On the Prospects of Intelligent Collaborative E-learning Systems

Miikka Miettinen [a,1] Jaakko Kurhila [b] and Henry Tirri [c]
[a] *Helsinki Institute for Information Technology, Finland*
[b] *Department of Computer Science, University of Helsinki, Finland*
[c] *Nokia Research Center, Nokia Group, Finland*

Abstract. Collaborative learning is question-driven and open-ended by nature. Many of the techniques developed for intelligent tutoring are applicable only in more structured settings, but fortunately there are other interesting opportunities to explore. In this paper we introduce a system called OurWeb, and use it as an exemplar framework for demonstrating some of these opportunities. We claim that effective participation in distributed and self-organizing collaboration requires sufficient awareness of the resources and dynamics of the community. The feasibility of implementing certain features of this kind is evaluated based on data from two university level courses.

Keywords. collaborative learning environments, collaborative learning, collaborative annotation, awareness, social navigation, information retrieval

1. Introduction

The objective of intelligent e-learning systems, as it is typically conceived, is to provide highly structured lessons that are to a large extent under automated control. Within this framework, the intelligence of the system often appears in the form of adaptive sequencing or personalization of the course material, adaptive guidance for navigation, or interactive problem solving support. All of these methods work the best in well-structured domains, and rely heavily on a fixed collection of pre-made course material.

While the prevailing approach has arguably proved to be appropriate in several contexts, there are good reasons to extend the perspective to other essential ways of learning. On the one hand, the theoretical assumptions implicit in the instruction method have received substantiated criticism. Learning has been claimed to be primarily a matter of participation [1] or collaborative knowledge building [2] rather than direct assimilation of facts from an authoritative source. The critics have suspected that excessive guidance places the students in a passive role, hampers the development of metacognitive skills, and results in an instructional setting that is too simplified and restricted to facilitate real-world problem solving [3,4,5]. These claims may or may not be justified, but in any

[1]Correspondence to: Miikka Miettinen, Helsinki Institute for Information Technology, P.O.Box 9800 FIN-02015 TKK, Finland; Tel.: +358 9 451 8123; E-mail: miikka.miettinen@hiit.fi.

This work was supported in part by the Academy of Finland under the Prima and Prose projects.

case they highlight the fact that some important aspects of learning do not fit well in the present framework.

On the other hand, collaborative learning has become a fairly common way of organizing education, and attempts to develop better tools for its particular needs are motivated in their own right. However, the needs turn out to be quite different from the ones that intelligent e-learning systems typically try to address. The collaborative learning process is highly unstructured and open-ended, and the activities of an individual student must be considered in a broader context. As a result, the most interesting opportunities to develop intelligent functionality are related to facilitating collaboration rather than adapting the learning material.

The next section introduces a system called OurWeb, which demonstrates the principles of collaborative learning and provides a suitable exemplar framework for the rest of the paper. In section 3 we present some general ideas of advanced features that might support the collaborative learning process, and continue with a preliminary feasibility study in section 4. Section 5 concludes with some general reflections of the issues involved.

2. Collaborative Learning with OurWeb

Collaborative learning takes place within the framework of joint activities. Rather than trying to master a fixed set of topics determined by the instructor, the students are engaged in an open-ended effort to advance their collective understanding [4]. Division of work and specialization are seen as opportunities, and the students are encouraged to rely on each other as sources of information and assistance. Genuine participation taking place in a meaningful social context is claimed to make learning a matter of personal development and result in deep intrinsic motivation [1]. In addition, interactions among the students facilitate learning directly by encouraging them to explain the subject matter to each other and revealing in a constructive way the inconsistencies and limitations in their knowledge [6].

OurWeb is an integrated set of tools for collaborative learning. The most essential principles underlying its design are *openness* and *transparency*. By openness we mean that the students should be enabled to utilize any available information sources with as few restrictions as possible. Transparency is pursued by attempting to provide tools that fuse seamlessly into the activities of the students, allowing them to benefit from the work of each other and participate in meaningful ways.

The OurWeb server acts as a proxy between the user's browser and the Web, capable of augmenting any page with additional content and functionality. Most features are located in a custom *popup menu*, which is opened with the right mouse button. Some of the menu items are used for manipulating the visible page and others for navigating between various parts of the system. This kind of a minimalist user interface is natural and appropriate when providing unrestricted access to heterogeneous Web pages.

OurWeb provides a shared *document pool*, which serves as a repository for both external resources and the students' own work. Any potentially useful Web page can be linked to the repository with the popup menu. The user simply opens the menu with the right mouse button and chooses the option labeled "Add to document pool". As a result, the document becomes visible to everyone on the various index pages and the internal

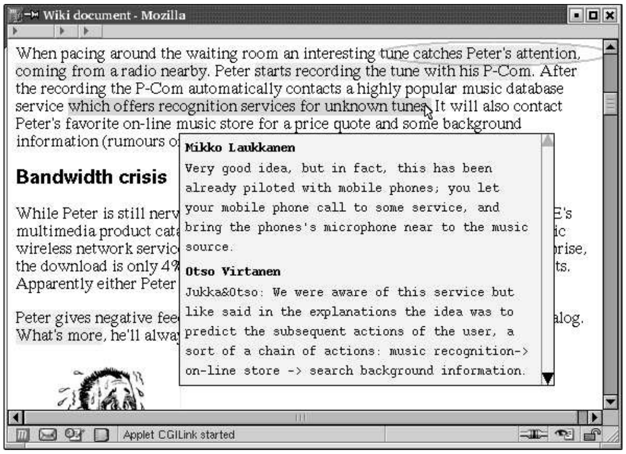

Figure 1. Comments in OurWeb.

search engine, and the full functionality of OurWeb (including e.g. annotations) can be applied to processing the contents effectively.

In collaborative learning, different groups of students are typically working on different topics, and the groups are organized by the students themselves instead of being assigned by the instructor. OurWeb supports the process by enabling the students to publish their ideas and suggestions as *projects*. The initial proposal consists of a title and a short description of the content, along with plans and schedules for organizing the effort in practice. Interested people can get involved by simply clicking a link labeled “Join team”. All ideas do not normally create sufficient interest, and the person who made the suggestion does not necessarily need to participate as an active team member. We want to avoid creating unnecessary barriers to collaboration, and encourage participation in all forms.

During the course of a project the team members are engaged in collaborative *process writing*. The idea is to produce a document incrementally, gathering feedback and ideas from the others along the way. In addition to supporting the work of each individual group, this enables cross-fertilization of ideas and fosters the sense of being part of a larger community.

OurWeb contains an integrated *Wiki*, which the groups use as a document editor. A Wiki (or WikiWikiWeb) is a tool for collaborative authoring of Web pages with an ordinary Web browser and a simple markup language [7]. At any point in time, the team has an internal “working copy” of the document being written. Intermediate versions can be published in the document pool with one mouse click, and are essentially snapshots of the continuously evolving document. The groups are encouraged to publish the first drafts already at the early stages of the work in order to get feedback and create opportunities for collaboration.

The primary means of collaboration are annotations and threaded discussions. Two different types of annotations are supported: *highlights* and *comments*. Highlights can be applied to marking important parts of the text, analogously to the way many people underline text on paper. In practice, adding a highlight involves selecting the text with the mouse, right-clicking the mouse to make the popup menu visible, and choosing the “Highlight” option from the menu. Comments are added the same way, except that the

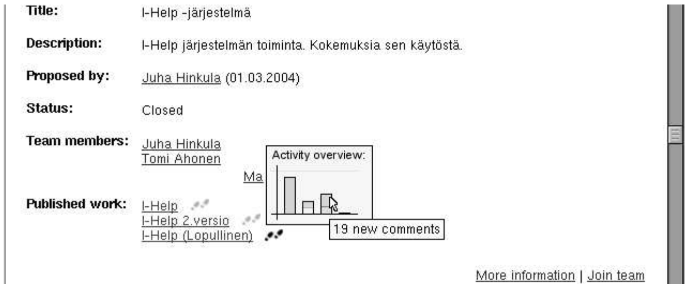

Figure 2. Footprint information of a project's published document version.

user types the input in a popup window. A comment appears as a tooltip when the mouse pointer is placed on top of the commented text fragment (see Figure 1). If several comments are attached to the same text, they appear one after the other as a dialogue. Longer reflections and remarks that may not be associated with any single passage of text can be posted in a *threaded discussion* located at the bottom of the page.

The annotation and discussion facilities of OurWeb allow the community to engage in *artifact-centered discourse* [8], in which the contributions appear in the immediate proximity of the relevant information. This has turned out to be very useful in practice. We have observed that especially comments are used extensively for short exchanges of feedback and ideas that would probably never have taken place in a detached discussion forum.

The number of documents in the document pool can grow large, and it is useful to provide several alternative views to the contents. These include e.g. lists organized by topic and the navigation history of the user, and a selection of documents that have recently received attention from the community. The system also contains an *internal search engine* covering the document pool as well as the annotations and threaded discussions. Google can be used through the OurWeb server for searching the entire Web.

Each link appearing on the index pages is followed by a *footprint icon*, which is either black or gray, depending on whether or not the document has received new activity since the user's previous visit (see Figure 2). When the user places the mouse pointer over the icon, a bar chart appears showing the relative amount of reading, highlighting, commenting, and threaded discussion activity associated with the document.

Other features of OurWeb include personalized desktop, automatic marking of new comments, and an interface for sending e-mail. The desktop serves as the entry point to the system, and contains recommended links to documents and discussion messages along with announcements from the instructor. Marking of new comments makes it easier to follow the gradual progress of asynchronous collaboration. The marks appear as ovals or lines around the commented text fragments (see the upper right corner of Figure 1). Finally, e-mail messages can be sent conveniently to an individual user or everyone in a particular project team by clicking links appearing in the project list.

3. Suggested Features for Intelligent Collaborative E-learning Systems

The "intelligent" functionality that is feasible and appropriate in the collaborative setting has to be quite different from a conventional intelligent e-learning system. The students are engaged in question-driven and open-ended inquiry, which would be very difficult to augment with automated guidance and problem solving support. In addition, it is not obvious that such facilities would be appropriate, even if they were feasible to implement. Identifying fruitful lines of inquiry and exchanging explanations in peer groups are essential elements of collaborative learning that should not be transferred away from the students.

Therefore, we propose a different approach. Rather than trying to guide the students directly, the system should support their activities with various kinds of supplementary information. Effective participation in distributed and self-organizing collaboration requires sufficient *awareness* of the resources and dynamics of the community. A suitable role for the system is to try to provide the right information at the right time, while the interpretation of the information and the associated decision making are best left to the user.

It seems plausible that several key activities involved in collaborative learning could be supported by better awareness. In this section we identify some relevant objectives and present general ideas of the additional functionality that would be needed for achieving them. The next section presents some data gathered from OurWeb in an attempt to assess in more detail the need for automated recommendation of collaboration opportunities.

3.1. Facilitating effective utilization of background material

At the age of the Internet, collaborative learning often happens at the edge of information overload. For almost any question the students might choose to examine, there is an endless supply of partially overlapping resources with additional details. The problem is not primarily technical by nature, but better tools could make it easier to locate relevant information and utilize the work of the others.

One obvious approach is to try to develop better facilities for information retrieval. In addition to the keyword search included in the current version of OurWeb, we have done some preliminary experimentation with *proactive search*. The idea is to observe the navigation and scrolling patterns of the user, and generate queries automatically to provide additional links to potentially relevant pages. Unlike the user, the search engine has a *global view* of the available contents and could (at least in principle) identify semantic relations between disparate sources of information. If successful, this would provide the user with improved awareness of the available contents, and reduce the cognitive demands involved in reading and constructing explicit queries at the same time.

Potentially relevant material can also be highlighted by presenting *recommended links*. In the absence of an explicit domain model, such recommendations are typically based on content-based or collaborative filtering. Both techniques rely on the notion of a *user profile*, which is assumed to be stable over long periods of time. In the present context this assumption is clearly invalid, because the usefulness of a document changes dynamically both as a result of learning and depending on the task that the student is working on at a particular moment.

Therefore, a better approach is to resist the temptation to give explicit recommendations, and focus on supporting the users' own judgment. For example, the kind of data underlying the "footprints" of OurWeb could be used as input for collaborative filtering, but presenting it directly to the users in summarized form is much more transparent and informative. Other examples of supporting cooperative processing of background material include OurWeb's shared document pool and annotations. Enabling the students to rely on the work of each other allows them to achieve a higher level of understanding than what would be possible if the same routines had to be repeated by each individual.

3.2. Making collaboration opportunities apparent

Informing the students about the activities of each other would also facilitate direct interactions. The shared workspace provides many *opportunities* for collaboration, and active encouragement from the system could make a significant difference in the engagement of the students. Ideally, the suggestions would be personalized and context-sensitive, adapted both to the needs of the individual and the overall status of the community. High precision would not be vital, however. Even if the suggestions were not especially pertinent, they might increase the amount of collaboration just by encouraging people to contact each other.

In order to form groups, the students need to be aware of the interests of each other. A suitable way of supporting such awareness would be to augment documents with information about people who have been actively utilizing them [9]. This would enable the students to identify potential collaboration partners when coming across interesting material.

When the groups are engaged in process writing, it is beneficial for their motivation and efficiency to get timely feedback. The system could encourage this by providing explicit notifications to potential reviewers. On the one hand, it would be appropriate to inform them whenever a new document version is published for review. Avoiding delays would ensure that the comments are valid and taken into account, as the document is often under continuous revision. On the other hand, the authors and the reviewers typically engage in asynchronous discussions, the status of which could be monitored and summarized automatically by the system. This would also help to eliminate delays by providing the users with better awareness of the progress of the discussions.

Real-time awareness of the presence of the others would facilitate *peripheral monitoring* of the workspace. When supplemented with synchronous communication tools such as chat and instant messaging, it would enable the students to engage in *spontaneous collaboration* motivated by momentary needs. This is claimed to be particularly useful in collaborative writing, which is characterized by frequent switches between independent work and focused group consideration of individual details [10].

3.3. Supporting coordination of group work

Effective group work also requires awareness of the activities of the other participants. Individual students need to coordinate the content and timing of their contributions with each other, and keep their efforts aligned with the overall objectives of the group. It is typical that the activities are reorganized repeatedly as new ideas and better understanding emerge [11]. Although continuous awareness can be maintained by means of explicit communication, utilizing data that accumulates automatically as a side product of the activities decreases the amount of routine communication and helps to eliminate unnecessary delays.

Different stages of the work call for different degrees of collaboration. Better awareness of the progress would enable flexible shifts between close and loose collaboration and make the interactions more fluid and natural [10]. Interestingly, it would also provide a basis for shared norms and conventions. The availability of relevant information would remove certain kinds of ignorance from the set of legitimate excuses, and foster stronger commitment the joint effort [9].

4. Feasibility Study

4.1. Study setting

Our empirical study assessed the need and feasibility of implementing automated recommendation of collaboration opportunities. We focused on three particular objectives:

1. *Supporting group formation by identifying students with shared interests*. As suggested in the previous section, a suitable way of supporting group formation might be to augment documents with information about people who have been actively utilizing them. The prerequisites for this would be the emergence of interest profiles from the activity patterns of the students, and sufficient overlap in the navigation of students with similar profiles.
2. *Increasing the timeliness of feedback*. The system could try to increase the fluidity of the review process by providing explicit notifications to potential reviewers. However, it would be useful to know specifically what kind of delays actually occur in the absence of this functionality.
3. *Providing real-time awareness of the presence of the others*. In order to cater for spontaneous collaboration, the system could inform the users about the presence and activities of each other. This is feasible only to the extent that there are several users logged in the system simultaneously.

The data was acquired from two university courses that employed the current version of OurWeb. During the first course titled “Computer Uses in Education”, 17 students were working on self-organized projects over a period of 10 weeks. The arrangements were extremely flexible, allowing the students to participate in projects of their own choice with roles and schedules negotiated among themselves. The second course involved doing a written and oral presentation on a free topic related to “Web Communities”. There were 16 students, and the work was done over a period of 7 weeks. The students of both courses were predominantly male and computer science majors.

4.2. Results

When a document is added to the shared document pool of OurWeb, it is assigned manually to one or more *topics*. As the first part of our analysis, we wanted to see if it would be possible to support group formation by identifying students with shared interests. We looked at the distribution of the students’ reading time with respect to the topics during the one week period preceding the formation of each group. Clear differences in the reading activity of the students were found. In 45% of the cases a single topic accounted for 50% or more of the student’s total reading time. There was also sufficient overlap in visits to individual documents. For example, for those with a clear interest profile on average 3 other students with the same profile had also visited a particular document associated with the dominant topic. Therefore, it seems that the suggested type of support for group formation could have been provided in practice.

There is also room for improvement in the timeliness of the feedback received by the project teams. On average only 42% of the feedback was received during the first two days after the publication of a draft, and 36% was received after 5 days or more. Turn taking in comment chains and discussion threads had an average delay of 38 hours.

Opportunities for synchronous interaction would have been limited. On average there were just 2.1 users online simultaneously, and the number went rarely above 5. Therefore, it seems that at least in small courses like ours the value of real-time awareness is questionable.

5. Conclusions

Collaborative learning is question-driven and open-ended by nature. Many of the techniques developed for intelligent tutoring are applicable only in more structured settings, but there are other interesting opportunities to explore. In this paper we suggested that trying to provide awareness of potentially relevant activities and resources is an appropriate direction for these explorations, and took some preliminary steps towards the implementation of such tools.

References

[1] Wenger, E. (1999). Communities of Practice: Learning, Meaning and Identity. Cambridge, UK: Cambridge University Press.

[2] Bereiter, C. (2002). Education and mind in the knowledge age. Mahwah, NJ: Lawrence Erlbaum Associates.

[3] Bredo, E. (1993). Reflections on the intelligence of ITSs: a response to Clancey's "Guidon-manage revisited". International Journal of Artificial Intelligence in Education, 4, 35-40.

[4] Scardamalia, M. & Bereiter, C. (1994). Computer support for knowledge-building communities. The Journal of the Learning Sciences, 3(3), 265-283.

[5] Deek, F.P. & McHugh, J. (1998). A Review and Analysis of Tools for Learning Programming. In Proceedings of the ED-MEDIA World Conference on Educational Multimedia, Hypermedia and Telecommunications, pages 251-256. AACE.

[6] Hatano, G. & Inagaki, K. (1993). Desituating cognition through the construction of conceptual knowledge. In Light, P. and Butterworth, G. (eds.), Context and Cognition: Ways of Learning and Knowing, 115-133. Lawrence Erlbaum Associates.

[7] Leuf, B. & Cunningham W. (2001). The Wiki Way: Quick Collaboration on the Web. Addison-Wesley Longmann.

[8] Suthers, D. & Xu, J. (2002). Kukakuka: An Online Environment for Artifact-Centered Discourse. In Proceedings of the Eleventh World Wide Web Conference (WWW 2002), 472-480.

[9] Erickson, T. & Kellogg, W.A. (2000). Social translucence: An approach to designing systems that support social processes. ACM Transactions on Computer-Human Interaction, 7(1), 59-83.

[10] Dourish, P. & Bellotti, V. (1992). Awareness and coordination in shared workspaces. In Proceedings of the 1992 ACM Conference on Computer-Supported Cooperative Work, pages 107-114. ACM Press.

[11] Gutwin, C. & Greenberg, S. (2002). A Descriptive Framework of Workspace Awareness for Real-Time Groupware. Computer Supported Cooperative Work, 11(3), 411-446.