

---

## **A haptic learner model**

---

**Tim Lloyd and Susan Bull\***

Electronic, Electrical and Computer Engineering,  
University of Birmingham,  
Edgbaston, Birmingham, B15 2TT, UK  
E-mail: TDL833@bham.ac.uk      E-mail: s.bull@bham.ac.uk

\*Corresponding author

**Abstract:** This paper describes an adaptive learning environment with an open learner model to promote learner reflection. The learner model can be presented using different techniques: textual and haptic. The paper investigates the potential for haptic learner models in terms of whether users can understand a haptic model, and whether they find it useful. Most liked both versions, but of those with a preference, more found the haptic method of interaction useful. However, some liked the textual version, suggesting the need for a variety of methods of accessing the model to be available in learning environments with an open learner model.

**Keywords:** open learner model; haptic learner model feedback; adaptive learning environment.

**Reference** to this paper should be made as follows: Lloyd, T. and Bull, S. (2006) 'A haptic learner model', *Int. J. Continuing Engineering Education and Lifelong Learning*, Vol. 16, Nos. 1/2, pp.137–149.

**Biographical notes:** Tim Lloyd recently obtained an MEng degree in Computer Interactive Systems from the Electronic, Electrical and Computer Engineering department at the University of Birmingham. The Computer Interactive Systems degree covers subjects ranging from mathematics and electronics to computer engineering and human factors in computer use. Tim implemented the haptic learner model, described in this paper, for his final MEng project.

Susan Bull obtained her PhD in Artificial Intelligence in Education from the University of Edinburgh in 1997. She is currently a lecturer in Electronic, Electrical and Computer Engineering, University of Birmingham. Her duties include teaching 'user modelling' and 'educational technology' subjects at both undergraduate and postgraduate level. Her research over the last ten years has focused mainly on adaptive learning environments with open learner models, where opening the learner model to the student is designed primarily to enhance learner reflection. Susan is cofounder of the 'Learner Modelling for Reflection (LeMoRe)' research network, one of the main aims of which is to promote understanding and use of open learner models to support student learning.

## 1 Introduction

Adaptive learning environments provide adaptive instruction, tutoring, scaffolding or other individualised learning support. To effect this individually tailored educational support, they require an understanding of the domain that is being taught, the learner who is using the system, and teaching strategies suitable for the particular subject and for a variety of individual users. Domain knowledge is defined according to what is known about the domain, its structure, interrelationships between concepts, etc., and teaching strategies are designed based on an understanding of relevant pedagogical considerations. The way in which adaptive learning environments maintain an understanding of the individual learner is to create a model of the user's beliefs (e.g., knowledge, difficulties and misconceptions), inferred during their interaction with the system. This may be derived from the student's answers to questions, the time they spend viewing information on different topics, their route through material, the help topics viewed, etc. This learner model is used in conjunction with the domain model to allow the learning environment to infer appropriate teaching or guidance strategies to suit an individual's current learning needs. Examples of such systems include the ELM-ART Lisp Course (Weber and Brusilovsky, 2001), which provides navigation support, recommending suitable topics for an individual to study according to the contents of their learner model; and INSPIRE (Papanikolaou et al., 2003), which suggests sequencing of materials and activities on computer architecture, according to a learner's preferred approaches to studying. While each of these environments makes individualised recommendations, they also allow the learner to have control over the interaction, i.e., in these systems, the learner is not forced to follow the system suggestions (as is the case with some adaptive learning environments).

Learner models are typically not seen by the user, but are used primarily to adapt the environment as appropriate for the individual learner, such as described above. However, allowing the student modelled to have access to their learner model data gives them additional information about their learning process and current understanding, providing them with another tool for learning. This can result in a greater insight into their achievements in a subject area, and where they need to focus more effort. A survey of 44 students in the Electronic, Electrical and Computer Engineering department at the University of Birmingham found that most would welcome the opportunity to view their learner model if it were open to them, as a means to prompt their reflection on important issues and problems (Bull, 2004).<sup>1</sup> This is despite this being an essentially reflective process, while many engineering students are more active learners (Felder and Silverman, 1988).

Some research has suggested that students will indeed view open learner models when given the opportunity (e.g., Bull and Pain, 1995; Dimitrova, 2003; Mitrovic and Martin, 2002), while others have found that students tend not to investigate their learner model when it is available optionally (Kay, 1995; Barnard and Sandberg, 1996). It is therefore too early to draw any general conclusions about the utility of open learner models. Nevertheless, there is increasing interest in open learner modelling in the research literature to show how inspectable student models can lead to another teaching aid, and provide valuable additional feedback for learners. Examples illustrate a variety of contexts where the focus is on the learner model, from schoolchildren as young as 8–9 (Bull and McKay, 2004), through 10–13 year olds (Zapata-Rivera and Greer, 2004), to university students (Bull and Pain, 1995; Dimitrova, 2003). Given this growing

enthusiasm for opening the learner model to users, coupled with the mixed results of studies to date, it is important to discover the conditions under which a student might benefit from accessing their learner model, and might be encouraged to view their model.

The manner in which a learner views their learner model is dependent on the particular system. For example, skill meters (partly shaded bars indicating the knowledge level of the student as a subset of expert knowledge) are used in some systems (Weber and Brusilovsky, 2001). Others use representations that permit greater detail to be shown to learners, such as conceptual graphs (Dimitrova, 2003) or textual descriptions (Bull and Pain, 1995). Usually there is no flexibility in the form of access to the learner model in a system, and there has been limited research into varying the interaction method used when inspecting the learner model. However, Mabbott and Bull (2004) found that engineering students have different preferences for how they choose to view their model, when multiple representations are offered.<sup>1</sup> These preferences were not related to individuals' learning style. It may be that, if learners are allowed to access their learner model in a manner that suits them, they may be more likely to take up the opportunity to do so, thereby gaining greater benefit from a system with an open learner model.

In this paper, we aim to extend the findings of Mabbott and Bull (2004) to investigate whether haptic feedback on learner model data might be beneficial to some learners. While there has been very little research to date into varying the way an open learner model is presented within a single system, no research has been undertaken into the potential advantages of a tactile interaction with it. This approach may particularly benefit those learners who prefer physical interaction within the learning process. We introduce an adaptive learning environment that builds a simple learner model, the contents of which are conveyed to the learner textually, or through a haptic experience facilitated by the use of different tactile properties for on screen objects representing domain concepts. The paper examines whether the haptic learner model is found to be useful, and whether students find the haptic or textual version most useful.

## **2 The learning environment**

While more commonly employed in subjects such as medical education, haptic devices have also been used to interact with virtual systems in engineering education, to allow students to learn through their sense of touch (e.g., Okamura et al., 2002; Williams et al., 2004). Rather than the main learning interaction, our focus, in this paper, is on the potential for a haptic learner model. This could be used in conjunction with a system providing a haptic learning experience through simulations or tutorials, etc., or with a more conventional approach to computer-based learning. As our concern is with learner modelling, our example system here is a more traditional (non-haptic) adaptive learning environment. This will enable us to evaluate whether any positive results regarding haptic feedback in the learner model are because of the haptic learner model rather than a combination of haptic learner model interactions and haptic experiences in the main system.

The system was designed to support an undergraduate computer graphics course in Electronic, Electrical and Computer Engineering at the University of Birmingham, as many engineering applications require computer graphics. The example in this paper is the graphics pipeline, which is fundamental to the understanding of computer graphics.

However, it is intended that this approach to open learner modelling should also be applicable in other domains.

A learner navigates the system by looking at individualised recommendations, as shown in Figure 1. These provide suggestions about what learning material should be viewed according to the contents of the individual's learner model, and what tasks should be undertaken next. The learner uses the tree structure to the left of the screen to access learning materials: slides (Figure 2), course notes, code examples, exercises, a discussion forum, recommended additional reading and the learner model (Figures 3 and 4, Section 3.2). This approach provides guidance for the learner in the form of individualised recommendations, while also allowing the learner the freedom to make their own choices, the importance of which was advised in Section 1 above.

Figure 1 Recommendations for study

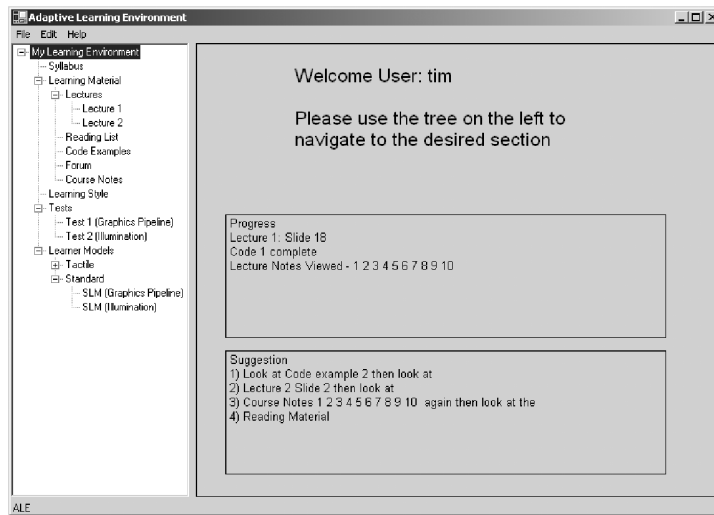
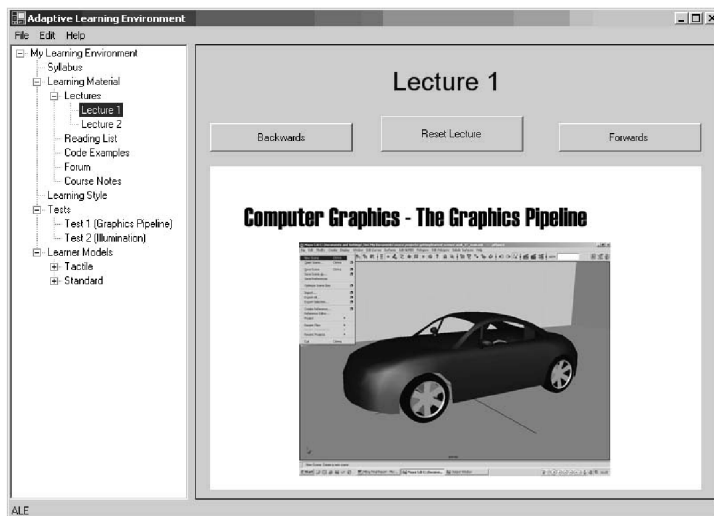


Figure 2 Example materials (slides)



### 3 Opening the learner model

This section describes the learner modelling process. We then present the user's view of their learner model.

#### 3.1 Building the learner model

A learner's understanding is tested by individually targeted questions, designed to identify a student's knowledge and elicit any misconceptions they may hold. The questions include ordering given items and multiple choice questions.

The learner modelling technique is relatively simple, sufficient for our purpose of evaluating the potential for a haptic learner model in engineering education. When a question is answered correctly, the student's score for that concept increases by one. If it is answered incorrectly, the score decreases by one. There are at least five questions used to test each concept in order that there is sufficient data from which to make a first attempt at inferring particular difficulties and misconceptions, where these exist. A misconceptions library is used to help identify misconceptions. If four or five questions are answered correctly out of five, then the concept being tested is said to be understood. If two or three questions are answered correctly, then the concept is represented in the learner model as partially known, but with some confusion; and when fewer than two questions are answered correctly, the concept is said to be unknown (but no misconception). Misconceptions are highlighted when there are three or more questions answered incorrectly, and the incorrect answers match information held in the misconceptions library. The Learner Model therefore holds the following type of information (a basic example from the start of the course) – knowledge level (percentage), known concepts and misconceptions:

- *known*: object and world coordinate systems (70%)
  - the order in which the two coordinate systems are defined
  - the OpenGL functions that support the transformation between coordinate systems
  - the transformation that takes place between object to world coordinate systems
- *misconception*: device coordinate system
  - misconception 1: the world coordinate system comes before the object, as the world takes into account all objects
  - misconception 2: the world coordinate system comes before the object, as the transformations convert world to object.

Areas of difficulty or confusion (which are not misconception related, or a simple lack of knowledge) can be inferred by comparing the number of questions attempted on a topic with the number of incorrect answers that have not been matched with the misconceptions library. The 'don't know' option allows learners to indicate a lack of knowledge.

### 3.2 Viewing and interacting with the learner model

Open learner models can vary in the level of detail and the complexity of information shown to the learner, as described in Section 1. As an initial investigation, the presentation of the learner model in our system has been kept quite simple. If the haptic model proves useful, its effectiveness can then be investigated further with information available at greater levels of detail.

Learners can choose the format of learner model they wish to use. The two methods of presenting the learner model are illustrated in Figures 3 (textual) and 4 (haptic). Figure 3 gives a textual description of the concepts the learner knows, areas of difficulty and the misconceptions that they may hold. This is accompanied by a weighting to show how strong or weak their knowledge of each of the concepts is, and the strength of evidence for the existence of misconceptions.

The setup in Figure 5 shows how a learner uses the SensAble Technologies PHANTOM haptic device (<http://www.sensible.com>) and Reachin Display unit (<http://www.reachin.se>) to interact with their haptic learner model illustrated in Figure 4. This provides force feedback in 3D environments, also allowing objects to have various tactile properties such as hardness, softness and friction. It allows users to feel the physical properties of 3D objects that are present on the screen.

**Figure 3** The textual learner model

The screenshot shows the 'Adaptive Learning Environment' (ALE) window. The title bar reads 'Adaptive Learning Environment'. The menu bar contains 'File', 'Edit', and 'Help'. On the left is a tree view under 'My Learning Environment' with the following structure:

- My Learning Environment
  - Syllabus
  - Learning Material
    - Lectures
      - Lecture 1
      - Lecture 2
    - Reading List
    - Code Examples
    - Forum
    - Course Notes
  - Learning Style
  - Tests
    - Test 1 (Graphics Pipeline)
    - Test 2 (Illumination)
  - Learner Models
    - Tactile
      - TLM (Graphics Pipeline)
      - TLM (Illumination)
    - Standard
      - SLM (Graphics Pipeline)
      - SLM (Illumination)

The main content area is titled 'Textual Learner Model'. It contains the following sections:

**Graphic Pipeline Concepts**

	Overall
Confusion between object and world coordinates systems	0%
Confusion between clipping and viewing coordinates systems	0%
Confusion between normalised device and device coordinates systems	0%

**Graphic Pipeline Misconceptions**

	Strength
Misconception about the world coordinate system	4 out of 4
Misconception about the clipping coordinate system	6 out of 6
Misconception about the device coordinate system	4 out of 4

**Results of the last test**

Graphics Pipeline	0 out of 10 attempted
Overall	0 out of 18
Questions Remaining	All Questions Attempted

The status bar at the bottom left of the window displays 'ALE'.

Figure 4 The haptic learner model

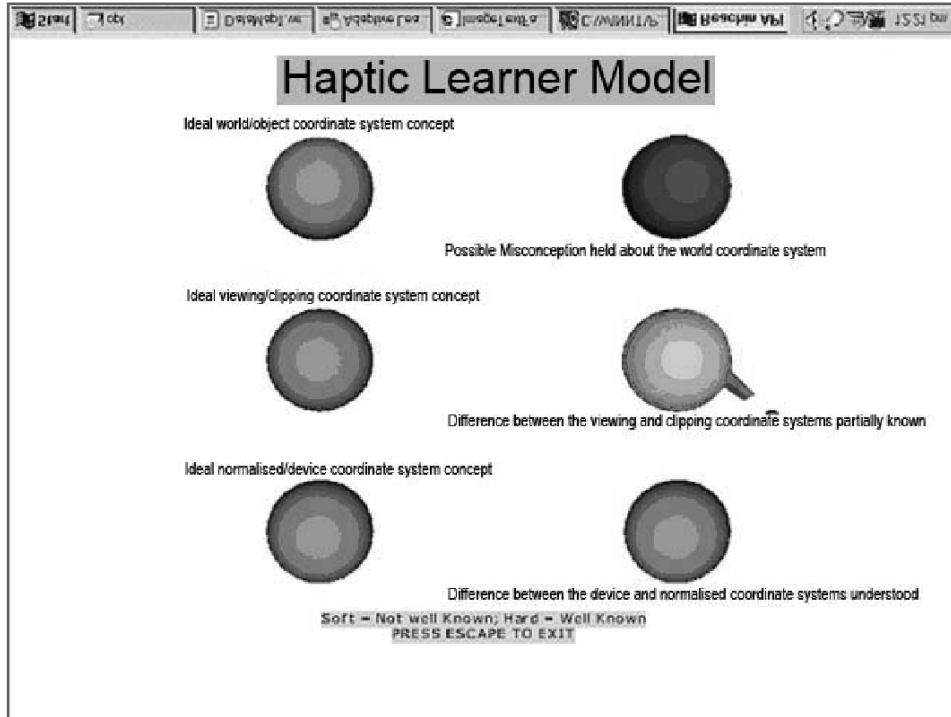


Figure 5 Setup for interacting with the haptic learner model




Figure 4 displays a 3D scene with haptic feedback where ‘concept spheres’ are presented alongside a textual description to allow the learner to both view and physically inspect their learner model. A control sphere is provided on the left, which illustrates to the learner, the state that they are aiming for at their current stage of learning. The spheres on the right show the learner’s current level of understanding of the relevant concepts. The better the learner understands a concept, the more closely matched the spheres on the right will be with those on the left, in terms of both visual and haptic properties. Concepts are displayed in various shades of green, depending on the level of understanding; and shades of orange if there are difficulties. Misconceptions are red. Haptic properties

are hard for known concepts, with spheres feeling softer if a concept is less well known. The property of magnetism is used for misconceptions to draw learners towards the problem, with the aim of encouraging them to investigate the issues. Thus misconceptions feel ‘soft and sticky’. The more severe the misconception is, the ‘stickier’ it feels.

The two methods of accessing the learner model provide the same information – knowledge, difficulties and misconceptions, together with an indication of the strength of knowledge or misconceptions. The provision of multiple ways of inspecting and interacting with the learner model allows the learner to choose between the two methods, to select the approach that suits them best. Learners have been shown to be able to select their preferred view of an open learner model when the choices are between various graphical and textual presentations of the learner model (Mabbott and Bull, 2004). This choice is expected also to be straightforward when learners are selecting between viewing a textual summary and viewing and interacting with a haptic learner model.

#### **4 Is a haptic learner model useful?**

An experimental study was conducted to determine whether users are able to use a haptic learner model, whether they consider it beneficial for their learning and whether they prefer investigating their model using a textual viewing or a haptic interaction. The haptic model necessarily combines graphics with haptic feedback, and also uses some text – to state what the concepts and misconceptions are (i.e., it uses three modes of presentation/interaction). We are therefore not aiming to make a direct comparison between the haptic and textual models in terms of whether one is better than the other. However, we are interested in whether there may be differences in learner preferences, which would suggest a need for providing quite different methods of accessing the learner model, for different learners.

##### *4.1 Participants*

Participants were 20 level 3 and 4 university students, studying either Computer Engineering or Computer Science.

##### *4.2 Materials and method*

A ten question paper-based pre-test, comprising multiple choice, item ordering and open questions, was administered to assess each student’s initial understanding of the target domain. This took approximately ten minutes. For the purpose of increasing the students’ appreciation of the tactile element of the session, a demonstration of the PHANTOM was carried out, using various interaction techniques. Participants then performed a simple task using the PHANTOM haptic device, in order to familiarise themselves with the unusual method of interaction. This took around five minutes. Each subject then undertook a learning interaction with the system, completed an exercise and finally inspected the textual and haptic learner model. Subjects were able to return to learning material or either version of the learner model at any time. This interaction occupied 30–45 minutes. The sessions were video-recorded. A ten question post-test (of the same format as the pre-test) was administered, which lasted about ten minutes. Subjects

then completed a questionnaire about their experience of using the system and the open learner model. In this paper, we focus on the open learner model results, in particular, the utility of the haptic learner model; and whether subjects preferred using the haptic model or the textual version. Responses to questionnaire statements were on a 5-point scale: strongly agree (5), agree (4), neutral (3), disagree (2) and strongly disagree (1). The post-interaction questions took around 10 minutes.

### 4.3 Results

The average score for the pre-test was 25%, and the average for the post-test was 58%. While we are not here evaluating the extent to which the system facilitates learning, it is nevertheless important to demonstrate that the evaluation of the learner model took place in the context of an effective system. The breakdown of test results is given in Table 1.

**Table 1** Pre-test/post-test scores (out of 10)

Results	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	S11	S12	S13	S14	S15	S16	S17	S18	S19	S20
Pre-test	3	4	2	2	3	5	3	4	4	3	5	3	5	4	6	3	2	2	7	3
Post-test	7	7	5	8	5	7	6	8	5	4	7	5	8	5	7	5	4	6	7	6

Of the 20 students, 19 improved between pre-test and post-test, with one student scoring equally on each. The latter student was the one with the highest pre-test score (S19). The range of scores for the pre-test was 2–7; the post-test range was 4–8. Improvement between pre-test and post-test ranged from a 0 point increase (i.e., same score, as mentioned above) to six points. The latter was one of the students who scored lowest on the pre-test, who subsequently obtained one of the highest scores on the post-test.

Table 2 shows the number of students giving each response on the questionnaire, with reference to whether the haptic and textual learner model access methods were understandable and useful. Responses are also given for whether students found the open learner model helpful for their learning in general (i.e., not with reference to either method specifically). For convenience of analysis, the mode and range values are also given.

**Table 2** Questionnaire responses

Questionnaire	Strongly agree (5)	Agree (4)	Neutral (3)	Disagree (2)	Strongly disagree (1)	Mode	Range
Haptic: understandable	2	10	5	3	0	4	2–5
Textual: understandable	0	7	10	3	0	3	2–4
Haptic: useful	3	9	7	1	0	4	2–5
Textual: useful	0	8	9	3	0	3	2–4
Learner model helped	2	14	4	0	0	4	3–5

Table 2 shows that most users found having an open learner model useful (final row), with only 4 neutral responses and no negative responses. Most users agreed that the haptic learner model was both understandable and useful. The scores for the textual learner model were lower, with more learners giving a neutral response. Nevertheless, there were still a large minority of students who found the textual version understandable and useful.

Table 3 provides a breakdown of the number of students preferring one form of the learner model over the other, in terms of how easy it was to understand, and how useful it was for their learning.

**Table 3** Preference for learner model presentation/interaction

<i>Preferred version</i>	<i>Understandability</i>	<i>Usefulness</i>
Haptic	8	8
Textual	4	2
Haptic and textual	8	10

While Table 2 illustrates the understandability and usefulness of the two versions of the learner model, Table 3 shows learner preferences for the two versions. Half the learners found both options for accessing the model equally useful, and 40% found them equally understandable. There was, however, a difference between results for students who had a preference for one version over the other, with a large minority (40%) preferring the haptic learner model – finding it more understandable and useful as a learning support, and 20% and 10% finding the textual model easier to understand and more useful, respectively.

Table 4 gives the level of knowledge at the start and level of knowledge at the end of the interaction, and the increase in performance between pre-test and post-test, for those who found one of the methods of accessing the learner model more useful than the other, and for those with no preference. For ease of visual interpretation, asterisks have been used to illustrate the scores. Thus, of those who found the haptic learner model more useful, two scored 2 on the pre-test, two scored 3, one scored 4 and so on.

**Table 4** Pre-test and post-test scores for students according to learner model preference

<i>Score</i>	<i>0</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>	<i>6</i>	<i>7</i>	<i>8</i>
Haptic: pre-test			**	**	*	**	*		
Haptic: post-test						**	**	***	*
Textual: pre-test				*				*	
Textual: post-test						*		*	
Both: pre-test			**	****	***	*			
Both: post-test					**	***	*	**	**
Haptic: improvement		**	***	*	*		*		
Textual: improvement	*		*						
Both: improvement		**	**	****	**				

It can be seen that there is a spread of starting level and final achievement of those preferring the haptic model, from scores of 2 to a score of 6 on the pre-test; and scores of 5–8 on the post-test. There are only two students who prefer the textual model; therefore it is not possible to look for patterns amongst the data. Nevertheless, it is interesting to note that the two students have quite different levels of performance, with scores of 3 and 7 on the pre-test and 5 and 7 on the post-test. Learners finding the haptic

model more useful also had a variety of degrees of improvement, with two learners improving only 1 point between pre-test and post-test, to one student improving six points. It is not possible to identify patterns with the two students who preferred the textual model, though it can be noted that they seem quite different, with one not improving between the two tests (this learner scored high (7) on both), and one improving two points (but amongst the weaker students at both pre-test and post-test stage). There are no obvious differences in the results for those learners with a preference for a particular version of the learner model, and those who found the two versions equally useful.

#### 4.4 Discussion

Table 1 shows that our learners had varying levels of expertise at the start of the interaction. The average improvement for a 30–45-minute learning session was 33%. All participants improved from pre-test to post-test, apart from one learner who already had a high level of knowledge, whose score remained the same. This suggests that the system as a whole (i.e., the learning interaction plus viewing the learner model) is educationally useful. We do not know, of course, to what extent this might be owing to the learner model interaction, beyond learner perceptions of the utility of the open learner model, as identified by questionnaire. This is an issue for further research. However, it is necessary to perform an evaluation of students' perceptions of an open learner model in the context of an *effective* learning environment.

The questionnaire results suggest that students found the existence of an open learner model as a learning resource, to be useful (Table 2). Thus an open learner model, in some form, is likely to be of benefit to learners such as those in our sample. Twelve of the twenty learners considered the haptic model understandable and a useful learning support, while a lower number found the textual model understandable and useful.

Our aim in this study was not to determine whether a haptic learner model *or* a textual learner model should be used in a system. We are also not claiming that an open learner model (of any kind) should necessarily be useful for all students. This is likely to depend very much on learners' preferences regarding the learning process. We simply wish to discover whether there are sufficient learners who would find a haptic model beneficial for their learning, to make it worthwhile implementing haptic models as an option for those who prefer to use these, over other presentations of learner model data. The results in Table 3 suggest that this is, indeed, the case. Although 50% of learners found the two versions of the learner model equally useful, and may therefore be satisfied with either, 40% felt that the haptic model better supported their learning. While this is a minority, it is a large minority. Fewer students preferred the textual version of the open learner model. Nevertheless, this still supports our view that the haptic option should be only one of the approaches to accessing the model that the learner can *select* to use, as it is clear that this was not the best method for all learners. This accords with Mabbott and Bull's (2004) recommendation that students be provided with different views on their learner model, extending that recommendation to include a haptic interaction, where appropriate equipment is available to students.

In order to discover whether preference for (and therefore greater use of) the haptic or textual version of the learner model might better facilitate learning, Table 4 shows the pre-test and post-test scores, and extent of improvement, of students preferring each version. The scores of those students who found the haptic model more useful, vary in

line with the group in general. Thus there appears no difference in learning gains between those who find the haptic model more useful, and those who found the two versions of the learner model equally useful. There are too few students who preferred the textual model to draw any conclusions based on that preference, though our two examples are quite different. Recalling our earlier observation that allowing students to access their learner model in the manner most suited to their preferences, might lead to more learners accessing their model and therefore gaining the associated learning benefits, seems valid in the context of our results. Neither version of the learner model was associated with greater learning gains, but some students did have a clear preference for the method by which to access their model. Thus it appears a matter of supporting student preferences to encourage learner model inspection, rather than prescribing an 'ideal' form of interaction with the learner model.

This was a 'single-shot' laboratory-based evaluation, and therefore does not necessarily reflect what would happen in practice. Students may well inspect their learner model when they initially use a system, but would this interest be maintained over time? Even short-term inspection of a learner model may be sufficient, in some cases, to raise learner awareness of their difficulties, and help students to develop into more reflective and autonomous learners. However, it would be interesting to discover the potential longer-term use of a haptic learner model. Our results have provided the essential starting point for further work. A major question for follow-up work is whether those who found the haptic approach useful, would continue to use it over a prolonged period. In such a longer study, we would want to investigate the use of the haptic learner model as is, and also a more detailed version – such as a structured learner model where, for example, prerequisite relationships between concepts can be displayed, or a more detailed breakdown of topics and sub-topics, or conceptual relationships.

Our learners were all computer-literate subjects. While valid for our engineering students, these results may not easily generalise to a broader learner group. This would be a useful further investigation.

A major issue, in terms of the likelihood of students having a haptic learner model available, is the cost of the equipment required. Our system was developed using existing equipment in the Electronic, Electrical and Computer Engineering department, at the University of Birmingham. However, it is not necessary that such expensive equipment be used. More affordable approaches have been employed in other interactive learning environments (e.g., Okamura et al., 2002; Williams et al., 2004).

## **5 Summary**

Textual and graphical open learner model presentations have already been investigated in a variety of contexts. This paper has suggested the potential for a haptic learner model to support learning for those students who prefer this approach.

The findings from this initial evaluation have provided some positive results. It was felt by the majority of users that the haptic learner model allowed them to reflect upon their learning. The majority of learners found it just as intuitive to inspect the haptic learner model as the textual one, and some preferred using it. Nearly all users found that inspection of their learner model was helpful to their learning. We therefore recommend providing haptic learner model interactions as options in addition to more standard methods of accessing the learner model, where such facilities are available.

## References

- Barnard, Y.F. and Sandberg, J.A.C. (1996) 'Self-explanations, do we get them from our students?', in Brna, P., Paiva, A. and Self, J. (Eds.): *Proceedings of European Conference on Artificial Intelligence in Education*, Lisbon, pp.115–121.
- Bull, S. (2004) 'Supporting learning with open learner models', *Proceedings of 4th Hellenic Conference in Information and Communication Technologies in Education*, Athens, pp.47–61.
- Bull, S. and McKay, M. (2004) 'An open learner model for children and teachers: inspecting knowledge level of individuals and peers', in Lester, J.C., Vicari, R.M. and Paraguacu, F. (Eds.): *Intelligent Tutoring Systems: 7th International Conference*, Springer-Verlag, Berlin Heidelberg, pp.646–655.
- Bull, S. and Pain, H. (1995) 'Did I say what I think I said, and do you agree with me?: inspecting and questioning the student model', in Greer, J. (Ed.): *Proceedings of World Conference on Artificial Intelligence in Education*, Association for the Advancement of Computing in Education (AACE), Charlottesville, VA, pp.501–508.
- Dimitrova, V. (2003) 'StyLE-OLM: interactive open learner modelling', *International Journal of Artificial Intelligence in Education*, Vol. 13, No. 1, pp.35–78.
- Felder, R.M. and Silverman, L.K. (1988) 'Learning and teaching styles in engineering education', *Engineering Education*, Vol. 78, No. 7, pp.674–681.
- Kay, J. (1995) 'The UM toolkit for cooperative user modelling', *User Modelling and User Adapted Interaction*, Vol. 4, pp.149–196.
- Mabbott, A. and Bull, S. (2004) 'Alternative views on knowledge: presentation of open learner models', in Lester, J.C., Vicari, R.M. and Paraguacu, F. (Eds.): *Intelligent Tutoring Systems: 7th International Conference*, Springer-Verlag, Berlin Heidelberg, pp.689–698.
- Mitrovic, A. and Martin, B. (2002) 'Evaluating the effects of open student models on learning', in de Bra, P., Brusilovsky, P. and Conejo, R. (Eds.): *Adaptive Hypermedia and Adaptive Web-Based Systems, Proceedings of Second International Conference*, Springer-Verlag, Berlin Heidelberg, pp.296–305.
- Okamura, A.M., Richard, C. and Cutkosky, M.R. (2002) 'Feeling is believing: using a force-feedback joystick to teach dynamic systems', *Journal of Engineering Education*, July, pp.345–349.
- Papanikolaou, K.A., Grigoriadou, M., Kornilakis, H. and Magoulas, G.D. (2003). 'Personalizing the interaction in a web-based educational hypermedia system: the case of INSPIRE', *User Modeling and User Adapted Interaction*, Vol. 13, pp.213–267.
- Weber, G. and Brusilovsky, P. (2001) 'ELM-ART: an adaptive versatile system for web-based instruction', *International Journal of Artificial Intelligence in Education*, Vol. 12, No. 4, pp.351–384.
- Williams II, R.L., He, X., Franklin, T. and Wang, S. (2004) 'Haptics-augmented undergraduate engineering education', *Proceedings of International Conference on Engineering Education*, University of Florida, <http://succeednow.org/icee>.
- Zapata-Rivera, J-D. and Greer, J.E. (2004) 'Interacting with inspectable bayesian student models', *International Journal of Artificial Intelligence in Education*, Vol. 14, No. 2, pp.127–163.

## Note

<sup>1</sup>These students were taking a course on either 'Educational Technology' or 'Interactive Learning Environments', and so had an interest in educational issues. Nevertheless, these courses are offered in the context of engineering degrees.