Knowledge matchmaking in Learning Networks: Alleviating the tutor load by mutually connecting learning network users

DOI:
10.1111/j.1467-8535.2006.00673.x

Document status and date:
Published: 01/11/2006

Document Version:
Peer reviewed version

Knowledge matchmaking in Learning Networks: Alleviating the tutor load by mutually connecting learning network users

Peter van Rosmalen, Peter Sloep, Francis Brouns, Liesbeth Kester, Malik Koné and Rob Koper

All authors are employed at the Open University of the Netherlands.
Address for correspondence: Peter van Rosmalen, Open University of the Netherlands, P.O. Box 2960, 6401 DL Heerlen, The Netherlands. Email: {Peter.vanRosmalen, Peter.Sloep, Francis.Brouns, Liesbeth.Kester, Malik.Kone and Rob.Koper} @ou.nl.

Abstract

Tutors have only limited time to support the learning process. In this paper we introduce a model that helps to answer the questions of students. It invokes the knowledge and skills of fellow-students, who jointly form an ad hoc, transient community. The paper situates the model within the context of a Learning Network, a self-organized, distributed system, designed to facilitate lifelong learning in a particular knowledge domain. We will discuss the design of the model and explain how we select and support capable peers. Finally, we will examine the calibration of the model and a simulation, which is intended to verify if the model is fit for use in experiments with students. The results indicate that, indeed, it is possible to identify and support capable peers efficiently and effectively.

Keywords

E-learning, Lifelong Learners, Tutor Workload, Peer Support, Ad Hoc Transient Community, Latent Semantic Analysis.

Introduction

In modern learning settings students typically spend a significant amount of time learning online. In this respect, these settings diverge from the classroom-based, face-to-face learning situations that we are all so familiar with. But they differ in more significant ways too. The advent of the knowledge economy and the individualisation of our society are two leading factors that underpin the increasing demand for flexibility: students want to be able to study at the place, time and pace of their own choosing (logistic flexibility); also, students are unwilling to submit themselves to pre-planned, rigid programmes but want their prior competences honoured and their specific study plans catered for (subject matter flexibility).

These developments called for a new perspective on learning that has become known as lifelong learning which upholds a central position for the learner. The lifelong learner is self-directed, and can perform different formal and informal learning activities in different contexts at the same time. Inherent to this is that learning activities take place in environments populated with learners in any given domain of knowledge with different levels of competence, varying from novices to top-experts, and different foci, varying from practitioners to researchers and developers. To accommodate the lifelong learners adequately, it is necessary to *maintain a record* of their growth in competency in a persistent and standard way to ensure that they can search for new learning facilities that fit and extend their current knowledge.

Networks for lifelong learning ('Learning Networks') embody these changes and at the same time seek to address the challenges they pose. A Learning Network (Koper et al., 2005) is a self-organized, distributed system, designed to facilitate lifelong learning in a particular knowledge domain. A Learning Network is special in that it follows a particular domain model (Koper, et al. 2006), that defines the concepts used and the overall architecture. A Learning Network is specific for a certain domain of knowledge (eg, an occupation) and consists of three entities:

a) Users (lifelong learners): people with the intent to learn and the willingness to share their knowledge in the specified domain.

b) Activity Nodes ie, a collection of learning activities that are created and shared in order to exchange knowledge and experience or to develop competences in the domain.

c) A set of defined learning outcomes, or 'goals' (eg, competence levels).

Learning and teaching in a Learning Network may have some unfortunate side-effects:

- Users are unlikely to arrive in groups, nor will they share their objectives or background. Missing the social structure of a class, students easily become socially isolated, ‘lone’ learners (Kester et al., 2006).
- The heterogeneity of the users and the lack of a readily available social structure that provides mutual support, makes a large demand on the tutors (Bartolic-Zlomislic & Bates, 1999; Bacsich & Ash, 2000; Koper, 2004). Tutors in an online learning context (Anderson, 2004) are no longer restricted to well-defined and pre-planned tasks but have to adopt to user needs on the fly. The tutor has to make provisions for the negotiation of activities to meet users’ unique learning needs, and equally well has to stimulate, guide and support the learning in a way that responds to common and unique user needs.

Moreover and of particular relevance in the context of the present paper, there is the additional challenge that Learning Networks are not meant merely to serve formal learning but also should cater for informal learning. For informal learning, there may not be any staff at all. However, also informal learners will have questions on where to start, how to proceed, how to understand and apply the available Activity Nodes, or will want to have their contributions assessed. As a consequence, there is a need to organize *and* support both formal and informal learning.

In this paper, we will concentrate on one element of this challenge, to wit answering questions related to the content studied. For a tutor this is considered a time consuming and disruptive task (De Vries et al., 2005). Yet, learning may improve if learners can ask questions and receive timely and relevant feedback (Howell, 2003). A number of models exist that address this particular problem. *Expertfinder* (Vivacqua & Lieberman, 2000) is an agent that classifies novice and expert knowledge by analysing documents created while working in the domain of Java programming. The model tries to distribute the question load evenly over several experts. It also prioritises not the best expert available but someone whose knowledge level is close to the questioner's level. This way, it is more likely to bring together people who share a similar mental model of the problem discussed. Interestingly, in 50% of the cases in which an answer was supplied, the expert did give an answer not directly but by looking it up. *Yenta* (Foner, 1997) a multi-agent, matchmaker system has been designed to find people with similar interests and introduce them to each other. The similarity of interest is based on the assumption that two users have similar interest if both possess similar documents (emails, newsgroup-articles, files). *FAQO* (Caron, 2000) relies on the use of *Latent Semantic Analysis* (LSA) (Landauer et al., 1998; Van Bruggen et al., 2004), a technology with a relatively wide-spread use in educational settings (Haley et al., 2005). LSA connects related words in a number of steps (eg, in documents in Computer Science the words human, computer and interface are related). In this way, although the actual keywords in the documents might differ, if there is sufficient similarity, documents are associated. FAQO allows the users to query questions in natural language in order to find relevant documents to solve their problems for specific technical problems.

In our model we combine a number of the characteristics of the above mentioned models. Crucially, we seek to solve content related questions by involving peers in answering them (peer tutoring). To that end, we identify appropriate and available users as well as documents and bring these together in a so called *ad hoc, transient* community. Such a community is ad hoc in that its only purpose is to solve a particular question; it is transient in that it vanishes the moment the question has been solved. In our view, ad hoc, transient communities are particularly well suited to assist peer tutoring (for a detailed discussion on the underlying theoretical aspects of our model on learning in communities and peer tutoring see: Kester et al., 2006; Sloep et al., in press). Obviously, one will have to heed the lessons learned on community building and peer tutoring.

First, for a social space to emerge, one should establish continuity of contact, recognisability of members, and a historical record of actions (Kollock, 1998). Furthermore, to assure the liveliness of a community, it should be populated with a heterogeneous group consisting of veterans and newbies; connectors, mavens, and salesmen; lurkers and posters (Preece, Nonneke & Andrews, 2004). Also, to facilitate co-operation in a community, clear boundaries and a clear set of rules that can be monitored and sanctioned within the community are required (Kollock & Schmidt, 1996). With respect to peer tutoring, we found out, among other things, that peer tutoring enhances the social embedding of students in a learning environment that facilitates social processes as engagement, commitment and a sense of belonging, and that peer tutoring does indeed help tutors and tutees to achieve higher learning outcomes (Fantuzzo et al., 1989).

At this point in time, we do not test any of these community formation conditions, but provisionally assume that we can sufficiently support the community with the help of e-portfolios, the expected heterogeneity of a Learning Network and by setting clear guidelines for the tasks supported. Similarly, although we will have to validate in future experiments that the expected benefits for learners and tutors will materialise, we provisionally assume them to be present. In the remainder of this paper, we concentrate on the main assumptions underlying our model, ie, that we can indeed identify appropriate and available peers and documents.

We will now explain our model by depicting the context of our model, a Learning Network, and describing its current implementation. In the sections to follow we will discuss the calibration of our model and the results of a simulation. The simulation will show how well we can map a set of pre-designed users' questions onto the Activity Nodes in a selected Learning Network. With this information we can identify capable peers and relevant textual resources in the Network.

Model implementation

A Learning Network

In order clearly to describe the context of our model implementation, we introduce a Learning Network example (Figure 1). Suppose we have a Learning Network in domain D, eg psychology, with a set of Activity Nodes A_1-A_{10}. Moreover, we have a Learning Network user P (Anne) who has formulated a goal that can be achieved by studying A_1, A_2, A_3, A_6, A_7, A_9, and A_{10}. Next, we know that Anne given her working experience and prior studies has exemptions for A_5 and A_6, and has already successfully finished A_7. Finally, let's assume that Anne while studying A_1 runs into problems. She has a problem understanding the relations between a number of concepts and as a consequence she is not able to complete an assignment. She studies some additional literature and searches the web, to no avail though. Anne studying on her own and thus out of touch with any peers, decides to pose a question to the 'on-line tutor'; she describes the general problem and her question.

Figure 1: A Learning Network for domain D

The 'on-line tutor' in our model consists of an ad-hoc, transient community populated with peer users who have, complementary content expertise. The goal of this community is to share knowledge and jointly come to an answer to the question in point. The central aim of our model (Table 1) is to set up and support the ad hoc, transient community that will help to answer the question within an agreed timeframe (eg two days) and to a mutually agreed quality (ie, the peer users decide together).

Table 1: The main steps of the model.

The prototype of the model (Figure 2) consists of five modules. For the users we have a Learning Network, its Activity Nodes and a question-interface. They are implemented in an instantiation of the Moodle environment (www.moodle.org). Additionally, each time a question is posed, a wiki is made available that includes the question and three documents selected from the Learning Network material. The wiki is populated with a selection of users who are invited to help. For the designer and for the runtime system we have three modules: a General Text Parser (GTP; Giles, Wo & Berry, 2001), a GTP calibrator (GTP Usability Prototype (GUP); De Jong et al., 2006) and a tutor locator (ASA Tutor Locator (ATL); Brouwers et al., 2006). We use GTP, an LSA implementation, to map the questions on the documents in the Learning Network. The GTP module returns correlations between the question and documents. The correlations are used to determine the Activity Node to which a question fits best and to select relevant text documents. The application of LSA, however, is not straight-forward. It depends on the corpus (the documents in the Learning Network) and its application. To assure optimal use, one has to calibrate a set of parameters. The GUP module has been built to ease the calibration. Finally, the ATL module takes care of the selection of the peer users who will assist. The selection is based on a weighted sum of four criteria that are derived from the users' background and performance. The designer can adjust the weightings.

Figure 2: The main modules of the model

The model covers three phases. In the first phase the working context is defined. The model is connected to the Learning Network. All text of this Learning Network is captured and put into a corpus for further analysis. This includes the calibration of a suitable set of parameters for LSA. The next phase starts when a user poses a question (Figure 3). First, the Activity Node(s) is (are) identified to which the question fits best. This is done by mapping the question with LSA on the documents of the corpus and to look for the three documents with the highest correlations. Later, the same three documents are given to the ad hoc, community to help the users get a quick overview of relevant documents in relation to the question. We chose three documents because it should be sufficient to distinguish and not too much to be read by the supporting peers. However, this number may be altered if experience suggests so. Next, knowing to which Activity Node the question fits best, the ATL module can identify peers who are competent in the pertinent Activity Node(s). ATL selects three to five users who, according to four different criteria, are best equipped to answer the question (Kester, 2006). The suitability ranking is a weighted sum of *tutor competency*, *content competency, availability* and *eligibility*:

- The *tutor competency* is the ability of a user to act as a tutor. The tutor competency is derived from a combination of data logging, i.e. from the frequency and size of the contributions, and ratings on answers given previously.
- The *content competency* indicates if a user has successfully completed the Activity Nodes related to the question.
- *Availability* is based on the actual availability as derived from the personal calendar of the users and on their past workload. Someone who has recently answered none or only a few questions should be preferred over someone who has answered many.
- Finally, *eligibility* measures the similarity of the users. It can be used to favour the selection of users with an almost identical competence level.

With all information available, ATL now attempts to form an ad hoc community. It creates a wiki and invites the selected users. The invitation includes the question, some guidelines and a small set of documents that have been identified as being relevant to drafting an answer.
Finally, in the last phase the users jointly formulate an answer to the question. After some time, the peer tutoring process ends and a response becomes available. Ideally, the process ends because the question asking user (tutee) is satisfied with the answer. However, if this is not the case, it may also end because a predefined period of time has elapsed or because the participants agree to end it. Whatever the reason, the tutee should rate the performance of the peer-tutors involved. If necessary, these data are used to alert the institution-bound tutor that there is an unresolved question or (in combination with other logging data) that some users perform sub-optimally.

Method

Before actual experiments with the model, involving real people, can be carried out, one has to prepare the required data structures (the text corpus) and calibrate the model, ie, determine a default setting for the LSA-parameters and for the weights of the peer selection criteria. In this paper we will concentrate on the corpus preparation and the LSA parameters. The selection of proper weights is out of the scope of the present paper, it will be determined in a future experiment with students. We will do a partial simulation of the model to ensure that the model operates according to its design. For a set of pre-designed questions we will look how well we can map them to the Activity Nodes of the Learning Network. This is of key importance for the selection of peer-users. Moreover, we will have the designers of the Learning Network rate the text documents that are selected for the users.

The corpus of the Learning Network

Fortunately, at the start of the work described, we had a Learning Network at our disposal developed for a study on navigation (Janssen et al., in press). The domain of this Learning Network is ‘Internet Basics’, a collection of texts, links and tasks which aim to instigate a basic understanding of the Internet. It contains 11 Activity Nodes, each of which introduces a different aspect of the Internet, ranging from 'Web searching', 'Chatting' to 'Worms and Horses'. The Activity Nodes consist of an introduction, exercises, references to external web pages for further study, and an assessment. The Learning Network matches our two start requirements ie (1) an accessible text corpus, a combination of the Moodle learning environments and external web pages, and (2) the users’ progress could be tracked by the data available from the Activity Node assessments. The corpus was extracted manually. It contained the Moodle pages and external web pages; assessment questions were left out, however. These questions were used to calibrate the model. The Activity Node of an assessment question is obvious and thus could be compared to verify the Activity Node determined with the help of LSA. The language of the corpus is Dutch – references to documents in English were ignored - admitting, though a considerable English Internet vocabulary. The documents were saved as ‘text only’, a quick way to get rid of all non-textual elements. The documents were used as raw input, this means no further corrections were applied such as removing irrelevant documents, diacritical signs or misspellings. The final corpus was relatively small. It consisted of 327 documents in size ranging from 50 to 23,534 bytes (41 documents smaller than 250 bytes; 50 documents above 3,000 bytes). The corpus contained a total of 82986 words divided over 10601 terms, 4440 of which occur in at least 2 documents.

The calibration of the LSA-parameters
Having created the corpus, our first action was to calibrate the LSA-parameters. A calibration is primarily focused on finding an optimal combination of parameters connected to a model. However, in our case it is equally important to find a way to define the parameters with a predefined, limited number of steps which can be easily repeated and automated at a later stage. In this way we ensure we can apply our model in real practice. An overview of applications with LSA (Haley et al., 2005) reveals that there is no straight-forward procedure to determine the LSA-parameters. The parameters are influenced by the corpus and the way LSA is applied. We selected the five steps (Giles, Wo & Berry, 2001; Wild et al., 2005) that should be the most important: the definition of a correlation measure and method, corpus pre-processing, normalisation, weighting and dimensionality. We will, however, not do an exhaustive test with different combinations of parameters. Instead, we started with an initial combination of parameters based on results reported (Van Bruggen et al., 2006; Wild et al., 2005) and in each step we tested one parameter in a limited number of test runs. Each time we continued to the next step we only used the best result(s) from the previous step.

Correlation measure and method
For our correlation measure we used cosine similarity. Our method directly follows from our model. We use LSA in two closely related ways. First, we use LSA to identify to which Activity Node(s) the question posed fits best. This information is used to identify peers that are competent in the pertinent topic. Second, we want to select the three documents (text fragments really) in the corpus most suited to assist the peers in answering the question. We combine the two by selecting the three best correlating documents and by assigning one point to each Activity Node a document originates from. This results in a maximum of three Activity Nodes the question relates to. We use the result of the mapping on the Activity Nodes to select the parameter combination with which to continue. In our case the questions, 16 in total, were chosen from the original assessment questions of the Learning Network. Therefore - in principle - each question should map to one known Activity Node.

Pre-processing the corpus can consist of stopping (removing 'meaningless' words) and stemming (reducing terms to their semantic stem). Since we did not have access to a stemming application for Dutch, we only considered stopping. Moreover, given the size of our corpus, we decided to follow a recommendation by Van Bruggen (Van Bruggen et al., submitted) to create our own stop lists based on the term frequency in the corpus. The stop list consisted of the terms that covered 33% (22 terms) respectively 50% (91 terms) of the overall term frequencies with the exception of terms that were judged corpus specific. By way of comparison, we also used a 'general' Dutch stop list (Oracle Text Reference: Release 9.2). For our corpus this resulted in a reduction of 188 terms. Finally, in each run (until the actual dimensionality step) we chose to limit the number of singular values (ie the number of dimensions) to 40% of the sum of the singular values (Wild et al., 2005). Next, as reported above, our corpus showed quite a spread in document lengths, while, at the same time the number of documents per Activity Node proved limited. Therefore it was decided to use *normalisation*. It makes the norm of each document vector equal to one. This has the effect that documents with the same semantic content are ranked equal in the question query. Next, we applied the three available types of *Global Weighting* and finally, in the last step we determined the best value for the *dimensionality* by comparing the initial value of 40% of the sum of the singular values to 30% and 50%.

A simulation of the model
After having studied the Learning Network and with a view to the simulation, we formulated a new set of 16 questions, each connected to one Activity Node. The questions were once again mapped on the Activity Nodes and the results compared with their known Activity Nodes. Please note, this time only the parameter combination that performed best in the calibration was applied. Next, we asked two of

the designers of the Learning Network to rate, on a five point scale (Figure 3), the suitability of the text fragments selected through the application of LSA. Obviously, a question may go beyond the content discussed in the Activity Nodes. In such cases the text documents identified by LSA have little bearing on the question, they can only serve to start off a discussion. Therefore, we instructed the designers to assess the suitability of the documents identified *relative to* the available text. This means that also a document that only starts off the discussion of a question, should be rated high in case there is no better alternative available. In addition, in case of a low rating we asked the designers to indicate a better alternative from within the corpus.

Figure 3: An example of a question and the way to assess the proposed text

Results

The first part of our study aimed to determine the LSA-parameters in a fixed, limited number of steps and a limited number of test runs. We achieved the following results. First, (Figure 4) we compared three stopping approaches: 33%, 50% and a general Dutch stop list (runs 1, 2 and 3). We were able to identify correctly the Activity Nodes of 5, 11, 11 questions, respectively. Second, as a result of this, we continued with normalisation for the 50% and the Dutch stop list (runs 4 and 5). The number of correctly recognised Activity Nodes remained 11. However, the questions with a single match increased, in particular in run 5 (Dutch stop list). We kept normalisation, continued with the Dutch stop list and compared global weights 'inverse document frequency', 'logarithm' and 'entropy' (runs 6, 7 and 8). This time the results improved to 12, 14, 15. For the last step, the dimensionality, we continued with the setting of runs 8 to run 9 (30% singular values) and run 10 (50% singular values). The overall results remained the same. The number of 100% recognitions increased by one. Finally, we carried out one additional run, which we had not planned beforehand; we used the 50% stop list in order to check if this would improve our results. The other parameters followed the settings of run 9. The result was good (15 out of 16) but not an improvement.

Overall, the results are encouraging. First of all, - at least for this corpus - it seems possible to determine such a combination of parameters that an important requirement of our model can be fulfilled: the mapping of a question to the appropriate Activity Node and, on the basis of this information, the ability to select appropriate peers. Second, the results suggest that the approach taken, to calibrate the parameters in a fixed setting with a limited number of test runs, is sound. Nevertheless, one should be open to retrace one's steps, in particular if the results are very close (as in our normalisation step) and improvements develop insufficiently. Since runs 9 and 10 had identical results, both were kept for the simulation.

Figure 4: The mapping of the questions on the Activity Nodes: (left) the assessment questions in the calibration runs; (right) the final questions

Having completed the calibration, we devoted the second part of our study simulating part of the model. We created 16 questions that we felt students may well have asked, mapped them on the Learning Network and invited two of the designers of the test Learning Network to rate the suitability of the proposed text fragments with respect to the questions. First, the model identified the correct Activity Node for 12 out of the 16 questions (Figure 4). Case one (the settings of run 9) did slightly better in the 100% recognition category. For this case (Figure 5), subsequently, the designers rated the supplied text fragments. Of the 16 questions:

- 6 (38%) and 4 (25%) respectively had at least 1 relevant text fragment (rating 4 or 5),
- 1 (6%) and 2 (13%) respectively had a text fragment that was of some use, and
- 9 (56%) and 10 (62%) respectively had no suitable text fragments connected to them.

Figure 5: The rating by designer 1 and 2 of the suggested text fragments

The results of the mapping are worse than in the calibration, but still quite accurate with a recognition of 75%. The suitability of the text fragments, approximately 40% of the questions receives one or more fragments rated 3 or above, looks far less accurate. However, they do answer our expectations very well, given the conditions we work with:

- We chose to only forward the first three fragments in order not to overload the users. Obviously, we thus run the risk that relevant fragments are left out. FAQO (cf. Introduction) for instance returns a top ten and, indeed, answers 4-10 do give a relevant contribution.
- The corpus is relatively small, this lowers the likelihood to find a relevant text for each question. Designer 2 confirmed that for 6 out of 10 questions (with a text rating of only 1 or 2) he could not identify a better alternative. In a real implementation, one can stepwise improve the likelihood of finding a relevant text by adding the answers of solved questions to the corpus.
- Finally, as with Expertfinder and Yenta (cf. Introduction) our intention is not so much to identify *the* answer. Our focus is on questions that are not readily answered by simply looking up the Learning Network-contents. But we do want to give the ad hoc communities a solid starting point to the extent that the corpus makes that feasible.

Conclusion

In this paper we described a model that intends to alleviate the support task of tutors. The model does so by invoking the knowledge and skills of fellow-students, who jointly form an ad hoc, transient community. We described how we calibrated LSA for an existing Learning Network. Subsequently and for the same Learning Network, we checked with a simulation whether the model is fit for experimentation with students. In our opinion, the results are promising. For 75% of the questions, we were able to identify to which Activity Node they belonged; for approximately 40% of the questions we could suggest one or more text fragments that could be useful when formulating an answer. Moreover, we were able to arrive at our results in a systematic way. The same steps can be followed for a new corpus or if the changes to an exiting corpus are relatively small, the known settings can be re-applied in just one additional run. Important characteristics of the procedure followed are that (1) it is relatively straightforward, there are no experts needed to apply it; and (2) it can be automated to a very large extent. Furthermore, the requirements to use the model are limited. They are restricted to having an accessible text corpus and accessible learner progress information. In a final system the first requirement, for instance, can be realised by adopting the widely accepted IMS-CP standard (IMS-CP 2003).

Obviously, there are a number of issues to be considered. First, the model has only been applied to questions that exactly match one Activity Node. It is fair to expect that, in real practice, some of the questions will cover not just one but more Activity Nodes. This may complicate the recognition and thus dilute the results. Next, as shown by some of the results, the approach is sensitive to the size (and content) of the available corpus. We do not know (yet) what the minimum size of a corpus should be. We also still have to determine a working combination of weights of the suitability ranking (*tutor competency*, *content competency availability* and *eligibility*). These issues, however, do not lend themselves to simulation and should be addressed in empirical tests.

The results indicate that the model is ready for use in experiments with students. A first experiment is planned for the second part of 2006. Here we will also investigate and optimise the community formation conditions discussed in the introduction. Ultimately, the experiment is meant to investigate our main hypothesis, to show that the task of staff in answering questions can be significantly alleviated by following our peer-tutoring model.

Acknowledgements

The authors thank the designers of the Learning Network 'Internet Basics', in particular José Janssen and Colin Tattersall, for their support; and Jan van Bruggen for his indispensable comments and advice on the use of LSA. The authors' efforts were funded by the OTEC Technology Development Programme in the project ASA (www.learningnetworks.org) and by the European Commission in TENCompetence (IST-2004-02787) (http://www.tencompetence.org).

References

Anderson T. (2004). Teaching in an Online Learning Context. In T. Anderson & F. Elloumi (Eds.) *Theory and Practice of Online Learning*, Athabasca University, 271-294. Retrieved online September 1, 2004 at: www.cde.athabascau.ca/online_book

Bacsich, P. & Ash, C. (2000). Costing the Lifecycle of Networked Learning: Documenting the costs from conception to evaluation, *Association for Learning Technology Journal (ALT-J)*. 8, 92-102.

Bartolic-Zlomislic, S. & Bates, A. W. (1999). Assessing the costs and benefits of telelearning: a case study from the University of British Columbia**.** *Network of Centers of Excellence (NCE)-Telelearning Project report*. Retrieved online November 8, 2005 at: http://research.cstudies.ubc.ca/nce/index.html

Brouwers, M., Brouns, F., Van Rosmalen, P., Sloep, P. B., Kester, L. & Koper, R. (2006). ASA Tutor Locator. URL: http://sourceforge.net/projects/asa-atl

Caron, J. (2000). *Applying LSA to Online Customer Support: A Trial Study*. Retrieved online April 2004 at: http://my.unidata.ucar.edu/content/staff/caron/faqo/faqoPaper1.pdf

De Jong, A., Brouns, F., Van Rosmalen, P., Sloep, P., Kester, L., & Koper, R. (2006). GUP GTP Usability Prototype. URL: http://sourceforge.net/projects/gup

De Vries, F., Sloep, P., Kester, L., Van Rosmalen, P., Brouns, F., De Croock, M., Pannekeet, C. & Koper, R. (2005). Identification of critical time-consuming student support activities that can be alleviated by technologies. *Research in Learning Technology (ALT-J).* 13, 219-229.

Fantuzzo, J. W., Riggio, R. E., Connelly, S., & Dimeff, L. A. (1989). Effects of reciprocal peer tutoring on academic achievement and psychological adjustment: A component analysis. *Journal of Educational Psycholog.y*. 81, 173-177.

Foner, L.N. (1997). Yenta: a multi-agent, referral-based matchmaking system. In *Proceedings of the first international conference on Autonomous agents***,** Marina del Rey, California, United States. Retrieved online October 13, 2004 at: http://doi.acm.org/10.1145/267658.267732

Giles, J. T., Wo, L., & Berry, M. W. (2001). *GTP (General Text Parser) Software for Text Mining*. Retrieved online January, 2005 at: http://www.cs.utk.edu/~berry/papers02/GTPchap.pdf

Haley, D. T., Thomas, P., De Roeck , A., & Petre, M. (2005). *A research taxonomy for Latent Semantic Analysis-based educational applications.* Technical Report No 2005/09. Retrieved online April, 2006, at: http://computing-reports.open.ac.uk/index.php/content/download/187/1136/file/TR2005_09.pdf

Howell, K. (2003). *Question Generation and Answering Systems R&D for Technology-Enabled Learning Systems Research Roadmap*. Retrieved online November 2004 at: http://www.thelearningfederation.org/qa.pdf

IMS-CP (2004). IMS Content Packaging v1.1.4 IMS Global Learning Consortium Inc. Retrieved online April, 2006 http://www.imsglobal.org/content/packaging/index.html

Janssen, J., Tattersall, C., Waterink, W.,Van den Berg, B., Van Es, R., Bolman, C., & Koper, R. (in press). Self-organising navigational support in lifelong learning: how predecessors can lead the way. *Computers & Education*.

Kester, L., Sloep, P., Brouns, F., Van Rosmalen, P., De Vries, F., De Croock, M. & Koper R. (2006). Enhancing Social Interaction and Spreading Tutor Responsibilities in Bottom-Up Organized Learning Networks. *Proceedings of the International Conference Web Based Communities* 2006, San Sebastian, Spain.

Kollock, P. (1998). Design principles for online communities. *PC Update.* 15, 58-60.

Kollock, P. & Smith, M. (1996). Managing the virtual commons: Cooperation and conflict in computer communities. *Computer-mediated communication: Linguistic, social, and cross-cultural perspectives* (pp. 109-128).(Ed. Susan Herings). Amsterdam: John Benjamins.

Koper, E. J. R., Giesbers, B., Van Rosmalen, P., Sloep, P., Van Bruggen, J., Tattersall, C., Vogten, H. & Brouns, F. (2005). A design model for lifelong learning networks. *Interactive Learning Environments.* 1-2, 71-92.

Koper, E. J. R. (2004). Use of the Semantic Web to Solve Some Basic Problems in Education: Increase Flexible, Distributed Lifelong Learning, Decrease Teacher's Workload. *Journal of Interactive Media in Education*, 6. Special Issue on the Educational Semantic Web. ISSN:1365-893X. Retrieved online September 2004 at: www-jime.open.ac.uk/2004/6

Koper, R. (2006). The TENCompetence Domain Model. *TENCompetence Project Report*. Retrieved online July 10, 2006 at: http://dspace.ou.nl/handle/1820/649

Landauer, T, Foltz, P. W. & Laham, D. (1998). An introduction to latent semantic analysis. *Discourse Processes.* 25, 259-284.

Oracle Text Reference: Release 9.2. Part Number A96518-01 Dutch (nl) Default Stoplist. Retrieved online April 2006 at:
http://www.lc.leidenuniv.nl/awcourse/oracle/text.920/a96518/astopsup.htm#37481

Preece, J., Nonneke, B., & Andrews, D., 2004. The top five reasons for lurking: Improving community experiences for everyone. *Computers in Human Behavior.* 20, 201-223.

Sloep, P., Kester, L., Van Rosmalen, P., Brouns, F., Koné, M., & Koper, R. (in press) Facilitating Community Building in Learning Networks Through Peer-Tutoring in Ad Hoc Transient Communities. *International Journal of Web Based Communities*.

Van Bruggen, J., Sloep, P., Van Rosmalen, P., Brouns, F., Vogten, H., Koper, R., & Tattersall, C. (2004). Latent semantic analysis as a tool for learner positioning in learning networks for lifelong learning. *British Journal of Educational Technology*. 35 , 729-738.

Van Bruggen, J.M., Rusman, E., Giesbers, B, & Koper, R. (2006). Latent Semantic analysis of small-scale corpora for positioning in learning networks. Manuscript submitted for publication.

Vivacqua, A. & Lieberman, H. (2000). Agents to assist in finding help. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, The Hague, The Netherlands. Retrieved online October 13, 2004 at: http://doi.acm.org/10.1145/332040.332408

Wild, F., Stahl, C., Stermsek, G., & Neumann, G. (2005) Parameters Driving Effectiveness of Automated Essay Scoring with LSA. In *Proceedings of the 9th International Computer Assisted*

Assessment Conference (CAA), 485-494. Loughborough, UK, July, 2005. Retrieved online March 2006 at: http://nm.wu-wien.ac.at/research/publications/b497.pdf

Figure 1: A Learning Network for domain D

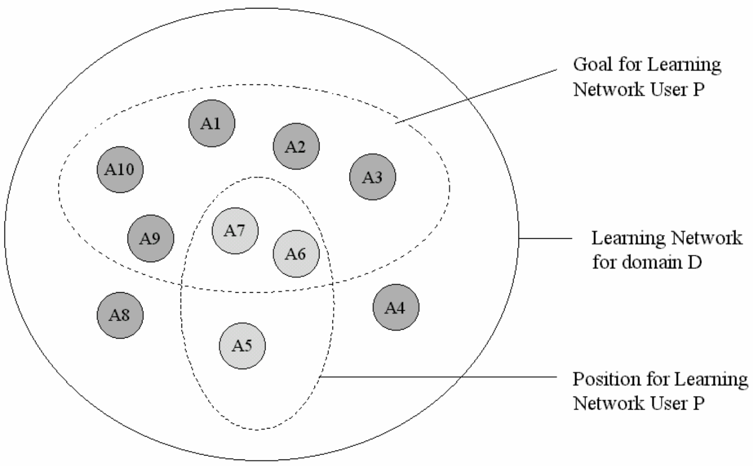

Figure 2: The main modules of the model

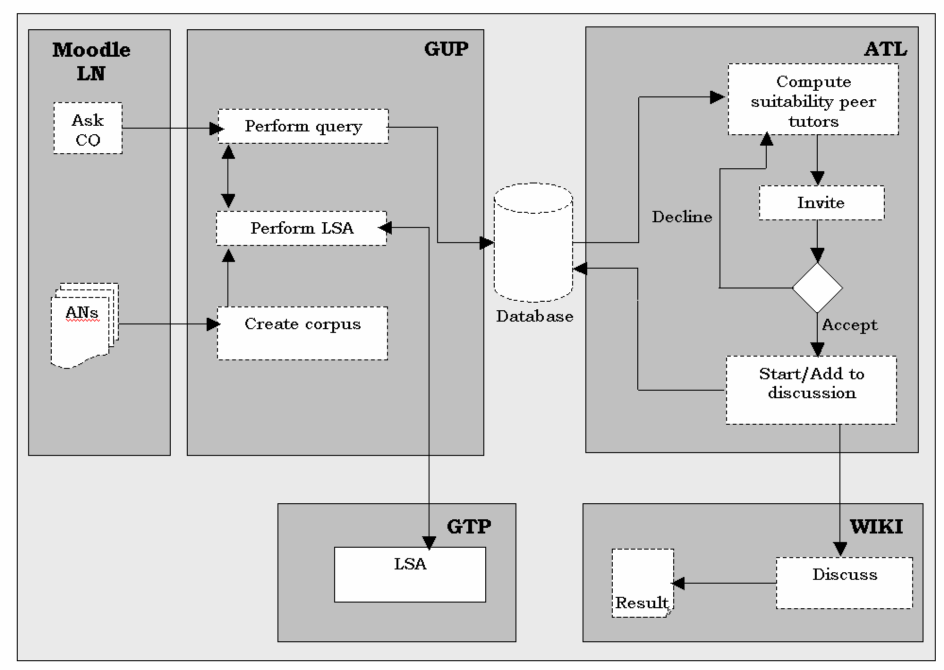

Figure 3: An example of a question and the way to assess the proposed text

Question 1
Wanneer ik me laat registreren om gebruik te kunnen maken van een chatroom kan ik dan met dezelfde registratie meerdere pseudoniemen gebruiken? [When I register for a particular chat room, does my registration allow me to use several pseudonyms?]

Proposed text fragments*	Activity Node						
237.txt	Net chatten en Netsletten (Net chat)						
Suitability	Not useful	1	2	3	4	5	useful
Your text fragment(s) in case of a rating of 1 or 2							
58.txt	Omgaan met ongewenste inhoud (Dealing with undesired content)						
Suitability	Not useful	1	2	3	4	5	useful
Your text fragment(s) in case of a rating of 1 or 2							
329.txt	Zoeken op het web (Searching on the web)						
Suitability	Not useful	1	2	3	4	5	useful
Your text fragment(s) in case of a rating of 1 or 2							

* Open the link to see the text

Figure 4: The mapping of the questions on the Activity Nodes: (left) the assessment questions in the calibration runs; (right) the final questions

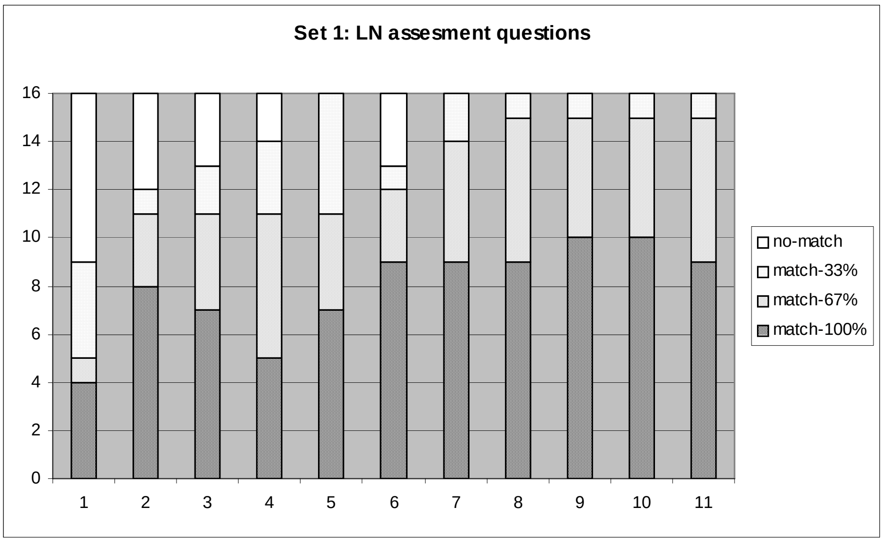

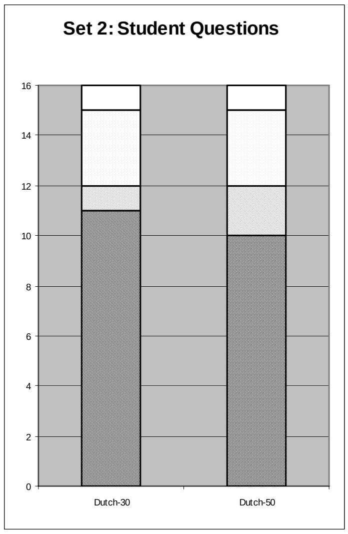

Figure 5: The rating by designer 1 and 2 of the suggested text fragments

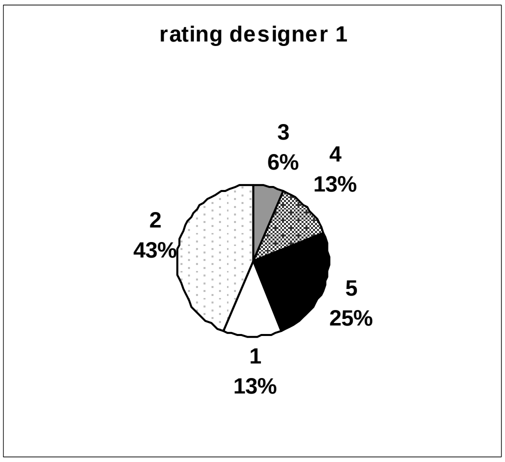

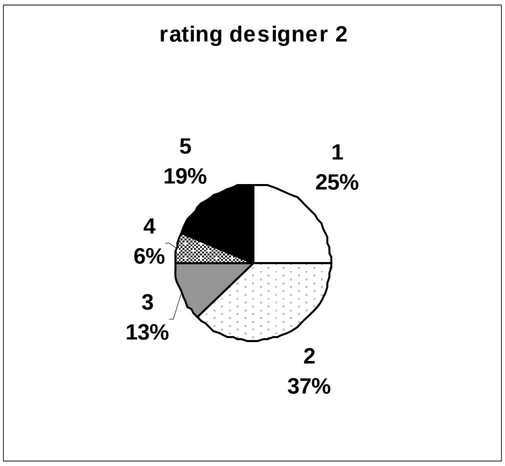

Table 1: The main steps of the model.

Pre-condition	A Learning Network with a set of Activity Nodes and a set of users with their profiles
Main steps	1) *Anne* poses a question. 2) The *system* determines: o the most relevant text fragments; o the concerning Activity Node(s); o the most suitable users. 3) The *system* sets up a wiki with the question, the text fragments and guidelines. 4) The selected *users* receive an invitation to assist. 5) *Anne and the users* discuss and phrase an answer in the wiki. 6) If answered (or after a given period of time) *Anne* closes the discussion and rates the answer.
Post-condition	The answer is stored.