

International Journal of Artificial Intelligence in Education

Teaching Database Design with Constraint-based Tutors

--Manuscript Draft--

Manuscript Number:	AIED-D-15-00008R2
Full Title:	Teaching Database Design with Constraint-based Tutors
Article Type:	SI - 25th Anniversary High-Impact
Funding Information:	
Abstract:	Design tasks are difficult to teach, due to large, unstructured solution spaces, underspecified problems, non-existent problem solving algorithms and stopping criteria. In this paper, we comment on our approach to develop KERMIT, a constraint-based tutor that taught database design. In later work, we re-implemented KERMIT as EER-Tutor, and extended its instructional domain. Several evaluation studies performed with KERMIT and EER-Tutor show that they are effective Intelligent Tutoring Systems (ITSs). We also comment on various extensions made to EER-Tutor over the years. There are several contributions of our research, such as developing effective problem-solving support for conceptual database design in terms of interface design. Our database design tutors deal with huge solution spaces efficiently by specifying constraints that capture equivalent solution states, and using ideal solutions to capture the semantics of the problem. Instead of requiring a problem solver, the ITS checks whether the student's database schema is correct by matching it to constraints and the ideal solution. Another contribution of our work is in guidelines for developing effective feedback to the student.
Corresponding Author:	Antonija Mitrovic, PhD University of Canterbury Christchurch, NEW ZEALAND
Corresponding Author Secondary Information:	
Corresponding Author's Institution:	University of Canterbury
Corresponding Author's Secondary Institution:	
First Author:	Antonija Mitrovic, PhD
First Author Secondary Information:	
Order of Authors:	Antonija Mitrovic, PhD Pramuditha Suraweera, PhD
Order of Authors Secondary Information:	
Author Comments:	
Response to Reviewers:	Dear Guest Editor We have revised our paper further to respond to the comments we received from the reviewers. We thank you and the reviewers on the efforts to improve the paper. Please find our replies below. Reviewer #1: This paper provides a useful references for IJAIED readers to examine the authors's series of researches. I think it satisfies the requests as a short commentary. If there is room to add pages, I hope the authors add more detailed explanation concerning the difficulty of the "ER modeling". This explanation is helpful to understand the motivation, approach and contributions of the series of the researches more

deeply.

Reply: We have expanded the introduction, and the section on KERMIT to address this comment, as well as the related comments of Reviewer 4.

Supplemental explanations to each interface would be also helpful for readers.

Reply: We have added more text to discuss the interfaces shown, and also separated Figure 2 into two figures.

Typo

- There is an extra parenthesis on page 4, line 45.
- There is a sentence expressed with different font size on page 5, line 41.
- "Educaiont" on page 11, line 17.

Reply: All of these changes have been done.

Reviewer #3: Numerous minor English errors required careful editing by a native English speaker compulsive about correct grammar and usage.

** = substantive

1.26: done -> made

46: common-sense -> common sense

2.7: / -> and

8: , lack -> , and lack

13: that -> , which: if claiming that SQL-Tutor was the first constraint-based tutor of any kind; delete ever: if claiming that SQL-Tutor was the first constraint-based tutor for database querying 27, 29: Section -> section

** 3.3 gains: on what measure?

Reply: The previous version of the paper stated that the effect size was based on gains. We have now expanded that part, to say "learning gains" and we added an explanation "(post-test score – pre-test score)".

3: the effect size of -> effect size d' = [assuming it's Cohen's effect size]

5: the feedback -> its feedback

15: in which we designed -> with [avoid awkwardness of 'we' following 3rd person on 3.14]

16..18: was... was... is... becomes: pick a consistent tense

45:).). ->).

49: -> defined by Elmasi and Navathe (2011),

50: (specializations) -> , specializations,

4.4: have -> has

14: difficult to anyone -> difficult for anyone

5.10: learnt -> learned

** 16: tracked based on what input?

Reply: we have added this information – the system used the student's facial features and also information about the student's problem-solving actions within the system

** 35: learning gains on same measure as before?

Reply: yes, the learning gains as defined previously

39: looked deeper -> looked more deeply

41-42: fix font size

52: Thank you -> We thank them [i.e. not the readers]

6.4: you -> them

Reply: all the changes required have been made

Reviewer #4: This is a nice paper and will make a good addition to the IJAIED special issue. However, it would be helpful to succinctly formulate a small number of key contributions of the work and include them in the abstract and discussion section.

Reply: done

Also, I would really like to get a better sense for how the particular type of ill-definedness that exists in this domain (no clearly defined process model) is dealt with by means of constraints and comparison of expert solutions. So it would help, for example, to clarify (a) how large the solution space tends to be for the type of ER models built with ER Tutor and (b) a bit more detail about how the tutor evaluates solutions.

Reply: Done. We have modified the introduction and the section on KERMIT to reply to this comment. We use our view of ill-definedness (presented in Mitrovic & Weerasinghe, 2009), in which we differentiate instructional domains (i.e. the domain theory) from instructional tasks. We hope the paper is now clear in that respect. We are aware that our views differ from that of other authors, but our view has been presented in the referenced AIED 2009 paper. We believe that adding further comparisons to the views of other researchers would be detrimental for the current paper, which is meant to be a commentary on our work.

Abstract

As mentioned in my original review, adding a clear and concise statement of the contributions of the work would help make a more informative abstract.

Reply: done

Introduction

"Additionally, the stopping criterion (i.e. knowing when the solution is reached) is unclear - how can a student evaluate the quality of the produced ER diagram?"

This raises the question: (How) can the tutor evaluate the quality of a student's ER diagram?

Reply: Done. We have added a couple of sentences to the Introduction, and also added a discussion of how the system analyses the student's solution to the section on KERMIT.

In my opinion, as far as ill-defined design tasks go, or ill-defined tasks more generally, ER design seems far over to the well-defined part of the spectrum. I get the sense from this paper that the essential correctness of solutions is not an issue in this domain - so in that sense, the domain is not ill-defined, unlike other domains, such as less well-defined design tasks, ethics, or even, legal reasoning. The current domain seems more similar algebra word problems, even if more complex.

Reply: We do not agree with this comment. There is no algorithm (problem-solving procedure) for database design, and therefore it is ill-defined. We agree that there is a whole spectrum of ill-defined tasks, but database design is definitely on that spectrum (as is programming and many other design tasks).

A few years back, IJAIED published a special issue (two, in fact) about AIED in ill-defined domains, including an article by Lynch et al. that tried to formulate some criteria for ill-definedness. It would be interesting to connect the current discussion to that prior work.

Reply: We have not done this. As stated in our previous reply, we are using our own criteria for ill/well defined instructional domains and instructional tasks, from the AIED

2009 paper that we refer to. Adding a comparison to other definitions would require a lot of space, and is not going to increase the value of our paper, which is meant to be a commentary on the original KERMIT paper.

"while in the case of ER diagrams, there is no prespecified structure at all."
Doesn't the diagram language itself provides lots of structure, as alluded to earlier in the paper? So this would seem to be an overstatement.

Reply: We have added additional discussion to the paper to address this issue. The ER model is well-defined, but the process of developing and ER diagram is ill-defined. There is no pre-specified structure for the solution, apart from it having to contain entities, attributes and relationships.

KERMIT: A KNOWLEDGE-BASED ER MODELING TUTOR

"and the constraint set ensured that other correct solutions would be recognized as such, by looking for equivalent ways of specifying solutions (Suraweera & Mitrovic, 2004)."

In other words, the system is able to evaluate the correctness of solutions? It is unclear however, how constraints are used to evaluate the correctness of solutions? Do these constraints essentially compare the student solution against the expert solution?

Reply: Yes, the system can evaluate the correctness of the solution. We have expanded this section to provide more detail on how that is achieved. The constraints do compare the student's solution to the ideal solution; in addition, they check for alternative ways of specifying ER components that are different from what is captured by the ideal solution.

1 International Journal of Artificial Intelligence in Education Volume# (YYYY) Number
2 IOS Press
3
4
5

6 Teaching Database Design with Constraint-based Tutors 7 8 9

10 **Antonija Mitrovic**, *Intelligent Computer Tutoring Group, Department of Computer Science*
11 *and Software Engineering, University of Canterbury, New Zealand*
12 tanja@cosc.canterbury.ac.nz
13

14
15 **Pramuditha Suraweera**, *Research Insights and Analytics, Chief Marketing Office, Telstra,*
16 *Australia*
17 pramudi@gmail.com
18
19

20
21 **Abstract.** Design tasks are difficult to teach, due to large, unstructured solution spaces, underspecified problems,
22 non-existent problem solving algorithms and stopping criteria. In this paper, we comment on our approach to
23 develop KERMIT, a constraint-based tutor that taught database design. In later work, we re-implemented
24 KERMIT as EER-Tutor, and extended its instructional domain. Several evaluation studies performed with
25 KERMIT and EER-Tutor show that they are effective Intelligent Tutoring Systems (ITSs). We also comment on
26 various extensions made to EER-Tutor over the years. There are several contributions of our research, such as
27 developing effective problem-solving support for conceptual database design in terms of interface design. Our
28 database design tutors deal with huge solution spaces efficiently by specifying constraints that capture equivalent
29 solution states, and using ideal solutions to capture the semantics of the problem. Instead of requiring a problem
30 solver, the ITS checks whether the student's database schema is correct by matching it to constraints and the
31 ideal solution. Another contribution of our work is in guidelines for developing effective feedback to the student.
32

33 **Keywords.** KERMIT, EER-Tutor, teaching design tasks, ill-defined task
34
35

36 INTRODUCTION 37

38 Work on a constraint-based tutor to teach database design started in 2000. The motivation for this
39 work was twofold. Firstly, at the time, Mitrovic was teaching (and still does) an introductory course on
40 relational databases. The course includes a large component on database design, including conceptual
41 database design using the Entity-Relationship (ER) model (Chen, 1976). Mitrovic observed that
42 students had many difficulties learning ER modeling. The second motivation for developing KERMIT
43 was to continue our work on Constraint-Based Modeling (CBM) (Ohlsson, 1992). We start the paper
44 by discussing these two motivations.
45

46 Mitrovic and Weerasinghe (2009) made a distinction between instructional domains and
47 instructional tasks when discussing ill-definedness. The ER model itself is well-defined; it consists of
48 only a handful of constructs with well-defined syntax. However, conceptual database design using the
49 ER model is an ill-defined task; it consists of mapping database requirements, usually provided to
50 students in the form of English text, to ER diagrams. Database requirements are usually
51 underspecified, and can be ambiguous. In order to understand them fully, students need a good
52 understanding of the application domain, common sense and background knowledge, which are often
53
54
55
56

57 1560-4292/08/\$17.00 © YYYY – IOS Press and the authors. All rights reserved.
58
59
60
61
62
63
64
65

1
2
3
4 missing. Furthermore, the outcome (i.e. the ER schema) is defined in abstract terms: the student is
5 required to develop a high quality schema that meets the requirements. Database design is a
6 demanding task, as there is no algorithm students can use to produce ER schemas. Additionally, the
7 stopping criterion (i.e. knowing when the solution is reached) is unclear – how can a student evaluate
8 the quality of the produced ER diagram? It is possible for the student to check the syntactic
9 correctness of the schema by making sure that all components satisfy the integrities of the ER model,
10 although even this task is complex for students new to database design. Checking semantic correctness
11 is much harder, as it requires the student to make sure the produced database schema covers
12 requirements completely and that it is of high quality.

13
14 All of these features of database design are typical for other design tasks. Several criteria have
15 been identified to describe design tasks in general (Goel & Pirolli, 1988; 1992; Reitman, 1964).
16 Design tasks require extensive domain expertise, use of artificial symbol systems and incremental
17 development. These tasks are characterized by underspecified start and goal states and problem-
18 solving algorithms, huge solution spaces, lack of operators for changing states, large solutions, and
19 lack of a definite test to decide whether the goal has been attained, and consequently, there is no best
20 solution, but rather a family of solutions for each problem.

21
22 The other motivation for developing KERMIT was to further research on CBM in terms of its
23 applicability in various instructional domains. Before KERMIT, Mitrovic and colleagues developed
24 SQL-Tutor, the first constraint-based tutor which teaches students how to query relational databases
25 (see Mitrovic & Ohlsson, this issue), and CAPIT, a constraint-based tutor which teaches elementary
26 school children about punctuation and capitalization rules in English (Mayo & Mitrovic, 2001). These
27 two tutors proved that CBM is an effective student modeling approach, and answered many questions
28 raised by Ohlsson when he originally proposed CBM in his 1992 paper (see also Ohlsson's
29 commentary in this issue). ER modeling differed sufficiently from the previous two instructional
30 domains to be interesting from this point of view. Punctuation and capitalization rules in English are
31 very straightforward in comparison to ER modeling. Writing queries in SQL is an ill-defined design
32 task, and therefore has some similarities with ER modeling. However, solutions are more structured in
33 SQL than in ER modeling; queries consist of six clauses, thus defining the elementary structure for the
34 solution space, while in the case of ER diagrams, there is no pre-specified structure at all.

35
36 We start the paper by briefly discussing the highly cited IJAIED paper on KERMIT in the next
37 section, and then discuss EER-Tutor, an enhanced Web-enabled version of the tutor. EER-Tutor
38 enabled many exciting research projects, which investigated meta-cognitive skills, affect-sensitive
39 pedagogical agents, and tutorial dialogues. We present those projects in the last section.
40
41

42 **KERMIT: A KNOWLEDGE-BASED ER MODELING TUTOR**

43
44
45 Similar to other constraint-based tutors, KERMIT was designed as a complement to database courses;
46 we assumed that the student has already acquired some knowledge via lectures and labs. KERMIT
47 allowed students to practice, by providing database requirements for which students developed ER
48 diagrams. The system, presented in Figure 1, was originally developed in Visual Basic as a stand-
49 alone application. The student was given the database requirements for the chosen problem at the top,
50 and could use various tools to draw ER components on the drawing pane. When the student submitted
51 a solution, the system analysed it, and provided feedback on identified mistakes (if any). Students
52 could request more detailed feedback messages depending on their needs. The animated pedagogical
53
54
55
56
57
58
59
60
61
62
63
64
65

1
2
3 agent (the Genie, implemented using Microsoft Agent) presented feedback verbally, using audio and
4 speech bubbles.
5

6 KERMITE analyses the submitted solution by matching it to the constraint set (which contained 92
7 constraints) and the pre-specified ideal solution (Suraweera & Mitrovic, 2004). There are two types of
8 constraints: the syntactic constraints check whether the student's solution satisfies the syntax of the
9 ER model, while the semantic constraints check whether the student's solution matches the problem
10 requirements by comparing it to the ideal solution. KERMITE stored only one correct solution per
11 problem (specified by the human teacher). For all but the simplest problems, there could be several
12 correct solutions and an infinite number of incorrect solutions, resulting in huge solution spaces. The
13 ideal solution captures the semantics of the problem without the need for a problem solver (which
14 would be difficult, if not impossible, to develop). Semantic constraints ensure that alternative correct
15 solutions would be recognized as such, by looking for equivalent ways of specifying ER components
16 (Suraweera & Mitrovic, 2004). Therefore, constraints and ideal solutions allowed KERMITE to deal
17 with huge solution spaces efficiently.
18

19 There were several challenges we faced during design and development of KERMITE. We wanted
20 to support learning, and therefore to lower the working memory load; to achieve that, we showed the
21 full problem text to the student, as well as the set of tools to draw diagrams. When a student creates a
22 new diagram component, he/she needs to specify its name; in KERMITE, the name must be a word or a
23 phrase from the problem text. Although there has been some criticism of this decision as being overly
24 restrictive, we still believe it is beneficial for students' learning, for several reasons. From our own
25 teaching experience, we observed that student often underline words or phrases from the problem text
26 using different colours when they study requirements – therefore, using a similar approach in the ITS
27 is not unnatural. Furthermore, using different colours for different component types makes it easier for
28 the student to go over the problem text and see how much of it is already covered, and what is still
29 outstanding. Another important reason in favour of such naming policy comes from Software
30 Engineering, where the use of the customer's language is strongly encouraged. By using parts of the
31 problem text, we prevent students from inventing their own labels, which might be difficult for anyone
32 else to interpret. Finally, this decision made the interpretation of students' solutions much easier; we
33 know exactly what each component represents, and avoid problems with natural language
34 understanding which we would face if the students named components freely.
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

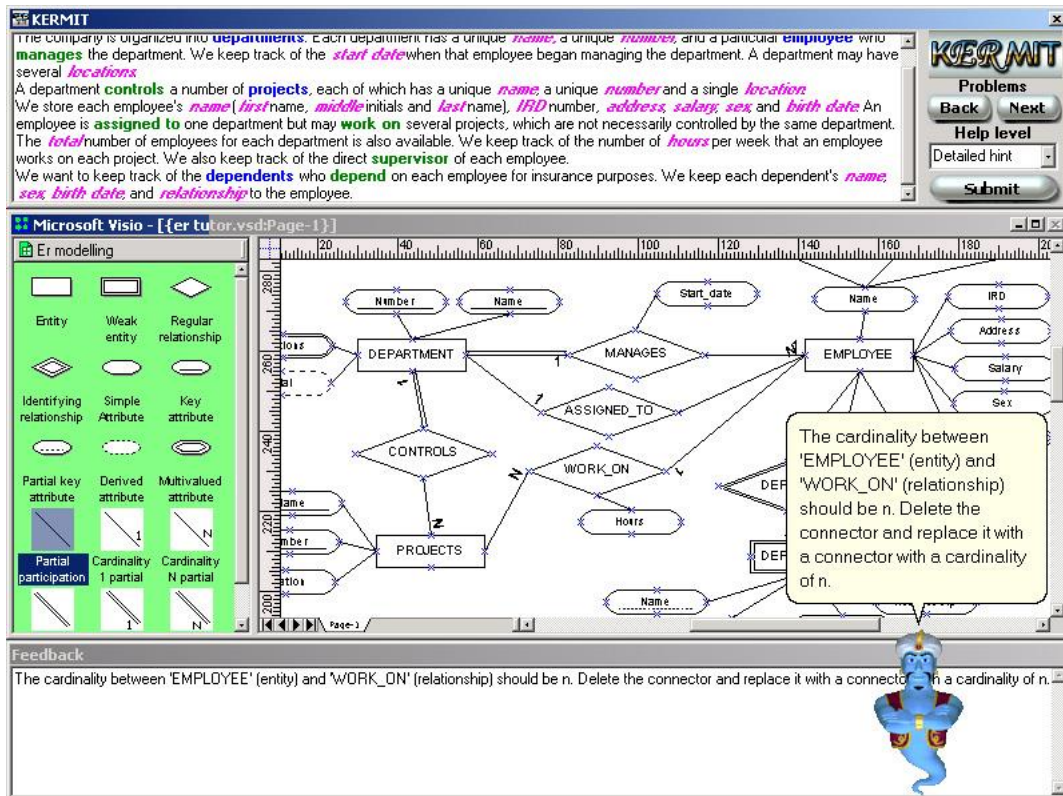


Figure 1. Screenshot of the KERMIT's interface

The 2004 paper presented the approach taken to develop KERMIT, as well as the findings from the pilot and evaluation studies conducted in 2001, which showed that KERMIT was highly effective. The full study compared the experimental group (26 participants using KERMIT) to the control group (31 participants) who used a version of the system with limited feedback (no feedback was given on the student's solution, only the full solution was provided after each problem). There was a statistically significant difference in the learning gains (post-test score – pre-test score) of the two groups, with the effect size d of 0.63, after students spent an average of only 66 minutes with the system. The students liked the system, and rated its feedback highly.

EXPANDING KERMIT

After showing that CBM is an effective approach to teach conceptual data modeling, we turned to other interesting research questions, such as modelling and supporting higher-order skills. One important metacognitive skill is reflection. Hartley and Mitrovic (2002) reported on an extension of KERMIT (named e-KERMIT) with two visualizations of the student model. We added skill meters to the interface (circled in Figure 2), which presented a high-level summary of the student model in terms of the student's knowledge of the ER notation and ability to identify entities, attributes and relationships. Skill meters presented continuous feedback on progress, and also served as an aid to remind and motivate students to further inspect their models. A more detailed, hierarchical open model (Figure 3) was available on request, when the student clicked on the *Show Me More* button. The hierarchical open student model presented a visualization of the student's knowledge on a finer granularity level; for each domain concept in the hierarchy, the model shows the student's progress in terms of the percentage that the student covered (i.e. the percentage of corresponding constraints the student used) and the percentage mastered (i.e. the percentage of corresponding constraints the student used correctly). The student could expand a concept to show concepts on the lower level of the category (e.g. *Type* under *Attribute identification* in Figure 3).

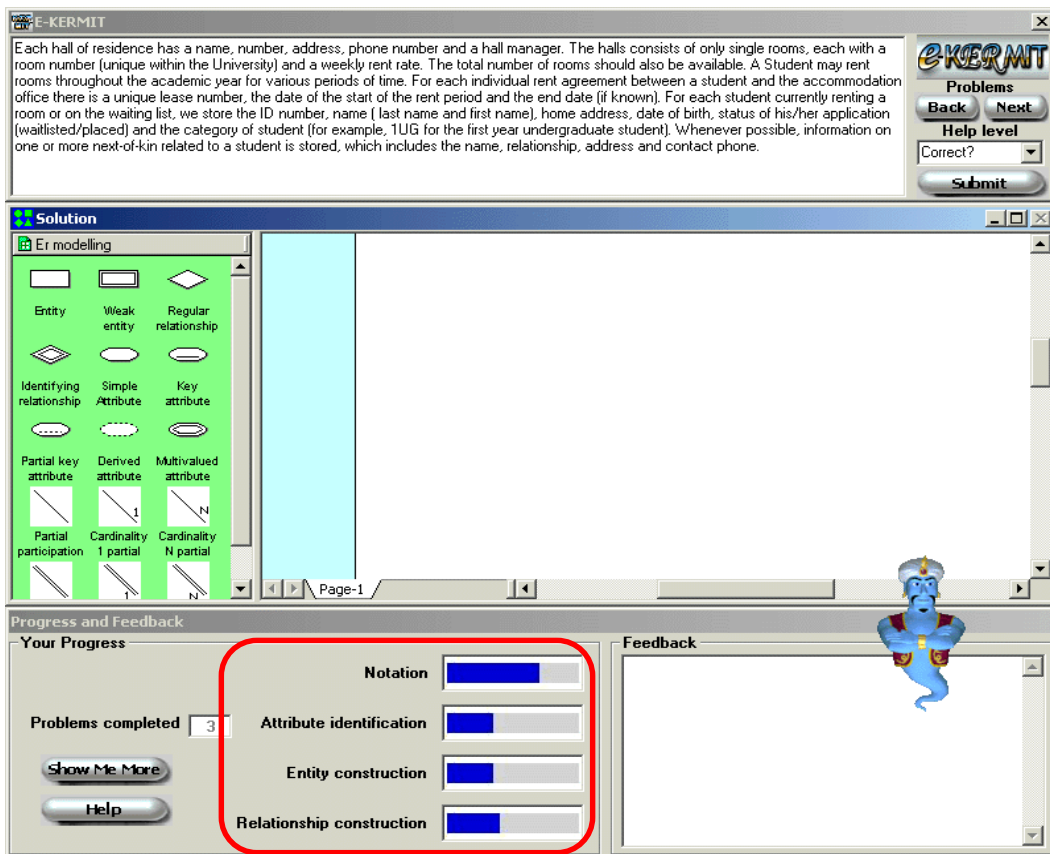


Figure 2. The interface of e-KERMIT

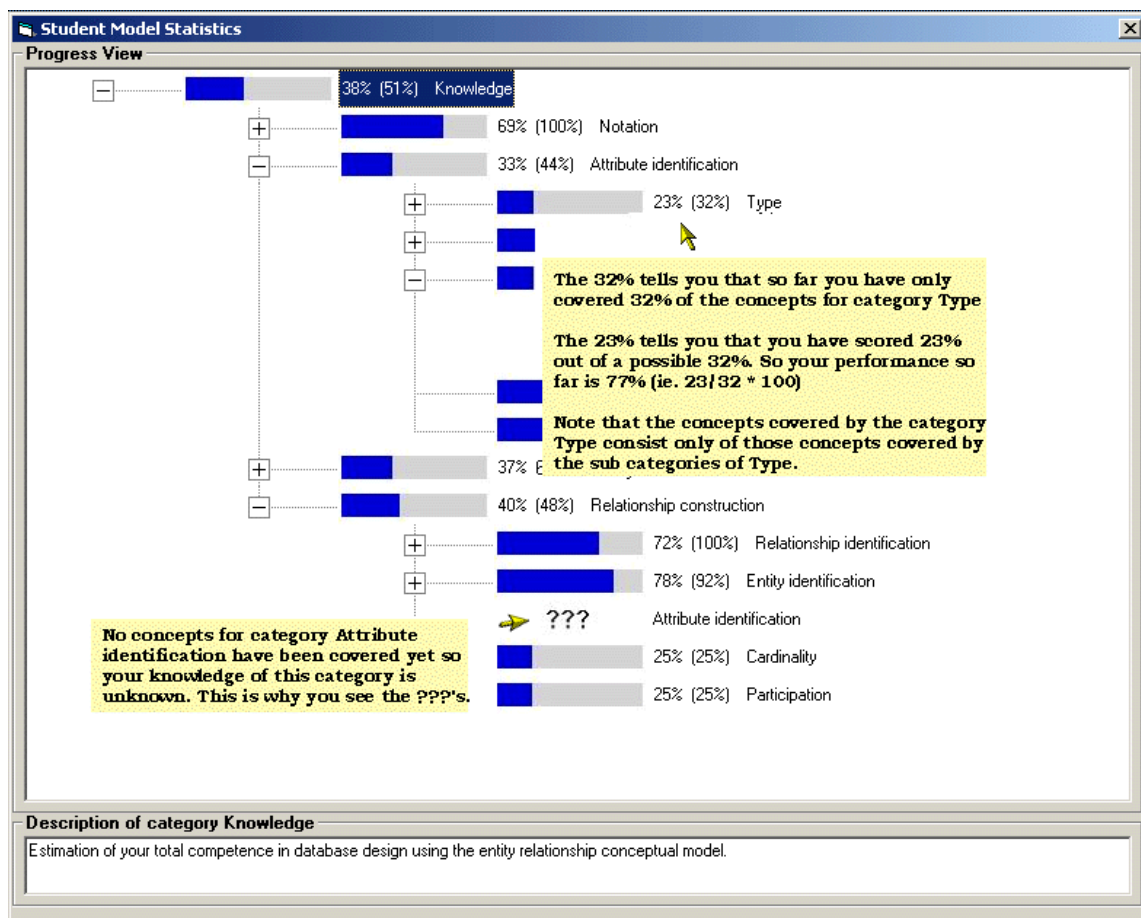


Figure 3. The expanded open student model

Besides summarizing the student's knowledge, the open model also presented a high level view of the instructional domain, which supported the student's understanding of the domain structure. When students inspected the open models, they could reflect on their knowledge and reconsider their beliefs about domain concepts. By providing such visualizations, the student model was not just a source of knowledge about the student valuable to the system, but became an important learning resource on its own. The study conducted with e-KERMIT showed encouraging results, with the majority of the students exploring the expanded open student model. Later on, we conducted studies on additional visualizations of student models (Duan, Mitrovic & Churcher, 2010; Mathews et al., 2012).

After the initial success with KERMIT, we developed EER-Tutor, a Web-enabled tutor. Figure 4 shows its interface, which has many similarities to the original version. The requirements are shown at the top of the interface, below which is the toolbar showing the tools corresponding to the components of the EER model. Students can submit their solutions whenever they want, after which they get feedback shown in the right pane of the interface. The system also highlights in red incorrect parts of the solution in the drawing area.

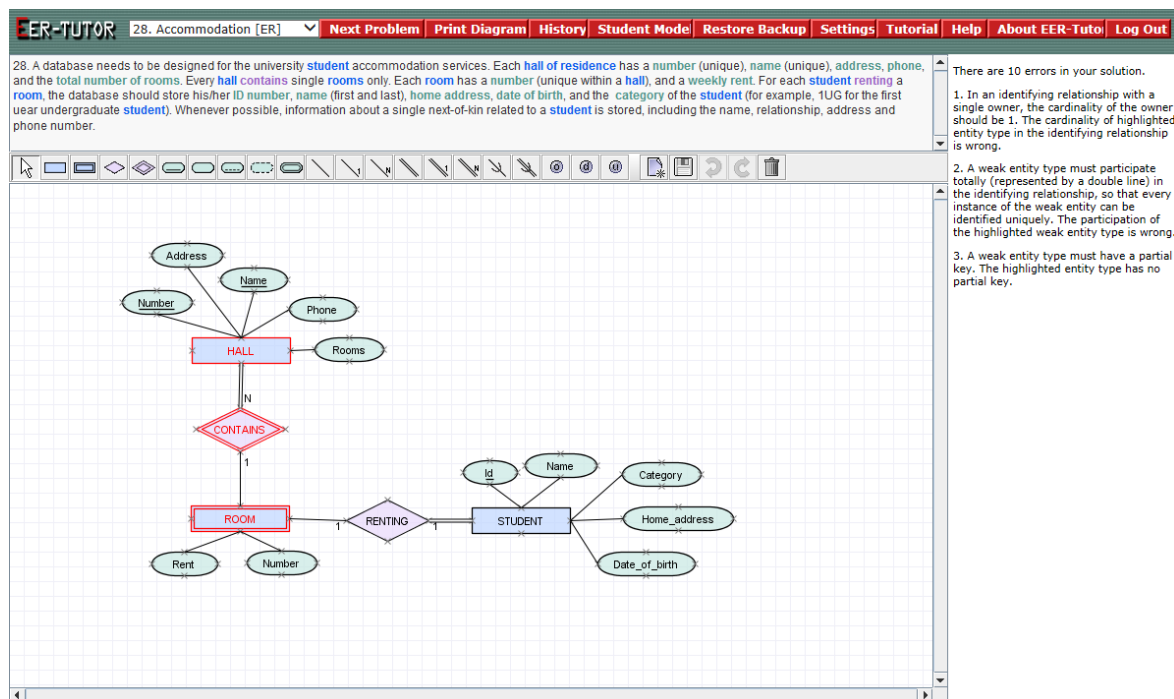


Figure 4. The interface of EER-Tutor

EER-Tutor was developed in WETAS, the authoring shell developed by Brent Martin (Martin & Mitrovic, 2002). It is one of our three tutors that have been available on Addison-Wesley's DatabasePlace Web portal, and used by more than 10,000 students worldwide. EER-Tutor supports the Enhanced Entity Relationship model as defined by Elmasri and Navathe (2011), which extends the basic ER model by introducing specializations and categories. The current version of EER-Tutor contains 57 problems, and over 200 constraints.

We conducted several other projects in the context of EER-Tutor. In one of them, we wanted to explore further the effect of feedback. Usually feedback for ITSs is defined by designers by using their intuition, which was the case with the feedback provided by KERMIT and the first version of EER-Tutor. However, the theory of learning from performance errors (Ohlsson, 1996), on which CBM is based, states that the role of the feedback should be to identify the error and explain what constitutes the error (blame assignment), as well as to re-iterate the domain principle violated by the student's solution (remediation). Therefore, we re-engineered feedback for EER-Tutor and conducted a study comparing EER-Tutor to a version of the system that provided theory-based feedback. The results of evaluation showed that theory-based feedback is more effective in supporting learning: students who received such feedback learned constraints faster than the students who received the original EER-Tutor feedback (Zakharov, Mitrovic & Ohlsson, 2005).

EER-Tutor paved the way for further research; its code was used as the foundation to develop COLLECT-UML, a constraint-based system that teaches UML design and supports pairwise collaboration between students (Baghaei, Mitrovic & Irwin, 2007). Kon Zakharov developed an affect-sensitive agent for EER-Tutor (Zakharov, Mitrovic & Johnston, 2008), which tracked the student's affective states from the student's facial features and problem-solving actions, and responded

1
2
3
4 to them by modifying its facial expressions and providing affect-sensitive feedback. For example, the
5 agent encouraged the student to continue with problem solving when there was only one error found in
6 the solution by saying “*Just a little more effort and you get there – it will make you feel great!*”, or
7 provided advice for a struggling student by saying “*It seems you are somewhat frustrated. Would it
8 help if you started working on one of the simpler problems?*”
9

10 Amali Weerasinghe developed tutorial dialogues for KERMIT (Weerasinghe & Mitrovic, 2003),
11 and later a model for providing adaptive tutorial dialogues for EER-Tutor, effectively providing
12 substep instruction (VanLehn, 2011). Recently, that model was extended to provide adaptive tutorial
13 dialogs (Weerasinghe et al., 2009, 2011). The dialogs are selected adaptively, on the basis of the
14 student model, by identifying the concepts that the student had most difficulties with, and then
15 selecting the tutorial dialogs corresponding to those concepts. Additionally, there are adaptation rules
16 which individualize the dialogs to suit the student’s knowledge, in terms of the length of the dialog
17 and the exact content of the dialog. In response to the generated dialog, learners are able to provide
18 answers by selecting the correct option from a list provided by the tutor. In a study conducted in
19 March 2010, the students who had adaptive dialogs outperformed their peers who only received non-
20 adaptive dialogs, with the effect size 0.69 on learning gains after approximately 100 minutes of
21 interaction with the system (Weerasinghe & Mitrovic, 2010). The obtained effect size is remarkable
22 because the only difference between the two groups was the adaptivity of the dialogues. Additionally,
23 Myse Elmadani has looked more deeply into how students interact with tutorial dialogues in EER-
24 Tutor by capturing and analyzing eye-tracking data (Elmadani, Mitrovic & Weerasinghe, 2013).
25

26 Our papers on teaching database design and on SQL-Tutor also contributed to popularization of
27 CBM, which has been used since by many researchers in addition to those coming from our research
28 group – see (Mitrovic & Ohlsson, 2015). Our future plans include enhancing EER-Tutor with other
29 forms of learning, such as learning from examples and providing motivational support to students.
30
31

32 33 **ACKNOWLEDGMENTS**

34
35
36 The work reported here could not have been done without the support of other members of the
37 Intelligent Computer Tutoring Group. We thank them all for the discussions and friendship over the
38 years. We thank Brent Martin, Konstantin Zakharov, Moffat Mathews and Jay Holland for their work
39 on EER-Tutor.
40

41 42 **REFERENCES**

- 43
44 Baghaei, N., Mitrovic, A., & Irwin, W. (2007) Supporting collaborative learning and problem-solving in a
45 constraint-based CSCL environment for UML class diagrams. *Computer-Supported Collaborative*
46 *Learning*, 2(2-3), 159-190.
47
48 Chen, P. P. (1976) The Entity Relationship Model - Toward a Unified View of Data. *ACM Transactions*
49 *Database Systems*, 1, 9-36
50
51 Duan, D., Mitrovic, A., & Churcher, N. (2010) Evaluating the effectiveness of multiple open student models in
52 EER-Tutor. S. L. Wong et al. (Eds.) Proc. 18th Int. Conf. Computers in Education (pp. 86-88)
53
54 Elmadani, M., Mitrovic, A., Weerasinghe, A. (2013) Understanding Student Interactions with Tutorial Dialogues
55 in EER-Tutor. L. H. Wong, C-C Liu, T. Hirashima, P. Sumedi, M. Lukman (Eds.) Proc. 21st Int. Conf. on
56 Computers in Education (pp. 30-40)
57
58
59
60
61
62
63
64
65

- 1
2
3
4 Elmasri, R., & Navathe, S.B. (2011) Fundamentals of Database Systems. Pearson Education, 6th edition.
5 Goel, V., & Pirolli, P. (1988) Motivating the Notion of Generic Design with Information Processing Theory: the
6 Design Problem Space. *AI Magazine*, 10, 19-36.
7 Goel, V., & Pirolli, P. (1992) The Structure of Design Problem Spaces. *Cognitive Science*, 16, 395-429.
8 Hartley, D., & Mitrovic, A. (2002) Supporting learning by opening the student model. In: S. Cerri, G.
9 Gouarderes, F. Paraguacu (Eds.) *Proc. 6th Int. Conf. Intelligent Tutoring Systems*, LCNS 2363 (pp. 453-
10 462)
11 Martin, B., & Mitrovic, A. (2002) WETAS: a Web-Based Authoring System for Constraint-Based ITS. In: P. de
12 Bra, P. Brusilovsky and R. Conejo (Eds.) *Proc. 2nd Int. Conf. on Adaptive Hypermedia and Adaptive Web-
13 based Systems*, LCNS 2347 (pp. 543-546)
14 Mathews, M., Mitrovic, A., Lin, B., Holland, J., & Churcher, N. (2012) Do your eyes give it away? Using eye
15 tracking data to understand students' attitudes towards open student model representations. S.A. Cerri and
16 B. Clancey (Eds.) *Proc. 11th Int. Conf. Intelligent Tutoring Systems*, Crete, Greece, Springer-Verlag,
17 LCNS 7315 (pp. 424-429)
18 Mayo, M., & Mitrovic, A. (2001) Optimising ITS Behaviour with Bayesian Networks and Decision Theory.
19 *Artificial Intelligence in Education*, 12(2), 124-153.
20 Mitrovic, A., & Weerasinghe, A. (2009) Revisiting the Ill-Definedness and Consequences for ITSs. Dimitrova,
21 V., Mizoguchi, R., du Boulay, B., Graesser, A (Eds.) *Proc 14th Int. Conf. Artificial Intelligence in
22 Education* (pp. 375-382)
23 Mitrovic, A., & Ohlsson, S. (2015) Implementing CBM: SQL-Tutor after fifteen years, *Artificial Intelligence in
24 Education*, Artificial Intelligence in Education, 25(2).
25 Ohlsson, S. (1992) Constraint-based student modelling. *Artificial intelligence in Education*, 3, 429-429.
26 Ohlsson, S. (1996) Learning from Performance Errors. *Psychological Review*, 103, 241-262.
27 Ohlsson, S. (2015) Constraint-Based Modeling: From Cognitive Theory to Computer Tutoring – and Back
28 Again, *Artificial Intelligence in Education* (this issue).
29 Reitman, W.R. (1964) Heuristic Decision Procedures, Open Constraints, and the Structure of Ill-defined
30 Problems. In M. W. Shelly & G. L. Bryan (Eds.) *Human Judgements and Optimality*. New York: Wiley.
31 Suraweera, P., & Mitrovic, A. (2004) An Intelligent Tutoring System for Entity Relationship Modelling.
32 *Artificial Intelligence in Education*, 14(3-4) (pp. 375-417).
33 VanLehn, K. (2011) The relative effectiveness of human tutoring, intelligent tutoring systems and other tutoring
34 systems. *Educational Psychologist*, 46(4), 197-221.
35 Weerasinghe, A., & Mitrovic, A. (2003) Effects of self-explanation in an open-ended domain. In: U. Hoppe, F.
36 Verdejo & J. Kay (ed) *Proc. 11th Int. Conference on Artificial Intelligence in Education AIED 2003*, IOS
37 Press (pp. 512-514)
38 Weerasinghe, A., Mitrovic, A., Martin, B. (2009) Towards individualized dialogue support for ill-defined
39 domains. *Artificial Intelligence in Education*, 19(4) (pp. 357-379)
40 Weerasinghe, A., Mitrovic, A., Martin, B. (2010) Evaluating the effectiveness of adaptive tutorial dialogues in
41 database design. S.L.Wong et al. (Eds.) *Proc. 18th Int. Conf. Computers in Education*, APSCE, (pp. 33-40)
42 Weerasinghe, A., Mitrovic, A., Thomson, D., Mogin, P., Martin, B. (2011) Evaluating a General Model of
43 Adaptive Tutorial Dialogues. In: Biswas, G., Bull, S., Kay, J., Mitrovic, A. (eds.), *Proc. 15th Int. Artificial
44 Intelligence in Education*, Springer, LNAI 6738 (pp. 394-402)
45 Zakharov, K., Mitrovic, A., & Ohlsson, S. (2005) Feedback Micro-engineering in EER-Tutor. In: C-K Looi, G.
46 McCalla, B. Bredeweg, J. Breuker (Eds.) *Proc. Artificial Intelligence in Education*, IOS Press (pp. 718-
47 725).
48 Zakharov, K., Mitrovic, A., & Johnston, L. (2008) Towards Emotionally-Intelligent Pedagogical Agents. B.
49 Woolf et al. (Eds.) *Proc. Intelligent Tutoring Systems* (pp. 19-28).
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65