Chapter 6
Learning to Learn and Work in Net-Based Teams: Supporting Emergent Collaboration with Visualization Tools

Peter Reimann and Judy Kay

Introduction

Collaboration in virtual teams and communities is becoming part of the K-12 curriculum ("twenty-first century skills" – see, e.g., iN2015 Education and Learning Sub-Committee, 2007; West Virginia Department of Education, 2007) and also increasingly well-established in university education (Resta & Laferriere, 2007). Not only do we see an increase in the frequency of employing group learning in educational settings, but also see a qualitative change taking place: the expectations as to the outcomes of collaborative learning have changed from a focus on improving individual learning (via motivational and cognitive processes) to also include gains with respect to shared knowledge (e.g., groups producing artifacts useful for others) and gains in social capital (e.g., students becoming integrated into social networks).

It is not that we have to force students to do things together. In their life outside of schools, today's young people spend an increasing amount of time using so-called social media such as Facebook, MySpace, or Second Life (in addition to using the equally social IM, SMS, phone, and e-mail services). create and share music, pictures, movies, homework, and experiences all the time.

This chapter is about creating and sharing knowledge and epistemic practices. With Scardamilia and Bereiter, we regard the knowledge challenge to be the central educational challenge of the twenty-first century: "how to develop citizens who not only possess up-to-date knowledge but are able to participate in the creation of new knowledge as a normal part of their lives" (Scardamalia & Bereiter, 2003). Due to the rapid adaption of information technologies in developing countries, tasks of

P. Reimann (✉)
Centre for Research on Computer-supported Learning and Cognition (CoCo),
University of Sydney, Sydney, NSW 2006, Australia
e-mail: peter.reimann@sydney.edu.au

J. Kay
School of Information Technologies, University of Sydney, Sydney, NSW 2006, Australia

many kinds can now be allocated to a globally distributed workforce, creating pressure in high-income countries to retain only knowledge- and service-intensive (nonroutine) economic activities. While information and communication technologies are for part "responsible" for the rapid growth of global competition for performing tasks (for pay) that require basic to medium skill levels, information and communication technologies are also part of the answer to the resulting knowledge challenge: "The same technologies that make innovative and creative thinking critical skills for the future also make it possible for students to prepare for that future" (Shaffer, 2008, p. 37). Computers and communication networks can connect students to resources, tools, and communities that are necessary to develop the skills, knowledge, values, mindsets (Shaffer speaks of "epistemic frames") required for complex, nonroutine, knowledge-rich problem solving.

Specifically, networked computers can be used to make collaborative learning easier to conduct, and more sustained and integrated with life in- and outside of formal educational settings. While it can be hard to manage collaborative learning from a teacher's perspective, at least for large classes (in university courses, often comprising hundreds of students), learning management systems have made it considerably easier to deal with the logistics, and tools such as LAMS (http://www. lamsfoundation.org/) that specializes in forms of team learning offer even more support. And while it is very hard in face-to-face groups to keep track of individual contributions (with problematic consequences for motivation and group dynamics), when collaboration is conducted through networked computers, students' individual contributions can easily be recorded and set in relation to each other. Thus, through technology we have fairly direct access to students' socially distributed practices, to the tools and artifacts (such as chats, forums, shared whiteboards, wikis) used in these practices, and to the products of their collaborative work, such as texts, models, programs.

The role that and product type artifacts play for knowledge and creation and learning has in particular been recognized in the "trialogic" framework (Paavola & Hakkarainen, 2005): "Trialogue means that by using various mediating artifacts (signs, concepts, and tools) and mediating processes (such as practices, or the interaction between tacit and explicit knowledge) people are developing common objects of activity (such as conceptual artifacts, practices, products, etc." (p. 546). This view that learning is not adequately described by the acquisition and/or the participation metaphor (Sfard, 1998), but that a third dimension needs to be taken into account: learning as artifact creation. Artifacts are always social in nature, not only because they are frequently created in a collaborative fashion, but also because they are intended for subsequent use by others. Once created, artifacts, in the form of concepts (such as a scientific theory), tools, and practices, enable people to engage in activities were not at their command before – Engeström (1987) speaks of learning as an "activity-producing activity." This third metaphor for learning is better suited than the acquisition and the participation metaphor to accommodate notions of innovation and knowledge building (Scardamalia & Bereiter, 2006).

Since knowledge can reside in peoples' minds, in their practices, and in the tools and artifacts they create and use, learning means not only increase individual

knowledge and skill, but can also take the form of improving upon social practices (such as forms of team work) and of creating or modifying tools and artifacts. Working together, for the purpose of learning, makes address all three forms, and indeed is a for practice improvement, such as improving on general and specific team skills.

Attending to "team skills" is not only important when the goal is to teach such skills directly, but a basic level of functioning as a group is a prerequisite for successful collaborative learning in general, as it is for working together (Arrow, McGrath, & Behrdal, 2000; Kreijns, Kirschner, & Jochems, 2003). While productive group collaboration can be "designed" to some extent from the outside, e.g., by setting up roles and workflows ("scripting" the collaboration, see (Kollar, Fischer, & Hesse, 2006), there are limits to this, in particular as we move to long-term collaboration (weeks and months instead of minutes and hours) and as collaboration skills themselves become the learning goal.

Long-term group work among students can take various forms. Examples are knowledge-building communities (Lee, Han, & van Aalst, 2006), problem-based learning (Zumbach, Hillers, & Reimann, 2003), and design-based learning (Kolodner et al., 2003). While knowledge-building communities can be seen as working with declarative knowledge directly by creating artifacts that represent ideas, theories, explanations and their relations, the pedagogy behind problem-based and design-based learning addresses learning and knowledge creation indirectly: students are engaged working on a task, the artifacts created are typically task-related (e.g., a design sketch), and learning is seen as occurring as a "side effect" of working on the task. In any case, for groups of students working together as a team for some time, not only do they need to get their task done (e.g., knowledge building, problem solving, design), they also need to manage their interaction, establish and maintain common ground, keep the group stable, and take care of individual members' concerns (McGrath, 1991). These collaboration management aspects mean that teams need to engage in ongoing learning about how to manage themselves.

To appreciate the complexity of teamwork, the conceptual framework suggested by Arrow et al. (2000) is illustrative. They describe groups as involving the elements members, tasks, and tools, and comprising six networks of relationships between these three elements: (a) the network between team members (social relationships, e.g., affiliation), (b) tasks (i.e., dependencies), and (c) tools (e.g., flow of data between various software tools), and furthermore (d) the role network between tools and members, (e) the labor network between members and tasks, and (f) the job network between tools and tasks. Each of these networks needs to be initially established, and then continuously elaborated, enacted, maintained (e.g., monitored), and modified. In the framework of Arrow and colleagues, this is called the "coordination cycle." An elaboration of one or more of the networks is often necessary because typically not everything a team needs to have and know in order to get started with its work (enactment phase) is provided from the outside. The task may be given to the team, but it may need identification of subtasks. Group membership may be specified from the outset, but roles may need to be identified by the team.

Elaborating, maintaining, and modifying (improving) these networks does not get easier when the team members interact (mainly) with each other and the task mediated by communication technology. "Virtual" teams face additional challenges resulting from too little information: the lack of social awareness (Bodker & Christiansen, 2006), lack of common ground (Clark & Brennan, 1991), lack of transactive memory (Wegner, 1986), lack of social control and, hence, increased tendency for free-riding (Albanese & Fleet, 1985), lack of experience with the communication technology, and so on. At the same time, virtual teams suffer from too much information: too many postings, e-mail messages, and parallel activities that are hard to monitor and to make sense of (Fussell et al., 1998).

These complexities of team work in general and of virtual teams in particular need to be addressed, certainly in cases where the goal is to develop team "skills" amongst students, but also for supporting groups for the benefit of subject-matter learning and knowledge development. Our research aims at developing and analyzing computational methods for visualizing aspects of team performance to help virtual learning teams with their production tasks and to help them to become better in their coordination tasks, i.e., to develop team skills.

We have been studying teams in the context of university courses (advanced undergraduate and graduate level) performance involves substantial collaboration, over several weeks if not months, where the primary goal for student teams is to create a shared artifact (such as wiki pages, programs, and models) and where the learning occurs in the context of working on the task and reflecting on one's performance. We use the term "team learning" to signal that under such circumstances students need to act very much like a "real" team, where this includes making substantial investments in managing team processes. In an educational context, such learning teams can be expected to (a) produce useful artifacts that constitute a contribution to socially shared knowledge (e.g., a problem solved), (b) to learn individually about the domain the problem is contextualized in, and (c) to learn individually about the team members and to develop knowledge and skill about collaboration management. On the group level, we can expect learning to occur (d) by improved team effectiveness, such as improved coordination of members' activities, and in general, changes in group work practices.

An example is the case where students work as a software team that needs to collaborate over several months to build a system for a client, using the Extreme programming method (Beck, 1999) for the broad software development process and Java as the programming language. In this case, by the end of their project students can be expected to (a) have produced a program that satisfies the requirements, (b) Java programming and know more about the domain that the program addressed, (c) have a better understanding of team members and team processes as well as improved collaboration skills, and (d) work better together as a software development team.

One cannot expect that the learning outcomes (b–d) will emerge automatically; providing students with a group task, some incentives for accomplishing the task, and collaboration tools are necessary, but not sufficient conditions for productive interactions and serious learning to occur (Kreijns et al., 2003). For learning

to take place along these dimensions, group members need to be, for example, motivated and supported. Motivation can be established, for instance, making the learning goals explicit and rewarding progress; in particular, it must be clear to the team members that they are supposed to learn about team management, leadership, online collaboration, and so on, and that this learning will be rewarded. Support can be established in many forms (see Reimann, 2003), but teams must be provided with the information required for learning along all of the dimensions and they must be provided with information on team processes, in addition to the task-related information. This, then, is our basic approach to "teach" team skills: We provide groups with a challenging task, including criteria for success, a suite of authentic collaboration technologies, access to information on group performance parameters, and expectations associated with indicators for effective team work.

In this chapter, we describe a number of mirroring and feedback approaches that we have developed to support collaboration management for teams. They make use of electronic communication tools, especially, but not only, wikis, in order to create a jointly constructed knowledge artifact, such as a program or a report. All these approaches were developed, and have been tested to varying degrees, in the context of university courses at undergraduate and graduate level.

We begin with an analysis of artifact-mediated, net-based collaboration, using wikis as a paradigmatic case. Demands that this mode of collaboration impose on students are identified, based on a review of the literature on computer-supported collaborative writing and on a semiotic analysis of wiki writing. This analysis yields two central areas for support: coordination of team members, and establishment of coherence in the shared document. We then describe our first approach to supporting artifact-mediated collaboration. It targets team member coordination by measures aimed at increasing task and member awareness. The approach exploits the database of student actions as they make long-term use of an online collaboration tool. We focus on providing visualizations of participation patterns. These are intended to support reflection by team members and, especially, team leaders.

We then turn our attention to the question of how students can be supported in creating coherence amongst their collaboratively developed ideas as reflected in the shared knowledge artifacts, wiki pages in particular. Just as a sitemap can be an invaluable aid for static web sites, a new tool, WikiNavMap, provides several ways to see the structure of a wiki and to "see" or visualize the most salient features, including the parts where a particular author made contributions, the parts that were edited at different times, and the rate of development of the wiki. Importantly, it also shows how parts of the wiki are linked to each other. We have trialed this in long-term group projects. This approach to mirroring information should help students reflect on many important questions for their collaboration and progress, including whether they have covered all the topics/aspects relevant to their task.

A second approach we have developed to help students with the task of creating coherent works on the level of individual wiki pages. Using computational text analysis methods, this approach identifies and visualizes a network of concept

relations. As in our first approach, we hope that mirroring information should help students to reflect on the knowledge contained in the pages and the learning with respect to domain concepts the group went through. Before we introduce these team support measures, we first the need for such support.

Collaboration Mediated by Knowledge Artifacts

We are focusing on wikis as the main collaboration medium not only because they are frequently used to accomplish work in (virtual) teams, but also because they play an interesting double role: they often operate as both the medium and the product of collaboration. As a product, they function as a knowledge object (Paavola & Hakkarainen, 2005), whereas as a medium, they function as a coordination device (Olson, Malone, & Smith, 2001).

In our research, we work with student groups that are using shared textual artifacts as a means to go about their work and to communicate their ideas, such as wiki pages and program code managed in a shared versioning file repository. This is different from the case where dedicated interaction technologies are used, such as chat, discussion boards, newsgroups, or email (e.g., Stahl, 2006). However, the use of shared artifacts of textual representations is quite typical for the communication that takes place between software developers (Ripochet & Sansonnet, 2006) or between authors of jointly written documents (Zacklad, 2006). In such groups, one finds collaboration typically being conducted through a combination of face-to-face meetings, synchronous remote communication such as phone conversations, and an asynchronous textual medium such as a wiki. The artifacts created on wikis and in version-controlled collaborative document repositories can be seen as combining work on the task with interaction and coordination functions, to the extent that such artifacts are used not only to document work, but also to coordinate team members' activities and to structure their interactions. Using such document-like artifacts is convenient because they are often part of the groups' work anyway and hence constitute little communication overhead (MacMillan, Entin, & Serfaty, 2004). For instance, software designers often use wikis to document use cases and to develop user manuals, and they use shared versioning systems to both manage the code and also to distribute tasks amongst the team members (Layman, Williams, Damian, & Bures, 2006).

This convenience factor can easily lead to problems. Due to the fact that interaction and coordination functions are not systematically separated from production tasks, and given that documents tend to grow quickly in size over a project's time, it can become hard for team members to keep track of tasks and commitments. One way to address this issue is to separate the coordination aspects from the production aspects but keep them within the same basic medium. This is, for instance, possible in systems such as Xplanner (http://www.xplanner.org) and in Trac (trac.edgewall. org), the wiki-based group support tool employed in our research.

Collaborative Computer-Supported Writing

Collaborative writing (CW) is widely performed in industry, academia, and government (Cross, 2001), and with the rise of Web 2.0 genres such as wikis and blogs, it has also become part of popular culture. Amongst the positive effects of writing documents collaboratively are learning, socialization, new ideas, and more understandable – if not more effective – documents (Phillips, Lawrence, & Hardy, 2004). However, outside of creative writing courses, writing in a collaborative manner is hardly taught and practiced in secondary and tertiary education (with exceptions, e.g., Lowry, Nunamaker, Curtis, & Lowry, 2005). The only forms of collaboration in the writing process that students might experience are typically variants of peer review (Topping, 1998), but even then the goal is still to improve upon an individually authored document.

Collaborative writing, defined by Lowry, Curtis, and Lowry (2004, p. 72) as "… an iterative and social process that involves a team focused on a common objective that negotiates, coordinates, and communicates during the creation of a common document" is a cognitively and organizationally demanding process. As a special form of group work, it involves a broad range of group activities, multiple roles, and subtasks. When performed by groups that communicate (partially or only) through communication media, the process typically involves additionally multiple tools (e.g., phone, mail, instant messaging, document management systems) that each have different use characteristics.

From a cognitive perspective, (individual) writing has been described as an "ill-structured" problem type, meaning that there is no single "correct" way to write a particular document, and that instead, the writing task has to be clarified by the writer(s) before engaging in any more targeted problem solving (Hayes & Flower, 1980). When performed in an educational context, a lecturer typically provides the writing task, writing and communication tools, and group composition, so that teams can focus on team planning and document production. Both of these are typically complex, and involve steps such as task decomposition, role definition, task allocation, milestone planning as components of team planning, and brainstorming, outlining, drafting, reviewing, revising, and copyediting as components of document production.

With respect to computer-supported collaborative writing, two areas of research, in particular, are relevant for our purpose: (a) research that analyzes CW in terms of group work processes, focusing on issues such as process loss, productivity, and quality of the outcomes (Erkens, Jaspers, Prangsma, & Kanselaar, 2005; Lowry et al., 2004); and (b) research that studies CW in terms of group learning processes by focusing on topics such as establishing common ground, knowledge building, and learning outcomes (Scardamalia & Bereiter, 1991). In the second line of research, writing is seen as a means to deepen students' engagement with ideas and the literature and for knowledge building (Scardamalia & Bereiter, 2006). In CSCL, in addition to knowledge building in asynchronous collaboration, synchronous collaborative development of argumentative structures and texts has received much

attention (e.g., van Amalesvoort, Andriessen, & Kanselaar, 2007). Recent research has had a particular focus on textual data-mining techniques to enhance student writing. For instance, Williams, Calvo, & Bell (2003) applied automatic classification techniques to classify student postings based on the topic of their content. Others have applied similar techniques to classify postings by the type of contribution they make to an argument (Dönmez, Rosé, Stegmann, Weinberger, & Fischer, 2005).

Collaborative Wiki Writing: A Semiotic Analysis

Wikis, as one type of the evolving class of online documents, can be seen as the current culmination point of three trends in document production: (a) from text media to multimedia: online documents often contain nontextual materials, such as images, sound, and video; (b) from a clear separation of production and use-time to a fusion as wikis are already "published" while they are written (in some cases, such as Wikipedia, they are constantly (re-)written, so that their end state is indeterminate); (c) from single authorship to collective authorship.

Wikis have rapidly become part and parcel of learning environments in Higher Education. They are now a component of most learning management systems, for instance, and are routinely used in courses and seminars. And with the rapid spread of blogs wikis on the Internet, wikis have become part of popular culture. Although the interest in wikis has been boosted by the success of Wikipedia, this specific use of a wiki engine to orchestrate mass cooperation around an online encyclopedia is not the one most typical for (higher) education. Wikis were originally developed for fostering writing and collaboration in small teams (software development teams in particular, Leuf & Cunningham, 2001), and it is this form of use that we will be analyzing.

We focus, in particular, on situations where wiki documents are used to mediate the coordination of distributed small teams of students who are working together toward a common goal. Wikis are frequently used by teams not only because they are easy to set up and use, they also typically provide excellent support for a deep notion of document versions. The wiki can be the immediate object of students' activities (e.g., a group writing assignment), or it can play a supporting function in creating other artifacts, such as computer programs or models. In the later case, wikis are often used to explore the problem, communicate with clients outside of the team, to support coordination within the group (such as meeting agendas, minutes and other joint planning pages) and to write documentation for the program or model. In all these cases, a wiki page (or set of wiki pages; we use the singular "wiki" for ease of reference) is appropriately seen as a "document for action" (Zacklad, 2006), as "… a set of fragments contributed by various authors, the final content of which remains largely indeterminate, while its fast dissemination makes it a useful tool for conveying information, assisting decision-making and probing situations" (p. 206). We will apply Zacklad's insightful analysis of electronic documents-for-action (DofA) to the special case of collaborative wiki text authoring.

Like blogs and other documents which exist mainly online, wiki pages are characterized by certain properties that make them very different from paper-based documents, amongst them: a prolonged state of incompletion, durability, high degree of fragmentation, diverse commitments of their authors, and the evolving nature of their content (Zacklad, 2006). The two main activities performed on wiki pages are adding fragments and adding annotations. Annotation activities can be defined as all those "…activities serving to link together the fragments of DofAs with a view to achieve the common goals of a distributed collective practice" (Zacklad, 2006, p. 6). Wiki pages, produced in the context of collaborative projects, mediate (potentially widely) distributed emerging communicational transactions. They pose unique challenges regarding the coordination of these transactions and the establishment of coherence amongst the text fragments.

Coherence amongst the fragments in such documents is an emerging quality rather than being due to a specific plan or outline. Establishing coherence becomes a particular challenge, as this has to be accomplished in and through the same medium – the wiki page – that contains the fragments. Coordination of the transactions on the perennial artifact "wiki page" is mediated through the wiki medium and at the same time geared toward the wiki page, in order to make it serve its purpose(s) and to keep it coherent.

Since the transactional practices around the construction of wiki pages are distributed over space, time, and actors, groups and organizations have to find ways to effectively manage these transactions. (Zacklad, 2004) distinguishes eight (not exclusive) methods to accomplish the coordination and distribution of communicational transactions in general, with the last four being of particular relevance for print and electronic documents, including wikis: (1) standardizing the transaction situation, (2) formalizing the modes of expression, (3) mnemotechnic ritualization (such as using rhymes), (4) abstraction, (5) substitutive mediation, (6) documentarization, (7) increased recourse to techno-informational equipment, and (8) substitutive coordination. We now discuss the last four.

Substitutive mediation (5) refers to the phenomenon of separating the production of a semiotic product (e.g., an utterance of a sentence by a speaker) from the reception of that product by a receiver (a listener) by one or more intermediate media. Substantive mediation is particularly effective to help with distribution of the semiotic product (and coordination of its production) if the intermediate medium is perennial, such as a writing substrate. Documentarization (6) is an extended version of substantive mediation as it refers to endowing perennial substrates "… with specific attributes making it possible: (i) to manage them along with other substrates, (ii) to handle them physically, which is a prerequisite to be able to browse semantically among the semiotic content, and lastly (iii), to guide not only the recipients, but also the producers themselves to an increasing extent, around the substrate by providing one or several maps of the semiotic contents" (Zacklad, 2006, p. 215). For instance, the creation of a table of contents, or of an index, is an example a documentarisation process. Techno-informational equipment (7) refers to things such as filing cabinets (and their digital "equivalents") that affect the use of documents (and/or their production) without being part of the documents, while substitutive coordination (8)

results from the automation performed by techno-informational equipment. The automatic indexing of a document, for instance, document retrieval technologies, and the coordination of a document production process with a workflow tool are examples of substitutive coordination.

In the case of wikis as the perennial artefact involved in groups' work, it is in particular the distinction between substitutive mediation (5) and documentarization (6) that is important. In order to make a wiki a good resource for team work, more is needed than just "writing things down"; the collectively generated fragments must be further reformulated and organized in a manner that reflects their shared meaning amongst the team members. If the artifact is supposed to be of use for others who are not part of the team, then even more documentarization work will need to be invested to turn notes into a knowledge resource.

Wikis share essential features with online newsgroups and discussion forums. Similar to those, wikis are used to conduct and store traces of a kind of polylogal (multiparty) communication. In this respect, wikis are also similar to face-to-face group communications and to chat communications, but unlike those, are not conducted synchronously. Different from online newsgroups and discussion forums, wikis lack an explicit representation for "contributions" and for the take-up of contributions in form of "responses." Thus, in terms of their affordances, wikis are more a medium for writing than a medium for discussing. However, because of their double role, at the same time the product of writing and the medium to coordinate writing activities across multiple authors, wiki pages will often contain "discussion" entries (for instance, discussions as to what to include and not include in the text) as well as substantial entries. These discussion entries have the characteristics of annotations since they refer to the text but are not (yet) a proper part of the text.

Annotations can be seen as a mechanism to deal with the main problem confronting members of groups coproducing DofAs in general and wikis in particular: the lack of information about the transactional context associated with a proposed fragment. While the transactional context is continuously and seemingly effortlessly established in face-to-face communication, mediated communication situations, and particularly those of the asynchronous kind lacking an immediate feedback/repair channel, make it much more difficult to establish context. Annotations are the main device to establish context in such situations.

Following Zacklad (2006), we can see an author contributing a free (i.e., not yet integrated) fragment to a wiki page as a basic turn, or transactional bid. The free fragment can be complete or incomplete. It is complete when participants (readers and coauthors) perceive it as a coherent micro-transaction, and incomplete otherwise. The relationship between a (complete) free fragment and the wiki page as the main semiotic product emerging in the framework of the transaction can be variable: They constitute accessory semiotic products as long as their status with respect to the page has not been clearly established. If taken up by coauthors, a fragment can either be rejected, or be (gradually) included in the main product by becoming subject to changes in the mode of expression and/or semiotic content. The uptake step is essential: "However, if free fragments are not properly articulated together as soon as they are inscribed on the substrate, the uptake process will

not be possible and the DofA will not be able to efficiently sustain the emergent distributed transactions involved in the cooperative activity" (p. 221). Hence, the uptake is conditional not only on the readers' processing of the text, but also on how the fragment is introduced by the author. In peer groups, uptake will be made more complex by the fact that there is not clear attribution of specific authority to accept/reject free fragments.

Some Conclusions About Wikis

Wikis constitute, from a semiotic perspective, a rather complex document category. They are complex not only because of the need to coordinate a multistep group writing process (Lowry et al., 2004), but also because of the need for extensive documentarization in order to create shared meaning for what is collaboratively written. Wikis, being essentially a writing medium, require substantial coordination efforts amongst the members of the team involved in creating documents in them. A particular challenge is the establishment of coherence, on the level of text (connecting fragments that take the form of sentences and paragraphs) as well as on the level of concepts (ideas, arguments). The same holds for other types of digital knowledge objects that can be authored by a number of people in a manner at least technically independent of each other, such as documents in shared repositories (e.g., Lotus Notes) or on the Web (e.g., Google Docs).

Based on our analysis of wikis as documents-for-action, they ought to be rather hard to use when the goal is to produce coherent documents. So far, systematic analysis of this issue has not been performed. We can take the increasing interest in semantic wikis (Souzis, 2005) and in visualization/navigation support for wiki sites (Reinhold, 2006) as indicative of the fact that today's standard wiki technology has recognized limitations. It is also informative in this regard to see the many social rules and the increasing user role differentiations that have been evolving around the mass-cooperation sites, notably Wikipedia.

In general, students who are not experienced in working in virtual teams and who are not experienced in authoring text or other knowledge artifacts together will need to be supported in order to produce knowledge objects of good quality, and to learn from this experience.

Below we report on our own work on providing such support, in three forms: (a) by monitoring and visualizing group members' interactions and contributions, (b) by visualizing wiki site structure, and (c) by providing information on wiki page content based on a text-statistical analysis. These measures aim at improving coordination of team members' activities and increasing document coherence. In an educational context, such as a university, students are ideally not only supported in creating knowledge artifacts collaboratively, but would also learn how to get better at this, in particular learning how to work in (partially or fully) virtualized teams and how to author collaboratively. Therefore, we will look into team skills development next.

Developing Team Skills

Providing groups with ill-structured problems (i.e., problems that require elaboration and negotiation to be defined, and for which no single correct solution exists) and mirroring/feedback information, instead of well-formed problems and strong scaffolds and guidelines, can be seen as a risky pedagogical strategy (Kirschner, Sweller, & Clark, 2006). However, when development of team skills is part of the learning goals it is necessary to provide sufficient space so that students can experience aspects of real teamwork such as breaking down a large task into subgoals, defining roles and allocating team members to roles and tasks, reacting adaptively to unforeseen obstacles, changes in the group environment as well as in members needs, etc. There might be no other way to learn to manage what sometimes looks like chaos than to experience the chaos. There is no question that students will sometimes fail to deal with the demands of dynamic group work, but failure can be productive (Kapur, 2006), in particular when followed by opportunities for reflection.

Because most of the research on collaborative learning pertains to forms of collaboration where many of the decisions concerning group composition, roles, responsibilities, and timing have been made for the group, not by the group, developing team skills as such has not been much of a research issue. There are exceptions, such as research on how training for group work in school settings affects individual learning (Yager, Johnson, Johnson, & Snider, 1986), group performance (Johnson, Johnson, Stanne, & Garibaldi, 1989), and collaborative language and behavior (Gillies, 2004; Gillies & Ashman, 1996) as well as more recent work in Higher Education (Prichard, Startford, & Bizo, 2006). But by and large, team (skills) development has received more attention in organizational psychology, in the context of work teams, and by far most of the research has been conducted on types of teams where communication and coordination breakdowns can quickly result in disaster, groups of soldiers and pilots in particular (e.g., Cannon-Bowers & Salas, 1998). In the organizational and military training research literature, we find a diversity of team-training approaches that have been developed, as well as reviews of their effectiveness (Dyer, 1984; Kozlowski & Ilgen, 2006).

Pedagogical Approaches

Team-skills training is a broad term, comprising numerous skills such as goal setting, interpersonal relations, and role clarification (e.g., Buller & Bell, 1986), and various training strategies, such as team building (Salas, Rozell, Mullen, & Driskell, 1999) and team self-correction strategies (Blickensderfer, Cannon-Bowers, & Salas, 1997). An important didactical decision is whether to teach such skills individually or practice them in teams. While research shows that team

skills can to a certain extent be taught individually, highly interdependent forms of team work requires practice in team form (Kozlowski, Brown, Weissbein, Cannon-Bowers, & Salas, 2000). Focusing on teams as the learning unit, a further decision has to be made instruction and training to develop spontaneously in the context of teamwork. The consensus seems to be that any kind of team skill can be enhanced by facilitation through training, if not always qualitatively then at least in the form of an acceleration of the development process (Prichard et al., 2006).

Skills Addressed by Training

In addition to the variation in didactical approaches, one finds a variety in the topics that are addressed in team-skills education, such as goal setting, interpersonal relations, and role clarification. As Prichard and Ashleigh (2007) summarize:

> Training directed at goal setting emphasizes the setting of goals and objectives, the identification of obstacles to achieving goals, and action planning to determine how goals are to be reached and obstacles overcome. The interpersonal model focuses on the development of open communication, mutual trust, and cohesion. Role clarification models emphasize the different interacting roles that people play in a group situation and aim to increase each person's knowledge about the roles played by others. (p. 703–704)

Increasingly, these elements are combined into integrated, generic training programs (e.g., Prichard, Stratford, & Hardy, 2004).

Analogous to using analysis of experts' competence and skills identify learning goals, we suggest the factors characterizing successful teamwork in order to decide on learning goals for team skill development. This is particularly relevant when supporting team learning in a computer-based manner. In this case, aspects of individual and group work to provide feedback, given that many aspects can be recorded easily, but only few can be visualized in the shared working environment on a computer screen.

The question of what processes and components comprise teamwork and how teamwork contributes to team effectiveness has received much attention in team research. A recent review of this body of research resulted in the identification of the "Big Five" components of teamwork (Salas, Sims, & Burke, 2005). This review identifies the elements that make up teamwork, independent of the task a team has to perform:

1. Team leadership: Ability to direct and coordinate the activities of other team members, assess team performance, assign tasks, develop team knowledge, skills, and abilities, motivate team members, plan and organize, and establish a positive atmosphere.
2. Mutual performance monitoring: The ability to develop common understandings of the team environment and apply appropriate task strategies to accurately monitor team-mate performance.

3. Backup behavior: Ability to anticipate other team members' needs through accu-
   rate knowledge about their responsibilities. This includes the ability to shift
   workload among members to achieve balance during high periods of workload
   or pressure.
4. Adaptability: Ability to adjust strategies based on information gathered from the
   environment through the use of backup behavior and reallocation of intrateam
   resources. Altering a course of action or team repertoire in response to changing
   conditions (internal or external).
5. Team orientation: Propensity to take other's behavior into account during group
   interaction and belief in the importance of team's goals over individual mem-
   bers' goals.

Teams that enact these five elements will enjoy improved performance. However,
in order to fully realize this performance improvement potential, research shows
that three additional coordinating mechanisms need to be in place (Salas et al.,
2005, p. 564): (a) Shared mental models: An organizing knowledge structure of the
relationships among the tasks the team is engaged in and how the team members
will interact; (b) Mutual trust: The shared belief that team members will perform
their roles and protect the interests of their teammates; (c) Closed-loop communication:
The mutual acknowledgment of the success or otherwise of an exchange of infor-
mation between a sender and a receiver, irrespective of the medium.

Transfer

An important issue is the extent to which team skills can be seen as generic, i.e.,
independent of a specific team and task. Cannon-Bowers, Tannenbaum, Salas, &
Volpe (1995) argue that team members must have both generic and specific team
competencies (knowledge, skills, and attitudes) because in order to be effective, the
generic aspects must be enriched by team-specific information. For instance, it is
not only important that team members share information (generic), but each team
member must have knowledge about the knowledge and skills other team members
have (specific). One implication is that regrouping of effective teams should lead to
performance loss, which has been widely demonstrated in team research (Prichard
et al., 2006).
    A fundamental challenge for developing team skills arises from the multilevel,
dynamic, and complex nature of groups. There is general consensus amongst
researchers that groups need to be conceptualized as embedded in a multilevel
system that has individual, team, and organizational aspects (Kozlowski & Ilgen,
2006) as well as a wider socio-cultural context (Bonk & Cunningham, 1998).
Furthermore, groups incorporate temporal dynamics involving active different time
scales, ranging from the episodic to the developmental (McGrath & Tschan, 2004).
And that groups share many features with other complex, open systems, in particular

the fact that many group processes are emergent phenonema and that so, there are limits to their predictability. If groups are essentially complex systems, than we cannot account for group phenomena by aggregating over the individual members, the effects of interactions on the member level will not directly and linearly show effects on the group level (but be highly state dependent), and group level phenomena (such as role distribution, power structure, shared knowledge) will affect member behavior, but cannot be reduced to it (Arrow et al., 2000; Kapur et al., 2007). Furthermore, complex systems are path-dependent: How they will react to a signal from their environment depends not only on the signal and the current state of the system, but also on the "history" of the system. Time matters (Reimann, 2007).

This does not only makes comparative (e.g., experimental) research on groups problematic (because it is not quite clear what we are comparing, given that any group will be different from any other after a couple of minutes into their group existence), it also make it hard to come up with general advise on group performance, given the differences in group behavior as well as the path-dependence of the effects of the advice. One way to avoid these conceptual issues is to provide groups with information that is specific to their history and their current situation, and is hence adapted to the potentially unique information needs for each specific group at each specific point in time. To accomplish this, computational means are needed, because it is not feasible that human tutors or facilitators can deliver this kind of information just in time.

Supporting Coordination by Visualizing Interactions

We have identified the need to help students working in teams to produce knowledge artifacts with respect to two aspects: coordinating member activities and establishing coherence amongst ideas (concepts) and of the documents produced. In this section, we focus on the first aspect, supporting coordination.

Systems that support the management of distributed collaborative learning processes can be classified as mirroring tools, metacognitive tools, or guiding/coaching systems (Soller, Martínez-Monés, Jermann, & Muehlenbrock, 2005). We are interested in the first two categories, as they place the locus of the processing into the learners' hands. Mirroring tools (see, e.g., Barros & Verdejo, 2000) impose few constraints on users, but are purely performance-based and do not offer semantic interpretation or analysis of the nature of the user's intervention. In particular, mirroring approaches do not compare current performance with target performance; they do not show a gap. Mirroring is particular appropriate when information about "good" or "optimal" behavior is lacking. Metacognitive approaches build on the notion that a model of good performance (if not theory-based, then at least "best practice") is available and hence the difference between current and target performance – the "gap" – to the learners. In other words, the metacognitive approach provides students with feedback.

Functionally, visualizations of the mirroring type should be effective for enhancing team work in general and member/task coordination in particular because they contribute to task awareness and/or social awareness, depending on what is visualized. Task awareness should be increased when the visualization displays information on tasks and the degree of task completion. A typical example is a Gantt chart, often used in project management. Social awareness should be increased the type of information displayed about team members (e.g., their expertise areas; available time), the relations between members and tasks (e.g., task load, roles), and relations between members (e.g., social networks diagrams).

Visualizations for Wiki-Mediated Collaboration: Wattle Trees

We have created a set of interrelated visualizations that display a useful overview of the vast amount of information stored in electronic traces such as log files. We designed the visualization to directly support team functions (Kay, Maisonneuve, Yacef, & Reimann, 2006). In our approach, we draw upon theories of group work to define dimensions of group operation that we wish to scaffold and where we can identify sources of evidence of the ways that these are operating within a group. We track students' interaction behavior along these dimensions and provide visualizations that are mirrored back to the groups. We believe that groups gain benefit from "just" mirroring information, provided that information speaks to the right issues.

Collaboration Environment

To support their tasks and communication, groups use Trac (http://www.edgewall. com), a tool designed for programmers build software. It has three tightly integrated parts:

- A wiki for collaborative editing of web pages for general group communication, and in our case, for collaboratively creating the major report for the project;
- An issue tracking system based on so-called tickets (see Fig. 6.1), where one creates a ticket when a task needs to be done and this is allocated to a team member and, when the task has been completed, the ticket is closed;
- A browsing interface to a repository based upon the version control system called Subversion (SVN), for storing documents like source code, including any versions.

We describe our visualizations in the context of trac as it has provided most of our experience to date. Moreover, it is an authentic tool that is widely used and is representative of a substantial class of tools that such groups use: it supports

Fig. 6.1 A ticket in Trac. A ticket represents a task issued by somebody (the reporter) to some-body else or to oneself. Tickets in Trac are associated with milestones (available on a different screen), and with the due dates specified in the milestones

task management and allocation, general communication and shared text space, and is the central repository for the work produced, three important pillars for team work.

Form of Team Work

In order to illustrate our approach, we use observations from a software development project where students work in groups of five to seven over 13 weeks. Team members tend to focus on the goal of producing a software product that meets their clients' needs, rather than the group management needed to achieve this. Following the Extreme Programming (XP) approach (e.g., http://www.extremeprogramming. org), each student takes one or more of the XP roles, such as team leader (who manages the group), tracker (who tracks people's work and ensure that things are progressing as planned), the programmer (who deals with technical issues), the tester (in charge of functional testing), the doomsayer and so on. Teams meet face-to-face in order to evaluate and coordinate their work, but the main work done through the wiki and the code versioning system (SVN), both accessible in Trac.

Wattle Trees

Given our focus on helping small groups (five to seven members) manage their group processes, including communication, interaction and workload management, what information can we extract from the use of these three mediums to inform team members about these? Wiki pages are not owned by any one person but by everyone who can access the site. When several people alter the same wiki page, we regard this as an interaction. The size of each contribution is taken into account. Similarly, when a team member assigns, or reassigns a ticket to another member or closes it, we also regard this as an interaction between these members. With SVN, when members work on the same files, we also regard this as an interaction and record the size of the contribution.

Figure 6.2 shows our main visualization, which we call the Wattle Tree. (We chose this name as the Wattle tree is an Australian native plant with fluffy golden-yellow round flowers, similar to this visualization). The design goal for this vizualisation was to create a single overview of the total activity of each group member over the 3-month project, with differentiation of the different media. Essentially, it is a bird's eye view of the thousands of actions of the team over a period of time.

Each member of the team is a single wattle tree, with its vertical green stem that grows up the page, each day from the start of the project activity on Trac. Wiki-related activity is represented by yellow "flowers," the circles on the left of the trees. SVN-related activity is similarly represented, as orange flowers on the right of the trees. The size of the flower indicates the size of the contribution. Ticket actions are represented by leaves – the green lines: a dark green leaf on the left indicates a ticket was opened by the user and a light green leaf on the right indicates the user closed a ticket. The length of the left leaf is proportional to the time it remained opened. Those still open are shown at a standard, maximal size (e.g., the ones around day 41 in Fig. 6.2). Often, a good team leader will take the responsibility for opening most of the tickets. We see that the leftmost person in Fig. 6.2 has opened many tickets while the closing of tickets is more evenly distributed across the team.

Although these visualizations are intended to be meaningful for the team, rather than the outsider viewing them, there are some features we can identify in Fig. 6.2. The student at the very left has many yellow circles reflecting high wiki activity until around day 40 and they have created many tickets, indicated by the dark green lines. The second student from the left has a similar level of wiki activity and has opened many short-lived tickets and closed many tickets. Overall, these two members appear to have been the most active in management aspects at the wiki and tickets: knowledge of the group bears this out. The fourth student from the left is particularly active on SVN corresponding to a larger role in the technical development. The fifth student from the left has a hiatus from about Day 27 corresponding to little activity. This group had times when several members were ill or had other difficulties and they could see the effect of these problems in the Wattle diagram. The team member with responsibility for tracking progress could use the Wattle
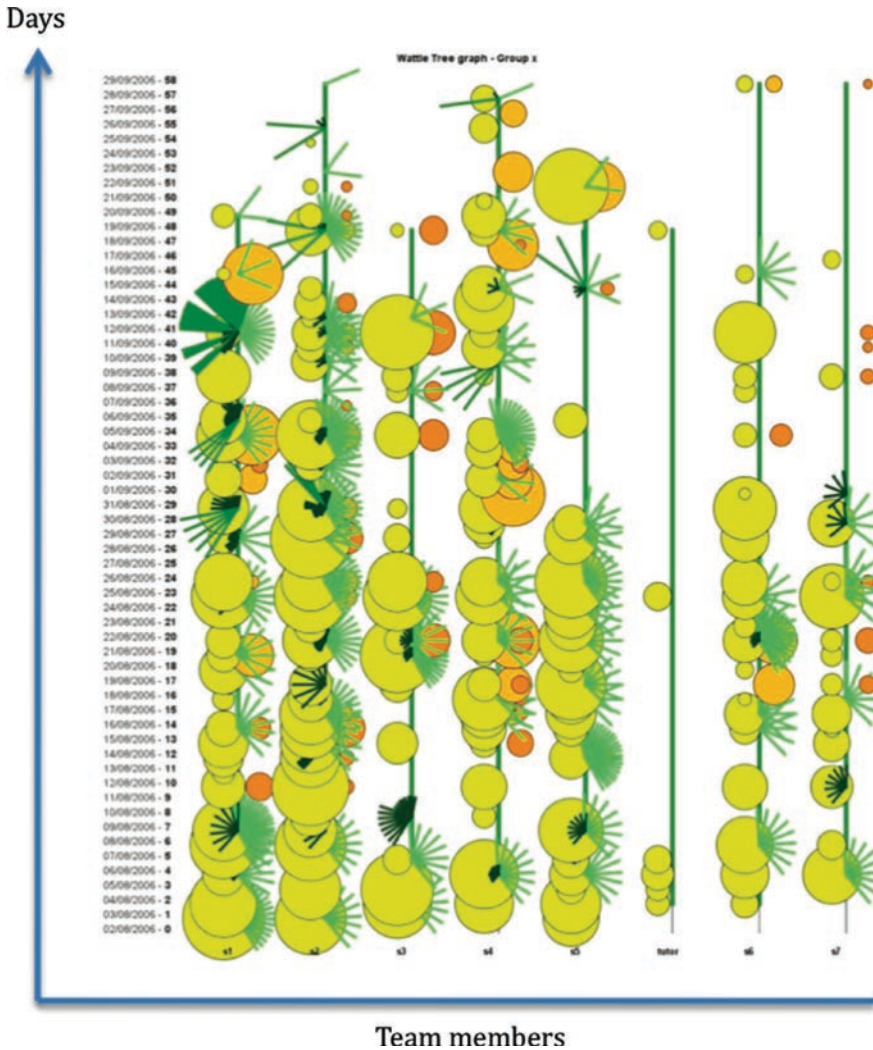
Fig. 6.2 Wattle tree diagram. Each person in the team appears as a "tree" that climbs up the page over time. The tree starts when the user first does an action on any of the three media considered. The vertical axis shows the day number and date

tree to get an overview of overall activity at a glance: they could also quickly see recent changes in activity by individuals, for example, those working on urgent or critical tasks. This serves as a starting point for delving into the details, as needed, by checking individual tickets, wiki pages and SVN documents and their histories. It also gives each individual team member a sense of how their level of activity compares with that of others in the team.

Social Network Diagrams

Wattle trees do not contain information on who issued tickets to whom, and who contributes to a wiki page. In order to visualize this kind of information, we use what we call an Interaction Network, inspired by the graphical notations used in Social Network Analysis (Scott, 1991), which aims to show relationships and flows between entities. The network is modeled as a graph, with each node representing a team member, always shown in the same, fixed position. So, for example, the person at 12 o'clock in Figs. 6.3 and 6.4 is the same in each of these visualizations. Lines between these nodes indicate interaction between these team members. We define interaction to occur when two people modify the same wiki page or SVN file or perform actions on the same ticket. The width of the edge is proportional to the number of interactions between them. For a given resource, the number of interactions is calculated as $n = \min(n1, n2)$ where $n1$ and $n2$ are the number of times user1 and user2 modified the resource.

So, for example, the interaction network for ticket interaction as depicted in Fig. 6.3 shows considerable interaction between most team members (but not the tutor). Note that some team members interact with more members of the team than others. For the tickets, we use color to indicate who initiates tickets, with blue at the node for a person who initiates more tickets. So, for example, a team leader often allocated tickets to all others and in this case, the lines from the leader are blue at the leader's end. Finally, the Interaction Diagram for the wiki shows that every member of the team interacts with every other one, including the tutor. These diagrams change over time. As we have already mentioned, this is intended to be meaningful for the team members who should know who was working on each aspect and who may have been interacting with others.

Table 6.1 briefly describes how the two visualizations relate to team success factors of the Big Five model. The aspects and behavioral markers are taken from (Salas et al., 2005). We show only those what are applicable to the visualization. When designing the visualizations, we had to balance the complexity of the display against the number of aspects presented. Importantly, we had to take account of which aspects could sensibly be inferred from the available data. So, for example, the first row indicates one role of the team leader, facilitating team problem solving. The next column briefly indicates how the interaction network can support this aspect. For example, one potentially pathological pattern occurred when one person could be seen interacting with every other team member on SVN: in this case, this person was fixing problems in all other team member's code, something that they should have been responsible for. This form of domination suggests a problem in the group. This can happen when several people have weak technical skills and they expect the top programmer to fix their code and do the difficult work. This pattern can, equally, occur when one person believes they are better than the others and that person leaps in and works on other people's code, not allowing them to complete it themselves, even though they are keen to do so and capable.

Fig. 6.3 Interaction network based on tickets

Group Interaction: ticket
Group x (Tue Aug 29 00:00:00 EST 2006)
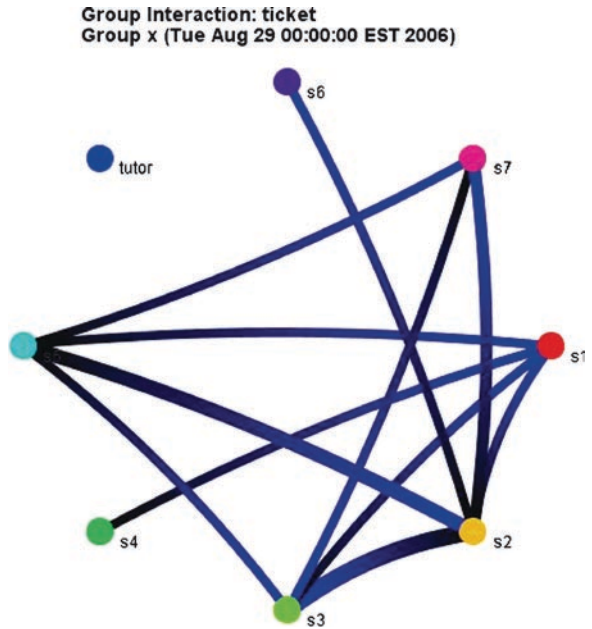
Group Interaction: wiki
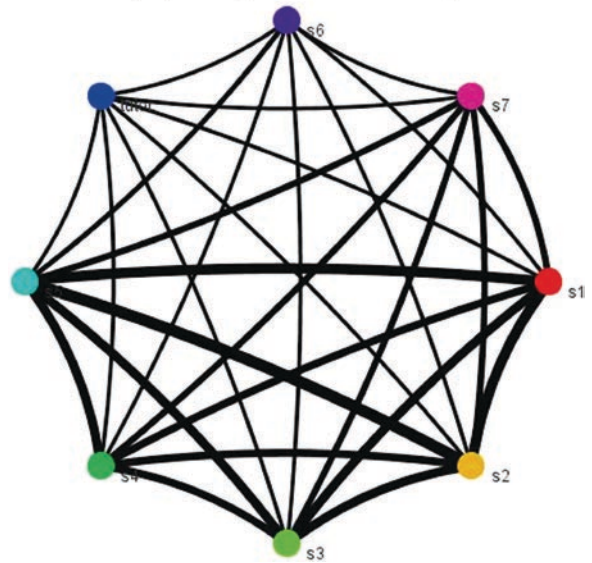**Group x (Tue Aug 29 00:00:00 EST 2006)**

Fig. 6.4 Interaction network based on wiki entries

Table 6.1 Relations between the visualizations and the Big Five framework

Aspect and behavioral markers	Interaction network	Wattle tree
Team leadership		
Facilitate team problem solving	Identify potential problems, due to unexpected interactions, such as one person altering work allocated to another. Starting point for leader to ask others about what appear to be problems	Missing activity where expected points to potential problems. This can be used as a starting point for conversations and checking details
Synchronize and combine individual team member contributions		Progress of each team member and type of activity can be seen and this can help with planning of integration
Seek and evaluate information that affects team functioning	Maintain awareness of who interacts with whom and, importantly clear lack of interaction when expected	
Clarify team member roles	Comparing expected interactions with the actual points to possible problems	Can see activity types and assess if they match expectations – and then act to address problems
Engage in preparatory meetings and feedback sessions with the team		Gain an overview of progress as part of planning the meeting and deciding whether there are problems to discuss at the meeting
Mutual performance monitoring		
Identifying mistakes and lapses in other team members' actions	Identifying lack of interaction when it was expected	Indicates the periods of inactivity
Providing feedback regarding team member actions to facilitate self-correction	Unexpected patterns of interaction may indicate problems, e.g., one person correcting/amending work allocated to another	Shows overall forms of activity, such as increased work on one medium
Backup behavior		
Recognition by potential backup providers that there is a workload distribution problem in their team		Can see distribution of workload, and potential for rebalancing of workload
Shifting of work responsibilities to underutilized team members	Can see interaction on same resources when one person begins work started by another	Can see both recent long-term contribution levels by those doing little at present
Completion of the whole task or parts of tasks by other team members	As above	Can see some rebalancing

Table 6.1 (continued)

Aspect and behavioral markers	Interaction network	Wattle tree
Adaptability		
Identify cues that a change has occurred, assign meaning to that change, and develop a new plan to deal with the changes	In new releases of the interaction overviews, can see new interactions, either confirming planned allocations or showing the unexpected	Can see unexpected patterns of activity and changes in level of use of each medium
Identify opportunities for improvement and innovation for habitual or routine practices	Can see that there is continuation of interactions that have been unproductive or start of new interactions	Can see changes/continuation of contribution levels on each medium and interpret this relative to plans
Remain vigilant to changes in the internal and external environment of the team	Quickly see unexpected changes in interaction	See unexpected patterns of activity, e.g., a person suddenly ill
Team orientation		
Taking into account alternative solutions provided by teammates and appraising that input to determine what is best	Partly visible as interaction on wiki and tickets	
Increased task involvement, information sharing, strategizing, and participation in goal setting	Each team member can see the interactions and link this to other knowledge of progress	Can see indications of chances for increased involvement, e.g., if one person was very active on the wiki, others may realize they need to check for the new information provided

This example illustrates some important aspect of our mirroring approach. Given the complexity of long-term group work on challenging tasks, the simple measures of interaction cannot possibly capture deep and subtle features of the group interaction. However, the team members have considerable knowledge of the nature of the tasks, the allocations as well as the personalities, skills and commitment and the particular circumstances of other team members. They can use this to interpret the visualizations, checking what these show against their expectations.

The third column in the table provides comments on the ways that the Wattle visualization can support the various aspects shown. Again taking the example of the first row, this visualization can either vindicate expected activity on each medium or, when a problem has occurred, it can give an early warning of this. For example, if a team member has been allocated a task, such as writing for part of the

task, one would expect to see progress on this as SVN activity. If there is none visible, it points to a possible problem. The leader can use the Wattle tree to begin the discussion about this: they can simply comment that they would have expected to see such activity.

First Experiences Using the Visualizations

We report here experiences from a semester-long project course (capstone project) where teams used Trac. There were seven groups of 5–7 students in each team, with 44 students making it to the end. We began the semester with a lecture about collaboration management, introducing the visualizations. We introduced weekly meetings with the leader of each team, to discuss concerns and challenges. There were also interviews with each team at three points in the semester to monitor their response to collaboration management.

The visualizations were made available to the students on a regular basis throughout the semester, on their wiki. In both the interviews and weekly meetings with the team leaders, there was a spontaneous response to the visualizations. Three of the seven groups showed great enthusiasm for them and asked to be able to generate them on demand. (This was not possible at that stage.) The students indicated that the visualizations were helpful for the tracker (the person who has to ensure that work is progressing as intended) and the manager (who distributes the workload). There has also been spontaneous reference to the visualizations in relation to some difficulties in groups, particularly in the case of seeming occurrences of social loafing, with an individual failing to carry their fair share of work in the group.

Students have also commented that the visualizations help individuals to see the amount of work they have contributed to the group, to compare it with that of others and to provide some quantitative measurement for balancing the group workload. Notably, six of the seven groups encountered serious problems with group members (absence due to sickness or travel, social loafing and technical weakness). Three groups asked for more frequent releases of the diagrams. Some students explained that they would like to see how the diagrams change after they have contributed a fair amount of work and see how this amount compares with the others. One group mentioned that the lack of contribution from a member showed up on the Wattle Tree. The group would have liked to see the evidence. The member said he took it as a wake-up call, and intended to participate more: importantly, he did so.

Overall, our experience has indicated that these visualizations are particularly useful for providing an early warning of problems. Since we have used them, we have had none of the most dysfunctional groups we saw in the past, especially for the case of social loafing that persisted through the semester. This operates for two classes of reasons. First, it was rather difficult for teachers to identify dysfunctional groups early enough to help them recover. Second, in the past, even when teachers could see indicators of dysfunction, it was very difficult to communicate this to the

students effectively. More recently, there have been spontaneous requests for the visualizations to be made available for other users of Trac.

The main negative feedback was related to the fact that the visualizations are based on simple counts of the amount of activity and there is no measure of quality. This is a very valid concern. There was also some negative feedback about the whole enterprise of monitoring activity and making this explicitly available in the visualizations. It was expressed most strongly by one student in these terms: "The virtual cybernetic monitoring of our work was counterposed by the need to set our own goals and this made for a fairly unalienated work environment but unfortunately there was also a resulting uneasy foreboding if we stopped working a while".

We certainly acknowledge the simplicity of the information presented. However, we also know that it is not really rewarding to play the system since we make no use of the visualizations for anything other than as a support for collaboration management. If a student played the system (and this did happened a couple of times, to a limited extent), perhaps creating many vacuous tickets or large quantities of low quality wiki or SVN content, the team members can readily see that if they scrutinize these. Every team member knows that this is the case. The students seem to generally be most positive about the Wattle tree: it summarizes temporal aspects and enables students to see and respond to changes. The interaction networks also seem to be important and they certainly did point to cases where an individual is isolated from the group.

From a questionnaire study, we gained the reactions as shown in Table 6.2. This shows two values, the average score on the Likert scale (from 1 to 7) first for the full cohort and then for just the seven managers. On average, most students found the visualizations somewhat informative and helpful (mean > 3.5). Notably, those in the manager role gave far stronger positive responses, around 1 point higher on the

Table 6.2 Students' reactions to the visualizations

Questionnaire item (each item scaled from 1 to 7 indicating increasing agreement)	Mean score (all/managers only)
I found the Wattle tree useful	4.7/5.8
I found the interaction network useful	3.8/4.8
I found that the combination of the diagrams helped me learn things that the individual diagrams did not…	3.8/4.5
The diagrams gave me a sense of "big picture" and/or revealed things I did not know about the members' contribution…	4.4/4.8
Without the diagrams, it would be difficult to get this overall sense of big picture…	3.9/4.8
The simplicity of the metrics used still conveys useful information…	4.1/5.0
I recommend the use of these diagrams for future offerings or other group project offerings	5.2/6.0
I/my group changed something during the semester in light of what I/we saw in the diagrams	Yes: 38% No: 62%

Likert scale for most aspects. This is reasonable, since the visualizations support
the role of the managers.

For each of the items, we also allowed for open answers. These are particularly
interesting for the question assessing the consequences of seeing the diagrams.
Students' open answers to this item are shown in Table 6.3. Only one student
(pessimists may say, the only honest one) remarked that the group began to
"play" the system ("cheat the system almost"). However, even this student's
behavior change may be seen as positive – they took more care to report their
work on the wiki and tickets in order to be seen to be working: this, in turn, meant
that the tracker could use these media without needing to wait for the next
meeting. In addition, if the student claimed to have done work, Trac makes it easy
to link this claim to the actual contribution, be it on the wiki or SVN. So, the
tracker should have been able to thoroughly check the work had been completed
satisfactorily. Many students referred to the visualizations for their reflective
statements in their final reports, pointing to features in the visualizations and
explaining the corresponding events.

Table 6.3 Students reporting consequences or implications of the visualizations

I/my group changed something during the semester in light of what I/we saw in the diagrams…	
Yes	No
I tried to show the work I was doing better, using more tickets	Did not shed light on something unknown, so no change was warranted
Make sure to contribute and interact on these levels (cheat the system almost)	Our approach didn't really need to be changed
We update our wikis more often and are pushed to commit stuff regularly	We only got them twice. Need weekly diagrams
	Not that useful, just states the obvious
I realized that I wasn't ticketing enough work so I stepped up	There was no need of change
	Because people were lazy
I took more effort on wiki and ticket after seeing the first diagram	Again, we were not surprised/upset by what they showed. Diagrams in contribution were already known and accounted for
Roles were changed in order to cater for aspects that group members were missing. Even though group members would gravitate toward original position	
I used the Wattle tree to highlight contribution of team members	
Some members find they weren't working enough and improved	
We worked harder. Created more tickets. Did more interactive group activities	
I became more aware of my interaction with the wiki/SVN	

New Developments Based on First Experiences

We were pleased to see how such visualizations are actually used by teams. We found that our teams need to be introduced to these tools. This goes hand-in-hand with the need to motivate team members to appreciate the importance of collaboration management. Compared to former semesters (not reported here), the students in this study worked much harder on their team performance, having been shown why they need to be concerned about group maintenance and having been shown how the visualizations might help. It may have also been important that we showed them how to interpret the diagrams. Equally important may be the value of showing examples of the varied forms these visualizations tend to take in highly successful groups as well as problem groups. This might make it easier to recognize the same visual pattern when it arose in their own group.

    We have subsequently been extending this approach on several fronts. When we began this work, we built the visualizations to be independent of any tool: so, for example, we built visualizations for discussion groups in WebCT and for groups using a set of different media in a Flash based online learning system. A disadvantage of this separation of the learning environment from the visualization software was that the visualizations had to be produced off-line and, hence, they were not available on demand. Rather they were generated and added to the Trac wikis at set times. We have been rebuilding the visualization software to be integrated into Trac, so that it can be used at any time, on demand.

    The second design goal was to generalize the visualization so that not only Trac components such wiki pages, tickets, and SVN interactions be visualized, but any combination of the many communication methods available in courses, and external to Trac, such as discussion board entries and chat entries. A third design goal was to link the visualizations with the underlying log data in an interactive manner: when selecting a specific part of the interaction visualization with the mouse, the log data behind that component of the visualization would be rendered on the screen. This implied a redesign of the visualizations itself, as shown in Fig. 6.5.

    The former Wattle Tree is now replaced by a set of "swim lanes," one for each student in a team (in Fig. 6.5, that is area A, with three students S1, S2, S3, and one tutor, T; time is in days, running from bottom to top). Color is used to represent the type of contribution (wiki, ticket, svn), per day (or other time units) and aggregated over the visualized time period (B). When the user clicks a point in one of the swim lanes that has an activity indicated (i.e., is colored), the underlying log data for that cell will be rendered on the screen (C). Since this visualization is now fully integrated into Trac, the user can further drill down by following the links to trac objects. For instance, in Fig. 6.5(C) ticket change events are shown, and clicking on each of the links will bring the user to the respective ticket.

    A second line of work is addressing the limitations of the visualizations is their use of very simple measures of the numbers of lines contributed to the wiki or SVN and the gross actions on tickets. We have been exploring ways to use machine learning, clustering, and data mining to identify patterns which could augment the simple line count (Perera, Kay, Yacef, & Koprinska, 2007). Notably, our clustering
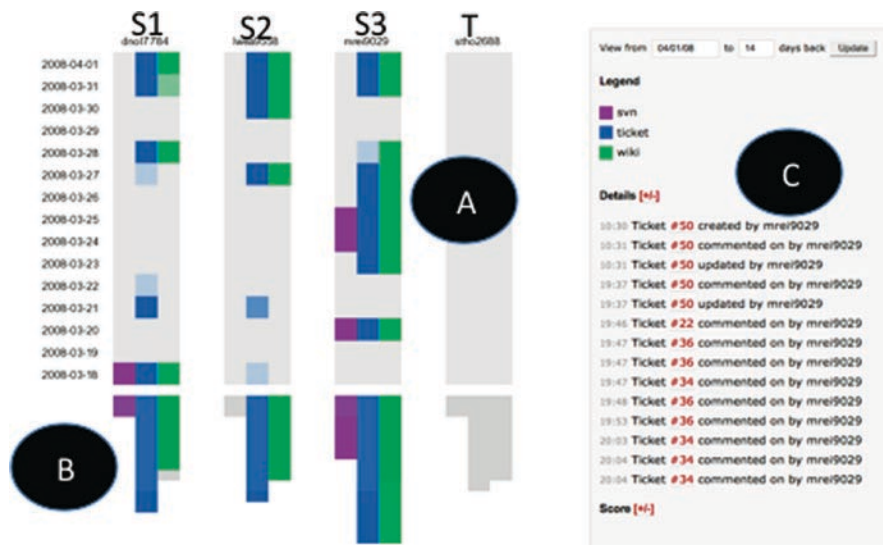
Fig. 6.5 Narcissus, an interactive form of interaction visualization. See text for explanations

approaches grouped individuals into some meaningful and valuable groups. For example, this gave one group that we would describe as having characteristics of a manager, with heavy use of ticketing and the wiki, based on a large collection of measures. Notably, successful groups had one such person and that person was the nominated manager. Problem groups had various other combinations, such as having several people with this cluster of behaviors with none being the nominated manager. We are exploring ways to use such data-mining approaches to provide additional mirroring information.

Data on students' interaction behavior are a rich source for mirroring and feedback, and particularly valuable when the learning goals comprise collaboration skills. But there is another source of information: the "product" that students generate in the course of their interactions. Mirroring, feedback, or guidance with respect to the group product will certainly be helpful for domain learning, but it is also important for the team process because, at the very least, teams need to know if/when the task is complete, and even better to what degree the task has been accomplished. For instance, if the goal is to develop a piece of software, then information on whether the software works and if the client is satisfied with the software is helpful. If the goal is to develop a written report, then information on the quality of writing and the satisfaction of the readers is helpful. In the following sections, we describe our first attempts to support teams with information on the features of their jointly authored wiki documents.

Visualizing Wiki Site Structure

The very nature of a wiki means that its structure and content will typically change over time. When the wiki is the main collaboration tool for learners, they need to be able to find the relevant parts of the wiki. On each return visit to the wiki, the learner needs to determine where to focus their attention. For example, they may need to determine where there have been changes made. If a person comes to a mature wiki project, with well-developed content, she or he needs to gain an overview of the wiki so that they can begin to get a sense of its extent and structure. A teacher is in a similar situation to the new visitor since they will typically visit intermittently and there may have been large changes in the wiki structure and content.

   Static web sites address some similar problems by providing a site map. Of course, for the case of a wiki, this would fail to account for the temporal issues. In a static site map, it is not usual to indicate the level of change that has occurred on parts of the site. Nor is there the need to indicate who made recent changes to parts of the site. However, in the case of a wiki, it will often be important to have such information. For example, if students use a wiki to write an essay collaboratively, a student may want to be able to see which parts have changed since they last visited the site. It may also be important to see which parts of the wiki were edited by a particular person, perhaps to monitor responses by others.

   WikiNavMap (Ullman & Kay, 2007) is an exploration of ways to support navigation in a wiki (see Fig. 6.6). It enables the user to customize the view of the wiki in terms of time and in relation to the authorship of activity on the pages. It aims to enable users to answer questions like these: Which are the pages that I have made contributions to? Which are the pages that another nominated person has made contributions to? Which are the pages associated with a certain task? Which are the pages with the most activity? Which pages changed in the last week? Which changed in a particular period of time, such as a particular month? What is the extent of the wiki? To do this, it provides a customization menu which enables the user to filter the pages, based on time and author. It also provides a complete overview, in a thumbnail and it presents a view of a larger version of the selected part of this, allowing the user to move this larger, viewed area. While these facilities give an overview, the user can see additional details via a mouse-over and then can click on the page to go to it, to see the full details of that page.

   WikiNavMap has been implemented as a plugin for Trac. As shown in Fig. 6.6, each rectangular box is a wiki page. The figure shows the page for a meeting. The user has their mouse over that box, causing the display of the information about recent actions on that page. The larger the box, the more activity there has been on it. The interface allows the user to control the shade of color to indicate the time. So, for example, the user could set the deepest color to show the pages which were last altered in the last week, the next deepest colour to show older pages, changed up to a month ago, and so on. Then, the largest boxes of a particular shade are the ones that had the most activity in the corresponding time period. Tickets are shown inside a box with a fine red line border, like the one at the bottom of the figure with
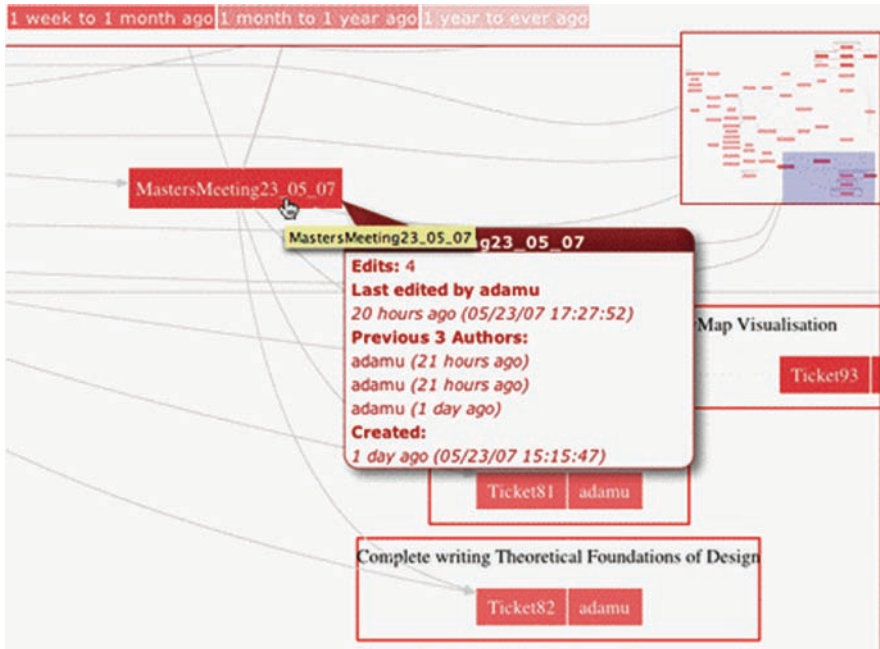
Fig. 6.6 WikiNavMap creates a dynamic visualization of a whole wiki site

Ticket 82. The tickets are grouped by their milestone, a notion supported by Trac to group tasks according to the higher-level goals of the group.

WikiNavMap plays both a navigational role, and also has the potential to increase member and task awareness (hence, affecting coordination), and helps to monitor coherence. If one filters the display to see the contributions of each team member in turn, it is easy to gain a sense of what each has contributed to the wiki. So, for example, one can set the colors to show work in the first, second, and third months of a 3-month project and then filter to show one team member. Then one can see which wiki pages and tickets this person completed work on in each month.

The display also shows lines from each page to other pages it links to. This are shown as very light lines in the display so as not to overwhelm the other information and because they can be very dense and complex. Such visualization of the hyperlinks between the wiki pages, which typically reflect semantic relations, provides the viewer not only with an idea of what topics are discussed on the pages, but also how they are related to each other. Thus, provided the links reflect content relations, the WikiNavMap visualization can be interpreted as a semantic network (Jonassen, Beissner, & Yacci, 1993) with directed, nonlabeled arcs.

Visualizing the Conceptual Structure of Wiki Page Content

In the third approach to supporting the use of wikis (as a paradigmatic collaborative writing technology) we focus directly on the conceptual content and the semantic relations. Providing information on what concepts are contained in a document and to what extent these are related to each other is helpful for individual writing, but it is particularly important for a collaboratively authored text. While an individual writer can be expected to know what concepts and ideas are in a document written by that author, on an active wiki site many changes will be made to a page, and it is hard for the writers to know at any point in time what is currently covered in the page.

In addition, both for the case of individual and collaborative writing, getting information on how the network of ideas and concepts contained in the text changes over time can be an important regulative for the writing process. In order to conceptualize such changes, one can build on a taxonomy suggested by Chi and Ohlsson (2005) for individual (declarative) learning:

- Knowing more on the same level of detail/abstractness;
- Increased density: new connection/relation identified;
- Increased consistency: errors/misconceptions identified and overcome, resulting in increased (local) consistence;
- Finer grain of representation: more details known, such as additional parts things are made of. Distinct from knowledge increment in as much as one moves down to identify which parts make up the thing/process described on the higher level;
- Greater complexity: integration of existing simpler ideas/theories/schemas;
- Higher level of abstraction, e.g., generalizing, conceptualizing;
- Shift in vantage point: a shift in perspective that allows us to see something in a new light, from a different angle;
- Identifying a dead-end line of inquiry/thought without being able to provide a better solution at this time.

Clearly, changes in conceptual knowledge are the hardest part to track and mirror automatically, even if we assume, as we do here, that text versions produced by students in the course of writing reflect changes in their declarative knowledge. Also, we assume that these types of changes are meaningful when applied to a jointly authored textual artefact – a document – not only when used to describe individual cognitive changes. Even when these assumptions hold, most of the forms of learning distinguished by Chi and Ohlsson require careful analysis on a semantic level. This cannot be accomplished computationally in full. However, techniques that are based on relations between text surface level and semantic level can be applied, and can support mirroring back to students some information about the knowledge contained in their (individually or collaboratively produced) texts. In the following, we illustrate how such an analysis can be performed, and what kind of information it yields.

Automap Analysis of Collaboratively Authored Wiki Pages

The automatic concept analysis method that we employed is based on Carley's map analysis technique (Carley, 1986, 1997) and the corresponding software, called Automap. This method is based on the assumption that peoples' mental models can be represented as concept maps (a variant of semantic networks), and that the mental models people have of a domain can be inferred from what people write about that domain. The map analysis method is predicated on the assumption that features of the text surface correspond to relations in the mental model: concepts that appear comparatively frequently in close proximity (window, e.g., within five words from each other) are treated as linked together in a statement (chain of interlinked concepts) in the mental model.

What counts as a concept needs to be defined by the analyst. Typically, and minimally, one would want to avoid treating certain words (such as the) as a concept, and one would want to treat singular and plural forms of a noun as the same concept. AutoMap provides the means to define concepts with (hierarchical) thesauri, thus accounting for concepts at multiple levels of generalization and abstraction.

Space prevents us from describing the technical details, so we confine ourselves to examples of how we use this method to provide information to students on the concept relations contained in their wiki contributions. The first example looks at a wiki page that has been coauthored by a number of students, on the topic of knowledge-building theory, based on their reading of Scardamalia and Bereiter (2006). This is one of the collectively authored wiki pages analyzed in Cai (2007), using a thesaurus of domain concepts for the learning sciences domain of approximately 300 entries. The length of the wiki analyzed was 663 words. Table 6.4 shows the basic parameters for the concept map based on the specific thesaurus.

There are 28 unique concepts and 110 unique statements in the map; the density of the map is 0.14; and the centrality of the map is 0.86. Density is calculated by dividing the number of identified links by the number of possible links. Centrality (of the map) reflects the extent to which a single concept has high centrality and the others low centrality, with a single concept's (in_degree) centrality defined as Total number of statements with concept in posterior position/Number of unique concepts per text. Figure 6.7 shows the map in a graphical format.

The value of such text statistics and displays becomes clearer when students can perform comparisons between wiki documents, as the absolute numbers do not provide much information in isolation. Comparisons are possible for instance between two pages on the same topic from different teams, or comparisons across versions. For instance, the same students produced a wiki page on another topic with more words (1,253), 55 unique concepts and 261 unique statements, with a density of 0.09 and a centrality of 0.62. In addition to comparing quantitative parameters of maps, visual inspection is useful, in particular for identifying information pertaining to specific concepts. For instance, a concept's relative centrality can be visually discerned in networks such as displayed in Fig. 6.7.

Table 6.4 Descriptive text statistics for a wiki page

Page name: "knowledge building"	
Page length	663 words
No. of concepts analyzed	
Total	105
Unique	28
No. of isolated concepts	
Total	2
Unique	1
No. of statements	
Total	162
Unique	110
Map density	0.14
Map centrality	0.86
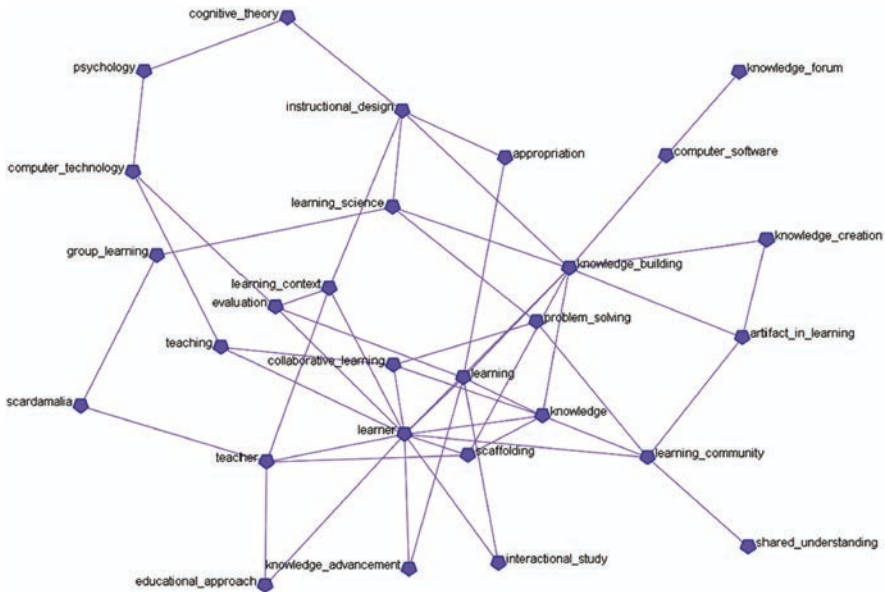
Fig. 6.7 A network view of the concepts identified in a jointly authored wiki page

Tracing a Document's Concept Structure Across Versions

An example for comparing across document versions, i.e., over time, is now illustrative as it sheds some light on the process of collaborative writing and knowledge construction. The wiki page analyzed here for the purpose of illustration went through 59 revisions, and reached a length of 4,128 words. The majority of the changes were done on the first day, 21st April, when the lecturer running the course set up a preliminary structure for the document. In the first few days the

majority of the versions were written (52), after which one change was written each on the 28th and 29th, following which was a break. On the 29th of May there was another change, and then there were four final changes on the 21st July. Most of the changes were written by the author A (12), followed by B (10) and C (9). The next five authors are responsible for around four to six changes, and three authors changed the wiki only once. This, however, does not correlate directly with the number of words contributed to the page. The largest number of words were, in fact, contributed by C (743) with the majority of students contributing an average of around 500 words. The mean word increase between documents is 70, with a standard deviation of 119. This illuminates the fact that a significant portion of the changes contributed less than 10 words, and a few others were very large, around 300–400 words increase.

It is illustrative to trace the development of this page across versions. As the wiki page grows, more concepts are added, and the concepts are used in statements with other concepts, this creating links in the concept analysis. However, with more concepts, the possible number of links between concepts grows, disproportionately to the actual number of links created. The density – defined as number of identified links/number of possible links – therefore drops as more versions added, despite the fact that some concepts have a large number of links to other concepts (see Fig. 6.8). This may be typical of a knowledge document, in which there is a large number of concepts used, and a limited opportunity to link them with other concepts. There were versions in which the density increased (between 5 and 6, 15 and 16, and 20 and 21). These revisions resulted in no or just one new concept
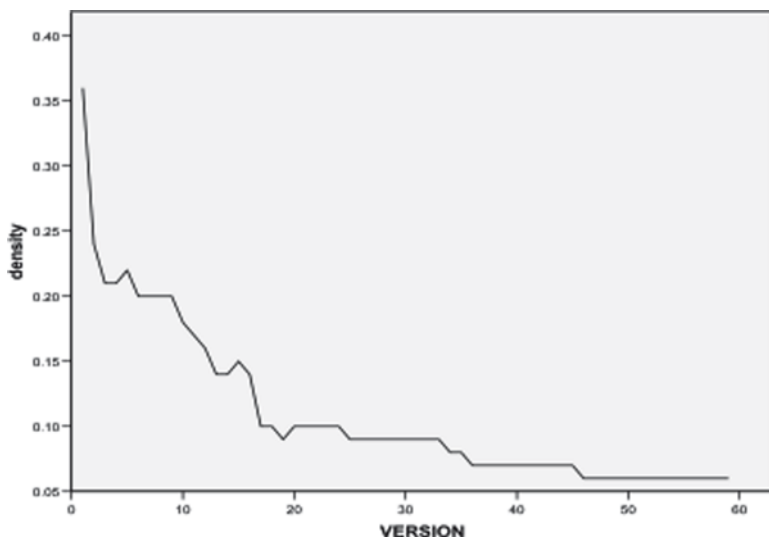

Fig. 6.8 Development of concept density over document versions

being introduced, but a substantial number of links was added. These could be instances of reevaluating, or enriching the understanding of, previously mentioned concepts. (Note that we speak rather loosely of "links added"; it needs to be kept in mind that it is not the students who are adding these links directly, but the numbers are based on an algorithm identifying such links given how students had been revising the text.)

Another useful indicator of coherence is the Variance of Concept Degree Centrality (VCDC). It measures the extent to which concepts are more central (more connected) than others:

$$VDVC = \frac{Sum(in\_degree - mean-in\_degree)^2}{No.\, of\, unique\, concepts}.$$

In looking at Fig. 6.9 we see that the VCDC value climbs sharply over time (versions). This is due to the fact that while a few concepts are very connected and central, others are very peripheral, particularly as more concepts added over time. Indeed, many of the concepts that have the most links in the last version were introduced in the first five revisions, while other concepts, with only one or two links, were introduced later in the document history. That being taken into account, this graph indicates that this particular wiki page is shaped on a few, key concepts, while mentioning a great many of peripheral ones.

It seems plausible that information such as provided in form of concept maps and/or concerning the development of concept map parameters such as density over
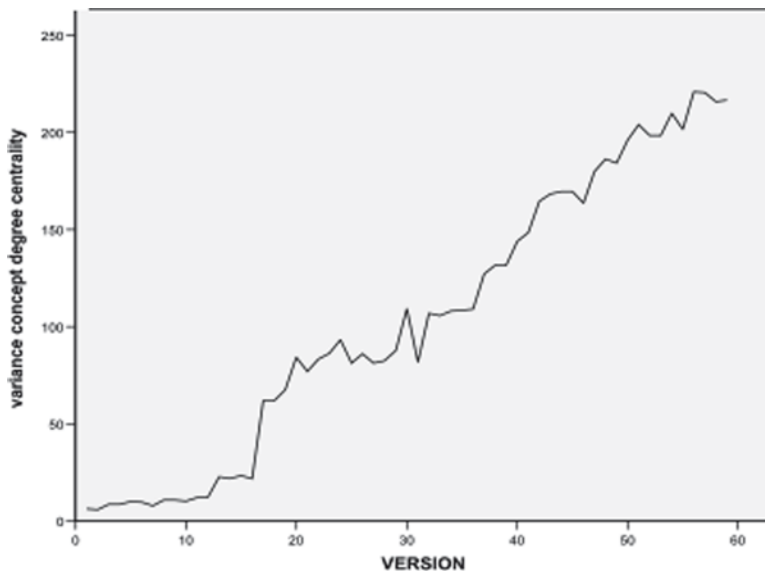

Fig. 6.9 Development of concept variance over versions

time can act as important regulators for a collaborative writing process, in particular concerning the coherence of the document. For individuals and groups to benefit from this information, the visualizations need to be made available for them, ideally on demand. Automap, as a stand-alone program, does not provide for this. We have, therefore, made the algorithm and visualization available online.

A Web-Based Program for Computing Concept Maps

An analysis similar to the Automap algorithm as described above is available on the Internet using Glosser, an online writing support environment developed at the University of Sydney (Villalon, Kearney, Calvo, & Reimann, 2008). Glosser uses text-mining techniques (based on Latent Semantic Analysis technique, Foltz, Kintsch, & Landauer, 1998) to provide student writers with information about their text on a number of dimensions, including conceptual coherence. A version of Glosser is integrated into Trac, another one can be used with Google Docs and Facebook (for more information, see http://www.weg.ee.usyd.edu.au/projects/glosser).

The Trac version of Glosser can be activated for any wiki page with a mouse click. The text of the page is then transferred to the Glosser server, which performs the text-mining analysis and renders the result to the students as shown in Fig. 6.10. In addition to the concept map visualization, Glosser currently provides information on text structure, text coherence, topics, as well as participation information. In addition, reflection questions are introduced for each of these aspects.

Discussion: Implications for Assessing Team Skills

In this chapter, we have described a number of approaches for supporting emergent collaboration, as distinct from designed collaboration, by providing mirroring and feedback information to groups in mainly graphical formats. Using wikis as a prototypical type of collaboration software, we have illustrated ways to visualize individuals' and groups' conceptual knowledge with automatically created concept maps and wiki site maps as well as various ways to visualize groups' work practices, e.g., with social network diagrams and the Wattle Tree visualization.

In the collaborative uses of wikis we studied, wikis sometimes played the role of the mediating artifact only – for instance when used by computer science students it was one of the means to develop the target activity (a software program) or by instructional design students to develop a course design – and sometimes wikis were both mediating artifact and target artifact at the same time –for instance when students create wiki pages that serve as research reports. In both cases, they were "activity-expanding" artifacts since they were intended for subsequent use: to guide

Fig. 6.10 Glosser web interface with reflection question on top and concept map on bottom

(later) action. In the case of students creating software, the software (and its design) was the tool for later use. In the case of students creating literature reviews and research reports, later use might be "knowledge building" activities (Scardamalia & Bereiter, 2003).

Reflecting further on the function of our concept map visualizations, created from students' individual and/or collective writing (see Fig. 6.7) and the wiki site

visualization (Fig. 6.6), one can see that these visualizations have two functions: a pragmatic and an epistemic one. Pragmatically, visualizations such as these can make it easier to manage the challenges of producing or maintaining documents with multiple authors (who can make changes in a quasi-parallel manner): they serve a techno-informational purpose and the purpose of substitutive coordination, in Zacklad's (2006) terminology. For instance, the graphical representation of a wiki site can be seen as an automatically created index of that site, where the index reflects modification activities by the authors.

At the same time, such visualizations have an epistemic function. The concept map in Fig. 6.7, for instance, can be seen as identifying the main ideas of the text that served as the source for the concept analysis, and of the connectedness of these ideas. It shows which concepts "go together" for the author or authors. In this sense, this kind of representation of text content is closer to Popper's (1972) "World 3," the world of ideas and concepts, than to the text surface (that would belong to "World 1," the physical world). This is not a representation that says anything about how well ideas are expressed in the text, but provides an answer to the question What is this text about? What do the authors use as the main concepts and how do they see them going together?

It needs to be mentioned that the concept maps created both with Automap (Carley, Diesner, & de Reno, 2006) as well as with Glosser (the web-based implementation) are not semantic nets: The links between the nodes are unlabeled and undirected. Analogous to Social Network Analysis (Wasserman & Faust, 1994), the basis for the visualization is cooccurrence information: the strength of the association between two concepts (expressed by the thinkness of the link, for instance) is solely determined by the frequency of the two concepts occuring in the same "window" (a certain number of words, a sentence, a paragraph, etc.). Nevertheless, this cooccurrence information can provide useful information about semantic relations as well, to the extent that the text surface reflects semantic relations – that things that go together are mentioned in close proximity to each other. In a more elaborated form, Shaffer and colleagues (Shaffer, Hartfield, Svarovsky, Nash, Nutley, Bagley et al., 2009) have used a similar approach to capture learners' "epistemic frames" and to track their development over time. The drawback with their method at this stage is that students' writings (and other textual data, such as transcripts from dialogs with mentors) need to be analyzed by trained human raters in order to identify if a certain element of the epistemic frame is realized or not. In contrast, our analysis works fully automatically.

The theory of Trialogic Learning (Paavola, Lipponen, & Hakkarainen, 2004) suggests that knowledge is to be found not only in peoples' head and the artifacts they create, but also in their practices: how they go about things. Learning in this respect means to become able to participate in practices, to become part of a community of practice for instance (Wenger, 1998). Visualizations of participation behavior (e.g., Fig. 6.2) can be seen as visualizing aspects of groups' practices. They thus can play a role in knowledge creation to the extent that such visualizations help groups to reflect on their practices, with a view to improving on them. We have at least some evidence that teams engage in such activities, from our

studies with software programming teams. To increase the frequency and depth of students' reflection on their team and work practices, we have begun to introduce example models of "good team work" into our groups, and work is under way to automatically identify students' work practices and to match them against these best-practice models.

Toward Assessing Team Practices and Artifacts

We have talked a lot about providing mirroring and feedback information, i.e., about formative assessment, but said nothing so far on summative feedback, on grading. While grading (and testing) may be of limited value to help with learning, they play an important role for evidence-based decision making on the level of schools and beyond, for placement and selection, and for large scale and long-term evaluations of curriculum reforms (Hickey, Suiker, Taasoobshirazi, Schafer, & Michael, 2006). Hence, any pedagogical or technological innovation needs eventually to be related to assessment in order to be integrated into an educational system (Fishman, Marx, Blumenfeld, Krajcik, & Soloway, 2004).

Again, computers can be used in various ways to help, both for assessment with a formative function (e.g., feedback on performance of a specific task) as well as a summative function (e.g., computer-based testing at the end of a school year). Computers are particularly well-suited for formative feedback, provided to the teacher and/or student directly contingent on performance – not test or exam performance, but problem-solving and decision-making performance, i.e., task performance. Such assessment is particularly important as it can affect learning while it is taking place (Shute, 2008).

Assessing group artifacts automatically is challenging, but can be done, in particular where the artifact has formal semantics. For instance, where students construct formal models such as Petri Nets (Reisig, 1985) as their products, it could be determined computationally if these nets are well-formed and able to produce the behavior required from the model. It is much harder to automatically evaluate artifacts of the computer program type with respect to their semantics (do they compute what they are supposed to compute?), but feedback on syntactic correctness can easily be provided. Moreover, current best practice, such as that advocated in Extreme Programming, require the programmer to create test cases before starting to write code. Such sets of tests can then be used for automated testing as the code development progresses, giving an ongoing form of formative feedback about the progress of the programming. If the artifacts take a textual form, a variety of methods for automatic essay scoring (Shermis & Burstein, 2003) can be employed. In particular methods that calculate similarities between documents can be used for both formative and summative feedback in a rather straightforward manner. For summative feedback, a reference solution needs to be provided in addition to students' essays. We can for instance use the text analysis methods described in the section on visualizing concept structures for assessment by calculating the similarity

between a student concept map and a map computed from a reference text (e.g., expert solutions).

Similarly, assessing group performance requires normative reference models of what constitutes "good teamwork," what processes characterize a good software team, for instance. In order to develop this line of thought a bit further, one can build on concepts developed in Intelligent Tutoring Systems (vanLehn, 2007) and in Evidence-centered Assessment Design (Mislevy, Steinberg, & Almond, 1999). Assessment here takes the form of updating a student model, a qualitative or quantitative representation of the skills and knowledge in terms of which one wants to make pedagogical or evaluative decisions. In the simplest, but most frequently used form, a student model is list of variables, of attribute–value pairs. The student model is constructed and maintained by calculating values (in the simplest case, quantitative values such as counts) for the variables in the student model based on observations about task performance. The relation between task performance and student model is mediated by an evidence model (Mislevy et al., 1999) that determines which aspects of students' performance to register (i.e., it defines event categories) and how to express the consequences of registering an event instance in terms of the student model (see Fig. 6.11). For instance, a variable may be increased when a certain student behavior is noted. Table 6.1 shows the relation between certain observable behaviors in the Trac collaboration environment and concepts of the Big Five framework for teamwork can be seen as forming the basis for a simple evidence model with the Big-5 concepts as the (latent) variables in the student model.

In state-of-the-art student-modeling approaches, task-related behavior is connected to student models using a Bayesian Net (Conati, Gertner, & vanLehn, 2002), thus accounting for the fact that the relation between latent variables and observable events is typically not a deterministic ("noise free") one. While this approach is elegant and computationally effective, it requires a careful analysis of the relation between events and variables, and it works best when the event categories and relations are not only known in advance, but also stay stable.

We have begun developing an approach that does not require such a detailed understanding and representation of the task domain. Instead of modeling students' capacities in a student model made out of variables, and calculating the value of variables based on performance observations, this approach works with a holistic, graphical model of team practices (Reimann, Frerejean, & Thompson, 2009). As illustrated in Fig. 6.12, a team practice (in this case, a decision-making process) can be represented as a formal process model (in this case, a transition diagram, see

Fig. 6.11 A conceptual assessment framework (adapted from Mislevy et al., 1999)
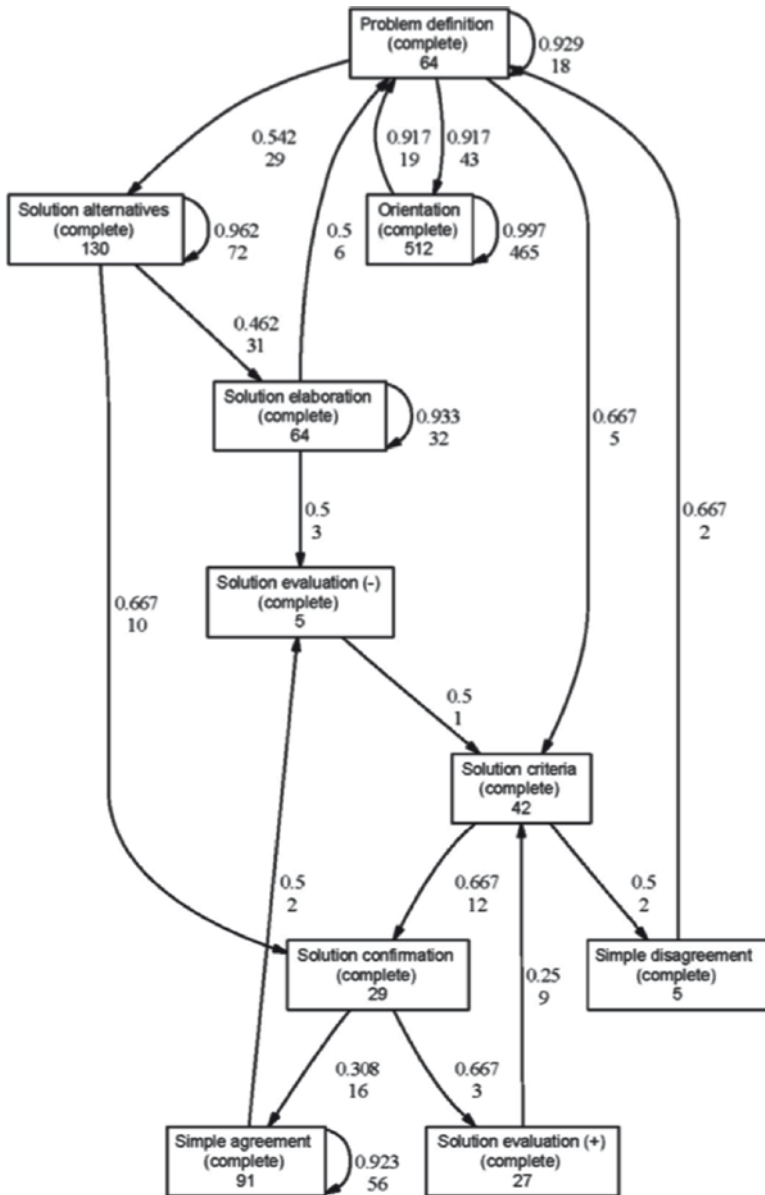
Fig. 6.12   A transition diagram model of a team decision process

Weijters, Van der Aalst, & Medeiros, 2006). To the extent that such models can be automatically identified based on observations (event logs), they become a basis for formative as well as summative assessment. Since methods of process mining (Van der Aalst & Weijters, 2005) provide the practical means for automatic process

modeling, this approach becomes practically feasible. For formative purposes, groups could receive process visualizations such as shown in Fig. 6.12 to reflect upon their practices. For assessment purposes, normative process models can be represented in the same format, and the similarity between the empirical and the normative models (calculated with standard algorithms for graph comparison) can form the basis for grading.

We end this chapter, which mainly dealt with issues of mirroring and feedback, with some thoughts on assessment because with few exceptions (notably Barros & Verdejo, 2000), assessment has not been in the focus of research on computer-supported interaction analysis (e.g., Bratitsis, Dimitracopoulou, Martínez-Monés, Marcos, & Dimitriadis, 2008). Given the role assessment plays on multiple levels of any educational system (Hickey et al., 2006), this led to the current situation where state-of-the-art methods and tools developed in research on computer-supported collaborative learning are disconnected from any mainstream educational assessment practices, despite the fact that collaborative learning forms an important part of educational policies and practices. Unfortunately, what students do in the course of their collaboration with peers does not relate to how they are assessed, and the outcomes of assessment rarely affect what they will do next. Further, considerable effort currently goes into developing individually administered tests, yet for addressing other twenty-first century "skills" such as those for collaborating and for communicating mediated by technology, there is seemingly little awareness of the potential of technology to capture students' team practices and to relating this information to dimensions relevant for assessment. It is our hope that professionals developing and implementing educational assessment methods will work much more closely with those researching technology-supported learning to move education toward twenty-first century assessment, which we believe will be (perhaps paradoxically) prerequisite for any meaningful realization of twenty-first century learning.

References

Albanese, R. & Fleet, D Dv. (1985). Rational behavior in groups: The free-riding tendency. *Academy of Management Review, 10*, 244–255.

Arrow, H., McGrath, J. E., & Behrdal, J. (2000). *Small groups as complex systems: Formation, co-ordination, development and adaptation*. Thousand Oaks, CA: Sage Publications.

Barros, B. & Verdejo, F. (2000). Analysing student interaction processes in order to improve collaboration. The DEGREE approach. *International Journal of Artificial Intelligence in Education, 11*, 221–241.

Beck, K. (1999). *Extreme programming explained: Embrace change*. New York: Addison-Wesley.

Blickensderfer, E. L., Cannon-Bowers, J. A., & Salas, E. (1997). Theoretical bases for team self-corrections: Fostering shared mental models. In D. J. M. M. Beyerlein & S. Beyerlein (Eds.), *Advances in interdisciplinary studies of work teams* (Vol. 4, pp. 249–279). Greenwich, CT: JAI.

Bodker, S. & Christiansen, E. (2006). Computer support for social awareness in flexible work. *Computer Supported Cooperative Work, 15*, 1–28.

Bonk, C. J. & Cunningham, D. J. (1998). Searching for learner-centered, constructivist, and socio-cultural components of collaborative educational learning tools. In C. J. Bonk & K. S. King (Eds.), *Electronic collaborators: Learner centered technologies for literacy, apprenticeship, and discourse* (pp. 25–50). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Bratitsis, T., Dimitracopoulou, A., Martínez-Monés, A., Marcos, J. A., & Dimitriadis, Y. (2008). Supporting members of a learning community using Interaction Analysis tools: The example of the Kaleidoscope NoE scientific network. In P. Diaz, Kinshuk, I. Aedo & E. Mora (Eds.), *The 8th IEEE International Conference on Advanced Learning Technologies (ICALT 2008)* (pp. 809-813). Washington, DC: IEEE Computer Society.

Buller, P. F. & Bell, C. H. (1986). Effects of team building and goal setting on productivity: A field experiment. *Academy of Management Journal, 29*, 305–328.

Cai, C. Y. (2007). *Extracting students' knowledge structures form their written texts through Automap approach.* Unpublished Master Thesis, University of Sydney, Sydney, Australia.

Cannon-Bowers, J. A. & Salas, E. (eds). (1998). *Making decisions under stress: Implications for individual and team training*. Washington, DC: American Psychological Association.

Cannon-Bowers, J. A., Tannenbaum, S. I., Salas, E., & Volpe, C. E. (1995). Defining competencies and establishing team training requirements. In E. S. R. A. Guzzo (Ed.), *Team effectiveness and decision making in organisations* (pp. 333–380). San Francisco, CA: Jossey-Bass.

Carley, K. M. (1986). An approach for relating social structure to cognitive structure. *Journal of Mathematical Sociology, 12*(2), 137–189.

Carley, K. M. (1997). Extracting team mental models through textual analysis. *Journal of Organizational Behavior, 18*, 533–558.

Carley, K. M., Diesner, J., & de Reno, M. (2006). *AutoMap user's guide*. Pittsburgh, PA: Carnegie-Mellon University. School of Computer Science. Institute for Software Research. Center for Computational Analysis of Social and Organizational Systems (CASOS).

Chi, M. T. H. & Ohlsson, S. (2005). Complex declarative learning. In K. Holyoak & R. G. Morrison (Eds.), *The cambridge handbook of thinking and reasoning* (pp. 371–399). Cambridge: Cambridge University Press.

Clark, H. H. & Brennan, S. E. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine & S. D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 127–149). Washington: American Psychological Association.

Conati, C., Gertner, A., & vanLehn, K. (2002). Using Bayesian networks to manage uncertainty in student modelling. *User Modelling and User-Adapted Interaction, 12*(4), 1573–1391.

Cross, G. A. (2001). *Forming the collective mind: A conceptual exploration of large-scale collaborative writing in industry*. Cresskill, NJ: Hampton.

Dönmez, P., Rosé, C., Stegmann, K., Weinberger, A., & Fischer, F. (2005, May 30-June 04). *Supporting CSCL with automatic corpus analysis technology.* Paper presented at the 2005 Conference on Computer Support For Collaborative Learning: Learning 2005: the Next 10 Years! Computer Support for Collaborative Learning, Taipei, Taiwan.

Dyer, J. C. (1984). Team research and team training: State-of-the-art review. In F. A. Muckler (Ed.), *Human factors review* (pp. 285–323). Santa Monica, CA: Human Factors Society.

West Virginia Department of Education (2007). *21st Century learning skills and technology tools: Content standards and objectives for WV Schools.* Available from http://wvde.state.wv.us/policies/p2520.14_ne.pdf

Engeström, Y. (1987). *Learning by expanding: An activity-theoretical approach to developmental research*. Helsinki: Orienta-Konsultit Oy.

Erkens, G., Jaspers, J., Prangsma, M., & Kanselaar, G. (2005). Coordination processes in computer supported collaborative writing. *Computers in Human Behavior, 21*, 463–486.

Fishman, B., Marx, R. W., Blumenfeld, P., Krajcik, J., & Soloway, E. (2004). Creating a framework for research on systemic technology innovations. *The Journal of the Learning Sciences, 13*(1), 43–76.

Foltz, P. W., Kintsch, W., & Landauer, T. K. (1998). The measurement of textual coherence with latent semantic analysis. *Discourse Processes, 25*(2–3), 285–307.

Fussell, S. R., Kraut, R. E., Lerch, F. J., Scherlis, W. L., McNally, M. M., & Cadiz, J. J. (1998). *Coordination, overload and team performance: Effects of team communication strategies*. Paper presented at the Proceedings of the 1998 ACM conference on Computer supported cooperative work, Seattle, WA

Gillies, R. (2004). The effects of cooperative learning on junior high school students during small group learning. *Learning and Instruction, 14*, 197–213.

Gillies, R. & Ashman, A. (1996). Teaching collaborative skills to primary school children in classroom-based work groups. *Learning and Instruction, 6*, 187–200.

Hayes, J. R. & Flower, L. S. (1980). Identifying the organization of the writing process. In L. W. Gregg & E. R. Steinberg (Eds.), *Cognitive processes in writing* (pp. 3–30). Hillsdale, NJ: Erlbaum.

Hickey, D. T., Suiker, S. T., Taasoobshirazi, G., Schafer, N. J., & Michael, M. A. (2006). Balancing varied assessment functions to attain systemic validity: Three is the magic number. *Studies in Educational Evaluation, 32*, 180–201.

iN2015 Education and Learning Sub-Committee. (2007). *Empowering learners and engaging minds, through infocomm*. Singapore: Ministry of Education.

Johnson, D. W., Johnson, R. T., Stanne, M. B., & Garibaldi, A. (1989). Impact of group processing on achievement in cooperative groups. *Journal of Social Psychology, 130*, 507–516.

Jonassen, D. H., Beissner, K., & Yacci, M. (1993). *Structural knowledge: techniques for representing, conveying, and acquiring structural knowledge*. Hillsdale, NJ: Lawrence Erlbaum.

Kapur, M. (2006). *Productive failure: A hidden efficacy of seemingly unproductive production*. Paper presented at the Cognitive Science Conference, Vancouver, Canada.

Kapur, M., Hung, D., Jacobson, M. J., Voiklis, J., Kinzer, C. K., & Victor, C. D.-T. (2007). Emergence of learning in computer-supported, large-scale collective dynamics: A research agenda. *Proceedings of the International Conference on Computer-supported Collaborative Learning (CSCL2007)*, New Brunswick, NJ.

Kay, J., Maisonneuve, N., Yacef, K., & Reimann, P. (2006). The big five and visualisations of team work activity. In M. Ikeda, K. D. Ashley & T. Chan (Eds.), *Intelligent tutoring systems*. *Proceedings of the 8th International Conference, ITS 2006* (pp. 197-206). Jhongli, Taiwan: Springer.

Kirschner, P. A., Sweller, J., & Clark, R. E. (2006). Why minimal guidance during instruction does not work: An analysis of the failure of constructivist, discovery, problem-based, experiental, and inquiry-based teaching. *Educational Psychologist, 41*(2), 75–86.

Kollar, I., Fischer, F., & Hesse, F. W. (2006). Collaboration scripts - a conceptual analysis. *Educational Psychological Review, 18*, 159–185.

Kolodner, J. L., Camp, P. J., Crismond, D., Fasse, B., Gray, J., Holbrook, J., et al. (2003). Problem-based learning meets case-based reasoning in the midlle-school science classroom: Putting learning by design into practice. *The Journal of the Learning Sciences, 12*(4), 495–547.

Kozlowski, S. W. J., Brown, K. G., Weissbein, D. A., Cannon-Bowers, J. A., & Salas, E. (2000). A multi-level perspective on training effectiveness: Enhancing horizontal and vertical transfer. In S. W. J. Kozlowski & K. J. Klein (Eds.), *Multilevel theory, research, and methods in organizations* (pp. 157–210). San Francisco, CA: Jossey-Bass.

Kozlowski, S. W. J. & Ilgen, D. R. (2006). Enhancing the effectiveness of work groups and teams. *Psychological Science in the Public Interest, 7*(3), 77–124.

Kreijns, K., Kirschner, P., & Jochems, W. (2003). Identifying the pitfalls for social interaction in computer-supported collaborative learning environments: A review of the research. *Computers in Human Behavior, 19*(3), 335–353.

Layman, L., Williams, L., Damian, D., & Bures, H. (2006). Essential communication practices for extreme programming in a global software development team. *Information and Software Technology, 48*, 781–794.

Lee, E. Y. C., Han, C. K. K., & van Aalst, J. (2006). Students assessing their own collaborative knowledge building. *International Journal of Computer-Supported Collaborative Learning, 1*(1), 57–88.

Leuf, B. & Cunningham, W. (2001). *The wiki way*. Englewood Cliffs, NJ: Prentice Hall.

Lowry, P. B., Curtis, A., & Lowry, M. R. (2004). Building a taxonomy and nomenclature of collaborative writing to improve interdisciplinary research and practice. *Journal of Business Communication, 41*(1), 66–99.

Lowry, P. B., Nunamaker, J. F., Jr., Curtis, A., & Lowry, M. R. (2005). The impact of process structure on novice, virtual collaborative writing teams. *IEEE Transactions on Professional Communication, 48*(4), 341–364.

MacMillan, J., Entin, E. E., & Serfaty, D. (2004). Communication overhead: the hidden cost of team cognition. In E. Salas & S. M. Fiore (Eds.), *Team cognition* (pp. 61–82). Washington, DC: American Psychological Association.

McGrath, J. E. (1991). Time, interaction, and performance (TIP). *A Theory of Groups. Small Group Research., 22*(2), 147–174.

McGrath, J. E. & Tschan, F. (2004). *Temporal matters in social psychology: Examining the role of time in the lives of groups and individuals*. Washington, DC: American Psychological Association.

Mislevy, R. J., Steinberg, L., & Almond, R. G. (1999). *Evidence-centered assessment design.* Available from http://www.education.umd.edu/EDMS/mislevy/papers/ECD_overview.html

Olson, G. M., Malone, T. W., & Smith, J. B. (eds). (2001). *Coordination theory and collaboration technology*. Mahwah, NJ: Lawrence Erlbaum.

Paavola, S. & Hakkarainen, K. (2005). The knowledge creation metaphor - An emergent epistemological approach to learning. *Science & Education, 14*, 535–557.

Paavola, S., Lipponen, L., & Hakkarainen, K. (2004). Models of innovative knowledge communities and three metaphors of learning. *Review of Educational Research, 74*(4), 557–576.

Perera, D., Kay, J., Yacef, K., & Koprinska, I. (2007). *Mining learners' traces from an online collaboration tool.* Paper presented at the Educational Data Mining Workshop at Artificial Intelligence and Education, Marina del Ray, CA.

Phillips, N., Lawrence, T. B., & Hardy, C. (2004). Discourse and institutions. *Academy of Management Review, 29*(4), 635–652.

Popper, K. (1972). *Objective knowledge: An evolutionary approach*. Oxford, UK: Clarendon Press.

Prichard, J. S. & Ashleigh, M. J. (2007). The effects of team-skills training on transactive memory and performance. *Small Group Research, 38*, 696–726.

Prichard, J. S., Startford, R. J., & Bizo, L. A. (2006). Team-skills training enhances collaborative learning. *Learning and instruction, 16*(3), 256–265.

Prichard, J. S., Stratford, R. J., & Hardy, C. (2004). *Training students to work in teams: Why and how?* York, UK: LTSN Psychology.

Reimann, P. (2003). How to support groups in learning: More than problem solving. In V. Aleven et al. (Eds.), *Artificial Intelligence in Education (AIED 2003). Supplementary Proceedings* (pp. 3-16). Sydney: University of Sydney.

Reimann, P. (2007). Time is precious: Why process analysis is essential for CSCL. In C. Chinn, G. Erkens & S. Puntambekar (Eds.), *Minds, minds, and society*. *Proceedings of the CSCL 2007 conference* (Vol. 8, pp. 598-607). New Brunswick, NJ: ISLS.

Reimann, P., Frerejean, J., & Thompson, K. (2009). Using process mining to identify models of group decision making processes in chat data. In C. O'Malley, D. Suthers, P. Reimann & A. Dimitracopoulou (Eds.), *Computer-supported collaborative learning practices: CSCL2009 conference proceedings* (pp. 98-107). International Society for the Learning Sciences.

Reinhold, S. (2006). *WikiTrails: augmenting Wiki structure for collaborative, interdisciplinary learning.* Paper presented at the 2006 international symposium on Wikis, Odense, Denmark

Reisig, W. (1985). *Petri nets. An introduction*. Berlin: Springer.

Resta, P. & Laferriere, T. (2007). Technology in support of collaborative learning. *Educational Psychology Review, 19*(1), 65–83.

Ripochet, G. & Sansonnet, J.-P. (2006). Experiences in automating the analysis of linguistic interactions for the study of distributed collectives. *Computer Supported Cooperative Work, 15*, 149–183.

Salas, E., Rozell, D., Mullen, B., & Driskell, J. E. (1999). The effect of team building on performance: An integration. *Small Group Research, 30*, 309–329.

Salas, E., Sims, D. E., & Burke, C. S. (2005). Is there a "Big Five" in teamwork? *Small Group Research, 36*(5), 555–599.

Scardamalia, M. & Bereiter, C. (1991). Higher levels of agency for children in knowledge building: A challenge for the design of new knowledge media. *The Journal of the Learning Sciences., 1*(1), 37–68.

Scardamalia, M. & Bereiter, C. (2003). Knowledge building. In J. W. Guthrie (Ed.), *Encyclopedia of education* (2nd ed.). New York: Macmillan.

Scardamalia, M. & Bereiter, C. (2006). Knowledge building: Theory, pedagogy, and technology. In R. K. Sawyer (Ed.), *The Cambridge handbook of the learning sciences* (pp. 97–115). New York: Cambride University Press.

Scott, J. (1991). *Social network analysis: A handbook*. London: Sage.

Sfard, A. (1998). On two metaphors of learning and the dangers of choosing just one. *Educational Researcher, 27*(2), 4–13.

Shaffer, D. W. (2008). Education in the digital age. *Digital Kompetanse, 3*(1), 37–50.

Shaffer, D. W., Hatfield, D., Svarovsky, G. N., Nash, P., Nutly, A., Bagley, E., et al. (2009). Epistemic Network Analysis: A prototype of 21st-Cenntury assessment of learning. International Journal of Learning and Media, 1, 33–53.

Shermis, M. D. & Burstein, J. C. (eds). (2003). *Automated essay scoring: A cross-disciplinary perspective*. Hillsdale, NJ: Lawrence Erlbaum.

Shute, V. J. (2008). Focus on formative feedback. *Review of Educational Research, 78*(1), 153–189.

Soller, A., Martínez-Monés, A., Jermann, P., & Muehlenbrock, M. (2005). From Mirroring to Guiding: A Review of State of the Art Technology for Supporting Collaborative Learning (pre-print). *International Journal of Artificial Intelligence in Education, 15* (4), 261–290.

Souzis, A. (2005). Building a semantic wiki. *IEEE Intelligent Systems, 20*, 87–91.

Stahl, G. (2006). Sustaining group cognition in a math chat environment. *Research and Practice in Technology Enhanced Learning, 1*(2), 85–113.

Topping, K. (1998). Peer assessment between students in colleges and universities. *Review of Educational Research, 68*(3), 249–276.

Ullman, A. & Kay, J. (2007). *WikiNavMap: A visualisation to supplement team based wikis Conference on Human Factors in Computing Systems, CHI '07 extended abstracts on Human factors in computing systems* (pp. 2711–2716). San Jose, CA: ACM Press.

van Amalesvoort, M., Andriessen, J., & Kanselaar, G. (2007). Representational tools in computer-supported colloborative argumentation-based learning: How dyads work with constructed and inspected argumentative diagrams. *Journal of the Learning Sciences, 16*(4), 485–521.

Van der Aalst, W. M. P. & Weijters, A. J. M. M. (2005). Process mining. In M. Dumas, W. M. P. van der Aalst & A. H. M. ter Hofstede (Eds.), *Process-aware information systems: Bridging people and software through process technology* (pp. 235–255). New York: Wiley.

vanLehn, K. (2007). Intelligent tutoring systems for continous, embedded assessment. In C. A. Dwyer (Ed.), *The future of assessment: shaping teaching and learning*. Mahaw, NL: Erlbaum.

Villalon, J., Kearney, P., Calvo, R. A., & Reimann, P. (2008, July 1-5). *Glosser: Enhanced feedback for student writing tasks.* Paper presented at the 8th IEEE International Conference on Advance Learning Technologies (ICALT), Santadar, Spain.

Wasserman, S. & Faust, K. (1994). *Social network analysis: methods and applications*. Cambridge, UK: Cambridge University Press.

Wegner, D. M. (1986). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of group behavior* (pp. 185–205). New York: Springer.

Weijters, A. J. M. M., Van der Aalst, W. M. P., & Medeiros, A. K. A. d. (2006). *Process mining with the heuristics miner-algorithm. BETA Working Paper Series WP 166*. Eindhoven, NL: Eindhoven University of Technology.

Wenger, E. (1998). *Communities of practice: Learning, meaning, and identity*. Cambridge: Cambridge University Press.

Williams, K., Calvo, R. A., & Bell, D. (2003, July 20-24). *Autmoatic categorization of questions for a mathematics education service*. Paper presented at the 11th World Conference on Artificial Intelligence in Education, Sydney, Australia.

Yager, S., Johnson, R. T., Johnson, D. W., & Snider, B. (1986). The impact of group processing on achievement in cooperative learning groups. *Journal of Social Psychology, 126*, 389–397.

Zacklad, M. (2004). *Processus de documentarisation dans les Documents pour l'Action (DopA): statut des annotations et technologies de la coope´ration associe´es.* In Le nume´rique : Impact sur le cycle de vie du document pour une analyse interdisciplinaire. Available from http://archivesic.ccsd.cnrs.fr/sic_00001072.html

Zacklad, M. (2006). Documentarisation processes in documents for action (DofA): The status of annotations and associated cooperation technologies. *Computer Supported Cooperative Work, 15*, 205–228.

Zumbach, J., Hillers, A., & Reimann, P. (2003). Supporting distributed problem-based learning: The use of feedback in online learning. In T. Roberts (Ed.), *Online collaborative learning: theory and practice* (pp. 86–103). Hershey, PA: Information Science Publishing.