Characterising knowledge construction through a process analysis of dialogues

Michael Tscholl, John Dowell
Department of Computer Science, University College London
Malet Place, WC1E 6BT London
e-mail: m.tscholl@cs.ucl.ac.uk, j.dowell@cs.ucl.ac.uk

Abstract: We present an analysis of discourse in case-centred learning. The analysis answers questions about how an abstract representation of a case is constructed through discourse, and about what cognitive products result from that construction. The analysis relies on a generic learning model that supports the coding and quantification of learning discourse. This analysis is demonstrated through its application to a set of dialogues taken from 2 groups of medical students who, as part of their professional training, were asked to explain the medical ethics engendered by two patient cases. The analysis shows that the learners strengthen their existing conceptual knowledge, rather than acquire new conceptual knowledge, and specifically that they make valuable new connections between structurally similar episodes, and between concepts and specific facts of the case. We assess the value of an analysis of cognitive processes for characterising collaborative case-centred discussions, and its use in showing differences of processing of information in different learning environments.

Introduction

In the discussion that lies at the core of learning by shared problem solving, learners construct a representation of the problem that is essentially virtual and occasionally real. Learners use their individual and shared knowledge, and they acquire knowledge through the construction. In this paper we examine how this virtual representation is constructed, and what knowledge learners acquire in the process.

We make the assumption that fundamentally the same cognitive processes occur in both individual learners and in groups of learners (cf. Dillenbourg, 1996). We therefore take the further step of suggesting that what is known about how individual learners learn from problem solving can be applied to a group of collaborative learners. Such processes as abstraction, generalisation or association that have been well established in studies of individual learners may equally occur with groups of learners. We present here a generalised model of learning from problems that integrates prior models and findings for application to both individuals and groups.

We make a further, methodological assumption that these cognitive processes of learning are reflected systematically in group dialogues and are amenable to analysis. In other words, we assume that people say what they think, and that what they think and know can be determined from what they say. The purpose of the generalised learning model is to direct this analysis of dialogues. Other recent work has similarly relied on an analysis of the cognitive processes evident in utterances to investigate collaborative learning. Moreover the general notion of co-construction, as one of the core concepts in constructivism, entails that utterances can be seen as constructing knowledge in a group learning situation: learners add knowledge, elaborate knowledge, analyse knowledge, etc. (e.g. Davies, 2003). However, this prior work does not make a commitment to *what is constructed* by these processes. Rather, the interest of this research is in investigating which processes occur in order to derive conclusions about the general quality of a discussion (cf. Hara et al., 2000; Newman et al., 1995) Other work has focussed on the knowledge and information processed in a group learning situation, exploiting the concept of critical thinking (Garrison et al., 2000); however the aim was to gauge how much a group learning environment supports a community of inquiry of which critical thinking is a core component.

Our approach posits a relation between processes and products, that is, it speculates about what knowledge is constructed by specific cognitive processes. We focus especially on how the correct abstract representation of the problem is constructed, and describe which processes are employed during the construction of that representation. This description allows us to characterise a learners' discussion on a case in terms of processes.

Learning as abstracting a structure

Learning with cases and problems is fundamentally the abstraction of structure; it is the recognition of deep features that give meaning to the situation. Cognitive science has consistently viewed learning in this way, particularly in relation to episode-based problem solving (Gick & Holyoak, 1980; Forbus, 2001). Our aim is to discover which cognitive processes occur when people attempt to solve problems and which processes give rise to the structure of a case. To this end we first compiled a generalised model of learning from cases,

characterising what is learnt and how. The model delineates a set of cognitive processes and a set of cognitive products (most notably the structure of the case) that are likely to occur as a result of the tension between the case (novel information) and existing knowledge. The model is summarised in the next section and is the basis for identifying learning processes occurring when people discuss cases.

A model of learning from cases

Learning with cases and problems occurs through two main processes: first, people use their existing knowledge to interpret and conceptualise the case, refining their knowledge in the process; second, they may abstract a structure from the case. These two processes are intertwined: using existing knowledge to interpret the case entails recognising new relationships within that knowledge, relationships that then represent the structure of the case. Even without acquiring structural knowledge, cases can promote learning by refining the existing knowledge because cases, consisting of specific concrete aspects, function implicitly as tests of existing knowledge. For example, people may ask whether some conceptual knowledge is really useful to interpret the case, and by verifying the applicability of that knowledge, they may learn more about when to apply that knowledge. Clearly, the process of verifying the applicability of knowledge is most pertinent with conceptual knowledge, but also models may be tested.

Research on learning from problems and examples has identified more specific processes of learning. In mathematics learning, for example, students acquire knowledge about the conditions of applicability of production rules by adding the specific features of an example to the condition part of that rule. A similar form of learning occurs when conceptual knowledge is used to interpret cases, leading to learning about the conditions under which specific concepts should be used. However, learning of concepts also entails recognising the structure of problem features, a pattern or constellation of concrete problem aspects that is representative for a specific type of problem. The literature on this form of learning is extensive, and a common theme is that experts possess a large set of domain-specific examples acquired through practice that permits them to see the structure of a problem situation beyond its surface features. Chi et al.'s (1981) research on experts vs. novices led to the claim about the role of a conceptual system (concepts related to other concepts) in problem solving. This research emphasised that learning occurs through practice which teaches how to apply concepts to examples, how to relate concepts and how to abstract structural (relational) knowledge from examples.

The research on example-based problem solving has identified generalisation over specific features of instances as the core learning mechanisms, by which those cognitive products found in the research cited above are build. Reusing an earlier example to solve a current one forces a generalisation over the two: the concrete features are abstracted and represented by more general knowledge (Ross & Kennedy, 1990). The role of earlier solved problems (memorised problem solving episodes) is hence crucial and explains why it is that expert-type knowledge is generated through practice and experience.

When earlier episodes are used, a problem may be solved through analogy, and people may abstract the structure common both to the current and the past problem (Gick & Holyoak, 1980). However, analogies are computationally complex, and demand some structural knowledge about the solved problem. As a result, though analogies may occur, they occur infrequently.

While the reuse of episodes is frequent in problem solving, novices are more likely to use conceptual knowledge because they still do not possess an adequate knowledge base of episodes. However when episodes are available and are used, conceptual knowledge is used at several stages of reuse: at retrieval, for the identification of similarities between the current and memorized episodes, and when transferring structural knowledge from the past to the current problem. In these stages, conceptual knowledge is used, but may also be acquired. Knowledge abstracted in this way constitutes domain-specific knowledge in the form of principles or 'points of the story'. Individual case aspects that have been found also provide indices for the organization of episodes in memory (Kolodner, 1993) and become part of the domain-specific vocabulary. In future situations, students will rely on this vocabulary to encode new cases. Table 1 presents the learning processes and associated learning products.

Identifying processes in dialogues

A small number of recent dialogue analyses have examined the latent content of utterances to answer questions about learning. Garrison et al. (2000) interpret utterances as processes to gauge the quality of learning. They explain their focus on this cognitive dimension of dialogues by contrasting it with "other dimensions, such as the participatory or social dimension, (that) lead to measures of quantity, not quality and tell us about motivations and conversations rather than the learning taking place" (Newman et al., 1995). Clearly this work assumes that learning manifests itself in dialogue in a systematic way amenable for analysis. By characterising learning, our model allows us to derive a distinct coding scheme allowing us to recognise learning in dialogues in terms of processes and products.

Table 1: processes and products of learning with cases.

	processes	description	products
1	relating aspects and concepts	emphasising that a concept applies because of a specific aspects	concept instantiation
2	relating concepts	implying that concepts are related (e.g. are similar, depend on each other, are in opposition)	domain-specific conceptual system
3	relating aspects	relating two or more aspects. Relations are normally causal, but can also be of other type (e.g. linear)	new structured episode
4	relating cases (from learning material or personal or public episodes)	saying that two cases are similar or different; or implying that a similarity or difference between cases has been identified	knowledge base of cases
5	abstracting similarities/ differences	emphasising an aspect that makes two cases similar or different	encoding knowledge
6	abstracting principles	providing a description of the case that captures its point, it story	encoding knowledge; problem types
7	analogies	using another case to interpret a current case	knowledge base of cases; encoding knowledge; domain-specific conceptual knowledge; new structured episode

A process coding scheme

The coding scheme deriving from the model is largely the description of what constitutes evidence for an element of learning. For example, to assign the category 'relating cases' to a proposition, specific indicators must be present in the proposition. The specification of the coding categories is a list and a description of these indicators. The presence of one or more indicators in a proposition is the basis upon which a categorisation of the proposition is made.

The process-coding scheme uses two kinds of indicators: *reference* and *keywords/key phrases*. To describe the indicators, a short review of the basic components of the model of learning is helpful. The model describes learning as a set of processes that generate a set of learning products. Learning processes operate on knowledge and information and hence the identification of information and knowledge is a first step towards the identification of the processes. For example, if a student draws attention to the similarity between two cases ("this situation is similar to [another one]", the proposition mentions the two "situations", and establishes a relationship of similarity between them ("…is similar to…").

However, to capture not only explicit mentioning of specific facts, concepts or cases, *references* to facts, concepts or cases are used to identify a unit where a learning process may be evident. A reference is a surface expression that points to a concrete aspect (fact), a case or a concept that employs a term that is different from the one used in the cases and lecture material.

The second type of indicator is keywords (and key phrases). Keywords are frequently used in content analysis schemes because they are reliable and objective indicators. Some research (cf. Chi, 1997) used keywords such as 'because' and 'so' to search for 'explanations'. The presence of such keywords is an indication that the proposition should be classified as an explanation. However keywords are found relatively rarely. That is, mental representations and mental phenomena can be identified only rarely with keywords.

The keywords relevant in this coding scheme are used to identify the processes of relationing. Specifically, keywords indicating a comparison ("is similar", "is different") are used to identify propositions where cases are related to cases. Keywords indicating an explanation ("because", "so") are used to identify the relationing of aspects with concepts ('this concept applies because of this fact'; e.g. "she is free to chose because she is of sound mind"). A last set of keywords and key phrases is used as an indicator for relationing concepts as 'contrasting' or 'being in opposition'. These are keywords indicating contrasts, such as 'but' (for example, "they cared about her welfare, but restricted her freedom"). Table 2 reproduces the cognitive processes and the indicators required for their identification.

Table 2: processes and indicators for their identification in utterances.

#	processes (coding category)	primary indicators	secondary indicators
1	relating aspects and concepts	reference to an aspect	reference to a concept; indication of relationing
2	relating concepts	reference to a concept	reference to a concept; indication of relationing
3	relating aspects	reference to an aspect	reference to an aspect; indication of relationing
4	relating cases	reference to other cases	reference to another case; keywords: "similar", "different", "is like"
5	abstracting similarities/ differences	reference to other cases	reference to other cases; reference to an aspect; keywords: "similar", "different", "is like"
6	abstracting principles	reference to whole case	reference to some abstract principle/type
7	analogies	reference to other cases	use of its structure (interpretation/decision/ solution)

Characterising small-group case-centred discussions

We observed two small groups of students taking a course on professional skills development within the undergraduate medical programme at UCL. The groups were given two clinical scenarios in each of which a profoundly ill patient wished to end their life. During the discussions, the students develop a shared explanation for each scenario that would take into account the relevant concepts and principles of medical ethics facing the clinicians treating the patients.

Two landmark cases (Mrs. B and Mrs. Pretty) were chosen for the sessions. The cases are similar in all aspects with the crucial difference, however, that Mrs. B is able end her own life, while the physical condition of Mrs. Pretty would require the active intervention of a physician (active versus passive euthanasia). Understanding the similarities and differences between the two cases is of major importance to understanding the domain of medical ethics.

Euthanasia cases are normally understood on the basis of several general ethical principles: autonomy (the right to chose), nonmaleficiance (the prohibition for physicians to harm patients) and beneficence (the duty to help and do good). Euthanasia cases raise ethical issues because these principles often are in conflict with each other. For example, a patient requesting euthanasia would emphasise that it is her right to do so under the principle of autonomy, while a physician would rely on the principle of nonmaleficiance to deny her request. However, if the quality of life of the patient is very low, then the principle of beneficence may apply because a physician may help end the patient's suffering.

The students were introduced to these basic principles in an earlier lecture.

The two groups comprised 3 and 5 members, respectively and their discussions were videotaped. Two tutors were present during the discussions, but did not act as facilitators. The groups were located in the same room, either around a table or sitting in a circle in chairs.

Results

In general, both groups identify the ethical dilemma in the cases and relate the case aspects correctly to the relevant concepts. The identification of the dilemma occurs gradually, with individual students bringing in different viewpoints that are then argued for and against by their peers. This characteristic is somewhat puzzling because the students' own opinion of the cases is based on the principle of autonomy, i.e. they place the 'freedom to choose' above all other considerations, and none of the students places much value on principles that conflict with the principle of autonomy. The arguments for and against the judges' decision are hence not made on the grounds that some students agree with the judges. Rather, it appears that some students take on temporarily the position of the authorities. In this sense, many exchanges can be characterised as arguments.

Between the groups, an average of 15% of the utterances could be coded. It should be pointed out that if the processing of a knowledge unit is distributed over several utterances (and different students) it is counted as a single positive instance (see figure 3 for an example).

Processes

The following sections present utterances that have been classified in one of the coding categories. That is, these utterances show how information and knowledge is processed. Below, the processes will be exemplified, and their significance for learning discussed.

The processes occur with different frequency. For example, the process of relating concepts to aspects is very frequent, while the process of relating two aspects is infrequent. The process of abstracting principles did not occur at all. A possible reason for the absence of this process is that the cases can be sufficiently structured in terms of the known ethical dilemma. Other learning material where less domain-specific conceptual knowledge is available may lead to the abstraction of a principle (e.g. the fortress problem and the convergence principle, cf. Gick & Holyoak, 1980).

Relating aspects with concepts

Multiple utterances show that the students check whether a specific aspect of a case warrants the application of a principle (see figures 1 and 2). We interpret these utterances as externalisations of a normally internally occurring verification process. That is, students may want to use a principle or think that a principle should apply, but verify the facts of the case about whether there are grounds for that application. There is a probable reason for that externalisation: aspects function as and are used as 'evidence' presented to the group in order to justify the use of a concept for the conceptualisation of the entire case.

The significance of the process for learning can be understood in terms of instantiation: the students learn about a concrete feature that warrants the application of a principle. Future cases characterised by similar features will prompt the application of the principle. It is clearly also possible that students generalise over the concrete feature of the cases, acquiring more general knowledge about its applicability.

145, E	but in this case [pointing to sheet] she was sick

Figure 1. Group 1: student E refers to the fact that the patient is suffering, in support of the principle of beneficence (good and gentle death).

75, I	[interrupting] I think she was {of/in} stable mind.

Figure 2. Group 2: student I refers to the ability of the patient to make informed decisions, a precondition for applying the principle of autonomy

Formulating the conditions of applicability of concepts

We observed that sometimes the verification of applicability of a concept leads to the explicit formulation of conditions of applicability. This learning process is highly valuable because, by explicitly formulating knowledge – specifically, the conditions of applicability of a concept – this knowledge is not tied to the studied case, and hence does not require the retrieval of the case to use it in a future problem solving situation.

The formulation of general knowledge occurs through the application of a concept to the case, and the subsequent verification of its applicability. Instead of terminating the verification of applicability, conditions of applicability may be formulated without explicit reference to the case (see figure 3). Similar to learning domain-specific factual knowledge, acquisition of this knowledge is not demanded by the task, and can hence be considered a valuable by-product.

146, I	she was, but you know in this case I think if they said that this is legal then it could potentially open up a lot of, ... a lot of blur kind of, a lot of grey areas
147, E	...where to draw the line...
148, I	exactly
149, E	... someone has to be able to say and someone has to be able to write...

Figure 3. Group 1: example of an exchange where a student formulates general conditions of applicability of the principle of autonomy (149, E)

Relating concepts

Case-centred learning promotes integration of conceptual knowledge and relations between concepts acquired in isolation. Ethical (as environmental, political, etc.) cases are useful to understand the conflicts

between principles, and the situations under which one principle may be favoured over another (e.g. Pata & Sarapuu, 2003). This educational value of the case method is recognised and accepted by many practitioners and educators and cases are indeed frequently employed in education to raise the awareness of potential conflicts, problem and trade-offs (Burgoyne & Mumford, 2001). The model of learning with cases presented here hence functions as a description of this utility in terms of cognitive processes.

While it can be assumed that the students possess some knowledge about contrasts between ethical principles, the discussion of the cases promotes the strengthening and clarification of the relationships between concepts. The discussion is here interpreted as providing an opportunity to voice the pre-existing conceptual structure (see figure 4).

103, E	because in this case they were operating within the law and, say, they wanted to preserve life, but the weren't... but they weren't ... really ... they weren't ... they weren't really relieving suffering were they? ... which some people would argue {it would be?} ** preserve life. Some people they * doctors *** relieve suffering * preserve life, but {when?} they are suffering more {I mean?}then they should be able to end their life

Figure 4. Group 1: student E formulates the dilemma between the principle of beneficence and nonmaleficence using a mundane terminology

Relating aspects

The process of relating aspects also entails a structuring of the case. However, here the structuring occurs on the basis of some mundane relation rather than a relation that is specific for the domain (such as a conceptual relation). Such relations are, in the simplest case, linear (the more…, the less…) or causal. Figure 5 exemplifies a relationing between the aspect of 'sound mind' (that is a necessary precondition for even considering the patient's request) and the patient's depression. In this specific case, authorities have argued that the patient's request may be tainted by her poor mental state. The student's comment criticises that argument by identifying a relationship between her physical and mental health.

It should be pointed out that relationing of aspects using relational knowledge other than domain-specific one is very infrequent in these discussions. We believe that in cases with more facts and in different domains (such as biomedical or business cases) the process is much more frequent.

170, I	They think she can't decide on her own because she is depressed? Well of course she's depressed

Figure 5. Group 1. Student I establishes a (causal) relationship between the patient's state of health and her state of mind. The student criticises with it the claim that the patient may not be able to decide for herself because of her depression.

Relating cases and abstracting similarities/differences

We present these two processes together because they often occur in a single utterance. That is, frequently when a reference to another cases is made, the similarity or difference between the cases is also explicitly stated. The process of relationing cases is highly valuable for learning. While memorisation of cases or episodes is by itself considered to be a form of learning (cf. Kolodner, 1993), the reuse of episodes is crucially dependent on their organisation in memory. Episodes must be organised so as to permit efficient access and retrieval. Episodes that share features are normally grouped under a single header, with features not shared providing a subsequent identification opportunity if the header matches the input. Relating cases becomes a crucial mechanism in adding a case to the knowledge base, since it entails an assessment of similarities and differences, and an assessment of the relevance of the features, i.e. the identification of that or those features that represent the core of the case.

As described, the two cases represent similar situations. Active versus passive intervention is the only specific difference between them. This difference assigns different weights to beneficence and nonmaleficence, but the validity of autonomy remains unchanged.

Relating the cases, i.e. seeing the differences and similarities (see figure 6), promote the reiteration of the application of the ethical principles, as well as the search for a solution that is applicable to both cases. It is a first step towards a deeper understanding of the principles.

64, C	(shakes head in disagreement) … there is the difference, because [looks at sheet], because, basically, Mrs. Pretty, no, Mrs. B … it is passive and they could just switch off the ventilator, because basically they are keeping her alive, without her will, …, whereas the other one [Mrs. Pretty] was actually actively given drugs to die. So that's why it's difference there

Figure 6. Group 2. Student C identifies the difference between the cases.

Analogies

References to several other cases are made in the discussions. These references serve the purpose of introducing knowledge, and in two forms. First, by drawing attention to a similar case, the terminology or concept used to describe that case is reused to describe the current case. As an example, in figure 7, the student introduces the concept of 'quality of life' through a reference to a similar case. Second, the students have knowledge about the verdicts in similar cases (figure 8). This knowledge can be used to reason about the possible consequences of a decision. This function of analogies is of little use in the present cases since the students are not required to give or reason about a decision. References to other cases appear to serve primarily as introductions rather than as prototype 'solutions' to the case. Nevertheless, references to other cases produce valuable learning gains: they involve an assessment of similarities, and hence cases become related. Further, the introduction of a specific terminology increases the likelihood that this terminology and its associated concept will become part of the students' domain-specific vocabulary, and that it will be used to encode future similar problem solving situation.

139, I	yeah, ... there was a huge case in Canada a few years ago, where a father assisted his daughter, she was really very disabled ... and no quality of life [C and E seem to know the case] and he just suffocated her, and I think he got 12 years manslaughter.

Figure 7. Group 1. Example of a reference to another case that prompts the introduction of the notion of 'quality of life'

147, I	I: because, I know someone who is a doctor and they had {referred?}, *** had a huge motorcycle accident he actually {had to have?} blood transfusions, * but when it actually came out ** [that he was a Jehovah's witness {who refuse blood transfusions}] , ** how can a doctor try to preserve life ***

Figure 8. Group 2. Example of a reference to another case that highlights the principle of 'autonomy'

Summary

Our analysis has identified the main cognitive processes occurring when students discuss cases. From the results of the analysis we conclude that learning is conservative in the sense that students strengthen their existing mundane conceptual system on ethics. However, some valuable learning occurs: the students increase their knowledge about when to apply ethical concepts and learn some relations between concepts. The cases fulfil their role as promoting the integration of conceptual knowledge with real-world specific knowledge, and also the integration of conceptual knowledge. Further, relationships with other cases play a significant role in learning: the students bring in cases from their personal experiences or public cases, and abstract similarities and differences from them. These cases permit students to recognise what is common in euthanasia cases and what are their core distinctive features.

Our analysis helps us to understand what were the cognitive products constructed during the discussions. The students identified correctly the ethical dilemma in each case, though it is likely that the discussions served mainly as opportunities to externalise existing knowledge about the ethical dilemma rather than the dilemma being identified for the first time. In any case, the discussion fostered a clearer recognition of the relationships between ethical principles that embody the dilemma; and further, these conceptual relations have become tied to two specific cases that may serve as the basis for conceptualisations of structurally similar cases.

The occurrence of relationing between cases is of major importance: when students discuss the similarities and differences between the cases, they develop a knowledge base of cases, with the aspects found to be different functioning as distinguishing features. These features are indices for future retrieval of cases (Kolodner, 1993). The abstraction of commonalties and differences between cases leads also to the acquisition of a domain-specific vocabulary that will be used to encode future cases of medical ethics.

Discussion

Construct validity is a paramount concern for dialogue coding schemes i.e. whether the construct that is sought in the data is actually found (cf. Murphy, 2004). We have shown that the coding scheme is able to discriminate between different processes and identify those processes valuable for learning, according to our model. Moreover, the coding scheme permits new insights into how an abstract representation is constructed and especially what processes are used when learners construct that representation.

The application of the coding scheme is not restricted to a qualitative characterisation such as the one described in this paper. Rather, characterisation of learner dialogues has found application especially in relation

to the use of technology in education. For example, Newman et al. (1995) used a content coding scheme to identify differences in the quality of learning (specified in relation to a measure of critical thinking) in traditional face-to-face versus asynchronous computer-based environment settings. A similar application may be envisioned for the coding scheme presented here.

The approach to analysing learner dialogues and deriving a characterisation from it is complementary to other approaches. Most research on learning in a collaborative setting relies on socio-cultural or socio-cognitive views of learning, where specific interactions are deemed of significance for learning. For example, much work has been devoted to understanding how argumentation affects learning, and correlations between argumentation structures and learning have been identified (Veerman et al., 1999). We argue that some of this work can be integrated with the approach presented here. Indeed, argumentation may entail some of the processes claimed to be conducive for learning. For example, learners may point to a specific fact to support a claim, an argumentation move that, in our approach, is seen as relating a concept to a fact. While we cannot substantiate a relation between dialogue moves and cognitive processes, we believe that analyses targeting one or the other aspects of a dialogue may overlap.

References

Burgoyne, J. & Mumford, A. (2001). Learning from the Case Method. Report to the European Case Clearing House (RP301).

Chi, M. (1997). Quantifying Qualitative Analyses of Verbal Data: A Practical Guide. *The Journal of the Learning Sciences*, 6(3).

Chi, M., Feltovich, P. & Glaser, R. (1981). Categorisation and Representation of Physics Problems by Experts and Novices. *Cognitive Science*, 5, 121-152.

Davies, E. (2003). Knowledge Integration in Science Teaching: Analysing Teachers' Knowledge Development. *Research in Science Education,* 34, 21-53.

Dillenbourg, P. (1996). Some Technical Implications of the Distributed Cognition Approach on the Design of Interactive Learning Environments. *Journal of Artificial Intelligence in Education* 7(2), 161-80.

Forbus, K. (2001). Exploring Analogy in the Large. In Gentner, D., Holyoak, K., & Kokinov, B. (Eds.) *The Analogical Mind: Perspectives from Cognitive Science*. Cambridge, MA: MIT Press.

Garrison, D. R., Anderson, T., & Archer, W. (2000). Critical Thinking in a Text-Based Environment. Computer Conferencing in Higher Education. *Internet in Higher Education,* 2(2), 87-105.

Gick, M. & Holyoak, K. (1980). Analogical Problem Solving. *Cognitive Psychology*, 12, 306-55.

Hara, N., Bonk, C. J., & Angeli, C. (2000). Content analysis of online discussion in an applied educational psychology course. *Instructional Science*, 28, 115-52.

Kolodner, J. (1993). *Case-based Reasoning*. San Matteo: CA, Morgan Kaufmann.

Murphy, E. (2004). Promoting Construct Validity in Instruments for the Analysis of Transcripts of Online Asynchronous Discussions. *Educational Technology Research and Development* (online).

Newman, D. R., Webb, B., & Cochrane, C. (1995). A content analysis method to measure critical thinking in face-to-face and computer supported group learning. *Interpersonal Computing and Technology*, 3(2), 56-77.

Pata, K. & Sarapuu, T. (2003). Development of Mental Models by Environmental Role-play in Synchronous Collaborative Virtual Workplace. *Proceedings of eLSEE, e-Learning in Science and Environmental Education*.

Ross, B. & Kennedy, P., (1990). Generalising from the Use of Earlier Examples in Problem Solving. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 16.

Veerman, A. L., Andriessen, J. & Kanselaar, G. (1999). Collaborative Learning through Computer-Mediated Argumentation. In Hoadley, C. & Roschelle, J. (Eds.) *Proceedings of CSCL99*. Mahwah: NJ, LEA.

Acknowledgements