring Systems* (pp. 432-441). Berlin: Springer Verlag.

Grigoriadou, M., Tsaganou, G., & Cavoura, T. (2004). Dialogue-Based Personalized Reflective Learning. In *Proceedings of the ICALT'04 - International Conference on Advanced Learning Technologies*, pp. 266-270.

Kay, J. (1997). Learner Know Thyself: Student Models to Give Learner Control and Responsibility. In Z. Halim, T. Ottomann & Z. Razak (Eds.) *Proceedings of the ICCE'97 - International Conference on Computers in Education*, pp. 17-24.

Monk, A., Wright, P. C., Haber, J., & Davenport, L. (1993). *Improving Your Human Computer Interface: A Practical Technique*. Prentice Hall.

Morales, R., Pain, H., & Conlon, T. (2000). Understandable Learner Models for a Sensorimotor Control Task. In G. Gauthier, C. Frasson & K. Vanlehn (Eds.) *Proceedings of the ITS'00 - 5th International Conference on Intelligent Tutoring Systems* (pp. 222–231). Berlin: Springer-Verlag.

Morales, R., Van Labeke, N., & Brna, P. (2006a). Approximate Modelling of the Multi-Dimensional Learner. In *Proceedings of the ITS'06 - 8th International Conference on Intelligent Tutoring Systems* (pp. 555-564). Berlin: Springer.

Morales, R., Van Labeke, N., & Brna, P. (2006b). A Contingency Analysis of Leactivemath's Learner Model. In *Proceedings of the MICAI 2006 - Mexican International Conference on Artificial Intelligence*.

Morales, R., Van Labeke, N., & Brna, P. (2006c). Towards a Learner Modelling Engine for the Semantic Web. In *Proceedings of the AH'06 - International Workshop on Applications of Semantic Web Technologies for E-Learning (SW-EL)*.

OECD. (1999). *Measuring Student Knowledge and Skills: A New Framework for Assessment*.

Porayska-Pomsta, K., & Pain, H. (2004). Providing Cognitive and Affective Scaffolding through Teaching Strategies. In J. C. Lester, R. M. Vicari & F. Paraguaçu (Eds.) *Proceedings of the ITS'04 - 7th International Conference on Intelligent Tutoring Systems* (pp. 77-86). Berlin: Springer-Verlag.

Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton University Press.

Smets, P., & Kennes, R. (1994). The Transferable Belief Model. *Artificial Intelligence, 66*, 191-234.

Suthers, D., Weiner, A., Connelly, J., & Paolucci, M. (1995). Belvedere: Engaging Students in Critical Discussion of Science and Public Policy Issues. In J. Greer (Ed.) *Proceedings of the AIED'95 - 7th World Conference on Artificial Intelligence in Education*. Charlottesville VA: AACE.

Tanimoto, S. (2005). Dimensions of Transparency in Open Learner Models. In *Proceedings of the AIED'05 - Workshop on Learner Modelling for Reflection, to Support Learner Control, Metacognition and Improved Communication between Teachers and Learners*.

Tchetagni, J., Nkambou, R., & Bourdeau, J. (2005). Supporting Student Reflection in an Intelligent Tutoring System for Logic Programming. In *Proceedings of the AIED'05 - Workshop on Learner Modelling for Reflection, to Support Learner Control, Metacognition and Improved Communication between Teachers and Learners*.

Toulmin, S. (1959). *The Uses of Arguments*. Cambridge University Press.

Van Labeke, N. (2006). *Student Model Dialogue* (LeActiveMath Deliverable No. D32). Technical Report, available from the authors.

Weibelzahl, S. (2005). Problems and Pitfalls in the Evaluation of Adaptive Systems. In S. Y. Chen & G. D. Magoulas (Eds.) *Adaptable and Adaptive Hypermedia Systems* (pp. 285-299). Hershey, PA: IMR Press.

Wright, P. C., & Monk, A. F. (1991). The Use of Think-Aloud Evaluation Methods in Design. *ACM SIGCHI Bulletin, 23*(1), 55-57.

Zapata-Rivera, J.-D., & Greer, J. (2003). Analyzing Student Reflection in the Learning Game. In S. Bull, P. Brna & V. Dimitrova (Eds.) *Proceedings of the AIED'03 - Workshop on Learner Modelling for Reflection*, pp. 288-298.

Zapata-Rivera, J.-D., & Greer, J. (2004). Inspectable Bayesian Student Modelling Servers in Multi-Agent Tutoring Systems. *International Journal of Human-Computer Studies, 61*(4), 535-563.