Proceedings of Workshop on Learner Modelling for Reflection,
International Conference on Artificial Intelligence in Education 2005, Amsterdam, pp. 81-90.

A Framework for Designing and Analysing Open Learner Modelling

Susan BULL[1] and Judy KAY[2]

[1] *Electronic, Electrical and Computer Engineering, University of Birmingham, U.K.*
[2] *School of Information Technologies, University of Sydney, Australia.*
s.bull@bham.ac.uk, judy@it.usyd.edu.au

Abstract. Recently learner models have been opened to the learners they represent. However, as yet there is no standard way of describing and analysing open learner models. This is in part due to the variety of issues that can be important or relevant in any particular model. Nevertheless, this lack of a framework to discuss open learner models makes it difficult to compare the features of open learner models in different systems. We believe this is a serious barrier to the effective use of open learner modelling. This paper presents such a framework and gives an example of its use to describe a system.

1. Introduction

Self [1] points out that various components of learning environments can be usefully made open, stating that "it is only the fact that learner models have traditionally been 'closed' that opening them up is such a big deal". However, this does not diminish the importance of open learner models.

Open learner models are models of the user that are available for viewing - usually by the learner, and sometimes also by others. The externalisation of the learner model can be in a simple form such as a skill meter showing learner progress as a subset of expert knowledge [2]; or the information presented can be more complex, such as a graphical externalisation of a Bayesian network [3]; a hierarchical tree structure [4]; a conceptual graph [5]; or textual descriptions of knowledge and misconceptions [6]. Students can even be involved in the maintenance of their learner model, for example by being able to edit it [4], or by negotiating the contents of the learner model with the system [5,6]. Furthermore, learner models can be opened not only to the learner modelled, but also to other users (e.g. peers or instructors) [7].

As indicated by the above, interactions with an open learner model, and presentation formats of open learner models, differ. Other differences include the extent to which a learner model is accessible to the learner; who may initiate access (learner and/or system?); whether the learner has access to information regarding uncertainty in the model; flexibility of access to the model. In fact, the seemingly simple matter of making a learner model open turns out to involve a quite complex range of choices, each with significant implications. This has important implications for the designers of new personalised teaching systems with open learner models. It also makes it harder to see the important differences between ways that existing systems have made use of openness of learner models. Furthermore, descriptions of open learner models do not follow any standard (or even any pattern). This paper addresses the need for an open learner modelling framework for use by researchers,

for the description and analysis of open learner models, and to support the design of open learner models in new systems.

2. The SMILI☺ Open Learner Modelling Framework

Table 1 shows our proposed SMILI☺ framework for *Student Models that Invite the Learner In*. The columns correspond to purposes or goals of the openness of the learner model: these are the reasons *why* a system would make the model open. The rows correspond to the elements and means of achieving openness. The cells marked with 'X' indicate particularly strong connections between purposes and elements. In the remainder of this section, we discuss several elements of the framework to give the flavour of each part.

The column labels indicate important *purposes for openness:* improving accuracy of the learner model by allowing learners to contribute information; promote reflection, an important metacognitive foundation for learning; help plan and/or monitor learning based upon the foundation of information available in the learner model; facilitate collaboration because partners can improve understanding of themselves and each other by gaining information from their respective learner model(s); afford learners greater control over learning through greater control over their learner model; and the privacy issue of the right to view data about oneself. A starting point in the design or analysis of openness of the learner modelling in a system is to treat this set of purposes as a checklist to review, carefully considering whether the system needs each (or, alternatively, carefully considering whether it can be argued that the system does *not* need each). Once the purposes, or goals, of openness are identified, one needs to determine how to achieve them. The rows identify issues to consider at this stage. They are numbered for reference. The 0-th is of a different character from the others, which all have choices indicated in the second column. We now briefly review each of the purposes, with discussion of the elements that are important, hence marked with a cross in the table. At this stage some of the discussion is necessarily speculative, and will need to be refined and extended as users of the proposed framework apply it to their systems.

Centrality of Openness of the Model:
Openness may be more or less central to a system's aims. Any of the listed purposes for opening the model could make the openness central. However, the learner's right to view their model is of a different character from the other purposes in that it would generally not affect the learner's normal interaction with the system.

Accuracy of the Model:
We discuss this in some detail to illustrate the interpretation of the figure and the framework. First, we consider the three aspects of element 1, the extent of openness. There is an especially strong connection between opening the model to increase its accuracy, and allowing complete access. If a learner is to contribute information to improve model accuracy, they need to be able to find out what is there. Particularly if the model is small or easy to understand, complete access may improve accuracy. However, partial access can also enhance accuracy, and partial access can therefore be helpful. Indeed, for complex learner models, the complete model may be so overwhelming that the learner could not easily correct problems. In this case, partial access may be more effective.

We now consider the second aspect of the extent (1), the types of modelled aspects: knowledge and difficulties. To improve model accuracy by allowing input from the learner, the learner should have access to representations of both their knowledge (content known

Table 1. The SMILI☺ Open Learner Modelling Framework

Open Learner Model	*Properties*	*Accu-racy*	*Reflec-tion*	*Plan / Monitor*	*Collab-oration*	*Control*	*Right to view*
0. Centrality of openness							
1. Extent of model accessible	Complete	X				X	X
	Partial	X					
	Knowledge	X	X	X		X	X
	Difficulties	X	X	X		X	X
	Learning issues	X				X	X
	Social issues	X				X	X
	Preferences	X				X	X
	Other	X				X	X
2. Presentation	Textual (i.e...)						
	Graphical (i.e...)						
	Summary						
	Overview						
	Targeted Details	X				X	X
	All Details	X				X	X
3. Similarity to underlying representation	Identical	X					X
	Similar	X					X
	Different						
4. Access to uncertainty	Complete	X				X	X
	Partial			X			
	None						
5. Role of time	Previous						X
	Current	X	X	X		X	X
	Future						X
6. Access method	Inspectable		X	X		X	X
	Co-operative	X					
	Editable	X				X	
	Negotiated	X				X	
7. Access initiative	System initiated	X					
	Learner initiated	X		X		X	X
8. Access to sources of input	Complete	X				X	X
	Partial	X					
	None						
	System	X	X	X	X	X	X
	Self	X				?	X
	Peer					?	?
	Teacher					?	?
	Other program					?	X
	Other					?	?
9. Control over accessibility (to others)	Complete				X	X	X
	Partial				X		
	None						
	System					X	X
	Self	X				X	X
	Peer					?	X
	Teacher					?	X
	Other program					?	X
	Other					?	X
10. Awareness of effect of model on personalisation	Complete					X	
	Partial			X	X		
	None						
11. Flexibility of access	Complete					X	
	Partial						
	None						

or knowledge 'level'), and about problematic areas. This may range from an indication of the areas in which they are having difficulty, to specific misconceptions. Only if the learner has access to positive and negative data (if both are modelled), can they make an informed decision on the correctness of the data. Of course, if the model is an overlay, only 'known' concepts, or an estimate of the knowledge level, will be available. Increasing the accuracy

of such a model will probably be focussed around knowledge only. (It is, however, possible that a learner could provide information that they do *not* know X, if X is not a precise detail that requires comprehension before articulation.) For accuracy, it is also important that the learner have access to all other model data.

The third aspect of extent shows four quite different types of modelling information. For similar reasons to the above, all these aspects of the learner should be available to improve the model accuracy. At this point, we note that it is quite feasible to have the model only partially open, but still to make those parts include all types.

We now consider the role of presentation (2) of the model for the goal of accuracy. The framework identifies only that the learner must be able to see at least targeted details, in the case of partial access, or the full model otherwise.

A model should be available in a form similar or identical to the underlying representation (3) for greater accuracy, as long as the presentation is understandable.

Complete access to uncertainty (4) is important for accuracy. For example, suppose a learner is aware they do not fully understand a particular concept. If they see an indication of the system's uncertainty of this aspect, they can make more sense of the model. By contrast, they may not accept that the model shows they do understand it (if the system withholds the information that it is uncertain).

The role of time (5) for accuracy is shown with current data being important.

The learner must have access (6) to the model to improve accuracy. They must be able to provide information to the model, or about the model, otherwise they cannot influence its accuracy. Usually this involves model inspection (not necessary where learner and system co-operate in the contribution of data, but the learner cannot view its contents). Co-operative models and negotiated models (where student and system discuss and agree the contents) are usually designed at least in part to improve accuracy. Similarly, this is usually an intended purpose of editable models (which allow the student to change the model directly), though here control is given to the learner.

Mixed initiative (7) is important for accuracy. The system should be able to request information when required. It is also important that the learner can choose to offer information; openness motivated by the goal of accuracy, is based on the expectation that the model may be wrong, and that the student may be able to help.

Similarly, the learner should have some degree of access to different sources (8) of input. If the model is constructed mainly by the system or jointly by system and student, the situation is straightforward: the learner needs access to both sources. However, if the model also contains data from others, e.g. teacher, this assessment may be more accurate. A similar argument could be made for peer contributions. Conversely, the student's self-assessment may be correct. However, as it is difficult to resolve discrepancies between peer- and self-assessment (and indeed, a learner may change a peer-assessment inappropriately), for the purpose of an accurate model, access to peer contributions may not be essential. Unless the system can accurately determine the reliability of evidence from different sources, the only point about which it can be confident is that the student's self-evaluations probably reflect their beliefs.

Control over who (or what) else can access the information (9) is largely not helpful for accuracy. Indeed, learner control over accessibility to others could result in decreased accuracy of the model. Certainly if the learner could withhold some data from the system, this would limit its ability to accurately model them.

Awareness of the effect of the model on adaptivity (10) is in itself not important for accuracy. However, such awareness may raise motivation to reach an exact model.

Finally, the flexibility of access (11) - i.e. are there choices about how to access the model? - is not important for accuracy, though it might be desirable that a learner can

access precisely those (and only those) contents required, presented in the manner they wish to view them, and perhaps interacting using a method they prefer.

This more detailed discussion of the elements that are important for accuracy of the model was intended to introduce the elements that characterise openness and support the various motivations for opening learner models. For the rest of this section, we will focus on the less obvious parts of the framework.

Reflection:

Where the aim of opening the model is to promote reflection, access to information about knowledge and difficulties (1) is important. A learner could reflect simply on information about their knowledge, as this could influence their confidence in their progress, while perhaps also helping to raise awareness of what they do not know. However, if the system models difficulties or misconceptions, reflection on these is important. Whether this information should be complete or partial will depend on the particular system. Similarly, in some systems it may be helpful to encourage reflection on other model contents, though this may not apply as a general rule.

It may seem surprising that we do not show that reflection, in general, requires a particular level of access extent (2) or that the model presentation to be similar or identical to the underlying representation (3), or that it is different. The critical factor is that the model is easily understood.

It may not be necessary to have access to uncertainty (4). If the learner notices disparities, this may lead them to reflect further! However, a question of trust in the utility of the model as a source of information for reflection may arise if they perceive it to be withholding information.

The role of time (5) for reflection on knowledge must focus on the current model, as this gives the learner the clearest indication of where they are. In some systems, reference forwards and backwards in time might be helpful.

To reflect on their learner model, the learner must be able to see the contents (6) – i.e. an inspectable model is assumed. This may be sufficient to encourage reflection. Other access methods are not critical.

While it is desirable that a system prompt (7) the learner to reflect on their model where this would be useful, and that the learner can undertake unsolicited viewing of the model, this combination is not essential. Conceivably, either approach separately could work well in a learning environment designed to promote learner reflection.

Of greatest importance, in most systems, students will need access to the system's inferences about their beliefs (8).

The remaining elements are not shown as important for reflection. In essence, the reflection column of the table indicates that the critical aspects for supporting reflection are that the learner be able inspect (6) the system's view (8) of their knowledge and difficulties (1) at the current time (5).

Planning/Monitoring Learning:

This goal is very similar to reflection, in that it involves the learner in assessing how they are doing, and how to use this information to decide on future learning goals. Accordingly, this discussion deals with the differences between this column and that for reflection. One difference is the greater importance of at least partial access to uncertainty (4). This supports decisions about the reliability of information in the model, which is used to help plan and monitor progress. It is likely that partial access will be sufficient, as it is only knowledge of the existence of uncertainty, or an indication of the extent, that the learner really requires for this purpose.

To monitor or plan learning, learners need to access the current model. This may be inspectable only. Users must be able to initiate inspection (7) to plan and monitor effectively, as autonomous learners. Most important is access to system inferences, as the system will have been designed to facilitate learning in an effective manner.

It is important the learner is aware of the effect of their model on the interaction (10), otherwise they will not be able to use their model data effectively in planning.

Collaboration:

Although they could be potentially relevant in a specific system, most open learner modelling properties are not essential to support collaboration - e.g. it is not crucial that users have complete or partial model access (1), as long as they have sufficient access according to the purpose of the open learner model.

Although access to representations of knowledge (1) is likely to be important in many environments, to support collaboration a system may instead be designed to match learners, or support them in some other way, according to social issues or preferences. Thus there is no pre-specified link with the extent of the model accessible.

Similarly, presentation method (2) is not strongly linked to purpose of collaboration (though it may be important in specific settings). A number of problems can arise, e.g. if a group model is presented to co-present peers, what format should the presentation take? Such issues may be resolved in different ways in different systems. Furthermore, in a collaborative context, users may still view their model individually. As in the cases described previously, it is difficult to prescribe a presentation format that would generally suit all students, all domains and all likely educational aims.

To support collaboration, it is not necessary that the presentation is identical or similar to the underlying representation (3). What is important is that the student can understand the model for the purpose that it was made accessible in the environment.

Access to uncertainty about representations (4), although potentially helpful in some settings, is probably not crucial in a collaborative environment.

With the other listed purposes of an open learner model, access to the current model (5) is important. In a collaborative environment this may also be so, but it may be, for example, that the learner can view the predicted future state of a peer, and then choose to work with that student once they have reached that state. Alternatively, a student might note that someone had previously had difficulties similar to their own, that are no longer problems according to their current model. This could indicate to the learner that this person may be particularly suitable to help them.

As learners may not only have access to their own model, issues such as whether an individual can edit another person's model (6) need to be considered; and whether all peers have the same kind of access to the same learner model.

Whether learner or system initiates model access (7) will depend on the system's aims. For example, in a system that matches partners, if it is a student's responsibility to find a partner by comparing learner models, they will need to access them. However, if the system does the matching, it may show learners the relevant models to help collaboration, once they are matched and maybe already working on a problem.

In a collaborative environment there are potentially many sources of data for the leaner model (8), as peers could be involved in providing data about each other. A learner model may be an individual model and/or it may be a model of all group members combined in some way from their individual models, or maintained according to group decisions. The extent of access to sources of input to the (individual or group) model is less critical in general, but as for other uses of open learner modelling, information modelled by the system is central to an open learner model to support collaboration. If peers or others contribute data to an individual's model, this may be accessed if it helps the collaborative

experience, but it is not usually essential that learners know that specific data comes from specific individuals (or peers in general).

It could be argued that the learner should have control over accessibility (9) of their model to peers (or others). This may be complete or partial - e.g. a learner might give permission for their data to be available to any peer, or they might restrict their individual model from (some) others, while it is still available as part of an aggregate model comprising data from all students. This, however, may be more to do with privacy than facilitating collaboration. From the perspective of *collaboration for learning*, at least in some systems, access to the models of others might be useful.

If the model is to be used to help students collaborate, to understand its relevance in a collaborative environment, students should have some awareness of its effect on personalisation (10). This is unlikely to be complete awareness, but it should be sufficient for students to understand the model's purpose. This should help motivate interaction with others if some contents available relate in some way to other learners.

Control Over the Model:

Control may refer to a learner's control over learning (related to learner autonomy); the interaction; or the model. These notions are linked - if the purpose of the model is understood, control over the model will lead to some degree of control over the interaction, resulting in increased control over learning.

Where learner control is central to a system (0), many properties of open learner modelling are key. If the aim is to allow users greater or full control over the model, interaction and/or learning, they should have complete access (1) to all contents.

Presentation mode (2), as other cases discussed, must be understandable with reference to controlling learning, the interaction, or the model. However, the learner should be able to access all details, or just those relevant for a particular purpose. The latter holds true especially if the model is large or complex - presentation of too much information may not be helpful, and may even be confusing. In practice it could result in little learner control being effected through the open learner model. A summary of contents or an overview of the model is not necessary (though it may be helpful).

Similarity of the presentation to the underlying representations (3) is not relevant, for similar reasons to those discussed concerning the potential complexity of the model being overwhelming if rendered at the wrong level or form.

It is important that learners can retrieve information about the uncertainty (4) of model data, since they can only have full control if they have all relevant information.

It is most important that the learner can access their current model (5), as it is their current learning or interaction over which they can have immediate control. Previous models, as a representation of previous states, will often not be manipulable. This may not apply, however, if the system designer accepts that a student's control over their model should be ultimate, or that the student may be able to correct something that was inaccurately represented in a previous model. Nevertheless, a question is whether a historical record should be alterable later. Arguments for access to future models are similar, with the additional difficulty that they are less certain.

The model should be at least inspectable (6). This may not allow students to alter its contents, and would therefore not provide control over their learner model or indirect control over the interaction, but it could offer greater opportunity to take control over learning through identification of the extent of their knowledge and potential problems (as long as the system allowed them some guidance of the interaction). A negotiated model would offer greater control, as the student could influence the model by argumentation. An editable model would give the greatest level of learner control.

A student must be able to initiate access (7) to their model to control it fully.

The student should have complete access to some sources (8) of input to their model. This applies in particular to the system's model and learner's own contributions. To help learners gain more control over their learning, it may be useful to have access to other sources. For example, an environment might allow the learner to view contents inferred by the system and also data contributed by peers. The student could inform the system of the information sources that should be regarded as most accurate to help shape the interaction. Against this is the possibility of others being more capable of making judgments about the student. Nevertheless, for the question of control, it is the learner's own control that is the issue. A different scenario might be a learner able to alter data supplied by another, if they believe it incorrect. For learner control this may be acceptable. However, for educational purposes it may not be appropriate to rely too heavily on the learner's views if they contrast with those of other people.

The learner should have complete control over the accessibility to other people (9), of their learner model, as well as to other programs using their learner model.

In order that a learner has control over their model, interaction and learning, they need full understanding of the effect of the model on personalisation (10).

Flexibility of access to the learner model should be complete (11).

Right to View the Model:

If learners should be allowed access to their data, as a right, this access should be complete (1). This includes representations of knowledge, difficulties, learning and social issues, preferences, and any other information (1).

The format of viewing is not prescribed, but learners should be able to access whatever details they wish - all, or selections (2). The latter issue of selecting data to view, is important to enable access in a manner which allows understandable presentation (information the learner wants, without being overwhelmed by additional data).

If opening the model as a right, the model should be viewable in the same (or similar) form to the underlying representations (3). This does not prevent other presentations being used in addition, to facilitate viewing and support learning. Moreover, if right to view implies right to understand, additional representations may be vital.

Since learners should have access to all data, this includes uncertainty (4). Similarly, access should be provided to all data held over time (5) - previous, current, future, if the system can predict that. The model need only be inspectable (6). Access must be able to be initiated by the learner (7).

Students should have access to complete information regarding sources of input (8). This applies to input from the system, explicitly from themselves, or other programs. In the case of input by peers - the student should know which data came from peers, but the privacy of the peer contributing data also needs to be protected. This is a difficult issue: whose right is greater? Ways around ethical problems might be to allow peers to contribute information only if they agree to their identity being revealed. However, this may stifle peer contributions; resulting in a decrease in the utility of learner models that can benefit from peer offerings. Whether the teacher's contributions should be named, is perhaps less controversial. However, does this violate teachers' rights? Conversely, if teachers are not identified, does this violate students' rights? Similar arguments apply where other people contribute information.

A learner should have complete control over access to their learner model by others (9). This is essential if considering open learner modelling from the perspective of the rights of individuals. This control applies to all potential accessors: the system, themselves, peers, teachers, other programs, and any other potential accessor.

The remaining aspects (10, 11) are less important for this goal.

3. An Example

The various aspects of open learner models are not usually described in sufficient detail to enable a reconstruction of Table 1 showing how some of the main open learner models to date, fit into the design space. We therefore use one of our own systems as an example.

Mr Collins [6] is a system with an open learner model for language learning, aimed at university level students. While several purposes of opening the model, and elements of openness, are relevant, only the most critical are given here as an illustration of the framework in use (Table 2). Mr Collins aims to improve the accuracy of learner modelling using a negotiated model, while also promoting reflection through the negotiation process Thus the openness of the model is vital. A negotiated model is necessarily inspectable, and the negotiation process can be initiated by either learner or system. Access to sources of input to the model from both system and learner, is complete. Full access is available to data about knowledge, difficulties, other learning issues (e.g. language transfer), preferences (learning strategies). The above are useful for both purposes of opening the model described here.

Table 2. Mr Collins in SMILI☺

Open Learner Model	*Properties*	*Acc*	*Ref*
0.Centrality of openness		X	X
1. Extent of model accessible	Complete	X	X
	Knowledge	X	X
	Difficulties	X	X
	Learning	X	X
	Preferences	X	X
2. Presentation	Textual	X	
	Overview		X
	Targeted		X
	All Details	X	X
3. Similarity underlying	Different		X
5. Role of time	Previous		X
	Current	X	X
	Future		X
6. Access method	Inspectable	X	X
	Negotiated	X	X
7. Access initiative	Sys initiated	X	X
	Stu initiated	X	X
8. Access to sources of input	Complete	X	X
	System	X	X
	Self	X	X

Presentation is by text and tables - detailed text being important to ensure that the learner understands the model contents before trying to negotiate changes. The level of detail to support reflection is up to the learner. Given that promoting reflection is a key aim of the open learner model, it is important that the learner can understand the model contents. Thus the representation is different from the underlying Prolog. To achieve an accurate model it is important that the learner can access the current model. To prompt reflection the learner can also track their progress over time, both backwards and forwards to anticipated future states.

Using SMILI☺, contrasting Mr Collins with another system is quite easy - for example, the editable learner model for the SAM text editor (also for university students) [4]. While, like Mr Collins, several purposes and elements of openness are relevant, the main aims are to increase model accuracy by permitting the learner to directly change

contents with which they disagree, thereby also allowing the learner to take greater control over the interaction. The two main 'purpose' columns would be *accuracy* and *control*. The main difference between the systems would be *access method* (editable rather than negotiated). Other differences include the *presentation format* (though similarly to Mr Collins, the format does differ from the underlying representation); and *role of time*. In both systems, the third most relevant column would be *planning/monitoring*, in particular the use of the open learner model to help students gauge their progress, and identify topics and concepts requiring more focussed effort.

4. Summary

The SMILI☺ Open Learner Modelling Framework provides a method to describe and analyse open learner models. The framework should also facilitate consideration of issues relevant to designing the openness of learner models in future adaptive learning environments. The framework also suggests dimensions of useful evaluations: for example, has the selection of the knowledge presented actually supported the goals of the openness? We expect that, for certain systems, some of the links between elements and purpose of modelling suggested as generally important, and those described as less crucial, may differ. The framework is intended to be flexible enough to allow for this currently, as unusual approaches to open learner modelling are described, and in the future as the field evolves.

References

[1] Self, J. (1999). Open Sesame?: Fifteen variations on the theme of openness in learning environments, *International Journal of AI in Education* 10, 1020-1029.

[2] Weber, G. & Brusilovsky, P. (2001). ELM-ART: An Adaptive Versatile System for Web-Based Instruction, *International Journal of AI in Education* 12(4), 351-384.

[3] Zapata-Rivera, J-D. & Greer, J.E. (2001). Externalising Learner Modelling Representations, *Proceedings of Workshop on External Representations of AIED,* AIED 2001, 71-76.

[4] Kay, J. (1997). Learner Know Thyself: Student Models to Give Learner Control and Responsibility, *International Conference on Computers in Education*, AACE, 17-24.

[5] Dimitrova, V. (2003). StyLE-OLM: Interactive Open Learner Modelling, *International Journal of AI in Education* 13(1), 35-78.

[6] 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, *Proceedings of AI in Education*, AACE, 501-508.

[7] Bull, S., Mangat, M., Mabbott, A., Abu Issa, A.S. & Marsh, J. (2005). Reactions to Inspectable Learner Models: Seven Year Olds to University Students, to appear in *Proceedings of Workshop on Learner Modelling for Reflection*, International Conference on Artificial Intelligence in Education, Amsterdam, 2005.