School of IT
Technical Report

The University of Sydney

TARO: TEACHING WITH AN AUTOMATICALLY RETRIEVED ONTOLOGY
TECHNICAL REPORT 596

JUDY KAY, ANDREW LUM, WILLIAM NIU
AND LICHAO LI
THE UNIVERSITY OF SYDNEY

SEPTEMBER, 2006

TARO: Teaching with an Automatically Retrieved Ontology

Judy KAY, Andrew LUM, William NIU, and Lichao LI
School of Information Technologies,
University of Sydney, NSW, 2006, Australia
{alum, judy, alum, niu, lli1 @it.usyd.edu.au

Abstract. Ontologies are used in an increasing number of roles within Intelligent Tutoring Systems. In this paper, we describe TARO (Teaching with an Automatically Retrieved Ontology), a system which explores a role for ontologies that are externalised and used, in conjunction with a learner model, as a basis for independent learning. A novel and important feature of our approach is that the ontology is built *automatically* from a dictionary designed for people to read: the same dictionary is used, in conjunction with the ontology built from it to provide both a mechanism for explaining the meaning of the parts of a learner model and as a new learning source. We report a small-scale qualitative evaluation with five users, who self-assessed their knowledge before and after using TARO and who used TARO to assist them in refining their knowledge, externalising this in a concept-map. The study gives promising results both in terms of learner perceptions of their knowledge and with all learners demonstrating improved understanding in the concept-mapping task. TARO offers promise as a tool that can be used in conjunction with any teaching system that builds a learner model which can be shared with the learner.

1 Introduction

Ontologies have many potential roles in intelligent teaching systems. This was recognised in the earliest teaching systems such as Carbonell’s geography tutor (Carbonell 1970). We are particularly concerned with the roles that ontologies can play in student modelling. They can play a critical role in supporting reasoning across granularity levels: for example, we can reason about the learner’s high level knowledge even when we can only access fine grained data about their knowledge (Kay and Lum 2005).

There is a growing body of evidence that there are real learning benefits from making learner models accessible to the learner. The series of workshops on Learner Modelling for Reflection (LeMoRe) have reported many systems that make various forms of learner model open to the learner so that they can monitor their own learning, plan and control future learning and contribute to the correctness of the model.

Ontologies play a role here, making it feasible to provide a learner with an overview of their knowledge as reflected in a large learner model for a substantial domain. The ontology is critical in providing a structure over the domain so that, when mapped to components in a learner model, we can give an overview visualisation of the whole domain with hundreds of learner model component values easily visible (Apted and Kay 2004). This can serve as a basis for finding and scrutinising details of particular interest. For example, where the learner knows much of an area well, they can focus on the elements that they do not know.

In this paper, we are concerned with two critical problems associated with effective use of learner models. Firstly, if we present a learner model to the learner, we want the learner to be able to determine the meaning of each component in the model. Taking an example where a learner model for *UNIX* shows that the student knows most concepts but does not know about *access control*. We want the student to be able to able to see what *access control* means in this context. This may be critical in enabling them to determine whether they want to learn about it. In life-long learner controlled learning and much workplace learning, this is important for encouraging learning. Essentially, this problem is important because we should be able to support learning about a domain by explaining the meaning of the learner model.

One crucial step in concept teaching is to determine a taxonomical structure of the content (Tennyson and Park 1980). This lead to our second problem where we want to address is the use of ontologies to support both the needs of the teaching system and the learner in understanding the domain structure. The latter is linked to our first problem because the ontology captures the relationships between the components of the learner model and these relationships are essential to understanding the domain. This second problem is important because we should be able to support learning about a domain by explaining the relationships between concepts in the learner model.

An associated goal of our work is to achieve low cost construction of ontologies at time as addresses the two problems above. Our approach is to make use of dictionaries and glossaries that are written for people to read. These are widely available on the web and are commonly available as adjunct resources for courses. We want to be able to automatically analyse these dictionary or glossary sources to build the ontology in such a way that we can then refer to the dictionary definition to address the first problem, providing learners with an explanation of any component of the learner model that they want to know more about. At the same time, we want to exploit the ontology, sharing it with the learner, to provide a richer description of the meaning of those components in terms of their relationship to all other concepts in the learning model.

Our solution to the two problems is called TARO, Teaching with an Automatically Retrieved Ontology. Section 2 gives an overview of the architecture and user view of TARO. Section 3 reports our evaluation where we examine how learners interact with TARO and whether they can use it to learn more about unknown components of their learner models. We then describe related work and finish with a conclusion of the contributions of our research.

2 Approach

This section describes components of TARO, as well as how they are integrated. Fig. 1 shows schematics of the architecture of TARO. There is a core system (TARO) as well as two major sub-systems (SIV and SAGO) which are denoted with round-square blocks in the diagram.

Teaching with an Automatically Retrieved Ontology

Teaching with an Automatically Retrieved Ontology (TARO) is a web-based system that takes inputs of a learner model and a glossary ontology that uses the Simple Knowledge Organisation System, SKOS (Miles and Brickley 2005), and provides visualisation of the ontology as well as its relationships and reasoning. Fig. 2 shows a screenshot of the system. The learner model indicates a learner’s familiarity over a set of concepts. This may help him or her spot the weaknesses and provide potential remedies by examining the less known concepts and their related ones to progressively build up new knowledge. The system is capable of displaying definitions in a glossary using SKOS, as well as (first-level) relationships and their reasoning, if provided.

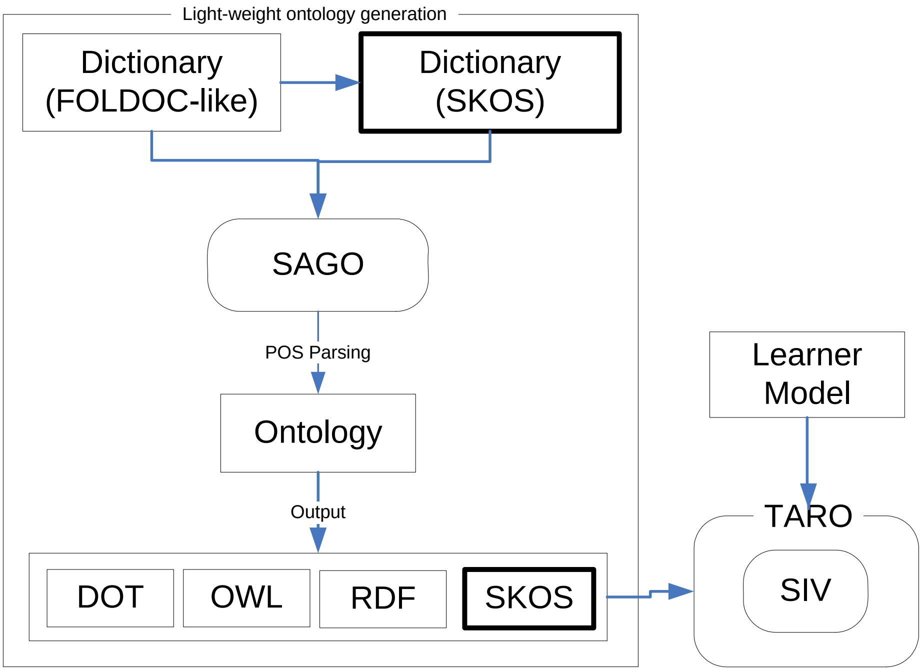

Fig. 1. Conceptual architecture of TARO

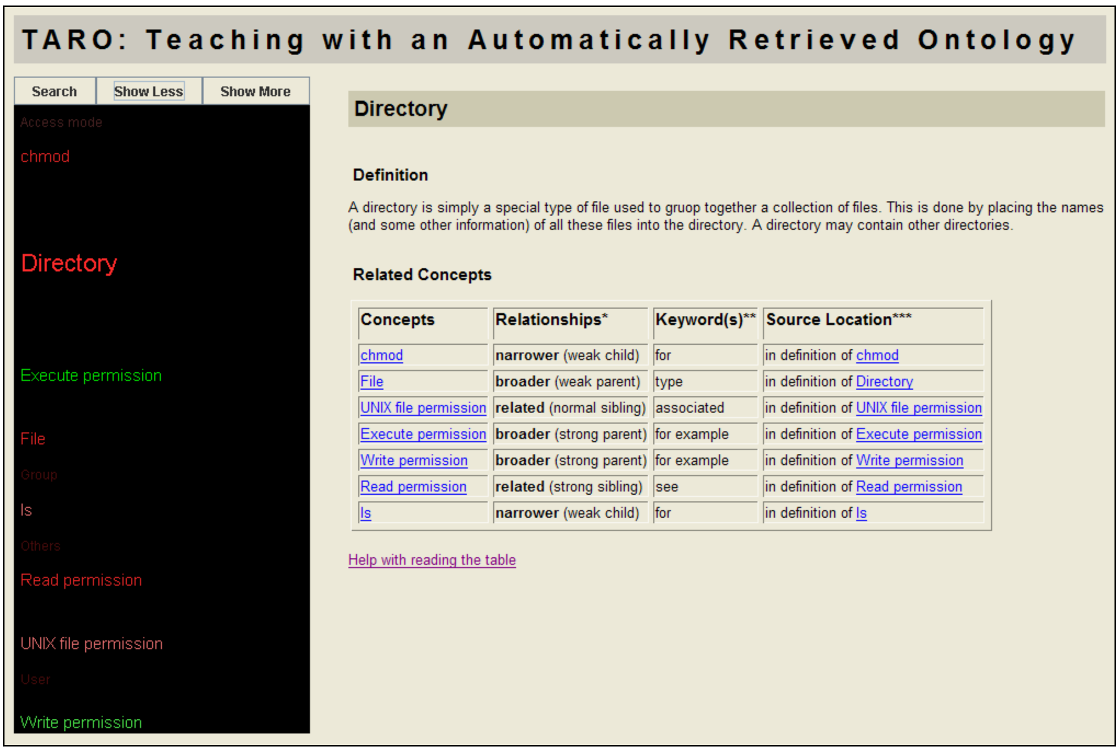

Fig. 2. A screenshot of the TARO interface, with SIV displayed on the left, and information about the selected concept on the right.

From a learner point of view, TARO enables a learner to scrutinise many aspects of a domain ontology, in particular, the relationships of each concept within the domain and their reasoning. In SIV, which is located on the left of a screen in Fig. 2, the colouring denotes a learner's understanding towards the concepts and font size and brightness present relatedness between the selected concept and the others. On the right side of the screen, the learner is presented with the definition of the selected concept and its related concepts with reasoning in an HTML table. Take the third row in the table with the related concept of *File* as an example, it reads "the concept *File* is a broader term (or a weak parent) of the concept *Directory*, and their relationship is found in the definition of *Directory* based on the keyword *type*". With careful look, it can be found that SAGO extracted this relationship from the first sentence in the definition in Fig. 2, "A **directory** is simply a special **type** of **file**…".

Although efforts are made to make TARO generic in terms of accepting SKOS glossary files, SKOS is still a developing standard, and there is instability in the use of classes, properties and attributes. There are, however, several projects that use SKOS, such as GEneral Multilingual Environmental Thesaurus[1] and W3C Glossary and Dictionary[2]. See Section 4 for details.

Scrutable Automatically Generated Ontologies

Scrutable Automatically Generated Ontologies (SAGO) builds on Mecureo (Apted and Kay 2004). Mecureo is a tool that automatically generates a light-weight ontology from a dictionary or a glossary in FOLDOC-like format. We would like to improve Mecureo in three aspects: conform both input (a dictionary/glossary source) and output (an ontology) to standards, refine the relationships, and add scrutability support.

We adopted the SKOS Core Vocabulary for both input and output format (Mizoguchi, Kozaki et al. 2000). SKOS is an application of Resource Description Framework (RDF)[3] and is designed to describe taxonomy and classification schemes which would be a plausible candidate for representing a subject domain ontology. Although this specification is still under development, it has been researched and used by several different research groups and communities (see Section 4). The fact that we use SKOS to represent both input and output of SAGO means we largely enrich a dictionary file by adding in inter-relationships, as well as greatly speed up the ontology building process by making it automatic.

Fig. 3 shows an example input and output files in SKOS format for SAGO. The part above the line is a `Concept` class, which corresponds to a concept, with the name of a concept (*Write permission*) and its definition. The other part represents the same concept with ontological information added by SAGO. Namely, `sago:keyword`, `sago:source`, `sago:strength`, `sago:type`, and `sago:linkWeight`. In SAGO, an ontology is a graph with each concept as a node and its relationships as edges. `Sago:linkWeight` is a reliability weighting for a link calculated by the system (Apted and Kay 2004). The other four properties correspond to the last three columns of the table in Fig. 2: *keyword(s)* = `sago:keyword`, *source location* = `sago:source`, *relationship* (inside the parentheses) = `sago:strength` and `sago:type`. With added SAGO vocabularies, a compatible parser may extract more detailed information about the ontology, but at the same time, a standard SKOS parser can get a simplified version of the ontology as well.

[1] http://www.eionet.europa.eu/GEMET
[2] W3C Glossary and Dictionary: http://www.w3.org/2003/glossary/
[3] Resource Description Framework (RDF): http://www.w3.org/RDF/

```
<skos:Concept rdf:about="write%20permission">
   <skos:prefLabel>Write permission</skos:prefLabel>
   <skos:definition>An access mode that allows a person to...</skos:definition>
</skos:Concept>
```

```
<skos:Concept rdf:about="Write%20permission">
   <skos:prefLabel>Write permission</skos:prefLabel>
   <skos:definition>An access mode that allows a person to...</skos:definition>
   <skos:narrower rdf:resource="#File">
       <sago:keyword sago:source="Write%20permission">for example</sago:keyword>
      <sago:strength>strong</sago:strength>
      <sago:type>child</sago:type>
      <sago:linkWeight>0.435500</sago:linkWeight>
   </skos:narrower>
   <skos:related rdf:resource="#UNIX%20file%20permission">
       <sago:keyword sago:source="UNIX%20file%20permission">
         associated</sago:keyword>
       <sago:strength>normal</sago:strength>
       <sago:type>sibling</sago:type>
       <sago:linkWeight>0.671950</sago:linkWeight>
   </skos:related>
   ...
</skos:Concept>
```

Fig. 3. SAGO input dictionary file (above the line) and output ontology file (below the line)

Table 1. Comparison of internal relationships between SKOS Core, Mecureo, Genetic Graph, and WordNet

SKOS Core	Mecureo	Genetic Graphs	WordNet
Broader	Parent	Generalisation	Hypernym
	Category	Simplification	Meronym
Narrower	Child	Specification	Troponym
		Refinement	Hyponym
			Holonym
Related	Sibling	Analogy	Morphological
	Synonym	Deviation	Synonym
		Correction	
	Antonym		Antonym

Another aspect of Mecureo we looked into was the internal representation of relationships: *parent*, *child*, *sibling*, *category*, *synonym*, *antonym*, and *unknown*. A couple of pairs of concepts have somewhat similar meaning. For example, *synonym* is a type of this. That may cause confusion in our situation when the ontology can be scrutinised by learners. Here we made comparison between four different relationship schemes: SKOS Core[4], Mecureo, Genetic Graph (Niem, Fugére et al. 1993), and WordNet (Miller 1995) in Table 1. SKOS Core vocabulary has a merit of simplicity, which is easier for learners to comprehend when used for ontology scrutiny, but it fails to address *antonym*s and may be too generalised.

Genetic graphs give smaller granularity in relationship classification, but it also fails to classify antonyms. One advantage of them over the others is that they were designed to address education relationships, such as *deviation* and *correction*. The *category* and *parent*, *sibling* and *synonym* relationships in Mecureo may cause confusion on learners' side, although internally they do cover different terms under each umbrella. WordNet is at the other spectrum of SKOS Core, which is often too lexical for education systems. We deliberately omit the SKOS Extensions specification because of its instability[5]. For example, relationships likes *broaderPartitive* and *relatedHasPart* have not been clearly defined.

Scrutability support for learners was the last aspect we examined. SAGO adds information generated in the ontology creation process to support user scrutiny. It retains the keyword(s) that Mecureo uses to identify each relationship as well as the location where the relationship is found, namely the file name and the concept whose definition contains the keyword(s).

Scrutable Inference Viewer

The Scrutable Inference Viewer (SIV) system (Kay and Lum 2005) builds upon a novel interface for visualising large user models (Uther 2001)and allows us to visualise learner models structured by ontologies automatically generated by Mecureo. We have extended SIV to enable the parsing and visualisation of the SAGO-generated ontologies described in the previous section. The remainder of this section, we describe the SIV interface in respect to the TARO system.

Concepts are listed vertically, and are shown in different font sizes and spacing. Clicking on a concept will put it into focus. It will have the largest font and spacing. In Fig. 1, the concept *Directory* is in focus. Concepts related to the focused one are shown in the next largest font.

In Fig. 4, the concepts *chmod*, *Execute permission*, and *File* are examples of concepts directly related to *Directory*. Clicking on one of these concepts will put it into focus, and the visualisation will change the font sizes of related concepts relative to the focused one. The amount of concepts displayed at any one time can be changed via the buttons described below.

We have modified the control elements above the visualisation to have three buttons. The **Search** button functions as normal, allowing users to type in a concept or part of a concept name, and SIV will display all relevant results. The **Show Less** and **Show More** buttons allow users to toggle the amount of concepts displayed on the visualisation.

Each concept on the visualisation is coloured in hues of red or green. The colouring represents information about the current user from their learner model. The system believes the user knows green concepts well, and red concepts are ones the user either does not know well or there is no evidence to support that they know it.

---

[4] Simple Knowledge Organisation System (SKOS) Core: http://www.w3.org/2004/02/skos/core

[5] This is stated in the SKOS Extensions Vocabulary Specification at: http://www.w3.org/2004/02/skos/extensions/spec/

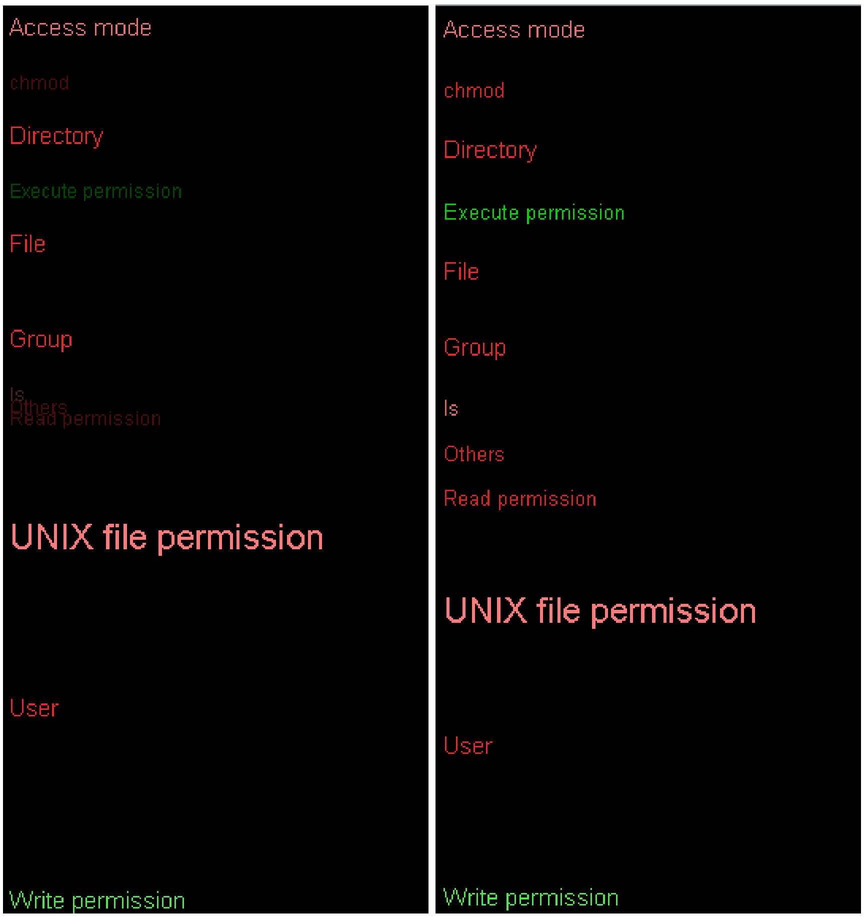

Fig. 4. The system defaults to showing just the immediately related concepts (left). We can show concepts related to the immediately related concepts by clicking on the **Show More** button (right).

3 Evaluation

We carried out a small-scale qualitative evaluation as our initial study for TARO to find out if the use of the system would enable learners to get a better understanding of the domain.

Hypothesis

Based on an assumption of a constructivist epistemology, human beings use past experience to construct meanings for events and objects (Novak 1996). It implies that people try to relate an event or an object to another one to help thinking. To depict this with a cognitive map in mind, it is conceivable that adding (or learning) new concepts one by one into the existing cognitive map is more manageable than flooding them all in at once.

Hence, one of our hypotheses is that it is effective to learn new concepts by progressively building on the existing ones. The fact that TARO only shows the directly related concepts to each concept ease the learning path and may help learners focus on smaller region of the subject domain.

The other hypothesis is that reasoning about the domain ontology and relationships between concepts helps learners comprehend how one concept related to another, and thereby they can more effectively understand the concepts.

Context

There are two requirements an evaluation domain has to meet: small enough to have a manageable number of concepts and large enough to contain a set of expressive inter-relationships between them. It has been shown that a concept map with more than 12 concepts may be difficult to compose with a reasonable time limit (Yin, Vanides et al. 2005).

After an iterative designing process, we went from the domain of a second-year course to a specific topic within it, namely *introduction to UNIX*, to a more specific area—*UNIX file permission*. We also trimmed the number of concepts down to 12:

read permission	write permission	execute permission
UNIX file system	chmod	ls
access mode	file	directory
user	others	group

Participants

The ideal participants were competent computer users with limited or no knowledge in UNIX. As this experiment was designed to be a qualitative initial study, we only recruited five participants; three females and two males. The age range was from 19 to 32. All of them were competent at using a browser to view Internet contents, which was what TARO user interface was based on. Three of them did not know UNIX operating system, one knew a limited set of commands, and one, being a regular user of Apple operating system, knew some basic commands. Four of them had very limited knowledge on file permission and one knew it fairly well. The bars with a lighter colour in Fig. 6 reflect this.

Experimental Design

There were two problems we faced when designing the study. First, we needed a way to quantify participants' knowledge gain over a set of concepts after using TARO. Second, as TARO is a passive teaching system, it needed a mechanism to start up the participant as well as guide him or her through the learning process. The theory of *concept inventory* (Tamir 1984) was adopted to overcome the first problem. Participants were asked to assess their own knowledge in a five-point scale towards the 12 concepts before and after using TARO. Fig. 5 shows a screenshot of the assessment. This allowed us to quickly quantify the participants' knowledge over the set of concepts.

The comparison of the two self-assessments gave the participants' potential knowledge gain after using the system. The first assessment also allows the generation of a simple learner model, which contributes to addressing the second problem. The learner model that was automatically constructed from the results of the first concept inventory was represented in five different colours, corresponding to the five-point scale for each concept, on SIV (see Fig. 4), with an exception being the concept of *UNIX file permission*, which was an average of the other concept ratings.

It is perceivable that learners would be more interested in either 1) finding out concepts they know less well or 2) confirming their understanding on the concepts they think they know. After started off, the participants would ideally try to expand their cognitive maps by finding out the related concepts that they know less well, thereby to progressively add new concepts into their cognitive maps as well as the concept inventories.

As described in Section 2, TARO is designed to increase scrutability on relationships and their reasoning between the concepts in a domain ontology. This allowed us to test our second hypothesis that reasoning for relationships between concepts helps learners to effectively understand the concepts.

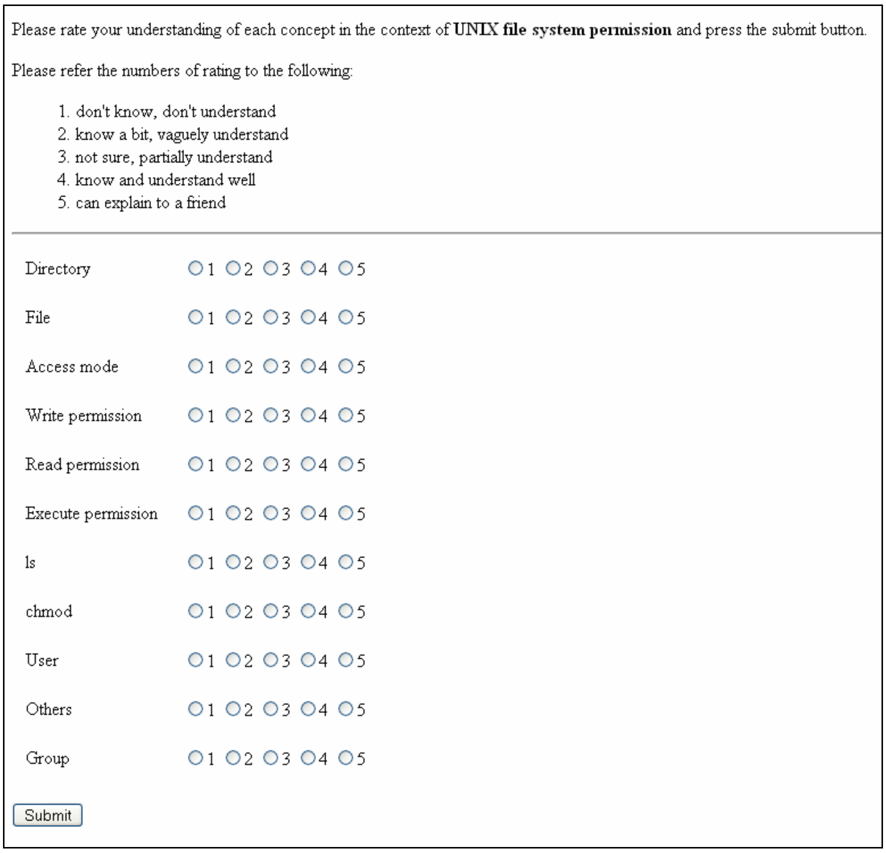

Fig. 5. A screenshot of the concept inventory for assessing learners' knowledge on UNIX file permission.

In addition, concept mapping was used to reflect how the participant organised thoughts with a cognitive map. Each participant was asked to construct an initial concept map before they started using the system. As some participants were not familiar with some of the concepts, they were not able to include all the concepts in their initial map. They then expanded the concept map as they went through TARO. They should arrive at a concept map with all 12 concepts included as well as have the relationships altered, added, or confirmed. Comparison between the initial and the final maps also served as another measurement for knowledge gain.

Results

We measured the potential knowledge gain of the learners based on three main factors: concept inventory before and after the learners used TARO (Fig. 6), proportion of "good" propositions in the concept maps constructed before and after the learners used TARO (Fig. 7) and average proposition accuracy for the concept maps constructed before and after the learners used TARO (Fig. 8). The results, as the initial qualitative study, were encouraging.

For the concept inventory, we simply compared the participants' perception of their own knowledge before and after using the system. Fig. 6 illustrates the average scores for each learner as well as their average before (the left, lighter colour bar) and after (the right, darker colour bar) the learner used the system. There was more than 40% of improvement (1.3 points).

As one of the feedback questions after the study, we asked "how accurate do you think a concept inventory task measures your understanding towards the concepts?" Four of the five learners agreed the concept inventory was "quite accurate" in terms of measuring their own knowledge. One thought it was "okay" because some concepts appeared too "vague".

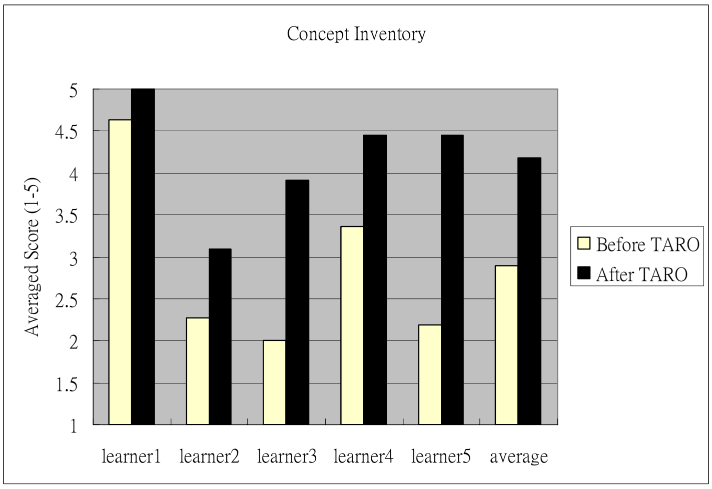

Fig. 6. Average scores of concept inventory tasks before and after using TARO

For scoring concept maps without a criterion map, we adopted a scoring scheme presented in (Yin, Vanides et al. 2005). We focused on the concept-map products (the resulting maps), as opposed to concept-map processes for the study. The paper suggests four criteria for comparing a concept map:

- Proposition choice (qualitative)
- Individual proportion score (quantitative)
- Structure complexity (qualitative)
- Total accuracy score (quantitative)

We chose not to use the criterion of total accuracy score because that was mainly used for evaluating equivalence between two mapping techniques and showing differences in students' knowledge structures, whereas we tried to compare the quality of two maps built by the same learner. For evaluating proposition choice, we used the same four-point scale scheme applied in (Yin, Vanides et al. 2005):

1. wrong or scientifically irrelevant proposition
2. partially incorrect propositions
3. correct but scientifically "thin" propositions
4. scientifically correct and scientifically stated propositions

This evaluates the quality of each proposition made. Fig. 7 shows the average proposition accuracy scores amount the learners and their average. The improvement was more than 45% in average (0.7 point). Even learner1, who already had fair knowledge in UNIX file permission before the study, acknowledged that the task helped refine their understanding.

Next we looked at the individual proportion score shown in Fig. 8. This is to divide the number of "good" proposition by the total proposition for each learner, so to see the proportion of the good propositions a learner made. We defined a good proposition to be the ones scored 2 or 3 on proposition choice criterion. There is more variation for this criterion. This is especially true for learner3 and learner5 who expressed that they had not heard of UNIX before. The average improvement was over 70%. Finally we looked at the structure complexity. There was only one person who had different map structures; from a tree proposition structure to a net one, which is more complex.

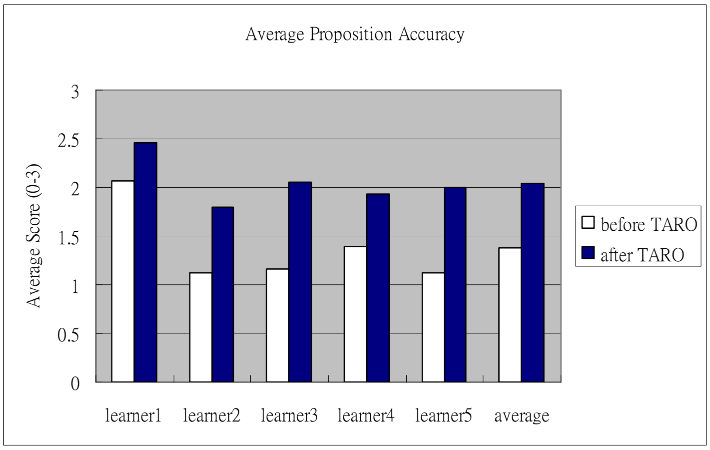

Fig. 7. Average proposition scores of the concept maps constructed before and after using TARO

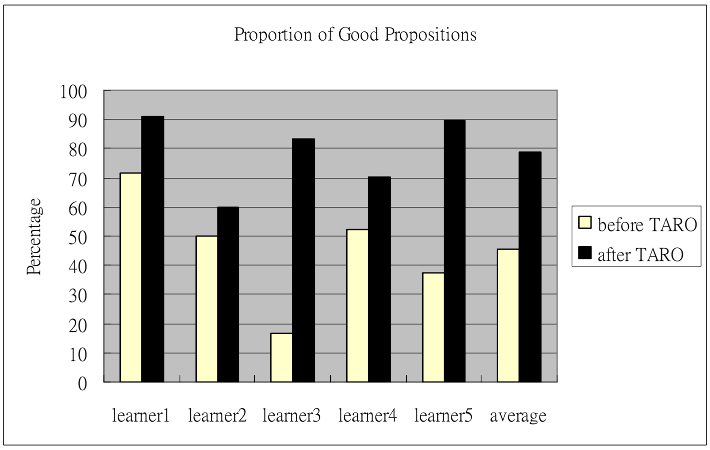

Fig. 8. Percentage of good propositions for the concept maps constructed before and after using TARO

Some general comments from the feedbacks:

- All participants agreed that the study helped refine their understanding towards some or all of the concepts.
- Four participants thought that concept inventory was quite accurate in terms of measuring their own knowledge in the domain. One thought it was okay and pointed out that some concepts were vague.
- All five agreed concept mapping somewhat helped them learn. Two of them pointed out the map got clogged up as all the concepts were put in.
- Only one person found the related concepts table helpful. The others either found it inaccurate and stopped using it or the table did not make much sense to them, although an example was included below the table on how to interpret it.
- One participant reported that she practically used the system like a glossary browser; that is, she ignored the font colour, size and brightness, as well as the related concepts, and only read the definitions.

The results set a promising start, although there were still many places we could have improved:

- user interface: as a user suggested, highlighting the keywords, and using hyperlinks may ease the scrutiny efforts at learners' side,
- more descriptive concepts: for example, use "owner group" instead of "group",

- reducing the number of concepts: two participants expressed the messiness of a fully-drawn map,
- improving the readability and perceptivity of the presentation for related concepts: we are unsure if putting them in English sentences would improve this situation, since we used that in the first version and thought it greatly decreased the readability when the number of related concepts got large,
- giving a better concept mapping tutorial,
- reducing user browsing options: for example, eliminating SIV, so they could only use related concepts to proceed,
- finally, the accuracy of SAGO needs to be improved.

Even though the table presentation of related concepts was not perceived well, the learners apparently improved their understanding in the subject domain by using the system. We would like conclude that ontological reasoning did assist in learning process.

4 Related Work

Ontological support has become one of the most discussed topics in ITS. Aroyo, Dicheva and Cristea (Aroyo, Dicheva et al. 2002) proposed semi-automate the process of courseware authoring as well as to support cooperative authoring by adding ontology-based layers into an existing prototype, AIMS. They used a subject domain ontology to represent the semantics of the course ontology which in term supports both the users (students) and authors (instructors). They lay the focus on teaching, whereas we emphasis on ontological support for learning.

A research group from the University of Saskatchewan proposed a student modelling architecture that used OWL[6], SKOS, and QTILite[7] specifications for portability as well as decentralises the original heavy-weight ontology into a student model ontology, a domain ontology, and a course-specific ontology for scalability (Winter, Brooks et al. 2005). They used SKOS vocabulary for the course topic ontologies, which is similar, except we propose automatic ontology construction and try to stick with the more stable version of the specification—SKOS Core—rather than SKOS Extensions.

Gasevic and Hatala (Gasevic and Hatala 2006) also adopt the SKOS specifications to define domain taxonomies. Their focus is on mappings between the local learning context and the remote ontologies, in order to facilitate learners to search for (remote) learning resources relevant to, say, a course curriculum.

There are several projects that are either using the SKOS specifications or have put efforts in testing their usability, such as the W3C Glossary and Dictionary project[8], GEMET[9], and SIMILE (Butler, Gilbert et al. 2004).

There are a number of concept learning studies reviewed in (Tennyson and Park 1980) that indicated importance of concept definitions in the process of concept attainment. In particular, studies showed that participants, when given a definition for each concept, performed significantly better on determining definitions and the interrelationships between concepts, than those who were not (Johnson and Stratton 1966; Anderson and Kulhavy 1972). In another study, it was discovered that a definition had almost the same effects as one rational set of examples and non-examples in terms of assisting concept attainment (Klausmeier and Feldman 1975).

---

[6] OWL Web Ontology Language: http://www.w3.org/TR/owl-features/
[7] IMS Question & Test Interoperability QTILite Specification:
http://www.imsglobal.org/question/qtiv1p2/imsqti_litev1p2.html.
[8] W3C Glossary and Dictionary: http://www.w3.org/2003/glossary/
[9] http://www.eionet.europa.eu/GEMET

5 Conclusion

We have described both the architecture and user interface of TARO, as well as conducted an evaluation on the system. The main contributions of this work are in the exploration of ways that an automatically constructed, built from an existing dictionary or glossary can serve as a support for learning in conjunction with an open learner model. We also make a contribution in exploring the potential of some of the standard ontological relationships from Genetic Graphs, SKOS and WordNet, both for representing the ontology and for providing details of relationships to learners. The evaluation points to areas of improvement in the user interface and the choice of ontological relationships that are more helpful for learners. At the same time, the evaluation suggests that the approach has promise for making open learner models more useful, both to make the models meaningful and to serve as a means for learners to build their understanding and learn more.

6 Acknowledgements

We would like to thank Smart Internet Technology CRC[10] and Hewlett-Packard[11] for parts of funding for this project. We also want to thank the participants for their valuable time spent on the user study.

References

Anderson, R. C. and R. W. Kulhavy (1972). "Learning Concepts from Definitions." American Educational Research Journal **9**(3): 385-390.

Apted, T. and J. Kay (2004). "MECUREO Ontology and Modelling Tools." WBES of the International Journal of Continuing Engineering Education and Lifelong Learning **14**(3): 191-211.

Apted, T. and J. Kay (2004). MECUREO Ontology and Modelling Tools. Special Issue on Concepts and Ontologies in WBES of the International Journal of Continuing Engineering Education and Lifelong Learning. **14:** 191-211.

Aroyo, L., D. Dicheva, et al. (2002). Ontological Support for Web Courseware Authoring. 6th International Conference on Intelligent Tutoring Systems, London, UK, Springer-Verlag.

Butler, M., J. Gilbert, et al. (2004). Data conversion, extraction and record linkage using XML and RDF tools in Project SIMILE. Bristol, UK, HP Labs.

Carbonell, J. R. (1970). "AI in CAI: an artificial-intelligence approach to computer-assisted instruction." IEEE Transactions on Man-Machine Systems **11**(4): 190-202.

Gasevic, D. and M. Hatala (2006). "Ontology mappings to improve learning resource search." British Journal of Educational Technology Special issue: The Semantic Web for E-learning **37**(3).

Johnson, D. M. and R. P. Stratton (1966). "Evaluation of Five Methods of Teaching Concepts." Journal of Educational Psychology **57**: 48-53.

Kay, J. and A. Lum (2005). Exploting Readily Available Web Data for Scrutable Student Models. 12th International Conference on Artificial Intelligence in Education, Amsterdam, Netherlands, IOS Press.

[10] http://www.smartinternet.com.au/
[11] http://hpl.hp.com/

Klausmeier, H. and K. Feldman (1975). "Effects of a definition and a varying number of examples and nonexamples on concept attainment." Journal of Educational Psychology **6**(7): 174-178.

Miles, A. and D. Brickley. (2005). "SKOS Core Guide." W3C Working Draft 2 November 2005. Retrieved 26 March, 2006, from http://www.w3.org/TR/2005/WD-swbp-skos-core-guide-20051102/.

Miller, G. A. (1995). "WordNet: a lexical database for English." Commun. ACM **38**(11): 39-41.

Mizoguchi, R., K. Kozaki, et al. (2000). Construction and Deployment of a Plant Ontology. Knowledge Engineering and Knowledge Management - Methods, Models and Tools **-:** 113-128.

Niem, L., B. J. Fugére, et al. (1993). "Defining the semantics of extended genetic graphs." User Modeling and User-Adapted Interaction **3**(2): 107-153.

Novak, J. D. (1996). Concept Mapping as a tool for improving science teaching and learning. Improving Teaching and Learning in Science and Mathematics. D. F. Treagust, R. Duit and B. J. Fraser. London, Teachers College Press**:** 32-43.

Tamir, P. (1984). "What do learning theories and research have to say to practioners in science education?" Research and Developments in Higher Education **7**: 166-162.

Tennyson, R. D. and O.-C. Park (1980). "The Teaching of Concepts: A Review of Instructional Design Research Literature." Review of Educational Research **50**(1): 55-70.

Uther, J. (2001). On the Visualisation of Large User Model in Web Based Systems. School of Information Technologies. Sydney, Australia, University of Sydney.

Winter, M., C. Brooks, et al. (2005). Towards Best Practices for Semantic Web Student Modelling. Artifical Intelligence in Education, Amsterdam, Netherlands, IOS Press.

Yin, Y., J. Vanides, et al. (2005). "Comparison of two concept-mapping techniques: Implications for scoring, interpretation, and use." Journal of Research in Science Teaching **42**(2): 166-184.

School of Information Technologies, J12
The University of Sydney
NSW 2006 AUSTRALIA
T +61 2 9351 3423
F +61 2 9351 3838
www.it.usyd.edu.au

ISBN 1 86487 867 6