20000 Inspections of a Domain-Independent Open Learner Model with Individual and Comparison Views

Susan Bull and Andrew Mabbott

Electronic, Electrical and Computer Engineering, University of Birmingham,
Edgbaston, Birmingham, B15 2TT, U.K.
{s.bull, axm891}@bham.ac.uk

Abstract. This paper introduces a domain-independent open learner model with multiple simple views on individual learner model data. Learners can also compare their knowledge level to their peer group, and to instructor expectations for different stages of their course. The aim is to help learners identify their knowledge, difficulties and misconceptions; prompt reflection on their knowledge and learning; and facilitate planning. We present a study of OLMlets in 4 university courses, with 114 users making over 20000 learner model inspections.

1 Introduction

There is an increasing trend towards opening the learner model of intelligent tutoring systems (ITS) to the learner that the model represents. The views of such 'open learner models' (OLM) may be simple overviews of knowledge level, or more detailed representations of knowledge, concepts, interrelationships between concepts, misconceptions, etc. In this paper we are concerned with students' use of simple learner model presentations that can be easily deployed into a range of courses.

Most simple open learner models are in the form of a skill meter, e.g. [1,2,3,4,5]. However, as yet there has been no investigation into whether learners find this the most useful presentation of simple-format learner model information. Students have clear individual preferences when multiple detailed views are available, but with none of the views standing out as most useful for most students, or generally less useful [6]. Given that skill meters are now becoming more common, it is important to investigate the use of multiple views of simple learner model presentations, to find out whether differences in learner preferences also exist with a simple open learner model, and hence whether the more widespread use of skill meters over other formats, is justified.

This paper introduces OLMlets (small OLMs - as in piglets or rootlets). OLMlets can be used in a range of courses for which multiple choice questions are appropriate. It is necessarily simple in order to encourage instructors to input the multiple choice questions required, and deploy the system in their courses. As a simple OLM, it is not currently part of an ITS, although the OLMlets approach could be harnessed for use in full systems. The aim of using OLMlets independently of an ITS is to prompt students to reflect on their knowledge (including lack of knowledge and misconceptions), facilitate planning of future learning episodes, and encourage learners to take greater responsibility for their learning. It has also been suggested that learners might like to

M. Ikeda, K. Ashley, and T.-W. Chan (Eds.): ITS 2006, LNCS 4053, pp. 422 – 432, 2006.

compare their knowledge to that of their peers, or to the expectations of the instructor for the current stage of their course [7]. OLMlets also has such comparison views.

This paper presents a study of the OLMlets logs with 114 users in 4 university courses. 13113 questions were answered; 17000 (exactly) inspections of the individual views of knowledge level were made; 520 additional inspections of misconception descriptions; 1637 inspections of the peer comparison view; and 1296 inspections of the instructor expectations view - a total of over 20000 model inspections.

2 A Domain-Independent Simple Open Learner Model

OLMlets has an interface through which instructors can enter multiple choice questions and responses; indicate which answer is correct; define misconceptions and assign these to incorrect answers; and indicate which responses are incorrect (but where no associated misconceptions are identified). Images, superscript/subscript and limited special characters (e.g. Ω,Π) may be used. Thus OLMlets is suitable for many courses for which multiple choice questions are suitable. Instructors may also define expected knowledge levels for each topic, at different stages of the course.

Based on a student's answers to the questions, OLMlets constructs a model of the learner's knowledge level of each topic - represented by a number between 0 and 1, with 1 representing probable mastery and 0 indicating no knowledge. Misconceptions are identified by comparing user input to a misconceptions library for a course (created when an instructor defines misconceptions). The probability of a misconception being held is also represented by a number between 0 and 1, with 1 indicating high probability that the learner holds the misconception. The last five attempts at questions for each topic contribute to the learner model of an individual, with successively heavier weighting on the most recent of those five attempts. An 'unsure' option is automatically included with the response choices created by an instructor, in order that a learner is not forced to guess if they do not know the answer, thus avoiding knowledge or misconceptions being represented in the model based on guessing.

Because OLMlets is domain-independent, relying on instructors to input questions and building the learner model as defined according to instructor input, the learner modelling is not complex. There is no domain model for the differing content of the various courses. The views of the learner model for an individual are therefore simple. There are five views, as shown in Figure 1: skill meter (the most common form of simple OLM); graph (the bars from the skill meter located over, or to one side of a 'neutral' axis to help visualisation of positive and negative data); boxes (coloured to indicate knowledge level); table (knowledge level ranked according to proficiency); text (a summary of knowledge in the order topics were sequenced by the instructor).

The large image shows the skill meter view in the full window. The smaller images show the other forms for comparison. The learner may change view by clicking on the links at the top of the page. Clicking on the misconceptions link next to a topic results in a textual description of the misconception being displayed at that location. Clicking on the compare link provides a comparison of the user's own knowledge for a topic, to that of the other users in their course, as shown in Figure 2. The star indicates the learner's own knowledge level for the topic, for ease of comparison against the rest of the group. Clicking on the numbers below the view links in Figure 1 (for week, day or lecture number - method selected by the instructor), displays the learner's knowledge level against the expectations for that stage of the course, as shown in Figure 3 for the skill meters and text views. Thus the learner can compare their current knowledge to the current expectations, and to expectations for previous or future stages of the course. This is useful, for example, if they are behind: the student may wish to try to 'catch up', and viewing expectations for different points of a course may help them to set realistic goals not only according to the current expectations, but also taking into account their position with respect to the expectations at different stages of learning.

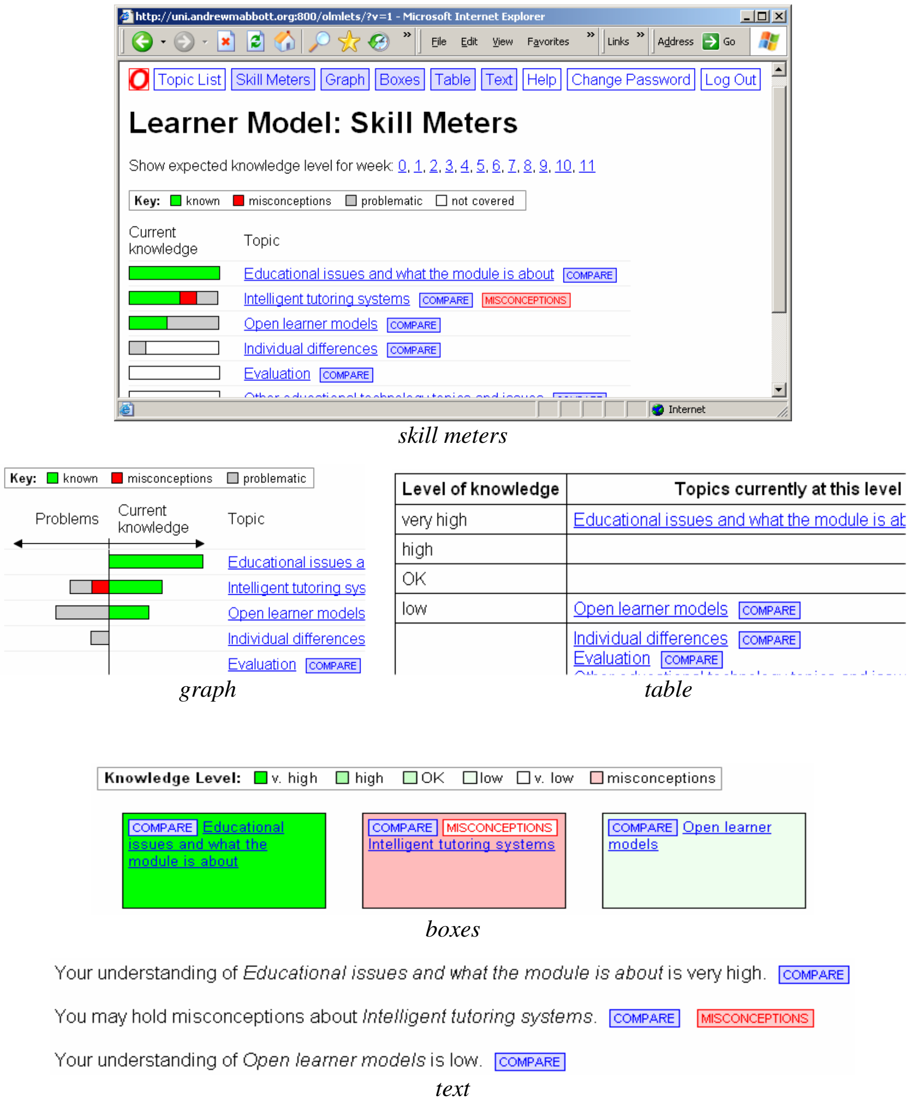

Fig. 1. The individual learner model views

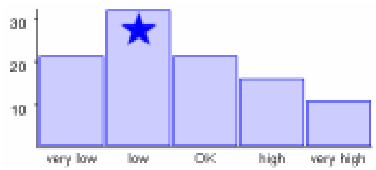

Fig. 2. Comparison to peers

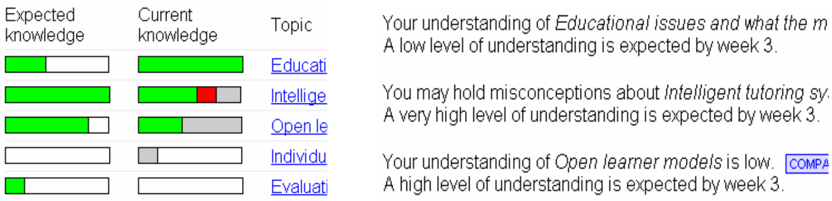

Fig. 3. Comparison to instructor expectations (skill meters and text views)

3 Use of the Individual and Comparison Leaner Model Views

While OLMlets could be incorporated into an ITS, it is currently deployed independently of a full system, alongside a range of courses. We present an overview of use in the first 4 courses in which OLMlets was deployed - the courses in which the teaching (but not necessarily the assessment) was completed at the time of the study.

3.1 Participants, Materials and Methods

Participants were 114 students in Electronic, Electrical and Computer Engineering at the University of Birmingham, using OLMlets in one of 4 courses as in Table 1.

Table 1. Courses in which OLMlets was used

Year	*Course*	*Taking Course*	*Used OLMlets*
1	Circuit Analysis	70	28 (40%)
1	Communication, IT & Lab Skills	108	47 (44%)
3	Interactive Learning Environments	29	29 (100%)
4/MSc/MRes	User Modelling	10	10 (100%)

Electronic, Electrical and Computer Engineering has 3 year BEng, 4 year MEng, and MSc/MRes degrees in subjects such as Electronic and Electrical Engineering, Communications Systems Engineering, Computer Systems Engineering, Computer Interactive Systems; and combinations with Computer Science, Languages, and Business Management. The 1st year Communication, IT & Lab Skills had students from the range of degrees as it is compulsory for all. The 1st year Circuit Analysis had students from the more traditional engineering degrees. The 3rd year Interactive Learning Environments had mainly students from the Computer Interactive Systems degree (which has a stronger focus on human factors and psychology in computing), but included some taking the Computer Systems Engineering degree, who chose it optionally. Students taking 4th year/MSc/MRes User Modelling were from a range of degrees.

The two 1st year courses introduced the OLMlets URL in a lecture, but had no further support. Those participating were therefore self-selecting. However, their data is still valid as we do not wish to argue that OLMlets should necessarily be used by everyone, if they have alternative successful learning strategies. Furthermore, the level of uptake (40% & 44%) is not especially low when considering that use was optional, and students had to familiarise themselves with OLMlets without support. There was no upcoming assessment for Circuit Analysis at the time of the study, thus students had no immediate assessment goal. The Communication, IT & Lab Skills course introduced OLMlets shortly before a multiple choice test was administered, the question types in OLMlets and the test being similar. The test contributed 25% to the final course mark. There may have been students who did not know about OLMlets in that course, as on the day that the URL was given out, there were only about 60 in attendance. Thus the percentage may be underestimating uptake from amongst those who knew about OLMlets. The 3rd year Interactive Learning Environments course introduced OLMlets in a lab at the start of the course, several weeks before the learner model of each student was assessed, the learner model counting for 10% of the course mark. The 4th year/MSc/MRes User Modelling also introduced OLMlets in a lab, a month before assessed reports were due, which comprised 100% of the assessment. The form of assessment (written report) was quite different from the OLMlets questions, though the questions addressed issues relating to the report requirements. Because these two courses introduced OLMlets in a lab, there was 100% uptake.

All groups had access to the five individual learner model views, and the group comparison - automatically generated based on all models. Expected knowledge was set in all courses. Misconceptions libraries were identified for 1st year Circuit Analysis and 3rd year Interactive Learning Environments. Students were not advised how to use OLMlets. Only in the 3rd year course did use contribute directly to assessment. Questionnaires were completed by the 3rd years at the end of the course. The analysis below is taken mostly from the system logs for all courses, but with some reference to the questionnaires. Questions required answers on a 5 point scale (strongly agree, agree, neutral, disagree, strongly disagree); and solicited open-ended comments.

3.2 Results

13113 questions were answered. 20523 learner model inspections were made. 1241 questions were answered in 1st year Circuit Analysis (average 44); 2051 in 1st year Communication, IT & Lab Skills (average 44); 9340 in 3rd year Interactive Learning Environments (average 322); 481 in 4th year/MSc/MRes User Modelling (average 48). Table 2 gives the breakdown for viewing the individual model views, which comprised 17000 of the model inspections. The final columns show the mean, median and range of inspections of individual views. The skill meter was the most frequently used by all groups. Apart from 3rd year Interactive Learning Environments, there was also usage, albeit at a lower level, of the graph view. There was occasional use of the other views (though these need to be accessed at least once before users can decide whether they are helpful). In the courses where the learner models were not assessed, the mean viewings per student was similar, at 31, 33 & 37, though the ranges were more varied. The median number of viewings was 16-28. The generally higher number of model inspections was in the more senior course, which had students with an interest in user modelling. In the course in which the learner models were assessed, the mean number of model inspections was much higher, at 491. Here the number of viewings ranged from 16 to 1112. However, only 3 learners made fewer than 100 inspections of their model, and only 5 made fewer than 200. The median was high, at 517. In this course more viewings were made of the learner model, than questions answered.

Table 2. Use of the individual model views

Course	*Skill M*	*Graph*	*Boxes*	*Table*	*Text*	*Mean: all views*	*Median: all views*	*Range: all views*
1 Circuit	534 59%	142 16%	92 10%	81 9%	62 7%	33	16	0 - 273
1 IT/Lab	997 68%	197 13%	90 6%	98 7%	95 6%	31	24	0 - 143
3 ILE	12110 85%	694 5%	590 4%	485 3%	357 3%	491	517	16 - 1112
4 UM	208 55%	63 17%	40 11%	37 10%	28 7%	37	28	14 - 82

45 of the 114 students used 1 view mainly, 44 of whom using the skill meter. 1 used mainly the boxes (usage defined by a view being selected at least 10% of times the individual model was accessed). 16 used 2 views, the most common combination being skill meter and graph, with 11; then other combinations, 2 of which did not include the skill meter, and 3 of which did not include the graph. All combinations included at least one of these views. 13 students used 3 views, each combination including the skill meter, with the second most frequent being the graph (9). 20 students used 4 views, with 2 not using the skill meter and 3 not using the graph; and 13 used all 5 views. Only 15 users did not use the skill meter as their most frequent view, and 72 were using the skill meter for over 50% of viewings of their individual model. Only the graph and boxes also had users accessing them at least 50% of the time, with 4 and 3 users respectively. 7 students did not access their individual learner model.

The most common strategy was to observe the model update after each question had been answered, with 5056 such occurrences, representing 30% of model viewings. 14 (12%) of the 114 users usually checked their model after each question, 11 of these in the 3rd year course in which the learner model was assessed. These users had attempted varying numbers of questions, the highest being a learner who attempted 626, where the model was checked after a single question had been answered on 449 occasions. 11 (10%) were checking their model on average after 10 or more questions, with the highest average being 16 (3 of these were from the 3rd year course). The remainder checked their model after answering 2-9 questions. The average for 1st year Circuit Analysis was 6.6; 1st year Communication, IT & Lab Skills, 4.5; 3rd year Interactive Learning Environments, 3.2; 4th year/MSc/MRes User Modelling, 7.8.

Misconception descriptions can be accessed in all views. Once opened, they remain visible until the user hides them. Thus there may be more occasions on which learners paid attention to misconceptions, than shown in the logs. Table 3 shows the number of times students opened a misconception in the courses in which misconceptions were modelled, and the mean, median and range. This includes only students who had misconceptions. Viewing misconceptions was higher in the 3rd year course. Only 1 student had no misconceptions. There was no clustering of viewing, with a spread over the range of 3-54 viewings - but with fewer at the higher end. In the 1st year course, 12 users had no misconceptions. Of the remainder, half viewed misconception descriptions once or twice, with most of the rest viewing fewer than 8 times.

Table 3. Inspection of misconception descriptions

Course	*Misconceptions*	*Mean*	*Median*	*Range*
1 Circuit Analysis	79	5	8.5	1 - 27
3 ILE	441	16	12.5	3 - 54

Table 4. Use of the peer comparison view

Course	*Peer View*	*Mean*	*Median*	*Range*
1 Circuit Analysis	217	8	5	0 - 68
1 IT/Lab Skills	326	7	5	0 - 34
3 ILE	1011	35	25	2 - 144
4 User Modelling	83	8	4	2 - 40

Table 5. Use of the instructor expectations comparison view

Course	*Expect. View*	*Mean*	*Median*	*Range*
1 Circuit Analysis	169	6	2	0 - 22
1 IT/Lab Skills	178	4	0	0 - 24
3 ILE	841	29	24	1 - 63
4 User Modelling	108	11	13	1 - 23

Table 4 shows use of the peer comparison view. There were 1637 inspections of the peer comparison. Some students were not interested in comparing their knowledge to that of others, with 19 of the 114 students not accessing this view. However, most accessed it several times - with means of 7 and 8; and medians of 4 and 5 for the courses in which the learner model was not assessed. In the course in which the learner model was assessed, students compared their knowledge to others' more often.

Table 5 gives the usage of the instructor expectations view. As with the misconceptions, the expectations comparison remains visible until the learner closes it. Therefore the logs may underestimate the attention users paid to this information. 1296 inspections were made. Most of these were in the 3rd year course where the learner model was assessed, although it was only the final model that contributed to the mark (i.e. assessment was *not* made at various points in the term to coincide with what students were expected to know). In the 1st year courses, inspections of the comparison to expectations were made occasionally, though 38 students were not interested in this information. In the Communication, IT & Lab Skills course, OLMlets was not introduced until shortly before the end of the teaching period (and so students had little opportunity to make use of this view). In the 4th year/MSc/MRes course, interest was higher, with 6 students making 12 or more inspections, but 4 making 5 or fewer.

The 3rd year group filled in a questionnaire at the end of the course. This was the only group that had completed all teaching and assessment at the time of the study. 23 questionnaires were returned, a 79% return rate. This paper has focussed on the logs rather than questionnaires. However, to demonstrate students' perceptions of the utility of OLMlets, and differences in preferences between individuals, we provide the results of 3 of the questions and excerpts from students' open-ended comments:

1. 22 students agreed or strongly agreed that using one or more of the individual views was useful in helping them identify their knowledge/difficulties; 1 user gave a neutral response.
2. 18 learners agreed or strongly agreed that the comparison to instructor expectations was useful; 5 gave a neutral response.
3. 20 students agreed or strongly agreed that the peer comparison view was useful; 2 gave a neutral response; and 1 disagreed.

- The graph view was less useful. The reason was because by having the grey and red areas separate it was difficult to see how much of the topic I had good knowledge of. In the skill meter view I could tell more easily as the 3 areas were all part of the same rectangle.
- The graph view is very much similar to the skill meter. They only differ as the graph has an axis. I found this view easier to interpret as there are results on both sides of the y axis.
- The only problem with the skill meters was that it was difficult to know what knowledge level I was in. This meant having to use other views such as the table for this information.
- I only used the table view towards the end when trying to get all my topics into very high.
- The boxes view was useful to me, as it gave me a graphical view of my progress as well as a textual description. This gave me much more information than the graph view on its own and a clearer description of how I was progressing than the skill meter. However, I did not find the boxes method useful for comparing my progress to what the lecturer wanted.
- [The boxes view] provides quite an easy way to compare current knowledge level to the expected ones, because it only contains a single colour tone rather than having to compare the proportion of good knowledge.
- The ability of checking my own progress against the rest of the students gave me a great deal of motivation, as I could actively view my rise in learning. Being quite competitive in nature this aid gave me a greater incentive to achieve.
- I found it reassuring to compare my results and see all students were struggling on a topic, also on the flip side it made me work harder on subjects in which I was below the average.

- Although I did prefer the skill meter view, the other views being available when I wanted them was of great use, as when I wanted to show my peers my model they could choose their favourite view to aid their understanding of my model.
- I did not think comparing my knowledge against what the lecturer expected me to know and to the knowledge of the rest of the group was helpful. When my knowledge of a topic was low, I felt low in confidence.

3.3 Discussion

Over 13000 questions were answered and over 20000 learner model inspections were made. It is perhaps not surprising that students viewed their learner model frequently in OLMlets, as it was the only feedback available. However, it was not expected that learners would make more model inspections than they answered questions, as occurred in the 3rd year course. In that course students were sensitive to the contents of their learner model, as it was assessed. Nevertheless, even in the other courses students were keen to see their learner model updating after only a few questions, as illustrated by the logs showing that only 11 of the 114 users answered 10 or more questions on average, before viewing their model. Most were viewing it more frequently, with the highest frequency (apart from the 3rd year course) being the course in which there was an upcoming test. Not surprisingly students felt motivated to perform self-evaluation that could lead to a higher result in the test. As stated above, learner model inspection was likely to be because students would otherwise have had no indication of their progress. The key question is, then, was inspection of the learner model made only because students would otherwise not receive feedback - i.e. is it simply that *some* information is better than *no* information? The questionnaire responses from the 3rd years suggest that students were indeed finding the learner model useful - both for viewing one or more of the individual views, and the comparisons. While in that course students were required to use OLMlets, use in the other courses - particularly the two 1st year courses where the URL was simply given out in a lecture, and the students had no particular interest in user/learner modelling - suggest that OLMlets was considered a helpful support alongside a lecture course. It would be interesting to see how use would compare if OLMlets was incorporated into an ITS.

Of the individual model views, the skill meter was the most commonly used, with most having it as their preferred view. Although the skill meter is the first in the series of links, and so more likely to be selected first, there was sufficient use of OLMlets that students tried all views. None is presented as a default, and so users have to make a choice about how to access their model. Skill meter usage provides hitherto lacking justification for the incorporation of skill meters ITSs or hypermedia systems that use this simple OLM format. Nevertheless, the graph was also used in combination with other views by 52 students, and 58 of those who used the skill meter were also using other views. Not all users had the skill meter as their first preference. Furthermore, questionnaire comments revealed that some students used different views for different purposes (e.g. to focus on their own knowledge and for comparison to the expectations; for a quick overview and to gauge knowledge level more precisely later). It is therefore suggested that, while it seems skill meters are a good choice for a simple OLM, additional views could be considered when designing the views of the OLM.

For the 2 courses in which misconceptions were defined, many students did view the misconception descriptions. In the 3rd year course where the learner model was assessed, the misconception inspections were higher as students had extrinsic motivation to ensure that their learner model was as good as they could get it. In 1st year Circuit Analysis, the mean and median viewings of misconceptions were 5 and 8.5 respectively (range 1-27). It should be noted that misconceptions remain visible once opened, until the learner hides them. Thus the actual attention paid to misconceptions may be higher than the logs suggest. Although not necessarily useful to all learners, it seems likely that there are sufficient who will use this information if it is available.

While some learners did not make use of the peer comparison view, most made several comparisons. Higher use was in the course that assessed the learner model. It is not clear whether this is higher simply because students used OLMlets more frequently, and so made the comparison more frequently, or whether they were really concerned about obtaining an equal or higher mark than the rest of the group. It would be interesting to pursue this question further. Several of the open-ended questionnaire comments certainly suggested that some students were motivated by being able to see their position in the group, and the fixed-response questions also suggested that students generally found this comparison useful. However, there was one student who found it demotivating to see that their knowledge was lower than that of other people.

The lecturer expectations view was used by many, most heavily in the course assessing the learner model. Nevertheless, there were students in the other courses using it. This may be because students find it useful to be able to make continued reference to their developing knowledge with regard to a target such as an upcoming test, as this provides external judgement rather than their own possibly inaccurate self-evaluation.

Where the learner model was assessed, there was higher use of the system and viewing of the model. More interesting is that students in other courses used OLMlets despite it not being directly assessed. Of course, the similarity of questions to the upcoming assessment in one of the 1st year courses no doubt contributed to students' motivation to use it. Nevertheless, it is unlikely that they would have continued if they had perceived no benefit. It appears, then, that a simple OLM can be a useful resource alongside a lecture course, even when not part of an ITS. However, our students are computer-literate, so the generality of this finding needs to be investigated in other course types, before recommendation for wider use of OLMlets can be made.

4 Summary

This paper has introduced OLMlets, a simple domain-independent open learner model with multiple individual and comparison views. Use of OLMlets in 4 university courses was described. Unlike with more detailed learner model information, where it has been found that students have preferences for a range of views [6], with this simple OLM the skill meter view was the most commonly used, with the graph in second place. It appears less critical, then, to provide a broad range of views if a simple format is used, and the growing use of skill meters in ITSs and educational hypermedia seems justified. Nevertheless, some users did use other views in preference to, or in combination with skill meters. Thus provision of a choice of view is still worthwhile. It was also found that students viewed a comparison of their own knowledge level to their peers; and instructor expectations of their knowledge level at different stages of the course. These features could therefore be usefully incorporated into systems.

References

1. Bull, S. & McEvoy, A. (2003). An Intelligent Learning Environment with an Open Learner Model for the Desktop PC and Pocket PC, *Artificial Intelligence in Education (Proceedings)*, IOS Press, Amsterdam, 389-391.
2. Corbett, A.T. & Anderson, J. (1995). Knowledge Tracing: Modeling the Acquisition of Procedural Knowledge, *User Modeling and User-Adapted Interaction* 4, 253-278.
3. Mitrovic, A. & Martin, B. (2002). Evaluating the Effects of Open Student Models on Learning, *Adaptive Hypermedia and Adaptive Web-Based Systems, Proc. Second International Conference*, Springer-Verlag, Berlin Heidelberg, 296-305.
4. Papanikolaou, K.A., Grigoriadou, M., Kornilakis, H. & Magoulas, G.D. (2003). Personalizing the Interaction in a Web-Based Educational Hypermedia System: The Case of INSPIRE, *User Modeling and User-Adapted Interaction* 13(3), 213-267.
5. Weber, G. & Brusilovsky, P. (2001). ELM-ART: An Adaptive Versatile System for Web-Based Instruction, *Int. Journal of Artificial Intelligence in Education* 12(4), 351-384.
6. Mabbott, A. & Bull, S. (2004). Alternative Views on Knowledge: Presentation of Open Learner Models, *Proc. Intelligent Tutoring Systems*, Springer, Berlin Heidelberg, 689-698.
7. Kay, J. (1997). Learner Know Thyself: Student Models to Give Learner Control and Responsibility, *Proc. International Conference on Computers in Education*, AACE, 17-24.