Visual Attention in Open Learner Model Presentations: An Eye-Tracking Investigation

Susan Bull, Neil Cooke, and Andrew Mabbott

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

Abstract. Using an eye-tracker, this paper investigates the information that learners visually attend to in their open learner model, and the degree to which this is related to the method of displaying the model to the learner. Participants were fourteen final year undergraduate students using six views of their learner model data. Results suggest some views of the learner model information may be more likely to encourage learners to inspect information about their level of knowledge, whereas in other views attention is directed more towards scanning the view, resulting in a lower proportion of time focussed on knowledge-related data. In some views there was a difference according to whether the learner model view was one of the participants' preferred formats for accessing their learner model information, while in other views there was little difference. This has implications for the design of open learner model views in systems opening the learner model to the learner for different purposes.

I Introduction

Open learner models - learner models that are accessible to users - are becoming more common in adaptive learning environments. Examples include simple presentations of the learner model data such as skill meters [1,2,3,4,5,6]; or more complex presentations such as a textual description of knowledge and misconceptions [7]; a text explanation of a fuzzy logic model [8]; a hierarchical tree structure [9]; a conceptual graph [10]; a graphical representation of a Bayesian model [11]. Reasons for opening the learner model to the user are varied (see [12]), but those particularly relevant to our current discussion include promoting learner reflection on their knowledge, facilitating planning, self-monitoring or navigation, and allowing the user to contribute information about their knowledge to improve the accuracy of their learner model in order to improve system adaptation.

While multiple representations to support learning or problem-solving have been investigated (e.g. [13,14]), the use of alternative presentations of the open learner model within a system has received less attention. However, in systems which do offer a choice of learner model presentations, it has been found that users have different preferences for how to access their model both when simple formats are used [15], and when more complex presentations are available [16]. Furthermore, students sometimes prefer or reject model presentation formats for what seem like similar

C. Conati, K. McCoy, and G. Paliouras (Eds.): UM 2007, LNAI 4511, pp. 187–196, 2007.

reasons, as illustrated by the following excerpts from open-ended questionnaire responses: "The concept map was the most useful as it shows the relationship between all subject areas and where my weaknesses lie"; "Concept map is a bit complex compared to the others, making it a bit difficult to understand". We therefore do not intend to recommend a single presentation or a specific set of presentation methods for open learner models based only on the type of information and level of detail that the learner model externalisations display. Rather, the purpose of this paper is to investigate the information that learners look at in their open learner model, taking into consideration whether the particular view is one of the user's preferred methods for accessing their model, and the structure of the model view. Based on this information we will consider requirements for the design of presentation formats for open learner models, which also allow for individual differences amongst users.

In line with previous studies of use of educational environments employing eye-trackers to record gaze, for example to help detect motivational factors [17] or enhance the interaction between a learner and software agent [18], we here describe a study of students' attention to the various components of the Flexi-OLM open learner model [19] revealed by use of an eye-tracker following students' inspection of six learner model views. While eye-tracking to determine gaze direction has been questioned as a way of identifying attention because it is often based on visual search (which is only one of the tasks a user may be undertaking) [20], here it is precisely visual search or scanning of the learner model that we are interested in measuring.

2 What Do Students Attend to in Their Open Learner Model?

In this section we investigate students' use of an open learner model with multiple views using the 'EyeLink 2' head-mounted binocular eye-tracking device which measured their gaze position to an accuracy of 0.5°-0.1° visual angle with a sample rate of 500 Hz, relative to a 22 inch computer monitor screen. The fact that the learner model information for an individual student was identical across learner model views (as these were representing the same underlying model), allows us to make comparisons between what learners visually attended to in the various views.

2.1 Participants, Materials and Methods

The Flexi-OLM open learner model was implemented for the domain of C programming. Participants were 14 final year undergraduate students who had completed a C programming course, and who had knowledge of what an open learner model represents, all having previously used a simple open learner model [1] in two of their courses. Students had not yet used more complex open learner model views as offered by Flexi-OLM, in any of their courses. However, they had participated in a two-hour lab session where their Flexi-OLM learner models were constructed, and their preferences for the learner model views investigated using questionnaires (data which was used in the current study). The eye-tracking study was therefore measuring their gaze in a learner model with which they were familiar.

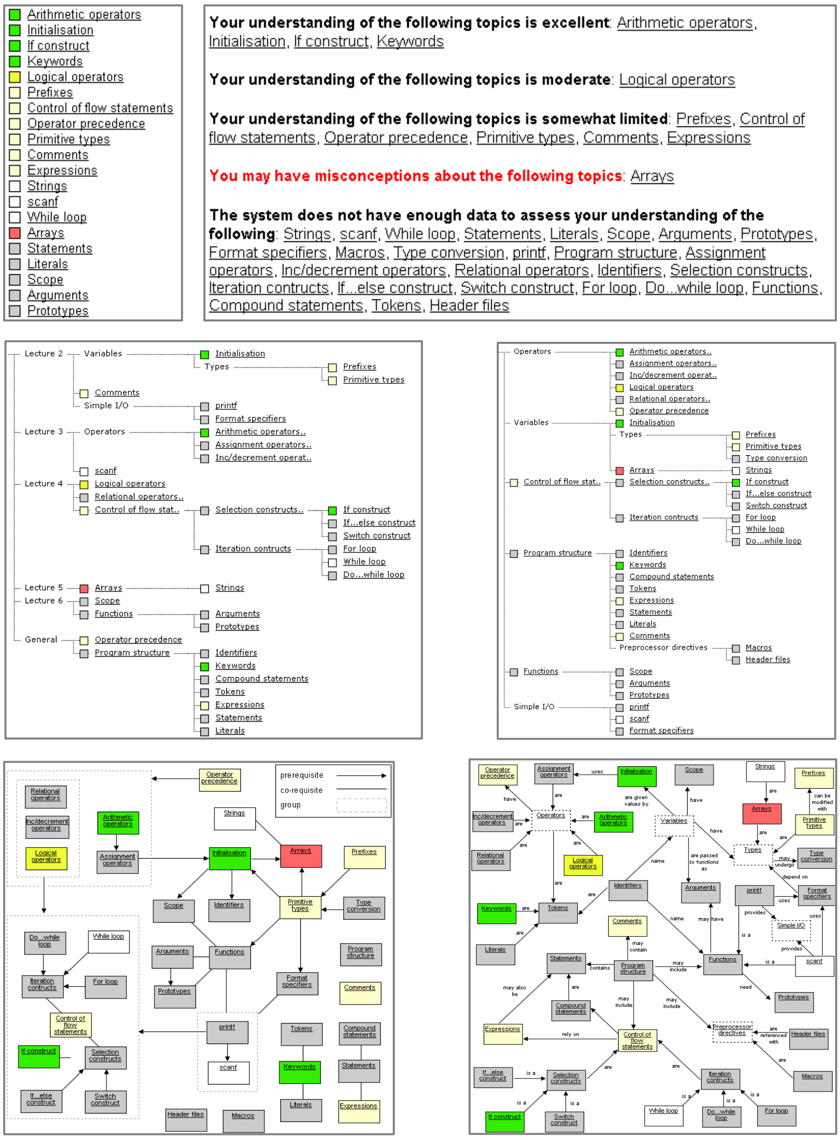

Fig. 1. The structure of the open learner model views: list ranked according to knowledge level, text summary, hierarchical lecture structure, hierarchical topic/concepts structure, prerequisite relationships, concept map

Flexi-OLM aims to help learners identify their learning needs for themselves, by permitting them access to the contents of their learner model constructed according to their answers to multiple choice and short answer questions. Flexi-OLM has seven views of the learner model data which are accessible at any time, to allow students to view the information about their understanding in the format or formats that suit them best: alphabetical index, list ranked according to knowledge level, text summary, lecture structure (i.e. structured according to the lecture course in which Flexi-OLM is deployed), related topics/concepts hierarchical structure, concept map, and prerequisite relationships. Previous studies have demonstrated that students differ in their preferences for the Flexi-OLM learner model views [16,19]. The eye-tracking was insufficiently accurate to determine the focus of visual attention confidently for the alphabetical index, so we here consider visual attention in the remaining six views. The differing layouts of these learner model views are illustrated in Figure 1.

In all but the text summary, coloured boxes or nodes indicate knowledge level (shades of green - for high levels of knowledge, through yellow, to white - representing low knowledge; and red for topics with misconceptions). Text statements of probable misconceptions are also given in each view, accessed by clicking on the links associated with the red nodes. In this study, given the size of the learner model and the fact that the previous lab session in which the learner models were constructed, lasted two hours only, all participants had more grey (insufficient data) areas in their learner models than coloured nodes. However, all data included here is from users who interacted sufficiently to allow a learner model with a variety of instantiated nodes to be built. This allowed investigation of visual attention in the context of a mix of instantiated and incomplete learner model information, potentially relevant to open learner model settings where students are encouraged to view their learner model to prompt self-monitoring or reflection, facilitate planning or navigation, or where they may contribute information directly to the model to increase its accuracy - thereby leading to improved adaptation.

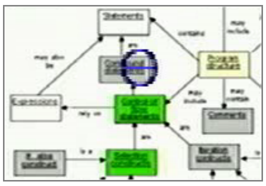

Fig. 2. Video excerpt showing part of screen illustrating visual attention

Figure 2 illustrates an excerpt from an eye-tracking video, showing the user's focus of visual attention. The 'EyeLink 2' system consists of a head-mounted camera system and two PCs for processing data and running experiments. On the head-mounted device, both left and right eye pupil position and the head position relative to computer monitor, were tracked. Combining the position of the head with the pupil movement relative to the screen, enabled recording of the gaze direction. Eye tracking data for both eyes was captured at a sample rate of 500Hz (2ms samples) with a resolution of 0.01° visual angle. The effective resolution, however, was limited by the calibration process, which provided an accuracy of typically 0.5° - 1.0° visual angle.

In addition to recording gaze direction, the eye tracker also detected eye movement types such as saccades, fixations and blinking. Fixation events, assigned to their nearest visual focus (i.e. nodes in the learner model), were used as the unit of analysis.

Sessions commenced by calibrating the eye tracker. This took 5-15 minutes. In 5 sessions, eye movement recordings showed poor calibration, so these have been omitted from this study. Rejection criteria were horizontal and/or vertical offsets in gaze direction, loss of gaze altogether, and general corruption in the eye signal, making the eye movement unintelligible. There are a number of causes for calibration failure. Overt indications of errors during recording occurred if a subject inadvertently moved the eye tracker when coughing, gesticulating or touching his or her face. As the human scalp is elastic, the head mounted device was also liable to moving. The eye tracker was also sensitive to the ambient light levels in its environment.

Once comfortable with the eye-tracker, students logged on to Flexi-OLM, accessing the learner model created in the previous lab session (see above). While students could move freely, they were requested to restrict their head movements as far as possible. They were instructed to continue using Flexi-OLM as they had previously, in order to avoid potentially biasing their choice of learner model views to access, and their behaviour within a view. Interaction may have involved answering additional questions, editing the learner model, or attempting to persuade the system to change the learner model contents (see [19]). This allowed the interaction to proceed as naturally as possible, observing attention in the context of using the system as a whole. The eye-tracking sessions lasted around 10 minutes to ensure that the calibration remained sufficiently accurate for the whole interaction. The eye-tracking data for each view was analysed from the videos, and compared against the data for the other learner model views and against the learner's preference for the view as indicated in their questionnaire responses (on a 5 point scale) in the previous lab session. As the participant numbers are relatively low, the figures were checked to ensure that they were not biased by unusual behaviour by a small minority of users.

2.2 Results

12 of the 14 participants stated in the questionnaire that the open learner model was useful (giving a score of 4 or 5 on a 5 point scale); 1 gave a neutral response (3); and 1 gave a negative response (2). 13 stated that they understood their learner model; the remaining student gave a neutral response. 4 students stated that they would regularly use 1 of the views only; 5 would use 2 views; 2 would use 3 views; 1 would use 4 views; and 2 would use 5 views.

During the 10 minute eye-tracking period, students spent an average of 145 seconds in their learner model (median - 143; range 70-261 seconds), spending up to 30 seconds in their learner model at any one viewing. The rest of the time was spent answering questions (all students), editing the model (7 students) and attempting to persuade the system that the learner model was incorrect (8 students). 5 students attempted to both persuade the learner model, and edit it directly.

Figure 3 shows the breakdown of visual attention to information in the learner model according to whether students were using one of their preferred views, as indicated in their questionnaire responses. 'High' indicates gaze focused on nodes showing a high level of knowledge for a topic or concept; 'medium' indicates a medium level of knowledge; 'low' indicates a low level of knowledge; and 'unknown' indicates that there is insufficient data on knowledge of that topic or concept for it to be modelled. The numbers in brackets show participants who claimed to be either users or non-users of a view (total = 84: 14 users x 6 views).

Figure 3 suggests that there is little difference in the distribution of visual attention according simply to whether a student is using one of their preferred views. However, there was a slight tendency for misconceptions to be accessed relatively more frequently in views that participants would be more likely to use.

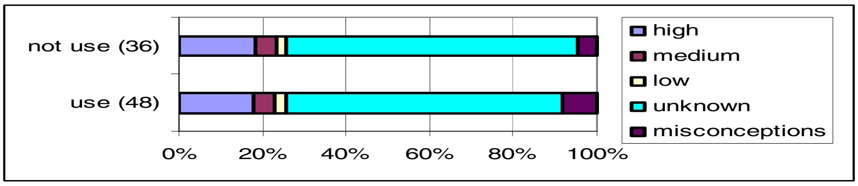

Fig. 3. Visual attention according to whether participants were users/non-users of the views

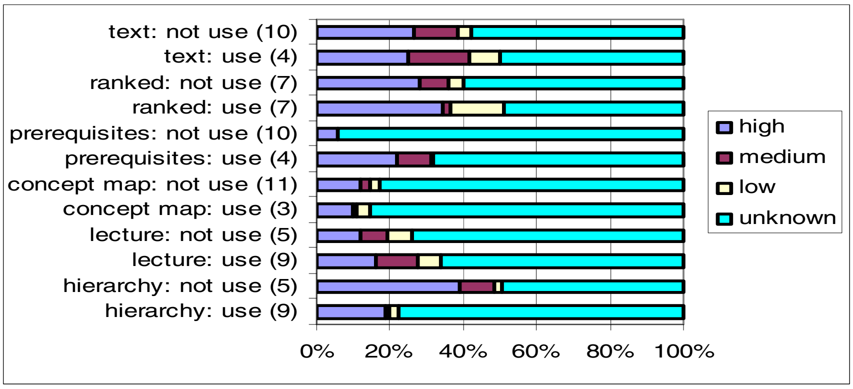

Fig. 4. Visual attention according to learner model view, and view users/non-users

Figure 4 gives the breakdown of the proportion of time participants spent viewing the various components of each of the learner model views according to whether they claimed that they would use the particular view in their questionnaire responses. As found previously [16,19], participants had different preferences for which views to use. Because not all participants had misconceptions, resulting in some of the categories in Figure 4 having no misconceptions to view, in order to compare the relative percentage of time directed to viewing knowledge level versus lack of data the figures for viewing misconceptions have been omitted. (In all but one category - non-users of the prerequisites view - in cases where students had misconceptions they tended to view them.)

Figure 4 shows differences between the learner model views with reference to the information users focus on. The text and ranked list appear to focus students' attention to a greater extent onto their level of knowledge than do the other views, regardless of whether these are the students' preferred views (40-51% of the time spent inspecting the view). The hierarchy view has a similar effect only for those who would tend not to use that view (50% of inspection time). The concept map focuses attention on knowledge to a lesser extent than the other views for both users and non-users (up to 18% of inspection time), as does the pre-requisites view for non-users, where focus on knowledge was particularly low (5%).

2.3 Discussion

Students generally claimed to find their learner model useful. Viewing the model prompted both selection of areas on which to answer questions, and interactions aimed at changing the learner model more directly. In this paper we are concerned with how students inspect their learner model rather than the resulting choice of task, as task selection may depend on what is found in the model. (Future work will investigate possible differences between browsing the model and using it for a specified purpose.)

When considering only whether students are using a learner model view that they regard as useful according to their individual preferences, there is very little variation in what they attend to - though there is a slightly higher tendency to access descriptions of misconceptions relatively more frequently in a view that they prefer. However, when separating out the data for the individual views it was found that variations between views resulted in gaze being directed towards different information. Using the concept map, regardless of whether participants considered this as one of their preferred views, they focussed less on their knowledge, exploring the areas with insufficient data in more detail. This may be due to the complexity of this view: students can easily see the differences in knowledge level at a glance due to the size of the nodes, but perhaps require more effort to gain an overview of the relationships between the concepts even when already familiar with the layout of the concept map; or, indeed, they may be particularly interested in gaining an understanding of such relationships. Of course, given the level of system usage at the time of the study there were still relatively many uninstantiated nodes in the learner model views, so this factor is likely to also contribute to the amount of time spent investigating this information. Nevertheless, the relative complexity of the model might still result in a broader spread of attention in a more complete model.

In contrast, using the text description and ranked list, although attention is still on the areas with insufficient data, 40-51% of the visual attention is directed at the nodes indicating knowledge level. This is likely to be because both the ranked list and the text summary sequence the information from known, to moderate knowledge, to limited knowledge, with the remainder of the concepts listed as having insufficient data to model them. Because of this sequencing of information this is not a surprising finding; however, we can now suggest that even when students are familiar with a learner model view, they may still follow it sequentially or focus more on the initial information if it is inherently a list format - regardless of their focus of attention in other layouts of the same information about their understanding. (It should be noted that preference for view has not been found to match a user's learning style (see [16]), as measured by Felder and Silverman's learning style classification [21].)

Interestingly, if our comment about ease of noticing knowledge level quickly in the concept map is correct, it is perhaps surprising that the same results were not observed in the pre-requisites structure. Clearly this was the case for the majority claiming to be non-users of this view; however, the 4 who found it a useful presentation format were looking at the nodes representing knowledge level to a greater extent. Perhaps for users who find that a structure showing pre-requisite relationships helps them to more easily identify the most appropriate areas to work on, are less distracted by the overall structure - a finding that did not seem to apply with the concept map.

Conversely, users of the hierarchical relationships structure seem to focus less on the knowledge-related nodes than do the non-users. It may be that users of this view are keen to understand the relationships and work out how their knowledge fits into the overall structure, whereas non-users do not find this helpful.

Although the number of participants is relatively low, there do appear to be some differences between the learner model views in terms of where gaze is directed - this sometimes being related to whether students find the particular view useful. However, it is rather early to be confident of the reasons for this - our suggestions are, at this stage, simply suggestions. But whatever the reasons, there does seem to be something happening. Congruent with one of the reasons suggested for offering multiple external representations in learning environments more generally [22], we have previously recommended that because students have different preferences for views to use when offered a choice, it is useful to offer multiple-view open learner models ([15,16]). Given the above findings we would now argue that consideration should also be given to whether it is important that students focus primarily on their knowledge level - for example to promote reflection and raise their awareness of their knowledge and difficulties, or to facilitate planning of learning episodes within or outside the system. Alternatively, a purpose of opening the model may be to allow the learner to provide information where there is, as yet, insufficient data, where focus specifically on those areas will be necessary - for example, where one of the aims is to use an open learner model to highlight gaps in knowledge to the learner; or if it will be important that the learner can review areas where their knowledge is low, as well as refer to areas with insufficient data in order to distinguish between these two types of information. An open learner model to facilitate navigation may also need to focus attention on both areas with low knowledge and areas with insufficient data - though other issues relating to the structure of the view may also be more relevant (for example, navigation might also be facilitated in particular by inclusion of information about pre-requisite relationships or the overall lecture structure of a course). In an open learner model designed primarily to allow the learner to correct inaccurate model data (or attempt to correct the data, but requiring system approval), it may be useful to use a presentation that encourages broader scanning. Thus, although we still maintain that individual preferences are important both to facilitate students' understanding of their learner model and to motivate them to access their model data, it may also be that different views of the learner model are better suited to different purposes for accessing the model. Students spent up to 30 seconds in their learner model at any one time. Therefore a format that encourages attention to be directed quickly at the components most relevant for the purpose of viewing the model in a particular system, would appear to be useful. This point does, however, need to be judged alongside consideration of a format that is appropriate for the domain. For example, if topics are defined very broadly in a system, it may be harder to design a useful concept map in sufficient detail; or if there is a range of topics and/or concepts that have few pre-requisite relationships, a pre-requisite structure will be less appropriate; or if the learner modelling is very coarse-grained, e.g. with concepts modelled simply as 'known' or 'not known', a ranked list may be less useful. Similarly, if a system expects users to understand relationships between different aspects of the domain, a learner model view that reflects these relationships may be more useful generally.

Of course, if students have only one or two learner model views available, they may focus more on aspects that they neglected in some of the views when they had a greater choice. Thus it may not be the case that a particular learner model view structure will necessarily encourage or inhibit attention to certain components. Nevertheless, given that students were often scanning their learner models quite quickly, it is still likely to be useful to consider what students would naturally attend to when designing an open learner model - if they find a particular structure difficult, they may be less inclined to use it.

Based on the above, our overall recommendation at this stage of the research in open learner modelling is to design open learner models that take into account: (i) the requirements of the domain as represented in the system; (ii) the educational aims of the system and the purpose of opening the model to the user; (iii) the individual preferences of the user; (iv) the user's focus of attention in the learner model.

Clearly there are limitations to this study, the most obvious being that using an eye-tracker may influence the natural choices of the user, given that they are aware that their gaze is being monitored. Furthermore, this was necessarily an experimental, lab-based study. It would obviously be impractical to measure use of Flexi-OLM in this way in a natural setting, where users are accessing the system regularly alongside a lecture course. However, given the results obtained, an interesting next step would be to follow up a set of users later on in a course to see whether their use of the learner model differs when they have had extensive previous interaction with it.

3 Summary

This paper has described an investigation into visual attention in an open learner model. Using an eye-tracker, students' use of a multiple-view open learner model was observed. Results indicate that, although in general students' preferences for learner model view do not affect the information that they visually attend to, some of the individual views do have a tendency to focus attention on information relating to knowledge level more than others. This suggests that, as well as considering individual differences in preference for the presentation of an open learner model, where it is important in a system that the learner model data is used for a particular purpose it may be useful to consider the information that students are likely to visually attend to in different model presentation formats.

References

1. Bull, S., Quigley, S., Mabbott, A.: Computer-Based Formative Assessment to Promote Reflection and Learner Autonomy. Engineering Education: Journal of the Higher Education Academy Engineering Subject Centre 1(1), 8–18 (2006)
2. Corbett, A.T., Anderson, J.: Knowledge Tracing: Modeling the Acquisition of Procedural Knowledge. User Modeling and User-Adapted Interaction 4, 253–278 (1995)
3. Linton, F., Schaefer, H.-P: Recommender Systems for Learning: Building User and Expert Models through Long-Term Observation of Application Use. User Modeling and User-Adapted Interaction 10, 181–207 (2000)
4. Mitrovic, A., Martin, B.: Evaluating the Effects of Open Student Models on Learning, Adaptive Hypermedia and Adaptive Web-Based Systems. In: Proceedings of Second International Conference, pp. 296–305. Springer, Heidelberg (2002)

5. Papanikolaou, K.A., Grigoriadou, M., Kornilakis, H., Magoulas, G.D.: Personalizing the Interaction in a Web-Based Educational Hypermedia System: The Case of INSPIRE. User-Modeling and User-Adapted Interaction 13(3), 213–267 (2003)
6. Weber, G., Brusilovsky, P.: ELM-ART: An Adaptive Versatile System for Web-Based Instruction. Int. Journal of Artificial Intelligence in Education 12(4), 351–384 (2001)
7. Bull, S., Pain, H.: Did I say what I think I said, and do you agree with me? Inspecting and Questioning the Student Model. In: Proceedings of World Conference on Artificial Intelligence in Education, AACE, pp. 501–508 (1995)
8. Mohanarajah, S., Kemp, R., Kemp, E.: Opening a Fuzzy Learner Model. In: Proceedings of Workshop on Learner Modelling for Reflection. In: International Conference on Artificial Intelligence in Education 2005, pp. 62–71 (2005)
9. Kay, J.: Learner Know Thyself: Student Models to Give Learner Control and Responsibility. In: International Conference on Computers in Education, AACE, pp. 17–24 (1997)
10. Dimitrova, V.: StyLE-OLM: Interactive Open Learner Modelling. Int. Journal of Artificial Intelligence in Education 13(1), 35–78 (2003)
11. Zapata-Rivera, J.-D., Greer, J.E.: Externalising Learner Modelling Representations. In: Proceedings of Workshop on External Representations of AIED, International Conference on Artificial Intelligence in Education 2001, pp. 71–76 (2001)
12. Bull, S., Kay, J. (in press). Student Models that Invite the Learner In: The SMILI Open Learner Modelling Framework, Int. Journal of Artificial Intelligence in Education
13. Ainsworth, S., Van Labeke, N.: Using a Multi-Representational Design Framework to Develop and Evaluate a Dynamic Simulation Environment. Dynamic Information and Visualisation Workshop, Tuebingen (July 2002)
14. Cox, R.: Representation Interpretation Versus Representation Construction: An ILE-Based Study Using switchERII. In: Artificial Intelligence in Education (Proceedings), pp. 434–441. IOS Press, Amsterdam (1997)
15. Bull, S., Mabbott, A.: 20000 Inspections of a Domain-Independent Open Learner Model with Individual and Comparison Views. In: International Conference on Intelligent Tutoring Systems, pp. 422–432. Springer-Verlag, Berlin Heidelberg (2006)
16. Mabbott, A., Bull, S.: Alternative Views on Knowledge: Presentation of Open Learner Models. In: Intelligent Tutoring Systems: 7th International Conference, pp. 689–698. Springer-Verlag, Berlin Heidelberg (2004)
17. Qu, L., Wang, N., Johnson, W.L.: Using Learner Focus of Attention to Detect Learner Motivation Factors. In: User Modeling 2005: 10th International Conference, pp. 70–73. Springer-Verlag, Berlin Heidelberg (2005)
18. Wang, H., Chignell, M., Ishizuka, M.: Empathic Tutoring Software Agents Using Real-Time Eye Tracking, Eye Tracking Research and Application. In: Proceedings of Symposium on Eye Tracking Research and Applications, pp. 73–78. ACM Press, New York (2006)
19. Mabbott, A., Bull, S.: Student Preferences for Editing, Persuading and Negotiating the Open Learner Model. In: Intelligent Tutoring Systems: 8th International Conference, pp. 481–490. Springer-Verlag, Berlin Heidelberg (2006)
20. Wood, S., Cox, R., Cheng, P.: Attention Design: Eight Issues to Consider. Computers in Human Behavior 22, 588–602 (2006)
21. Felder, R.M., Silverman, L.K.: Learning and Teaching Styles in Engineering Education. Engineering Education 78(7), 674–681 (1988)
22. Ainsworth, S.: The Functions of Multiple Representations. Computers and Education 33, 131–152 (1999)