Learner Modelling in Microworlds: conceptual model and architecture in MiGen

Manolis Mavrikis, Sergio Gutierrez-Santos, Darren Pearce-Lazard,
Alexandra Poulovassilis, and George Magoulas*

London Knowledge Lab, 23-29 Emerald Str, London, UK.

Abstract. The design of user-modelling components for Exploratory Learning Environments (ELEs) presents a significant challenge, particularly because of the ill-structured nature of the tasks that students are asked to undertake. This paper focuses particularly on Microworlds and argues that following the usual approach of modelling a learner's knowledge just in relation to concepts is not straightforward in such cases. We present a conceptual model that, as well as epistemic concepts, incorporates epistemic and unproductive 'ways of thinking' and their operationalisations through the affordances of the microworld. The paper also presents the architecture of the user model of the MiGen system, a microworld-based system aiming to support 11–14-year-old students in their learning of mathematical generalisation. We conclude the paper with a brief discussion of the generality of our approach and its applicability to other microworlds and ELEs.

1 Introduction

The work presented in this paper focuses on modelling learners as they are undertaking tasks within microworlds (MWs) [1]. A recent review of ill-defined domains [2] refers to such environments as 'model building systems' and identifies them as belonging to a particular genre of discovery learning whereby users (learners in our case) are provided with a suite of model-building tools and are encouraged to 'test their own intuitions about a domain' [2]. Most microworlds provide non-adaptive scaffolds designed to help students explore the domain. However, as with other constructivist approaches (c.f. [3]), related research suggests that explicit support is an important determinant of learning [4].

The onus thus falls on the learner modelling component of the system to provide the substrate that enables the provision of adaptive feedback, to facilitate self-regulation through open learner models, and to assist teachers in their efforts to integrate MWs into the classroom. However, following the usual approach in the field, whereby a learner's knowledge or skills are modelled only in relation to domain concepts, is not straightforward in MWs. As the paper presents in detail, even in the cases where the subject domain underlying a MW is well-defined (e.g. mathematics), the knowledge that students are expected to develop and

* The authors would like to acknowledge the rest of the members of the MiGen team and financial support from the TLRP (e-Learning Phase-II, RES-139-25-0381).

the tasks they undertake within a MW are ill-structured (c.f. [5] that highlights the need to consider the domain and the task as two orthogonal diamensions when discussing ill-definedness).

In this paper, we present our approach to tackling this problem. Section 2 provides a brief background on microworlds and on the challenges that their characteristics introduce for user modelling. It also describes the eXpresser microworld — the main user-facing component of the MiGen system — which is used throughout the paper to present our approach. Section 3 presents our conceptual model, the architecture of the learner model entities in MiGen, and the process by which these are updated as students undertake tasks. Section 4 concludes the paper by presenting how our approach can be applied to modelling learners working in other similar environments, and presents our future work.

2 Microworlds and Learner Modelling

2.1 Microworlds

Microworlds provide learners with opportunities to develop their own models within the domain of interest, using appropriate objects, interactive tools and/or a formal language. By directly representing and giving access to the domain, students are able to not only explore the structure of accessible objects in the environment, but also to construct their own objects and explore the representations that make these objects accessible [1]. Perhaps the most well-known microworlds are the ones based on LOGO or variants of it, such as the recent Scratch. Microworlds have mostly been used in geometry (see Section 4) and other topics in mathematics education [4] or inquiry learning [6].

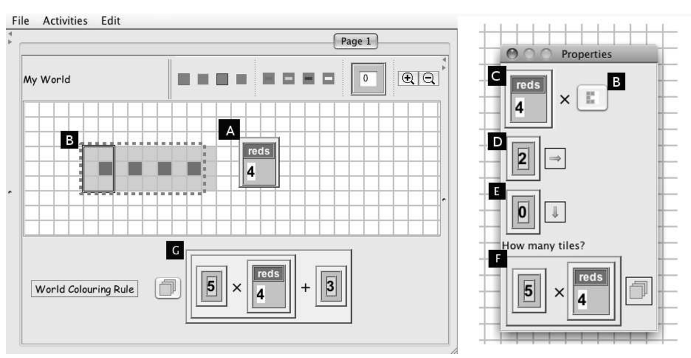

Fig. 1: Constructing a pattern in the eXpresser and describing it with a rule. Letters highlight the main features: (A) An ‘unlocked’ number that acts like a variable is given the name ‘reds’ and signifies the number of red (dark grey) tiles in the pattern. (B) Building block to be repeated to make a pattern. (C) Number of repetitions (in this case, the value of the variable ‘reds’). (D,E) Number of grid squares to translate B to the right and down after each repetition. (F) Units of colour required to paint the pattern. (G) General expression that gives the total number of units of colour required to paint the whole pattern.

As an example, Figure 1 shows the eXpresser microworld developed in the context of the MiGen project, with which the authors are currently involved. eXpresser encourages students to build constructions for patterns and to find general expressions (rules) underpinning such patterns. In order to do that, students can use building blocks of square tiles to make patterns. In order to represent the generalities they perceive, they can use numbers that can be 'unlocked' to become variables. Locked and unlocked numbers can be used in expressions. This microworld gives a lot of freedom to students, who may construct their patterns in a multitude of different ways. For a detailed description of the eXpresser the interested reader is referred to [7]. Figure 1 illustrates some of the core aspects of the eXpresser — in particular, the number (A) in the figure shows a variable and the expression (G) a general expression.

2.2 Related work in Learner Modelling

Research in the User Modelling (UM) and Intelligent Tutoring Systems (ITS) fields, has for many years highlighted tractability issues relating to learner (or student) modelling [8]. This is particularly challenging in ill-structured (or ill-defined) domains such as the ones that underlie microworlds. These difficulties have motivated many approaches to learner modelling, starting from the traditional overlay models and bug libraries to approaches that ground a learner model in cognitive theories of learning (for reviews see [9]).

Particularly relevant to our work is learner modelling in Open or Exploratory Learning Environments (ELEs) where, as well as the common challenges of user modelling, additional ones make the problem even more challenging. These include unstructured interaction, limited explicit information about the learners' knowledge, the need to achieve a balance between freedom and control, and the fact that there is no clear separation between 'right' and 'wrong' answers or a clear definition of 'good' or 'bad' explorations — we refer the reader to [10, 11] for more detailed analyses of these and related problems.

Regardless of the precise techniques used to trace students' actions or to update a learner model, representing the knowledge required for the learner model is an important prerequisite, particularly when we need to expose these environments to stakeholders such as teachers (e.g., through open learner modelling approaches [12]). In relation to domain and diagnostic knowledge, previous related research proposes various approaches. First, it is possible to avoid representing explicitly the domain and to model high-level processes instead. For example, to support students working in a virtual laboratory for physics, the work in [13] employs heuristics to assesses students' scientific experimentation processes in relation to their ability to conduct experiments that confirm hypotheses. Similarly, rather than determining deviations by tracing a student's model, which requires time-consuming representation of an expert's reasoning, constraint-based student modelling [14] advocates employing violation of constraints in order to guide the system's interventions.

Another approach, employed by ACE [10], is to assess the effectiveness of learners' exploratory behaviour. This is based on a Bayesian Belief Network that models salient cases for the problems under exploration. This approach seems particularly appealing in environments that behave more like simulators (model exploration) rather than microworlds (model building [2]). This functional difference is significant as it affects the anticipated type of student interaction. In the first case, it requires only exploration of the effect of, for example, different inputs to the pre-constructed model of the simulator, whereas in the case of microworlds it allows new constructions and affords creative and innovative approaches to solving a task (c.f. [2]). It is this distinction, in particular, that guides the conceptual model we present in the next section.

3 Learner Modelling in MiGen

3.1 Conceptual Model

Learner Modelling in microworlds requires particular attention to their epistemology, i.e. how a microworld provides scope for knowledge, or limits it; in other words, how it transforms the nature of the domain it represents [15].

First of all, interaction with a microworld does not necessarily provide direct evidence for assessing students' understanding of concepts. In particular, such interaction cannot be separated from the pragmatic context in which it occurs. It is this fact that introduces the need for explicit support, beyond the non-adaptive scaffolds that microworlds provide by design. In this sense, microworlds are not intended as a means of communicating just concepts. In contrast, by allowing direct interaction with a representation of the domain, they aim to provide opportunities for learners to develop subject-specific 'ways of thinking' (WOTs) [16] in the particular domain e.g. to think as a scientist, a mathematician, a programmer. For example, in mathematical microworlds, students are expected to learn not just the mathematical concepts (e.g., what a 'variable' is) but also mathematical ways of thinking (e.g. 'thinking by employing variables'). Accordingly, this constructionist orientation, determines that the tasks that students undertake are effectively ill-structured, since not only they have multiple solutions but also students are encouraged to explore the environment, and follow a variety of strategies, not all of which can be sequenced or pre-defined. Therefore, the cognitive analysis (and consequently the learner model) should include, apart from the traditional epistemic concepts and possible invalid conceptions for particular contexts, epistemic and 'unproductive' WOTs. This approach is in line with [17]) which argues for extending the scope of the learner model with aspects outside the domain boundary [1].

Second, it is important to make a distinction between the overall subject domain that the microworld is designed for, and what is referred to as 'epistemological domain of validity' of the microworld [18] i.e. the knowledge domain as it has been transformed by the affordances and interface of the environment.

[1] We recognise that the learner model should include affective and motivational characteristics. Currently, such information is encapsulated under our modelling of ways of thinking, as a fully-fledged affective modelling is beyond the scope of this research.

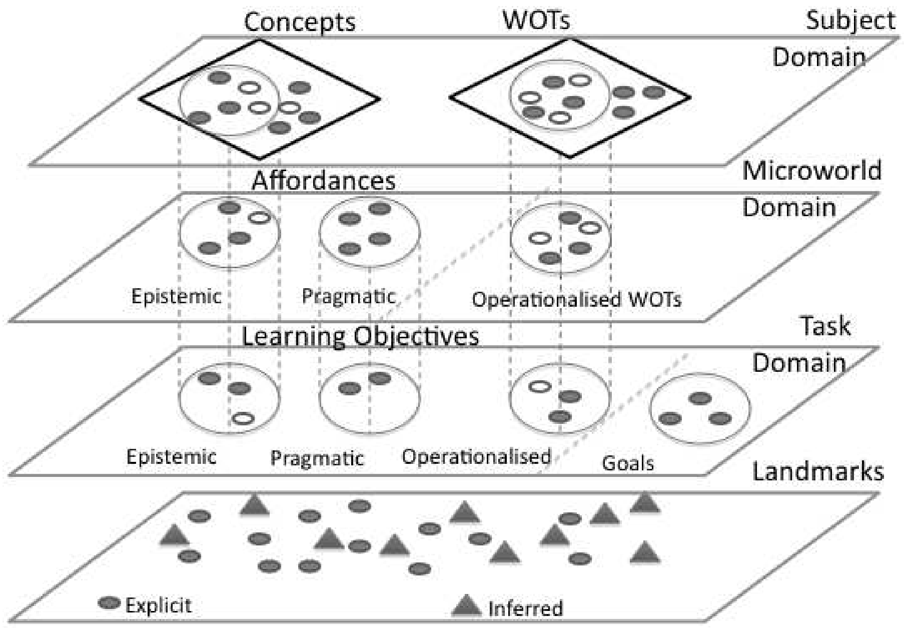

Fig. 2: Conceptual model for learner modelling in microworlds. Concepts in the top layer (including 'invalid' conceptions for a particular context — blank ovals) are operationalised to epistemic affordances in the microworld layer. The same applies to productive and unproductive ways of thinking (WOTs). The microworld layer also includes pragmatic affordances corresponding to actions independent of any epistemic basis. It is projected to the task layer comprising concrete goals and learning objectives. Lastly, landmarks indicate the completion of goals and attainment of learning objectives

Taking into account the situated nature of learning, this difference is important especially when one considers that the design of the microworld transforms or operationalises the knowledge domain and can make a difference on what is learnt and how. For example, the notion of a variable in the case of MiGen is linked with the view of a variable as a 'generalised number' in the eXpresser's affordances and operationalised as an 'unlocked number' (see Fig 1(A)).

If the objective behind the modelling process were just to develop a model of the 'user', then this distinction between the subject domain and the microworld domain would not be so important. However, our goal is to model learners and their learning progress, outside the boundaries of the microworld. Accordingly, we need to represent explicitly the relationship between the subject and the microworld domains. This becomes even more important when microworlds are integrated into a classroom curriculum. Not only do we need to take into account that teachers require a correspondence between learners' interactions with the microworld and the subject domain, but also that teachers are often required to identify and work towards specific learning objectives.

We address these requirements by considering a 'layer' of knowledge that involves microworld-specific concepts and that operationalises both the concepts of the subject domain and the ways of thinking. The means by which the subject domain concepts are operationalised in the microworld are the possible actions available using the objects and tools of the microworld, i.e. its affordances. Because of their direct relationship to knowledge, we refer to these as 'epistemic' affordances. In order to distinguish those actions in the microworld that are independent of any epistemic basis (e.g. 'knowledge of creating a building block' — see Fig 1(B)), we refer to these as 'pragmatic' affordances. In addition, this layer includes what we refer to as 'operationalised' ways of thinking. The two top layers of Fig. 2 present schematically the relationships between the subject domain and the microworld domain. More examples of the concepts and WOTs involved in the eXpresser and other microworlds are given in Table 1 later.

This conceptualisation has the additional advantage of enabling us to take into account that learners' previous knowledge about the domain, and their intuitions, play an important role and can influence the way they perceive and interact with the environment. By representing both the domain and its operationalisation through the microworld, it is possible to use such information (if available) to guide the adaptation process.

Finally, we recognise that modelling in the absence of any context is a very daunting challenge. Therefore, we impose a requirement that specific tasks are designed to contextualise the, otherwise unbounded, interaction. This provides high-level goals that the learner is required to achieve during a task and makes the diagnostic problem much more tractable. Such goals include tangible objectives such as 'find a general expression to colour a [certain] pattern' — see Fig 1(G). To further simplify the modelling problem, and in order to expose the learner model to teachers, our approach requires that learning objectives are assigned to each task. For example, an introductory task could have the simple objective to 'explore how a [certain] tool behaves'; a more complex task could have the objective to 'appreciate the power of unlocked numbers'. As the task domain layer in Fig. 2 shows, there are three types of learning objectives: pragmatic learning objectives correspond to pragmatic affordances of the microworld and are independent of the subject domain e.g. 'knows how to drag numbers on the canvas'; epistemic affordances (e.g. 'unlocking a number') are mapped to epistemic learning objectives; and WOTs to operationalised WOTs (e.g. 'validates the generality of construction by animation'). In MiGen, learning objectives, tasks and goals are currently co-designed by the research team with teachers (see also [19]); though we intend that in the future it will be possible for teachers (or appropriate task designers) to define their own.

Ensuring that tasks have tangible goals and are associated with learning objectives, enables a much more tractable diagnostic aim: measuring beliefs in relation to learning objectives and not abstractly in relation to a student's state of mind. This is achieved through the inference of *landmarks* while students interact with the microworld (bottom layer of Fig. 2). In particular, Explicit landmarks occur when specific actions are undertaken by the student, e.g. 'clicking the animate button to validate their construction', while Inferred landmarks are derived from occurrences of combinations of actions, e.g. 'the student has started to construct generally', or 'the student is exploring in a systematic way'.

3.2 Learner Modelling Architecture

This section formalises the main entities of the learner model architecture implementing the above conceptual model in the context of the MiGen system. It also gives details of the relationships between these entities and the way they are updated as learners interact with the eXpresser component of MiGen. An earlier paper [20] described the conceptual and architectural design of the overall MiGen system. This paper extends that work by focusing in particular on details of the learner modelling aspects, and the way in which the eGeneraliser component of MiGen updates the various learner model entities as learners interact with the eXpresser microworld (the eGeneraliser is the MiGen component that provides the core intelligent support functionality of the system — see [20] for further details of this, and the other, MiGen components and their interactions).

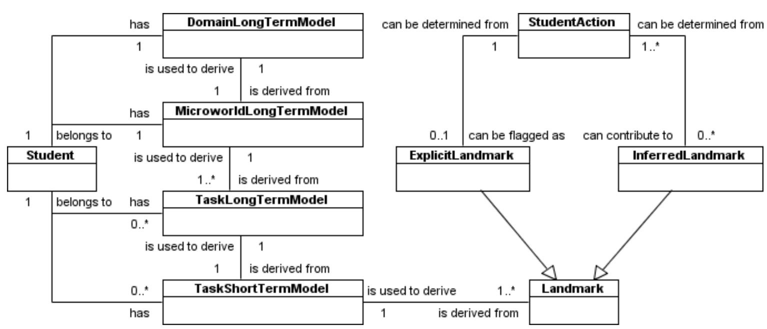

Fig. 3: MiGen learner model. At each end of an edge linking two entities is an indication of the cardinality of that end of the relationship and a verb phrase, e.g. a StudentAction can contribute to zero or more InferredLandmarks. A single-headed arrow indicates a sub-class relationship between two entities.

Figure 3 shows the major entities comprising the MiGen Learner Model and the relationships between them, as well as some associated entities. In particular, for each eXpresser task undertaken by a Student, the eGeneraliser maintains information on the student's ongoing progress through the task within a TaskShortTermModel. This information is derived from the occurrence of Landmarks as the student undertakes a task, as explained in Section 3.1. The eGeneraliser uses the TaskShortTermModel to derive a longer-term model of the student's strategies and outcomes in relation to a task — the TaskLongTermModel

— every time the student finishes a task. This TaskLongTermModel in turn is used to derive a model of the students' attainment in relation to learning objectives that pertain to the whole microworld — the MicroworldLongTermModel. Finally, this is used to derive a model of the learner's attainment of learning objectives related to the domain of mathematical generalisation as defined in the student's DomainLongTermModel. Thus, overall, a student's learner model consists of their current TaskShortTermModels, TaskLongTermModels, MicroworldLongTermModel and DomainLongTermModel.

For inferring the occurrence of landmarks and to update the learner model, we are employing a combination of different techniques. For example, an adaptation of case-based reasoning is used to determine if the students are employing appropriate structures for their constructions [21], rule-based reasoning is used to determine if the student has coloured their pattern in a general or a specific way, and a combination of a sliding window algorithm and string metrics are used to detect rhythm in the actions of the student. Rule-based reasoning is also used to update the TaskShortTermModel based on its current values and information about the occurrence of new landmarks [22]. Once a student completes a task, the higher layers of the learner model are updated: another rule-based component is applied to infer updates to the TaskLongTermModel based on its current values and the current values of the TaskShortTermModel. Similar processes are then applied to the MicroworldLongTermModel and DomainLongTermModel layers, driven by their corresponding rule-based components.

The student's DomainLongTermModel is consistent with the subject domain model of MiGen, which includes concepts such as 'constants', 'variables', 'constructions' and 'expressions' and the corresponding learning objectives mapped from the U.K. National Maths Curriculum, e.g. 'visualise and draw on grids of different types where a shape will be after a translation', 'understand and use the rules of arithmetic in the context of positive integers', 'explore number relationships and propose a general statement involving numbers'. Similarly, the MicroworldLongTermModel is consistent with the second layer of Figure 2, and the TaskShortTermModels and TaskLongTermModels with the third layer.

In practice, a teacher may initialise some attributes in a student's DomainLongTermModel, or may make explicit modifications to them over time. Such a change may result in an update of the student's MicroworldLongTermModel, and this, in turn, may result in an update of the TaskLongTermModels. So there are additional relationships between DomainLongTermModel, MicroworldLongTermModel and TaskLongTermModels capturing this behaviour which, for readability, we have not shown in Figure 3. This "top-down" inference process triggered by a teacher's explicit update is the counterpart of the "bottom-up" inference process described earlier (which is triggered by students' actions).

Figure 4 shows the major types of Learning Objectives in MiGen and the relationships between them, as well as some associated entities. DomainLearningObjectives are separated into epistemic objectives (shown as ConceptualLearningObjective in the diagram) and objectives related to mathematical ways of thinking within the domain of mathematical generalisation, e.g. 'appreciation of the

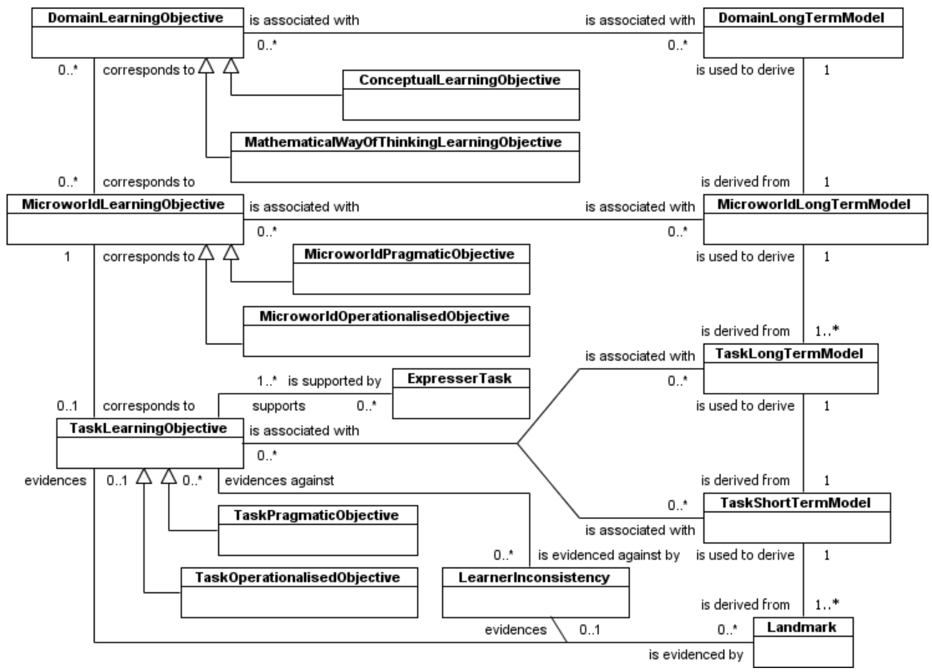

Fig. 4: Learning Objectives

use of variables'. Each task undertaken by a student supports one or more TaskLearningObjectives, as explained in Section 3.1.

Each TaskLearningObjective may be associated with a number of TaskShortTermModels and TaskLongTermModels, in the sense that a student's degree of attainment of that learning objective is recorded as an attribute value within their short or long term learner model. Likewise, each DomainLearningObjective may be associated with a number of students' DomainLongTermModels. Consistent with Section 3.1, landmarks provide evidence for TaskLearningObjectives, as well as for LearnerInconsistencies — these are system-specified stumbling blocks, e.g. 'using more variables than needed'.

Each TaskLearningObjective corresponds to a MicroworldLearningObjective; though there may be additional instances of the latter that have no counterpart TaskLearningObjective. There is a looser correspondence between sets of MicroworldLearningObjectives and sets of DomainLearningObjectives. For example, the MicroworldLearningObjective 'identify variants and invariants of constructions and expressions' corresponds to the DomainLearningObjective 'identify variants and invariants', as mapped directly from the National Curriculum; while 'select an appropriate building block to construct a pattern' corresponds to 'select appropriate strategies in problem solving situations'.

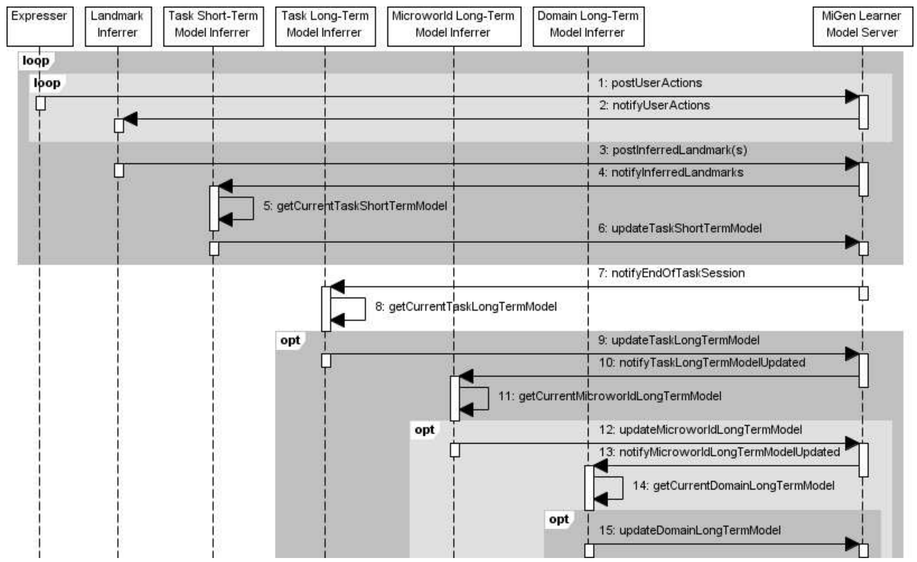

Fig. 5: UML sequence diagram

Figure 5 gives a UML sequence diagram showing the interactions between the main user-facing MiGen component, the eXpresser, and the system's infrastructure for managing the iterative update of the various learner model entities of Figure 3. The 'Inferrers' shown along the top of Figure 5 are sub-modules of the overall eGeneraliser component of MiGen.

As a user undertakes a task within the eXpresser, information on their actions is posted to the MiGenLearnerModelServer (message $\boxed{1}$). The latter is part of the overall MiGen Data Server which is responsible for managing all client-server interaction and interfacing with the system's database (see [20]); it provides a generic mechanism for other modules of the MiGen architecture to register their interest in being notified of updates to various database entities. Message $\boxed{2}$ shows that the Learner Model Server notifies the LandmarkInferrer of the user's actions as these happen (more specifically, it notifies the LandMarkInferrer of the occurrence of those user actions for which the latter has registered an interest in being notified). The LandMarkInferrer may infer the occurrence of landmarks from these actions and, if so, posts these back to the server $\boxed{3}$.

The TaskShortTermModelInferrer is then notified $\boxed{4}$. Using this information on new landmark occurrences, and also the current state of the TaskShortTermModel $\boxed{5}$, it determines whether any attributes in the user's TaskShortTermModel need to be updated. All the changes (if any) are posted back to the server $\boxed{6}$. The TaskLongTermModelInferrer is notified when the user's current session ends $\boxed{7}$ and it uses the updates made to the TaskShortTermModel during this session, and the current state of the TaskLongTermModel $\boxed{8}$, to update the latter if necessary $\boxed{9}$. The MicroworldLongTermModelInferrer receives notifica-

Microworld	Subject Domain		Microworld Domain		
	Epistemic		Operationalised		Pragmatic Affordances
	Concepts	WoTs	Epistemic Affordances	WOTs	
eXpresser [4]	- Variables (generalised numbers) - Algebraic Operations	- Working generally through the particular	- Unlocked numbers - General expressions	- Validation of generalisation by animation	- Drag numbers on the canvas
Geometry microworlds: Cabri[2], GS[3]	- Isosceles triangle - Congruent angles	- Identifying geometric invariants	- Construction tools - Angle tool	- Create and manipulate constructions	- Change properties of shapes
Programming: Scratch[4], Alice[5], ToonTalk[6]	- syntactic correctness - loop	- break complex tasks into sequence of actions	- Same shape blocks - Yellow C shapes-controls	- Connect blocks from a provided set	- Drag block from toolkit to the scripts area
Mathematical simulators: ACE [6],applets in WaLliS[7]	- Function - Variables	- Exploring extreme values to understand the behaviour of a system	- Graphs - Editable text boxes	- Different values in a function machine	- Type values for variables - Move slider
Science labs: SimCalc[8] SimQuest[5]	- Speed, PH, Voltage - Chemical equations	- Designing and conducting scientific experiments	- Gauges - Mixing reactants	- Conduct scripted experiments	- Place reactaant vials to mixer

Table 1. Examples of microworlds and other ELEs in the light of our model.

tion of these changes [10] and it updates the MicroworldLongTermModel [11, 12]. Finally, the DomainLongTermModelInferrer is notified of such updates [13] and it updates the DomainLongTermModel entity held on the server [14, 15].

4 Discussion and Conclusions

Motivated by the nature of the interaction in microworlds, and the intention underlying their design, this paper has argued that learner modelling in microworlds introduces the need to extend the standard approach of representing the domain as concepts or skills with additional information representing learners' 'epistemic' ways of thinking. The conceptual model we have presented here takes into account the transformative nature that microworlds have on the nature of knowledge and on the domain they represent. Our layered approach to knowledge representation enables simplification of the modelling problem by contextualising the domain to its operationalisation in the microworld, and subsequently to tasks with goals and particular learning objectives.

Several microworlds have been reported in the educational technology literature. In most cases, their integration into the classroom has been largely hindered because of the extensive requirement on teachers for helping students both with pragmatic and epistemic aspects. The modelling approach presented in Section 3.1 generalises to microworlds other than eXpresser. This is illustrated in

[2] http://www.cabri.com/ [3] http://www.dynamicgeometry.com/
[4] http://scratch.mit.edu/ [5] http://www.alice.org/
[6] http://www.toontalk.com/ [7] http://www.maths.ed.ac.uk/wallis
[8] http://www.kaputcenter.umassd.edu/products/software/smwcomp/

Table 1, which illustrates that it is possible to model learners undertaking tasks in other microworlds by making the distinction between the subject and microworld domains, and considering how the latter operationalises the epistemic concepts and ways of thinking.

This paper has also presented the learner model architecture of the MiGen system and the process by which learner model entities are updated as students undertake tasks in the eXpresser microworld. Details of the inference mechanisms and of how the information in the learner model is used to personalise the generation of feedback for students will appear in a forthcoming paper. Developing a conceptual model and architecture for the MiGen learner model was a necessary first step in the development of these mechanisms, and has involved several detailed iterations of research, analysis and design involving the authors from their multiple disciplinary perspectives, resulting in the model and architecture presented here. Currently, our approach requires explicit definition of the interaction and task models, as well as the relationships between domain concepts, mathematical ways of thinking and their operationalisations in the microworld. In the future, as learners interact with the microworld, and teachers are exposed to our conceptual model in operation, we plan to design tools to allow for the iterative refinement of our learner model architecture.

References

1. Thompson, P.W.: Mathematical microworlds and intelligent computer-assisted instruction. In: Artificial intelligence and instruction: Applications and methods, Boston, MA, USA, Addison-Wesley Longman Publishing Co., Inc. (1987) 83–109
2. Lynch, C.F., Ashley, K.D., Aleven, V., Pinkwart, N.: Defining ill-defined domains;a literature survey. In: Intelligent Tutoring Systems (ITS 2006): Workshop on Intelligent Tutoring Systems for Ill-Defined Domains. (2006)
3. Mayer, R.E.: Should there be a three-strikes rule against pure discovery learning? - the case for guided methods of instruction. American Psychologist (2004) 14–19
4. Noss, R., Hoyles, C.: Windows on mathematical meanings: Learning cultures and computers. Dordrecht: Kluwer (1996)
5. Mitrovic, A., Weerasinghe, A.: Revisiting ill-definedness and the consequences for ITSs. In: Proceeding of the 2009 conference on Artificial Intelligence in Education, Amsterdam, The Netherlands, The Netherlands, IOS Press (2009) 375–382
6. Joolingen, W.R., Zacharia, Z.C.: Developments in inquiry learning. In Balacheff, N., Ludvigsen, S., de Jong, T., Lazonder, A., Barnes, S., eds.: Technology-Enhanced Learning. (2009) 21–37
7. Noss, R., Hoyles, C., Mavrikis, M., Geraniou, E., Gutierrez-Santos, S., Pearce, D.: Broadening the sense of 'dynamic': a microworld to support students' mathematical generalisation. Special Issue of The International Journal on Mathematics Education (ZDM): Transforming Mathematics Education through the Use of Dynamic Mathematics Technologies **41**(4) (2009) 493–503
8. Self, J.: Bypassing the intractable problem of student modelling. In: Proceedings of the Intelligent Tutoring Systems conference, , ITS'88. (1988) 18–24
9. Woolf, B.: Building Intelligent Interactive Tutors, Student-Centered Strategies for Revolutionizing E-Learning. Elsevier & Morgan Kaufmann (2008)

10. Bunt, A., Conati, C.: Probabilistic student modelling to improve exploratory behaviour. User Modeling and User-Adapted Interaction **13** (August 2003) 269–309
11. Cocea, M.: Learner modelling in exploratory learning for mathematical generalisation. In: International Conference on Adaptive Hypermedia (AH'08). (2008)
12. Dimitrova, V., McCalla, G., Bull, S.: Open learner models: Future research directions. Int. Journal of Artificial Intelligence in Education **17**(3) (2007) 217–226
13. Veermans, K., van Joolingen, W.R.: Combining heuristics and formal methods in a tool for supporting simulation-based discovery learning. In: Intelligent Tutoring Systems. Volume 3220., Springer (2004) 217–226
14. Ohlsson, S., Mitrovic, A.: Constraint-based knowledge representation for individualized instruction. Computer Science and Information Systems **3** (2006) 1–22
15. diSessa, A.A.: Epistemology and systems design. In diSessa, A., Hoyles, C., Noss, R., Edwards, L., eds.: Computers and Exploratory Learning. Springer (1995) 15–29
16. Papert, S.: Teaching children to be mathematicians vs. teaching about mathematics. Artificial Intelligence Memo Number 249 (July 1971)
17. Bull, S., Brna, P., Pain, H.: Extending the scope of the student model. User Modeling and User-Adapted Interaction **5**(1) (March 1995) 45–65
18. Balacheff, N., Sutherland, R.: Epistemological domain of validity of microworlds: the case of logo and cabri-géomètre. In: Proceedings of the IFIP TC3/WG3.3 Working Conference on Lessons from Learning. (1994) 137–150
19. Mavrikis, M., Gutierrez-Santos, S.: Not all wizards are from oz: Iterative design of intelligent learning environments by communication capacity tapering. Computers & Education **54**(3) (2010) 641–651
20. Pearce, D., Poulovassilis, A.: The conceptual and architectural design of a system supporting exploratory learning of mathematics generalisation. In Cress, U., Dimitrova, V., Specht, M., eds.: Learning in the Synergy of Multiple Disciplines. Volume 5794. Springer Berlin Heidelberg, Berlin, Heidelberg (2009) 22–36
21. Gutierrez-Santos, S., Cocea, M., Magoulas, G.: A case-based reasoning approach to provide adaptive feedback in microworlds. In: Intelligent Tutoring Systems, ITS'2010. (to appear)
22. Gutierrez-Santos, S., Mavrikis, M., Magoulas, G.: The layered development and evaluation for intelligent support in exploratory environments: the case of microworlds. In: Intelligent Tutoring Systems, ITS'2010. (to appear)