A knowledge-based coach for reasoning about historical causation

Liz MASTERMAN
Learning Technologies Group, Oxford University, UK

Abstract. The ability to explain the causes of historical events is a key skill for learners to acquire, but the ill-structured nature of the task means they cannot be guided through a problem-space of well-defined moves to reach a correct answer. This paper investigates whether a knowledge-based computer coach can provide effective guidance to learners as they construct diagrammatic explanations of the causes leading to a particular event. The design of the coach was based on a model of expert reasoning synthesised from the historiographical literature and on an analysis of teacher-learner interactions observed during classroom activities. Coaching was provided at two levels: a) generalised (decontextualised) guidance and b) guidance directly relevant to the topic of study. Where appropriate, learners could choose to disregard the coach's advice. The knowledge-base underlying the coach could also be made available as a scaffolding aid. An evaluation with three groups of students aged 12-13 showed that i) maximal scaffolding and content-specific coaching resulted in diagrammatic explanations of greater accuracy and superior structural quality to those produced either with generalised guidance or with no guidance at all, and ii) learners' appreciation of the subjective nature of historical explanations was not compromised by the coaching interventions.

Introduction

Causation is one of a set of key concepts that provide both experts and learners with a structure for understanding and thinking about history [11]. However, reasoning about historical causation—that is, identifying and explaining the relationships between a particular event and those antecedent events that may be said to have brought it about—poses a substantially greater challenge than reasoning about causation in everyday life or in the natural sciences or law. The uniqueness of historical events and their remoteness in time mean that the historian can neither conduct experiments that make the effects of such events perceptible nor interrogate historical personages about their motives and intentions. Instead, he or she must rely on a vast knowledge-base of evidence derived from sources which may be incomplete, inconsistent and difficult to interpret. Thus, an answer to the question "why?" about history can never be definitive; rather, it is a matter of personal interpretation influenced in part by the historian's general knowledge and perspective.

A key challenge in fostering learners' reasoning about causation, therefore, is to induce them to appreciate the admissibility of alternative solutions to problems of historical causation. This paper outlines a study which investigated whether a knowledge-based coaching system can provide effective support for learners' emergent reasoning as they construct historical explanations on the computer. It begins by outlining the characteristics of expert reasoning about causation before examining how teachers introduce their learners to the task and scaffold the solving of "causation" problems through verbal interactions with learners and through different forms of external representation. The paper then describes how the information obtained from this research provided the basis for a computer program in which learners constructed diagrammatic explanations of the causes of the English Civil War. Finally, it reports and discusses the findings of an evaluation comparing the outcomes of learners' use of the program with and without support from a computer-based coach.

1. Characteristics of expert reasoning about historical causation

A model of expert performance in a particular domain can give an indication of what the outcome of successful learning should look like [1]. However, constructing such a model for historical causation is a challenging task, since there is neither an agreed terminology nor an agreed set of procedures among the experts, with historiographers arguing the case for and against causal reasoning as a deductive, inductive, adductive or associative process. It is, however, best characterised as an *informal* logic, governed by internal principles which have more to do with rhetoric than with propositions of formal logic [17] or estimates of probability. In order to identify those concepts and procedures most commonly associated with this logic, the author undertook an extensive synthesis of the historiographical literature on causation. Figure 1 summarises the outcome of this task.

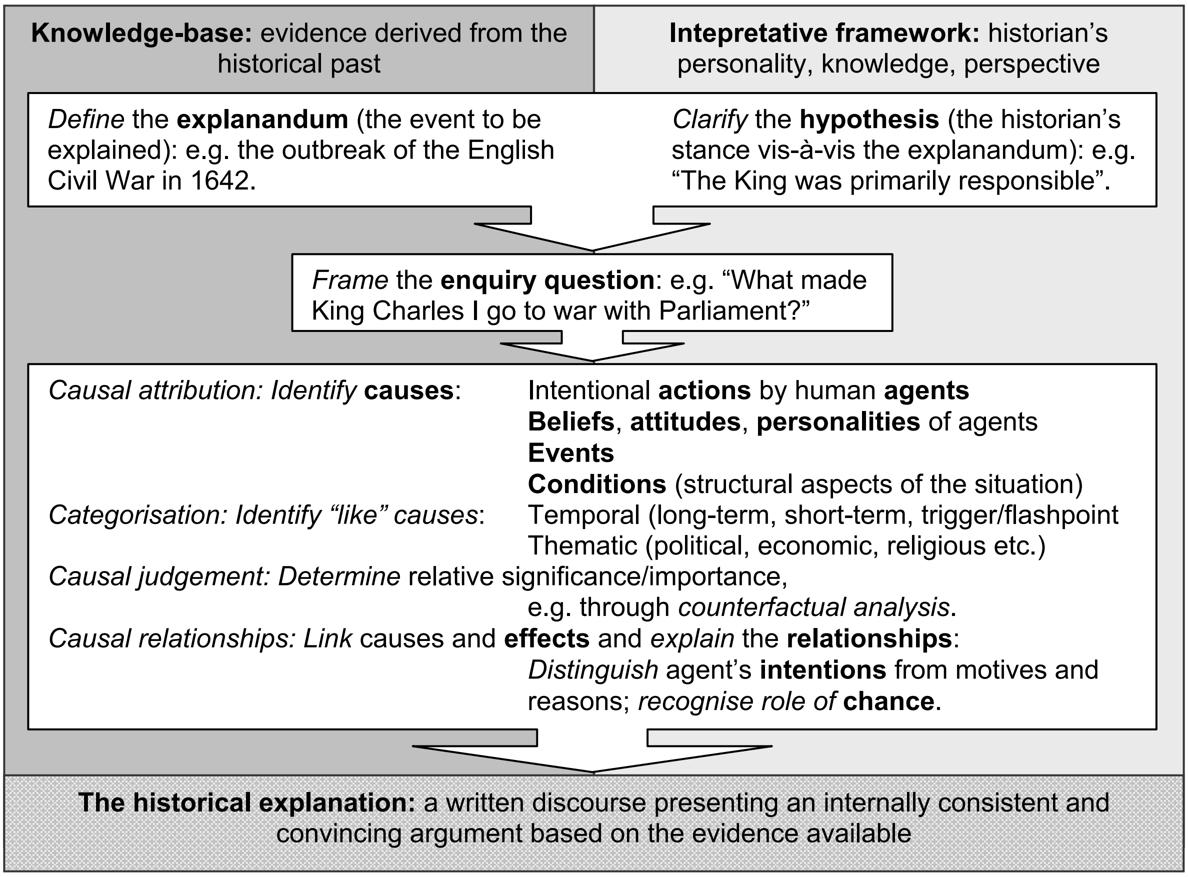

Figure 1. Reasoning about historical causation: summary of the principal **concepts** and associated *procedures*. Synthesised from numerous sources cited in [9]

To arrive at a historical explanation, the historian applies his or her interpretative framework to the knowledge base of sources in order to identify the relevant causes and to categorise, and/or judge the significance of, these different causes as desired. Establishing a cause-effect relationship is relatively straightforward in the case of conditions and events, especially where the agent is a force of nature. However, determining causal relationships where human actions are involved necessitates inferring the agent's conscious intentions, as distinct from their motives (often unconscious) and reasons (how the agent might justify the action). It should be recognised, of course, that the procedures in Figure 1 are iterative rather than sequential; that is, an attempt to establish a causal relationship between two factors might trigger a return to the evidence to search for a third, intermediate, factor.

Perhaps the cognitive model to which reasoning about historical causation is closest is that for the solving of ill-structured problems [18]. In line with this class of problem, historical causation is distinguished by i) an initial state (the explanandum) and a goal state (the historical explanation), each of which may be open to multiple interpretations; ii) the presence of a large number of open constraints (i.e. gaps and inconsistencies in the evidence) which different members of the problem-solving community may fill in different ways, thereby leading to iii) differing solutions, the quality of which is largely a matter of pragmatic judgement. Furthermore, as with other ill-structured problems, constructing a historical explanation involves selecting relevant information from a considerable body of data and decomposing the main problem into multiple relatively well-structured problems.

The ramification of this model for history teachers is clear in that, unlike problem-solving tasks in maths, science or logic, they cannot direct learners through a problem space of well-defined moves where specific constraints must (and can) be satisfied in order to arrive at the "correct" answer. Indeed, they must actively avoid creating the impression that problems of historical causation are solved in this way.

2. Promoting reasoning about historical causation in the classroom

While providing clues to the nature of expert reasoning, historiographers give little guidance about how to guide learners towards the desired performance [19]. Therefore, in order to determine how far the model of reasoning presented in Figure 1 is reflected in the classroom, what sorts of misconceptions learners have, how teachers guide learners through the problem-solving task and what forms of representation they use to mediate this process, the author combined a survey of recent research on the development of learners' causal reasoning [e.g. 8] and a review of the literature on teaching causation in the UK [e.g. 6] with classroom observations. The observations covered 46 lessons on a range of "causation" topics, involving students aged from 11 to 17 in three mixed-ability co-educational schools and one school for medium- to high-ability girls. The aim was to establish, from these multiple sources of data, generalisations applicable to the design of the proposed program.

2.1 Introducing learners to the concepts and procedures involved in causal reasoning

There was no overt teaching of any "global" logic for reasoning about historical causation in any of the schools observed. Rather, concepts and procedures were introduced gradually, according to the demands of the subject matter and the teacher's perception of the students' readiness for tackling new concepts or familiar ones at a higher level. Nevertheless, the principal elements of Figure 1 were discernible in the observations, albeit in a somewhat simplified form. For example, students, with their initially naïve interpretative frameworks, were not expected to generate their own hypotheses and so were given enquiry questions which had been pre-defined by the teacher. Overall, therefore, teachers may be seen as fostering a model of *competent*, rather than expert, reasoning which students in the UK might be expected to acquire before they end their compulsory study of history at age 14.

The observations also validated the equation of reasoning about historical causation with the solving of ill-structured problems, in that teachers laid stress on the multiplicity of possible solutions; provided students with a subset of sources (usually from a textbook);

and subdivided the topic into manageable phases: information-gathering and interpretation, knowledge-construction (categorising, judging significance, identifying causal relationships) and knowledge communication (usually as a written historical explanation).

2.2 Teacher-learner interactions

The analysis of observed interactions revealed a high level of input by teachers during the information-gathering phase in helping students to interpret evidence, alerting them to their misconceptions and explaining archaic terms and the abstract concepts associated with historical causation. In the knowledge-construction phase, when students were engaged in semi-independent problem-solving activities, the teacher would move around the class and engage with students individually, quickly reviewing their work and offering advice and feedback (i.e. coaching). The observations yielded four *styles* of coaching intervention:

- *Directive:* unsolicited advice and hints at the outset of an activity.
- *Responsive:* guidance in response to a student's request for help.
- *Reactive:* immediate feedback on an action by an individual student.
- *Retrospective:* holistic feedback either to an individual learner or to the whole class when the activity has reached an advanced stage or has been completed.

2.3 Representations of knowledge

The outcome of a causation enquiry is normally a written explanation, an exercise which is fraught with difficulty for many learners [4]. Therefore, teachers tended to scaffold this process by helping students to formulate their ideas through constructing representations which made fewer demands on students' linguistic abilities; for example, tabulated charts, directed graphs and "cause cards" (that is, labelled slips of paper, each corresponding to a cause, which pupils sorted into different configurations in categorisation activities).

3. Design of the computer-based coach

The next stage in the study was to feed the findings from the observation into the design of the program, titled 20/20. This design hinged on three interrelated decisions: i) the phase(s) in a causation enquiry which the program would support; ii) the role of the computer coach vis-à-vis the teacher; and iii) the form of representation to be supported at the interface. These decisions were made by marrying observational data with a theoretical framework which places teacher-learner interactions and learning activities within a *modelling-supporting-fading* paradigm [9, 10], where the teacher adopts the role of more able partner. The observational data suggested that there would be almost insurmountable difficulties in implementing a system in which the computer assumed the role of *replacement* teacher, since teachers often used topical references or their personal knowledge of students when explaining abstract concepts. However, it was also noted that, during classroom activities, the teacher did not always have time to provide guidance to individual students. Hence, it seemed that the computer could fulfil the role of *adjunct* to the teacher by coaching students when the latter was unavailable. However, the teacher would remain responsible for diagnosing learners' levels of ability and deciding the amount of support to be provided by the computer.

The representational form, a diagram akin to a concept map, was chosen because of its simplicity (consisting of two basic elements: boxes and arrows) and because it combined two forms already used in the classroom: namely, directed graphs and cause cards. The guiding principle in devising the notation was the need for a perspicuous scheme which did not impose an additional cognitive burden on students and made it possible to represent multiple perspectives simultaneously (e.g. temporal classifications plus thematic groupings). Table 1 maps the key concepts associated with historical causation supported by 20/20 to the notation used. Figure 2 shows the notation in context: a student's diagram.

Table 1. The key concepts associated with causation and their representation at the interface

Concept	**Examples**	**Notation used**	**Rationale for notation**
Causal factor	Action, event, condition	Box with text label	Visual similarity to cause card
"Temporal" classification	Long-term, short-term, flashpoint/trigger	a) Dark cloud, yellow or red lightning flash *or* b) Binoculars, spectacles	a) Meteorological metaphor is associated with build-up to cataclysmic events e.g. wars. b) Ocular metaphor is more suited to undramatic events.
Thematic grouping	Political, economic, religious, military	Colour-coding in boxes	Cultural associations (where possible); e.g. red = military
Significance	Major cause, minor cause	Variations in thickness of box borders	Visual salience
Chance factor		**!** in cause box	Used on "hazard" road signs
Causal relationships	N/A	Arrow connecting cause to effect	Arrow is suggestive of causal stimulus

The core system consisted of a "workspace" where learners explored and experimented with their ideas, creating and manipulating configurations of cause boxes and links to build a diagrammatic representation of the causes of the event in question (see Figure 2).The procedures involved in causal reasoning were mostly carried out through "point-and-click" operations using buttons in the toolbar.

The central challenges in designing the coach which was to be overlaid on the core system were primarily pedagogical; viz. i) how to guide learners towards a plausible solution to the question while simultaneously reinforcing an appreciation of the subjective quality of that solution, and ii) how to diagnose the misconceptions behind their actions. To meet both challenges, moves that could prompt coaching interventions were divided into:

- *Strong issues:* illogical moves (e.g. linking an effect to its cause instead of vice versa), in which the coach would always intervene to enforce correction.
- *Weak issues:* matters that were open to interpretation. Here, the coach would display a pop-up message alerting the learner to the discrepancy between their diagram and its own view, but give the learner the freedom to ignore its advice.

The frequency of interventions was defined by a set of rules derived from the WEST system [2] and by experimentation. The style of interventions by the computer coach was determined both by observational data and by technological constraints. For example, the object-oriented behaviour of the interface (i.e. select object→perform action) precluded directive coaching for almost all moves. Also, to avoid processing natural-language input, responsive coaching was implemented as a list of frequently-asked questions under the heading "Help me to decide".

The design allowed for two levels of coaching (as well as none at all), with the teacher predetermining the level to be used with any one group of learners. "Generalised" coaching gave broad guidance only (e.g. decontextualised definitions of concepts). "Content-specific" coaching offered additional guidance relevant to the situation in question, although this meant restricting learners to choosing causes from three pre-defined lists: actions and events (the "Time-Line" in Figure 2), beliefs and attitudes of the agents involved ("People"), and the underlying conditions ("Big Issues"). These lists could also be made available as optional scaffolding aids for learners receiving generalised (or no) coaching.

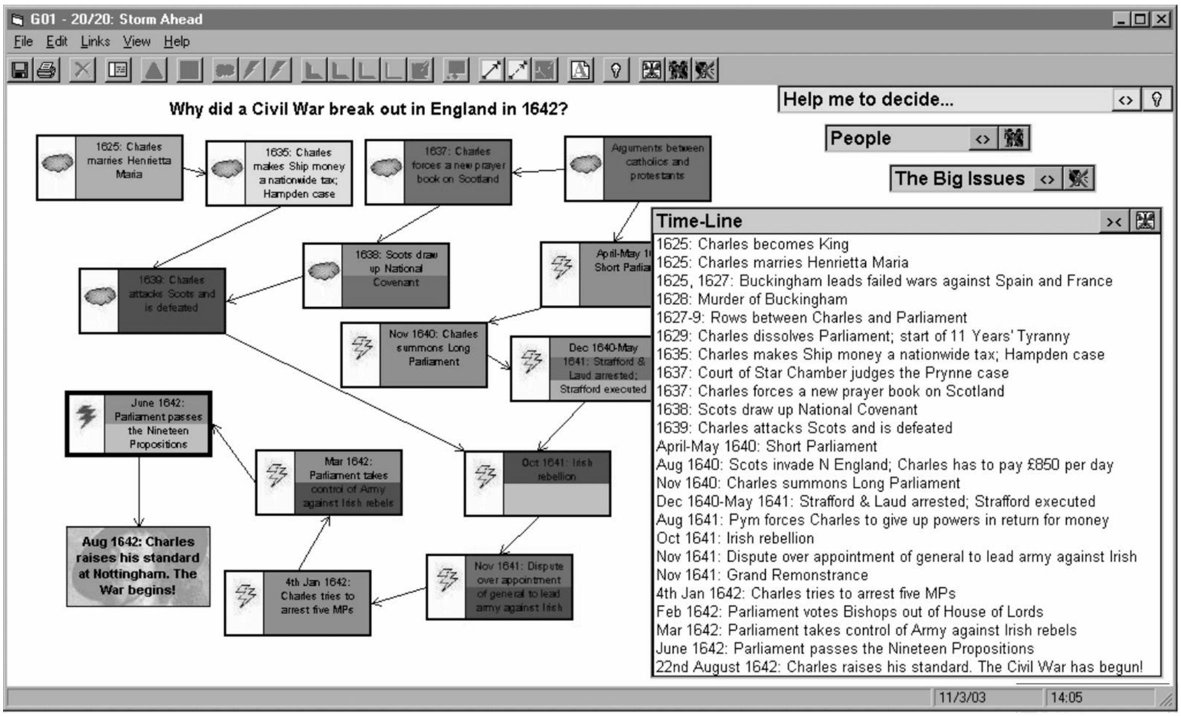

Figure 2. Workspace of 20/20, here subtitled "Storm Ahead" because the meteorological icons are in use. Two of the lists of pre-defined causes are closed, as is the set of issues for which responsive coaching is available

The coaching system was implemented as a combination of a) immutable rules embedded in the program code along with generalised coaching messages applicable to all situations, and b) a knowledge base of pre-defined causes specific to a particular historical situation, stored in a database along with their attributes (including relationships with other causes) and the coaching messages relating directly to them. This knowledge base thus served two purposes: to scaffold learners' tasks and to function as an "expert version" with which the coach could compare the outcome of learners' actions and give fully contextualised guidance.

4. Evaluation

The 20/20 coach was evaluated with three mixed-ability classes of students aged 12-13 at one of the co-educational schools involved in the observations. The hypothesis proposed that students who received higher levels of computer-based support would produce diagrams that were i) more accurate (i.e. closer to the expert version) and ii) of superior structural quality (i.e. containing more causes and links) than students who received less support. Each class constituted a separate experimental condition (see Table 2). They had already studied the causes of the English Civil War and spent two one-hour sessions using 20/20 to construct a diagram explaining why, in their view, the war broke out.

Table 2. Experimental conditions in the evaluation of the 20/20 coach

	Group T (26 students)	**Group G (22 students)**	**Group N (20 students)**
Composition	26 students; Teacher A	22 students; Teacher B	20 students; Teacher B
Scaffolding	Use causes from pre-defined lists only	Select causes from pre-defined lists + optionally devise their own causes from researching in their textbooks	
Coaching	Content-specific	Generalised	None

Analysis of the records of learners' actions in 20/20 confirmed that group T did receive more coaching: one reactive intervention per 5.33 actions and one retrospective intervention per 72.36 actions, compared with 28.14 and 65.56 for group G. Both groups appeared to act on the computer's coaching of "weak" issues roughly two-thirds of the time, suggesting that they were not completely in thrall to the computer coach. Recourse to responsive coaching was minimal, with a total of 21 requests from the two groups.

The completed diagrams of all three groups were scored using formulae based on [5, 13] and described in detail in [9]. Accuracy scores could range from 1 (maximum) down to 0 (minimum), and scores for structural quality could range from 1 (maximum) down to values below 0. Table 3 summarises the mean scores and the results of statistical tests performed on them.

Table 3. Mean scores (and standard deviations) obtained by the three groups, and results of statistical tests

Criterion	**Group T**	**Group G**	**Group N**	**Kruskal-Wallis test**
Accuracy: cause boxes	0.57 (0.15)	0.19 (0.13)	0.26 (0.09)	$\chi^2 = 25.386$, p = .000
Accuracy: links	0.03 (0.03)	0.01 (0.02)	0.02 (0.02)	$\chi^2 = 5.705$, p = .058
Structural quality	0.53 (0.04)	0.14 (0.04)	0.17 (0.28)	$\chi^2 = 9.605$, p = .008

Differences among the groups were significant at p<=.05 except for the accuracy of links, where the differences approached significance. However, it is notable not only that group T's scores were well ahead of the other two groups, but also that group G actually scored slightly lower than group N. Hence, the hypothesis was only partly supported.

5. Discussion

The question investigated in this paper is whether a knowledge-based coach can provide effective support for learners' emergent reasoning about historical causation. Since causal reasoning is a skill which requires several years to develop, it was possible to evaluate only a short-term intervention. Findings from the 20/20 evaluation showed that varying the amount and content of computer-based support could result in differences in performance in a single task without excessively compromising learners' independence of thought. However, it appeared that diagrams of significantly higher quality were produced only where a) the level of scaffolding was sufficiently high as to minimise the risk of students' voluntarily making unacceptable moves, and b) the coaching delivered was relevant to the topic of study. Coaching which offered only generalised advice and feedback often resulted in diagrams that differed little from those produced without any coaching at all—perhaps because such advice provided insufficient clues as to how learners should act in a specific situation. Although group T had a different teacher, observational notes from the evaluation sessions suggest that differences in the two teachers' styles were insufficient to account for such large variations in scores. Nevertheless, the investigation would benefit from a) a longitudinal study to determine, inter alia, whether learners can generalise from the advice received in relation to one historical situation and apply it, after an extended period, to a novel situation, and b) more rigorous control of variables such as teaching styles.

The program 20/20 is innovative in that it supports a domain traditionally under-represented in artificial education research, viz. history (an exception is Disciple [14]), but it also continues a well-established tradition of intelligent graphical reasoning tools that includes Belvedere, Convince Me and Reason!Able [15, 12, 16], as well as the more recent Reasonable Fallible Analyser (RFA) [3]. The option of a content-specific knowledge-based coach has commonalities with Belvedere; however, 20/20 does not currently support learners' construction of a substantiated argument like Belvedere, Convince Me and Reason!Able or allow learners to argue in favour of their position, as does the RFA. It would be worthwhile, therefore, to consider adding either or both of these facilities to 20/20.

Ultimately, further developments to the 20/20 coach must recognise the central tension between that which can be achieved technologically and that which is acceptable historiographically and, hence, pedagogically. At present, a major limitation of 20/20 is the lack of coaching for the key procedure of explaining causal relationships. Yet it is not only impossible to formulate the universal rules that might underlie a coach for this task (e.g. "people of disposition X faced with situations of type Y are likely to act in manner Z"), but such rules would negate the very essence of historical causation: namely, to explain why particular individuals acted as they did in specific situations [7]. History may be full of ill-structured problems with diverse solutions, but its internal logic must be strictly observed.

With acknowledgements to Mike Sharples for his invaluable support during the study.

References

[1] Bransford, J.D., Brown, A.L. & Cocking, R.R. (Eds.) (1999). *How People Learn: Brain, Mind, Experience, and School*. Washington, DC: National Academy Press.

[2] Burton, R.R. & Brown, J.S. (1982). An investigation of computer coaching for informal learning activities. In D. Sleeman & J.S. Brown (Eds.), *Intelligent Tutoring Systems* (pp. 79-98). London: Academic Press.

[3] Conlon, T. (2004). 'Please Argue, I Could Be Wrong': a Reasonable Fallible Analyser for Student Concept Maps. AACE Journal, 12(4). Available: http://dl.aace.org/15571 [Accessed 25/01/05]

[4] Curtis, S. (1994). Communication in History—A process based approach to developing writing skills. *Teaching History*, 77, 25-30.

[5] Funke, J. (1985). Steuerung dynamischer Systeme durch Aufbau und Anwendung subjectiver Kausalmodelle. *Zeitschrift fur Psychologie*, 193(4), 443-465.

[6] Husbands, C. (1996). *What is history teaching? Language, ideas and meaning in learning about the past*. Buckingham: Open University Press.

[7] Lee. P.J. (1984). Why Learn History? In A.K. Dickinson, P.J. Lee and P.J. Rogers (Eds.), *Learning History* (pp. 1-19). London: Heinemann.

[8] Lee, P., Dickinson, A. and Ashby, R. (1998). Researching Children's Ideas about History. In J.F. Voss and M. Carretero (Eds.), *Learning and Reasoning in History* (pp. 227-251). London: Woburn Press.

[9] Masterman, E.F. (2004). *Representation, mediation, conversation: integrating sociocultural and cognitive perspectives in the design of a learning technology artefact for reasoning about historical causation*. Unpublished doctoral thesis, University of Birmingham, UK.

[10] Masterman, L. and Sharples, M. (2002). A theory-informed framework for designing software to support reasoning about causation in history. *Computers & Education*, 38, 165-185.

[11] Nichol, J. (1999). Who wants to fight? Who wants to flee? Teaching history from a "thinking skills" perspective. *Teaching History*, 95, 6-13.

[12] Schank, P. & Ranney, M. (1995). Improved reasoning with Convince Me. *CHI '95 Proceedings: Short Papers*. Available: http://www.acm.org/sigchi/chi95/Electronic/documnts/shortppr/psk_bdy.htm

[13] Seel, N.M. (2001). Epistemology, situated cognition, and mental models: 'Like a bridge over troubled water'. *Instructional Science*, 29, 403-427.

[14] Tecuci, G. & Keeling, H. (1999). Developing an Intelligent Educational Agent with Disciple. *International Journal of Artificial Intelligence in Education*, 10, 221-237.

[15] Toth, J.A., Suthers, D. & Weiner, A. (1997). Providing Expert Advice in the Domain of Scientific Enquiry. In B. du Boulay & R. Mizoguchi (Eds.), *Artificial Intelligence in Education: Knowledge and Media in Learning Systems. Proceedings of AI-ED 97 World Conference on Artificial Intelligence in Education* (pp. 302-308). Amsterdam: IOS Press.

[16] van Gelder, T. (2002). Argument Mapping with Reason!Able. *American Philosophical Association Newsletter on Philosophy and Computers*, 85-90.

[17] Voss, J.F., Perkins, D.N. & Segal, J.W. (1991). Introduction. In J.F. Voss, D.N. Perkins & J.W. Segal (Eds.), *Informal Reasoning and Education*. Hillsdale, NJ: LEA.

[18] Voss, J.F. & Post, T.A. (1988). On the Solving of Ill-Structured Problems. In M.T.H. Chi, R. Glaser & M.J. Farr (Eds.), *The Nature of Expertise* (pp. 261-285). Hillsdale, NJ: Lawrence Erlbaum Associates.

[19] Wineburg, S. (2001). *Historical Thinking and Other Unnatural Acts: Charting the Future of Teaching the Past*. Philadelphia: Temple University Press.