Proc. Int'l Workshop on Learner Modelling for Reflection, to Support Learner Control, Metacognition and Improved Communic. between Teachers and Learners, in conj.with AIED2005, Amsterdam, July, 2005. Copyright 2005.

Dimensions of Transparency in Open Learner Models[1]

Steven Tanimoto [2]

Department of Computer Science and Engineering
University of Washington, Seattle, WA, 98195, USA.

Abstract. The design of learner models that are open to the student's perusal is challenging, because a variety of competing objectives must be reconciled: comprehensiveness versus comprehensibility, and student control versus model validity. This paper suggests one approach to meeting the challenge that begins by identifying three dimensions of transparency in learner models. These provide a framework in which to design appropriate views of the data and limited controls, that best satisfy the requirements for the model. Part of one solution is to logically and physically distribute the records that comprise the model. The current status of learner modeling in the INFACT system is described.

Keywords. learner model, student model, user model, transcript, computer-based learning environments, transparency, tutoring system, open learner model, metacognition, dimension.

1. Introduction

A *learner model* is a computer-based data management component or system that contains information about a person's learning activity. It typically forms a part of a larger system such as a learning management system or an intelligent tutoring system. An *open* learner model is one with specific provisions for the learner to have one or more views of the information in the model[2]. Furthermore, these views are relatively comprehensive, hiding relatively little from the learner. An open model can be contrasted with a closed model in which the student has no direct view of the model's contents.

An important potential advantage of open learner models is that they may encourage valuable metacognitive activity and thereby help a learner to learn more effectively[7,2]. Another possible benefit is that they may permit modelling errors to be detected more readily so that they can then be corrected[3,1].

Influences promoting open learner models include the following: the trend toward accountability in artificially intelligent systems, exemplified by explanation features in expert systems of the 1970s and 1980s; movements to create standards for learning technology; efforts to help learners take better advantage of the growing array of resources

[1] Research supported in part by the US National Science Foundation under grant EIA-0121345.

[2] Correspondence to: Steven Tanimoto, Box 352350, Department of Computer Science and Engineering, University of Washington, Seattle, WA 98195, USA. Tel.: +1 206 543 4848; Fax: +1 206 543 2969; E-mail: tanimoto@cs.washington.edu.

on the World Wide Web; gaining students' trust[14]; and recognition by psychologists of the value of metacognition by students. Technology to support proactive learners[9] and more detailed transcripts for learner-to-resource matchmaking[10,11] also call out for open learner models.

There are several challenges in designing open learner models. These challenges arise primarily from conflicting demands for which good compromises must be found. One axis of conflict is openness versus confusion; learner models may be complex and contain components that are difficult to explain, and presenting their details to unsophisticated learners may not only confuse them but upset them. Another challenge for open modelling is keeping the model valid. If a model shows a learner's warts, he or she may be tempted to "fix" it by hacking into it to alter or remove unflattering parts.

To meet the challenges, three aspects of transparent models are identified. These form a basis that spans a space of abstract designs for learner models. A good design for a sophisticated open learner model will tend to be at an extreme point of the space.

2. Transparency in Learner Models

A *transparent* system is one that makes some of its inner workings visible to the user[12]. It can be contrasted with a "black box" that accepts inputs and returns outputs but hides the transformation mechanism. Transparency is often a desirable feature in software systems, because it can reveal to users how the system works, helping to engender trust, permit error detection, and foster learning about how software systems work.

Let us consider the application of transparency to some learner models. To start, let's look at some relatively simple kinds of models: overlay models and locus-within-strand models. Then, we'll examine more challenging ones such as facet-based models and student achievement scores computed using complex inference processes.

2.1. Student-Understandable Learner Models

An overlay model typically takes advantage of a taxonomy of skills and concepts within an academic subject and represents the student's state of knowledge using a set of numerical scores, with one score per skill or concept. It is easy to make such a model transparent by simply providing to the learner, either on demand, or through the initiative of the program, a list of skills and concepts with the student's current scores for them alongside them. If the list is organized hierarchically, then section (or other subdivision units) scores and a total score can also be given. Such a model and its view do not communicate to the learner the reason for any of the basic scores. The learner might assume, for example, that the scores are the results of answering questions correctly or incorrectly.

A strand-based learner model organizes skills and concepts into linear sequences called strands. A strand typically consists of a progression of subtopics in which one subtopic may have prerequisite subtopics that appear earlier in the progression. A strand-based tutoring system might start a session with a student at the beginning of a strand and only allow the student to progress to the next subtopic in the strand when its preceding topics have been mastered sufficiently. The student's state within the strand is considered to be completely represented by the index of the last subtopic mastered (the student's "locus" within the strand.). This kind of learner model is somewhat like a coarse-grained overlay model, because a student may have separate loci in a number of different strands at the same time. However, the meaning of the locus is somewhat different from the meaning of the overlay model's mastery value for a single topic.

It is easy to open up a strand-based learner model by identifying the subtopics in the strand and showing the learner the locus, either in real-time, on-demand, or at the end of a session.

If learner models were always this simple, it would be fairly easy, in principle, to keep them transparent and fully open to learners. There would still be issues to resolve regarding best modes and styles of presentation, though. Complex models, however, pose significant challenges to designers of open learner models.

2.2. Difficult Learner Models for Students

Let's consider two somewhat more problematical kinds of learner models. One is called a "facet-based" model, and it is based on the work by Minstrell to categorize common misconceptions in physics[8]. Like an overlay model, a facet-based model contains a construct for each skill, concept or subtopic within an academic subject. However, the construct is not simply a number that represents degree of mastery but a probability distribution over a set of conceptions and misconceptions. In the simplest case of such models, the probability values are restricted to being 0 or 1, so that only one facet (conception or misconception) is indicated for a given subtopic for a given student at a given time.

Making a facet-based learner model transparent to the student is problematical for two reasons. First, facets are defined in pedagogical terms as well as subject-area terms, and students cannot be expected to understand either of these sublanguages; generally, students are not teachers, and students are not experts in the subjects they are currently studying. Second, the probability values assigned to the various conceptions and misconceptions may be the result of processes that are incomprehensible to the student. Achieving a useful kind of transparency for facet-based learner models therefore 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. Without such a mechanism, an open facet-based learner model would at least have to have accompanying documentation that says as much as "students are not expected to completely understand the classifications of their states of knowledge of particular topics or the inference processes by which those classifications are made."

A kind of learner model complexity different from the explicit incorporation of misconceptions is complexity which occurs in the inference processes that produce model values. A good example of an inference process that is difficult for students to understand is Latent Semantic Analysis (LSA), which is a technique that uses numerical matrix computations to compare textual documents. It has been used for essay grading and student written-answer classification. LSA typically is applied by converting a student's essay into a vector of word occurrence counts. That vector is compared with vectors representing, say, good and bad essays, or essays in various topical or stylistic categories. Providing an accurate description of this process to, say, an English major, would probably not be very helpful.

2.3. Gaming the System

There is another, very different, problem for achieving transparency with a learner model based on a technique such as LSA. Although the numerical method itself is not something that a student could understand without having studied linear algebra, there are properties of LSA that a student can easily understand that teachers and testing agencies may wish to keep hidden from students. LSA-based scoring ignores word order in the input and bases its results only on the frequencies of occurrence of the words. A student who knows this might be tempted to game the grading system, find out what sorts of words will be required to be used for a particular kind of essay, and submit gobbledygook on a writing assignment that nonetheless satisfies the LSA-based assessment system. This is, of course, not specifically a problem for learner models, but for educational assessment in general. However, learner models are products of educational assessment, and making them open may expose parts of the assessment process that will no longer work properly if exposed. An open learner model should not be one that says, “Here’s how to fool the system.”

3. Validity in Learner Models

In the preceding section, we developed the issue of transparency versus learner confusion — some models may contain information that students cannot be expected to understand. Now let us consider another challenge for transparency in learner models: maintaining the validity of the information in the model.

There are two threats to the validity of data in the model. The more serious of these is that students, once aware of what their learner models represent, and perhaps inadvertently empowered to edit those models, make changes to the models’ data that corrupt the models and that make them unfit for the purposes for which they were designed, e.g., to improve learning. The other threat may arise as a design bias.

3.1. The Risk of Tampering

Transparency has some obvious benefits. However, as a means of raising students’ awareness of their weaknesses as well as their strengths, transparency can beget a desire to control. The extent to which a learner model becomes a representation of success or failure in any learning domain can lead to a lessor or greater wish on the part of the learner to change the representation (rather than the phenomenon it represents). Just as some students have hacked into school computer systems to change their grades, some students may attempt to gain editing access to details of their learner models and change them, instead of trying to change their own knowledge.

Transparency without appropriate validity protections would pose a substantial risk. The modern-day epidemic of cheating in school in countries such as the United States could be the downfall of open learner models.

3.2. The Risk of Design Bias

The other threat is unlikely to be catastrophic, but may have a subtly negative influence on the effectiveness of the models in guiding educational software systems. That threat is that designers, avoiding the incorporation of model components that are problematical for transparency, weaken the models' pedagogical value. For example, a designer might decide that it is too difficult to have the software explain to the learner what a particular facet assessment means and therefore decide that facet-based model components should not be included in the model. The components of the model may tend to be limited to lowest-common-denominator constructs: those that can easily be understood by all members of the expected learner population. Constructs with any degree of pedagogical sophistication would be avoided.

In the worst case, the model would consist of nothing but scores on tests and questionnaires, and would not contain any diagnostic results involving expressions of uncertainty.

4. Dimensions of Transparency

The fundamental challenge of designing transparent systems is answering the question of what to show and how to show it. The following three questions about this challenge lead us to corresponding "dimensions" of transparency:

- How much of the learner model is made available to the learner? The answer, which we could somewhat simplistically represent as a percentage, lies along an axis that we can call the *quantitative dimension* (or alternatively it could be called the "data dimension.")
- How much support is provided in explaining and interpreting for the student the learner-model data that is made available? This is the *interpretive dimension.*
- To what extent are authentication facilities visibly integrated into the model? This is the *validation dimension*.

In each of these dimensions the designers of a learner model may strive, through an appropriate mechanism, to achieve a balance between the needs of individual students and the need for consistent, reliable systems.

4.1. Quantitative Dimension

The question of what data within the model is made available can be considered as an elaboration on the question of how much is made available. Answering the *what data* question is clearly necessary in the design of a learner model, but it appears to be less fundamental, from a philosophical point of view, than the question of whether the model will be 100 percent open or not, or whether it will be open at all.

4.2. Interpretive Dimension

The interpretive dimension calls attention to the fact that nontrivial software development may be required to make complex models understandable by students. There is probably no limit on the extent to which explanation facilities might be taken in order to help students appreciate pedagogical judgments represented in the learner model.

There may appear to be some ambiguity between these first two dimensions, because an interpretation may call upon additional information that might seem to be a part of the model but is not explicitly recorded there. To explain a diagnosis ("student has misconception 2 on the topic of additive color mixing"), some context must be provided, which seems to be increasing the quantity of information disclosed (first dimension). However, if someone made a list of the items in the model and then chose one item to portray, one could argue that any interpretative information presented along with this item is attributable to this item and not the other items in the model. Thus the interpretive dimension can be kept distinct from the quantitative one.

Hansen and McCalla[6] have argued for “active learner models" that are constructed on-demand for particular purposes. With an active model, if a student wanted to know something about herself, she might query the model, and it would synthesize an explanation for her from the more-or-less raw evidence it encapsulates. It is more difficult to place such an active model on the interpretive dimension, because in a sense, an active model is itself a sort of interpretive process, and to interpret it would be to interpret and interpreter. In order to make such an active model transparent, it then becomes necessary to reveal the mechanism that generates the new model components.

4.3. Validation Dimension

The validation dimension arises directly in response to the challenge of maintaining model validity while putting it directly into the hands of students and opening it up. A learner model that falls at the high end of this axis has explicit validation facilities built into it in a way that both the learner and any agents accessing the model can see and activate. Much as a job resume lists personal or job references that can be contacted for information about the applicant, validation links in the open learner model identify agents such as web servers that are prepared to corroborate or elaborate upon elements in the student’s copy of the model. It is important not only that these exist in the model, but that they be visible to the learner. The transparency of the validation facilities can help the learner to put faith in the model and encourage the learner to maintain the integrity of the model.

One approach that addresses the validity issue in a novel way is the “negotiated learner model"[1]. With this method, a student may inspect the system-generated part of her model and try to change it through a kind of negotiation process. If the system refuses to change its evaluation on the basis of new evidence provided by the student, the student may still succeed in registering her disagreement in the model and thereby recording the objection. This by itself doesn’t prevent hacking, but it may reduce the temptation to edit values inappropriately by providing a legitimate “channel” for objections.

In addition to the issue of whether or not parts of a model may have been forged, in some situations a model could be incomplete; it could be a partial copy of the full model, or it could be a new fragment that has not yet been joined to or synchronized with its main model[4,11]. The possibility that alternative versions of a learner model may exist means that additional identifying information is required to specify, as clearly as possible, the relationship between the given version of the model and a reference version of the model.

5. Learner Models in INFACT

INFACT is a suite of tools for online communication, construction, and assessment[13]. At its center is a database containing evidence of student learning and evaluations of learning. The evidence consists of students' textual discussion postings, graphical sketches (made with an online applet), and user-interface events posted by tools for computer programming and image processing. The assessments consist of "facet-assessment records" each of which comprises an evidence identifier, a facet (conception or misconception) identifier, a student identifier, an assessing teacher or agent identifier, a timestamp, and a certainty value. All the facet assessment records for a given student, together with the evidence pointed to by these records, constitutes a learner model.

Learner models in INFACT have some transparency features but are not yet very open. Rating them roughly on each of the three dimensions gives an abstract idea of what kind of system we have—*Quantitative: medium; interpretive: medium; validation: medium*. With regard to the quantitative aspect, facet assessments are, by default, private to the teacher at this time, although most of the evidence on which they are based (messages and sketches) remains available to the students. INFACT provides two interpretive mechanisms that teachers can configure to give instant feedback or delayed feedback to students. Configuration requires a significant effort, because rules and/or Bayes nets must be authored (using INFACT's assessment development tools) for each facet within the topic being assessed.

Along the validation dimension, INFACT provides two features that contribute to maintaining the validity of each learner model. One of these is simply that all models reside on a protected server, and students cannot directly edit a model. The other is that each facet assessment record includes an assessor ID field, and thus a human (if the assessor was a person) could theoretically be asked to verify an assessment description that happened to be in doubt, or an agent (if the assessment was performed by software) could have its rules inspected by a teacher to check for consistency with the assessment.

In the future, we wish to add new transparency features to INFACT, both in order to evaluate their effects on learning, and to provide more options to teachers and students in organizing the online learning experience.

6. Discussion

Good strategies are needed for designing an open modeling system to take the three dimensions into account. Here are some possible aspects of a strategy.

On the quantitative dimension, the question of what parts of the model to show or not show can be divided into four parts: the easy and innocuous, the easy and more valuable, the harder but meaningful, and the inappropriate. The easy and innocuous material consists of the raw evidence that is collected in the model. If a student wants to see this, it can easily be shown, although some of it might not be very intelligible. The possible benefits to the student of offering this view include a feeling that the system is being open, the possibility of increased trust, and motivation to perform better due to the knowledge of being watched.

The easy and more valuable material consist of items that can be directly shared with students or shared with only a small amount of interpretation. If a student achieves a notable milestone in the course of a lesson, as this fact is recorded in the model, it can be shared with the student as it is recorded. Also in this category are skill meter readings[5]; they provide displays derived from values in the model using a relatively small amount of interpretation.

The harder but meaningful components of the model are those that require either or both of (a) nontrivial interpretation, or (b) judgment in real time about the appropriateness of showing the material. An example of this is an explanation that a student has a misconception about a concept. Such an explanation may well need to be synthesized just-in-time, shaded in such a way as to take into account the computed probabilities of the student's holding various alternative beliefs about the concept. Making a judgment about whether to provide such an explanation might take into account both the strength of diagnosis achieved and what is known about the student's willingness to accept such a report.

Material that is inappropriate for the student to see, if any, might include any of the following: (a) model variables for which no satisfactory explanation mechanism has yet been implemented in the system, (b) elements intended only for the system, such as checksums or other details, and (c) negative or even derogatory comments from a system or an exasperated teacher.

The dimension that affords the greatest challenge is probably the interpretive dimension. It is important because students cannot make use of complex assessment information if it is not presented in ways that are meaningful for them. Interpretation may be increasingly difficult as the sophistication of the assessment methods grows.

We can hope that the validation dimension might turn out to be just a technical implementation issue, and that once solved, it need not be something that students, teachers, and designers need to keep thinking about. However, mechanisms to engender trust in the validity of documents typically require the existence of trusted agents and trusted signatures, and there is always the potential for trouble there.

7. Acknowledgements

Thanks to Adam Carlson, Earl Hunt, Daryl Lawton, William Winn, the student developers of INFACT, the anonymous reviewers, and the National Science Foundation.

References

[1] P. Brna, J. Self, S. Bull, and H. Pain. Negotiated collaborative assessment through collaborative student modeling. In R. Morales, H. Pain, S. Bull and J. Kay (eds), *Proc. Workshop on Open, Interactive and Other Overt Approaches to Learner Modelling*, held in conjunction with AIED 99, Lemans, France, 1999, 35-42.

[2] S. Bull. Supporting learning with open learner models. *Proc. 4th Hellenic Conference with International Participation: Information and Communication Technologies in Education*, Athens, 2004.

[3] S. Bull, and H. Pain. '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 World Conference on Artificial Intelligence and Education*, AACE, Charlottesville VA, pp.501-508. 1995.

[4] S. Bull, A. T. McEvoy, and E. Reid. Learner models to promote reflection in combined desktop PC / mobile intelligent learning environments. *Proceedings of the Workshop on Learner Modelling for Reflection*, held in conjunction with AIED2003, Sydney, 2003.

[5] A. Corbett, and A. Bhatnagar. Student modeling in the ACT programming tutor: Adjusting a procedural learning model with declarative knowledge. In A. Jameson, C. Paris, and C. Tasso (eds), *User Modeling: Proceedings of the Sixth International Conference*. New York: Springer, 1997, pp.243-254.

[6] C. Hansen and G McCalla. Active open learner modelling. *Proceedings of the Workshop on Learner Modelling for Reflection*, held in conjunction with AIED2003, Sydney, 2003.

[7] J. Kay. A scrutable user modelling shell for user-adapted interaction. Ph.D. dissertation, Dept. of Computer Science, University of Sydney, Australia, 1998.

[8] J. Minstrell. Facets of students' knowledge and relevant instruction. In Duit, R., Goldberg, F., and Niedderer, H. (eds.), *Research in Physics Learning: Theoretical Issues and Empirical Studies*. Kiel, Germany: Kiel University, Institute for Science Education, 1992.

[9] J. Schoening. Creating lifelong learners through quality assurance. *Educational Technology and Society*, Vol. 5, No. 1, 2002. ISSN 1436-4522.

[10] S. Tanimoto. Beyond the naivety of grades: Educational record keeping for the twenty-first century. Technical Report 92-07-09, Dept. of Computer Science and Engineering, Univ. of Washington, Seattle WA, USA, July, 1992.

[11] S. Tanimoto. Distributed transcripts for online learning. *Journal of Interactive Media in Education,* 2001 (2) [www-jime.open.ac.uk/2001/2]. Published 10 Sept., 2001. ISSN:1365-893X.

[12] S. Tanimoto. Transparent interfaces: Design issues. In *Proceedings of the Workshop on Invisible and Transparent Interfaces, ITI 2004*, Gallipoli, Italy, 2004.

[13] S. Tanimoto, A. Carlson, J. Husted, E. Hunt, J. Larsson, D. Madigan, and J. Minstrell. Text forum features for small group discussions with facet-based pedagogy. *Proceedings of CSCL'02,* Boulder, CO., 2002.

[14] J. Vassileva, J. Greer, G. McCalla. Openness and disclosure in multi-agent learner models. In R. Morales, H. Pain, S. Bull and J. Kay (eds), *Proceedings of the Workshop on Open, Interactive, and Other Overt Approaches to Learner Modelling*, held in conjunction with AIED 99, Lemans, France, 1999.

[15] J.-D. Zapata-Rivera, J. Greer. Inspectable Bayesian student modelling servers in multi-agent tutoring systems. *International Journal of Human-Computer Studies*, Vol. 61, 2004, pp.535-563.