QUESTION-ANSWERING THROUGH SELECTING AND CONNECTING PEER-STUDENTS

Peter Van Rosmalen, Francis Brouns, Peter Sloep, Liesbeth Kester, Adriana Berlanga, Marlies Bitter, Rob Koper (Open University of the Netherlands)

Abstract: Tutors have only limited time to support the learning process. In this paper, we introduce a model that helps answering the questions of students. The model invokes the knowledge and skills of peer-students by bringing them together based on the combination of questions posed and their study progress; it supports the process with text fragments selected from the material studied. We will explain the model and the use of Latent Semantic Analysis to select and support the peers. Finally, we will discuss the results of a calibration and simulation of the model and present the first results of an experiment with over a 100 students using a prototype of the model for a period of 8 weeks.

Keywords: Elearning, Lifelong Learning, Tutor workload, Peer support

Introduction

In modern learning settings, students typically spend a substantial amount of time learning online. 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). However, as in regular settings, students will have questions on where to start, how to proceed, how to understand and apply the available study material or will want to have their contributions assessed. In this paper, we will concentrate on one element of this challenge, to wit, answering questions related to the content studied.

We will do this in the context of Networks for lifelong learning ('Learning Networks'). 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, 2006), that defines the concepts used and the overall architecture. A Learning Network is specific for a certain domain of knowledge (e.g. an occupation) and consists of three entities:

- Students (lifelong learners): people with the intent to learn and the willingness to share their knowledge in the specified domain,
- A set of competences to be achieved,
- Activity Nodes i.e. a collection of learning activities that are created and shared in order to exchange knowledge and experience and to develop the competences defined.

Learning networks can be seen as communities for learning. In Learning Networks users can and should take on any role, including the one of tutor. For a tutor, answering questions is considered a time-consuming and disruptive task (De Vries et al., 2005). Yet, learning may improve if students can ask questions and receive timely and relevant feedback (Howell, 2003). The flexibility required by students nowadays combined with the new forms of learning required by the knowledge society asks for a new vision on education. Students taking on the role of tutor and supporting each other can be part of this new form.

In our model (Van Rosmalen et al., 2006; Kester et al., 2007) we seek to assist staff-tutors in solving content-related questions by involving peers in answering them (peer tutoring). To that end, we identify appropriate and available students as well as relevant documents (text fragments) from the material studied, 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. The model distinguishes (Table 1) six main steps of which the second step depends on Latent Semantic Analysis (LSA) and portfolio information regarding personal information, profiles, goal, and competences to be able to select appropriate students. The model relies on competence measures and data to be present in portfolios, such as proof of prior learning, evidence of competences obtained. In the following section we will introduce the current implementation, next we will discuss the results of a calibration and simulation of a prototype of the model and finally we will conclude with the first results of an experiment with approximately 100 students in a Learning Network, the topic of which is basic Internet skills.

Table 1: The main steps of the model

Pre-condition	A Learning Network with a set of Activity Nodes and a set of students with their profiles indicating their progress with regard to the Activity Nodes
Main steps	1. Anne poses a question. 2. The system determines: a) the most relevant text fragments, b) the appropriate Activity Nodes, c) the most. 3. The system sets up a wiki with the question, the text fragments and guidelines. 4. The selected students receive an invitation to assist. 5. Anne and the peer-students discuss and jointly formulate 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.

Model implementation

A prototype has been developed to test the model. As depicted in Figure 1, the system consists of five modules. Moodle (moodle.org) has been used to provide the Learning Network (in the figure indicated by LN) and the Activity Nodes (in the figure indicated by ANs) to the students. A question interface (AskCQ) has been added. Additionally, each time a question is posed, a wiki is made available that includes the question and the three most fitting text fragments selected from the Learning Network material. The wiki is populated with a selection of peer-students who are invited to help. In addition, in the background, we have three modules: a general text parser (GTP; Giles et al., 2001), a GTP calibrator (GTP Usability Prototype –GUP-; De Jong et al., 2006) and A Tutor Locator (ATL; De Jong et al., 2007). We use GTP, an LSA implementation, to map the questions onto the documents in the Learning Network. The GTP module returns correlations between the question and documents. The GUP module has been built to ease the calibration. Finally, the ATL module takes care of the selection of the peer-students who will assist. The selection is based on a weighted sum of four criteria that are derived from the students' background and performance (Van Rosmalen et al., in press). The designer can adjust the weights as required.

The criteria are:

- The *tutor competency* is the ability of a student to act as a tutor. The tutor competency is derived from among other things ratings on answers given previously,

- The *content competency* indicates if a student 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 students 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,
- *Eligibility* measures the similarity of the students. It can be used to favour the selection of students with similar competence levels.

The model relies heavily on data and information in the student's portfolio; in particular evidence on tutor and content competence; and the model is designed to take competence levels into account. Content competence on prior learning can be determined on evidence present in a user's portfolio; for new competences the grade obtained in the competence assessment could be taken into account together with other evidence in the portfolio. The current implementation of the model contained only a limited portfolio with no information on prior learning and provided no opportunity for complex competence assessment. This forced us to use a limited form of the tutor and content competence criteria: tutor competence is based on rating, and content competence is indicated as none, started and completed. Future implementation will consider competence levels and use more advanced methods to assess competence.

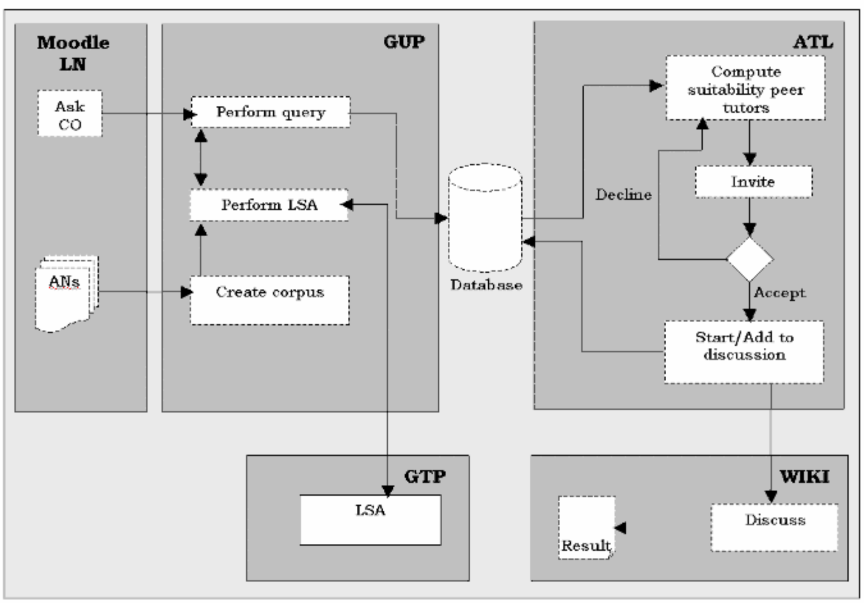

Figure 1. The main modules of the model.

The system basically consists of three phases. In the design-phase, the working context is defined. The text of all resources of the Learning Network is captured and put into a corpus for further analysis and; all parameters, the LSA and the peer selection parameters, are set. The question-phase (Figure 2) starts when a student poses a question (e.g. "when I register for a particular chat room, does my registration allow me to use several pseudonyms?"). First, those Activity Nodes which the question fits best are identified. 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 students get a quick overview of the documents most related to the question. We chose three documents because three should suffice to be of assistance to help answering the question and should not be too much to be read by the peers that will be selected to assist in answering. Next, knowing which Activity Node the question fits best, the ATL module can identify peer-students who are competent in the pertinent Activity Node(s). ATL selects 2 to 5 students who, according to a weighted sum of the four criteria, are best equipped to answer the question. Finally, in the answer-phase the students invited discuss and formulate an answer.

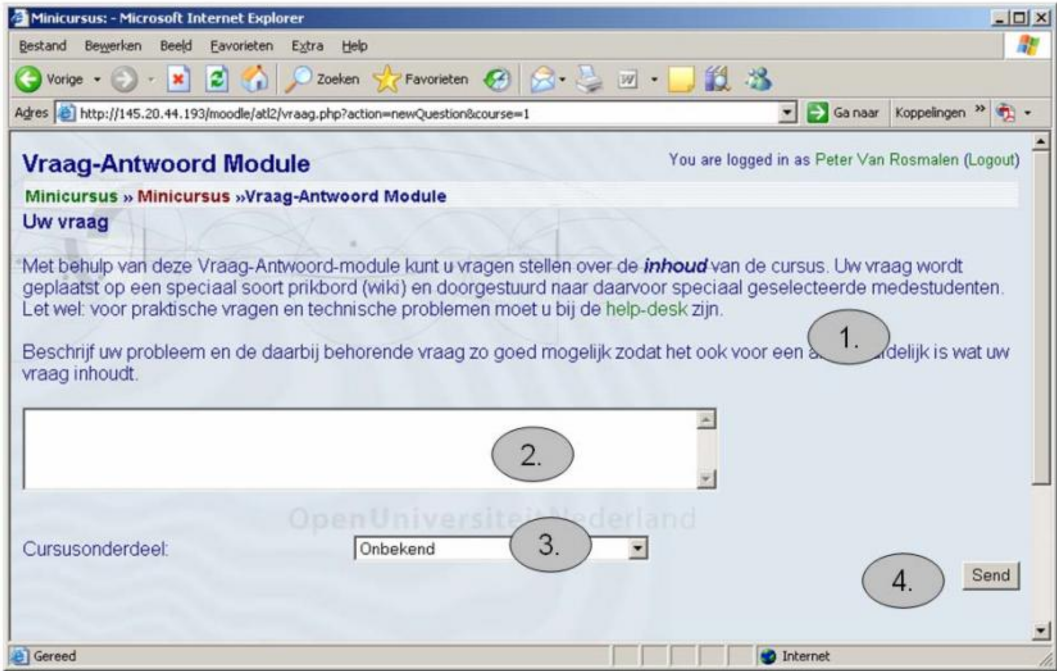

Figure 2. The interface to pose a question containing an instruction (1), a text field to enter the question (2), a pull down menu to select -if known- which Activity Node a question fits best (3), and a send button to submit the question (4).

Calibration and a first simulation

To ensure that our system is viable we calibrated the LSA-parameters, and simulated and tested two key aspects of the system. First, we checked how good LSA is at identifying the topic of a question (i.e. to which Activity Nodes a question belongs) and at selecting text fragments useful for answering the question. Second, we checked if the peer selection criteria met our expectations. The domain of the Learning Network we used is 'Internet Basics', a collection of texts, links and tasks that aim to instigate a basic understanding of the Internet (Janssen et al., 2007). It contains 11 Activity Nodes, each of which introduces a different aspect of the Internet. The Activity Nodes consist of an introduction, exercises, references to external web pages for further study, and an assessment.

For the simulation, we formulated a set of 16 test questions, each related to exactly one Activity Node. The system correctly identified the Activity Node for 12 out of the 16 questions (75%). Moreover, two developers of the Learning Network in question, evaluated the fit of the text fragments proposed by the system, three for each question. The developers indicated that for 7 of the 16 questions one or more of the text fragments were suitable for answering those questions. The developers also indicated that 5 of the 16 questions posed, were beyond the scope of the contents of the Activity Node studied. Taking this into account, the score is 7 questions with useful text fragments out of a total of 11 (about 60%, for details, see Van Rosmalen et al., 2006).

To test the peer selection criteria we created five student accounts (Table 2) and we assigned a set of test values to the parameters of the peer selection formula (for details see Van Rosmalen et al., in press). Next, student L_1 twice 'asked' one of the 16 questions mentioned above. The question was related to Activity Node 2. The first time the student asked the question, the peer-student with the highest rank was selected.

The results of the test showed however, that we could balance the selection of peers by taking availability and eligibility into account. For the first question the value of eligibility favoured student L_2 over student L_3, i.e., it prioritised the selection of a peer-student in the same study- phase. (Note: L_2 and L_3 have content competency 1 and availability 0.5. However, only L_2 and L_1 finished Activity Node 1, therefore the eligibility of L_2 with regard to L_1 is higher) However, if we pose the same question again the balance will be shifted due to the decreased availability of Student L_2. (Note: Because of the first question the availability of L_2 will become 0).

Table 2. Content competency of student L1 - L5 for Activity Node 1 and Activity Node 2, and their availability score.

	Content competency Activity Node 1	**Content competency Activity Node 2**	**Availability (at the start)**
L1	1 (= successfully completed)	0.3 (= in progress)	0.5 (=moderately available)
L2	1	1	0,5
L3	0,3	1	0,5
L4	0 (= not started)	0	0,5
L5	0	0	0,5

Experiment

The results discussed in the section above suggested that the system delivers as expected. Therefore, we set up an experiment to verify the hypothesis that the task of staff in answering questions can be facilitated and significantly alleviated by following the peer-student model proposed. The assumption to be validated is that it should be possible to solve at least 50% of students' questions without staff support.

For the experiment, we organised a course in the Learning Network on Internet Basics, lasting 8 weeks; 111 students registered. The students were divided at random over two groups. This was done to study the effect of different parameter settings of the student selection criteria. In group 1 we used a weighted selection of 1.0 x content competency, 0.5x availability and 0.5x eligibility. In group 2 we only made use of the availability criteria to select peer-tutors. Students received general instructions related to the Learning Network and a specific instruction on how to use the AskCQ-module for all their content-related questions. To avoid any unclear dependencies, it was decided that for this first experiment the students would not receive any incentives to use the AskCQ-module. We also decided that the staff-tutors would not assist during the study of the Learning Network with answering content- related questions; they would only rate the result of each question-answer pair.

We distinguished four types of data to collect. The first two, logging data and student ratings, were collected during the experiment; the last two, staff-tutor ratings and a questionnaire, will be collected after the experiment:

- Logging data. The progress data of the students and the data related to each question,
- Student ratings. For each question, the students that accept the invitation rate their own peer-tutor' suitability to answer the question; the student that posed the question rates the answer received,
- Staff-tutor ratings. At the end of the course, two staff-tutors will rate the answers of all finished questions-answer pairs,
- Finally, at the end of the course the students received a questionnaire focused on the usability

aspects of the AskCQ-module.

Results

At the moment of writing, the 8-week experiment just finished. Only the first results, based on the logging data and student ratings, are available:

- In total 111 students registered for the course, one of them withdrew officially after week 1, leaving a total of 110 students. Of these 110, 80 students were active, i.e. 30 students showed no or limited activity. This means they did not complete any of the 11 Activity Nodes,
- 101 questions were posed,
- 82 questions were resolved; 10 were still under discussion when the experiment was terminated, and 9 questions failed because the invited peer-students did not react or refused the invitation to contribute. Of the 9 questions that failed, 4 were posed on the very last day ofthe experiment,
- the average answer rating was 3.8 on a 5-point scale (see Figure 3 for the distribution).
- 48 students posed one or more questions,
- 65 students assisted in answering one or more questions,
- In total 69 students have been actively involved either posing or answering questions.

The rating of the answers

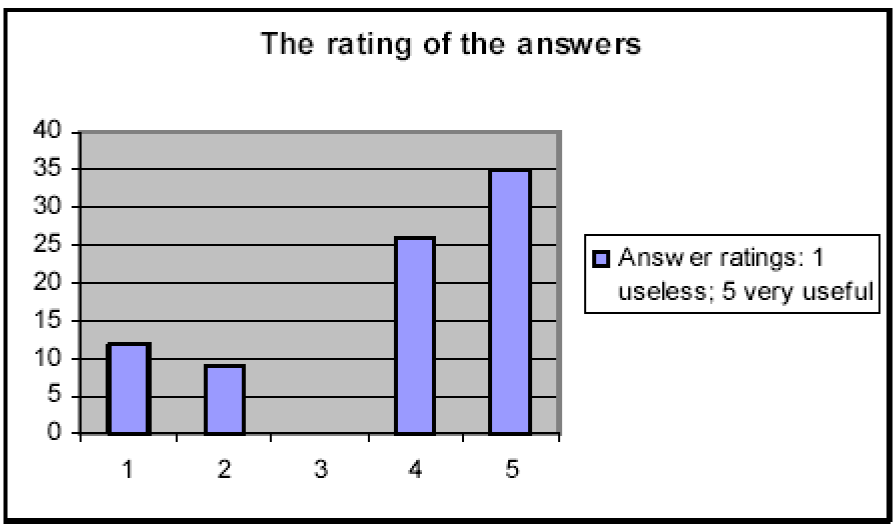

Figure 3. The rating of the answers by the students that posed a question.

Based on the available data and the first analysis so far, it is fair to conclude that the results are promising. The overall participation level with regard to AskCQ module is high, 69 out of 80 active students used the module to ask a question or to assist in answering in the role of peer-tutor. Moreover, according to the students which posed the questions at least 60% of the questions were resolved: the answer to 61 of the 101 questions was rated 4 or higher on a 5-point scale.

As soon as the staff-tutor ratings and the results of the questionnaires are available we will analyse whether they match the first results. The most important aspects we will look at are:

- whether or not the questions posed are also solved in a way that it is satisfactory from the perspective of a staff-tutor,

- whether or not there is a difference in the number of questions resolved successfully between groups 1 and group 2,
- what is the overall opinion of the students on the usability and the use of the AskCQ module.

Conclusion

In this paper, we described a model that intends to help students with questions that arise while studying. We described how we tested the model on two of its key aspects and presented the first results of an experiment with students. The test results indicate that we were able to identify the relevant Activity Node for a question, to select text fragments useful for answering the question, and to apply our peer selection formulas to the extent that it warrants carrying out an empirical study with 'real' students. The first results of the experiment suggest that the task of staff-tutors in answering questions can be facilitated, even with the limited portfolio data present in the experiment. It is likely that more advanced algorithms for tutor and content competence, based on competence assessment protocols and portfolio information, will result in a more effective implementation of the model. The results also indicate that the approach taken is a suitable model for learning in a network. Obviously, without a full set of data, i.e. including the staff-tutor ratings and the questionnaire results, and a detailed analysis of them it is too early to draw any final conclusions.

Acknowledgements

The authors' efforts were partly funded by the European Commission in TENCompetence (IST-2004-02787) (http://www.tencompetence.org).

References

1. De Jong, A.; Brouwers, M.; Brouns, F.; van Rosmalen, P.; Kester, L.; Sloep, P. B. (2007). Atl: A Tutor Locator. [Avaliavble from: http://hdl.handle.net/1820/960].
2. De Jong, A.; Brouns, F.; Van Rosmalen, P.; Sloep, P.; Kester, L.; Koper, R. (2006). Gup Gtp Usability Prototype. [Avaliavble from: http://sourceforge.net/projects/Gup
3. 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), Vol. 13 (3), pp. 219-229.
4. Howell, K. (2003). Question Generation and Answering Systems R&D for Technology-Enabled Learning Systems Research Roadmap. [Avaliavble from: http://www.thelearning federation.org/qa.pdf, accessed: 11.2004].
5. Janssen, J.; Tattersall, C.; Waterink, W.; van Den Berg B.; van Es, R.; Bolman, C.; Koper, R. (2007). Self-Organising Navigational Support in Lifelong Learning: How Predecessors Can Lead the Way. *Computers & Education*, Vol. 49 (3), pp. 781-793.
6. Giles, J. T.; Wo, L.; Berry, M. W. (2001). Gtp (General Text Parser) Software for Text Mining. [Avaliavble from: http://www.cs.utk.edu/~berry/papers02/gtpchap.pdf, accessed: 01.2005].
7. Kester, L.; Sloep, P. B.; van Rosmalen, P.; Brouns, F.; Koné, M.; Koper, R. (2007). Facilitating Community Building in Learning Networks Through Peer-Tutoring in ad hoc Transient Communities. International Journal on Web Based *Communities*, Vol. 3 (2), pp. 198-205.
8. 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*, Vol. 13 (1-2), pp. 71-92.
9. Koper, R. (2006). The Tencompetence Domain Model. Tencompetence Project Report.

[Avaliable from: http://hdl.handle.net/1820/649, 10.07.2006].

10. Van Rosmalen, P.; Sloep, P.; Brouns, F.; Kester, L.; Koné, M.; Koper, R. (2006). Knowledge Matchmaking in Learning Networks: Alleviating the Tutor Load by Mutually Connecting Learning Network Users. *British Journal Of Educational Technology*, Vol. 37 (6), pp. 881-895.
11. Van Rosmalen, P.; Sloep, P.; Kester, L.;Brouns, F.; De Croock, M.; Pannekeet, K.; Koper, R. (In Press). A Learner Support Model Based On Peer Tutor Selection. *Journal Of Computer Assisted Learning*. [Available online at: Doi: 10.1111/J.1365-2729.2007.00245.X.]

Authors

Peter van Rosmalen,

Dr ir. Francis Brouns,

Dr Peter B. Sloep,

Dr Liesbeth Kester,

Dr Adriana J. Berlanga,

Dr Marlies E. Bitter-Rijpkema,

Prof. dr. Rob Koper

Open Universiteit Nederland, Educational Technology Expertise Centre

P.O. Box 2960, 6401 DL Heerlen,

The Netherlands.

Peter.vanRosmalen@ou.nl,

francis.brouns@ou.nl,

peter.sloep@ou.nl,

liesbeth.kester@ou.nl,

adriana.berlanga@ou.nl,

marlies.bitter@ou.nl,

rob.koper@ou.nl