

Providing Support for Creative Group Brainstorming: Taxonomy and Technologies*

Hao-Chuan Wang¹, Carolyn P. Rosé¹, Tsai-Yen Li², Chun-Yen Chang³

¹Carnegie Mellon University
Pittsburgh PA, USA
{haochuan, cprose}
@cs.cmu.edu

²National Chengchi University
Taipei, Taiwan
li@nccu.edu.tw

³National Taiwan Normal
University
Taipei, Taiwan
changcy@ntnu.edu.tw

Abstract. This paper describes our plans and current work towards developing a user modeling-based learning environment for creative group brainstorming in qualitative domains. This research framework aims to guide the investigation and integration of different brainstorming activities, theoretical foundations and supportive technologies in order to benefit learning most. Intervention studies and system development are proposed to take place in parallel to ensure that the design of system behavior is informed by research results.

Keywords: Group Brainstorming, Creative Problem Solving, Collaborative Learning

INTRODUCTION

Research in the area of Intelligent Tutoring Systems (ITS) has achieved impressive results in improving student learning especially in well defined problem solving domains, such as basic algebra (Koedinger et al., 1997) and quantitative physics (VanLehn et al., 2005), just to name a couple. In this paper we begin to consider how to expand the frontiers of this success into areas less touched by ITS research. The research objectives of this work are: (1) to provide more open ended qualitative scientific problem solving opportunities for students, also known as creative problem solving (CPS), and (2) supporting productive small group dynamics for collaborative idea generation, joint problem solving, and learning, respectively. Recent work begins to approach this area, such as exploratory problem solving (Kumar et al., 2006), natural language-based tutorial dialogue systems (Grasser et al., 2005; Kumar et al., 2006), and adaptive support for collaborative learning (Gweon et al., 2006). In this paper, we present a tutoring system that brainstorms with students, called VIBRANT (Virtual Brainstorming), which we propose as an instructional tool for science education.

Supporting CPS poses different challenges than the types of problem solving domains more frequently studied in the ITS community. Consider the following sample CPS question: “*What are the possible factors that might cause a debris-flow hazard to happen?*”, and subsequently, “*How could we prevent it from happening?*” One salient characteristic of this task is the unlikelihood of modeling the process of answering these questions in a procedural and goal-directed manner as in more traditional ITSs. Notice that the goal for students here is not to select and then apply a known procedure for solving this problem. In contrast, the main purpose here is to let students actively generate the candidate problem solving steps/options by themselves. Also, because problems such as these do not have a single right answer, another important feature of this type of task could be the necessity of recruiting human judges in evaluating the goodness of students’ problem solving behaviors. In fact, beyond offering students the opportunity to generate possible solutions to problems, these tasks offer students the opportunity to weigh and balance trade-offs between alternative solutions. In our previous work, students were required to answer CPS questions independently without accessing external resources or peer support. Human graders were recruited to score student answers quantitatively using a rubric devised by domain experts (Chang & Weng 2002). However, this operationalization of CPS has previously been used primarily in assessing students’ problem solving abilities, and the focus was not specifically on how to scaffold students’ idea-generation and creative problem solving, and/or to investigate the embedded instructional opportunities with an additional factor of the presence of peers for collaboration. In contrast, the purpose of VIBRANT is to support students in the process of CPS, with the goal of improving their problem solving skills.

* A short version introducing the technologies discussed in this workshop paper is included as a poster in the main conference of ITS 2006 (Wang et al., 2006).

Thus, our current work focuses on data-driven design and implementation of system behavior for supporting individual/collaborative brainstorming activities. Three types of brainstorming activities are considered, including (1) brainstorming for idea generation, (2) brainstorming for creative problem-solving, and (3) brainstorming for inquiry learning, which are all related. VIBRANT offers user modeling-based supportive technologies for activities including *cognitively oriented support providing brainstorming feedback* and *socially oriented support for discussion group formation*. The combination of *activities, the degrees of socialization (e.g., group vs. solitary brainstorming), and technologies* results in a very rich design space for researchers to investigate. Our research goal is to explore how to combine findings from the social psychology community on how to maximize productivity in brainstorming with findings from the learning sciences community on how to maximize the learning benefit from exploration. Rather than arguing that we have the answer, instead we present VIBRANT as a framework in which researchers may begin to address the many research questions that stand in between our current understanding and the ultimate solution.

In the remainder of the paper, we describe our theoretical foundation, drawn from both communities of research. We then discuss our hypotheses and the research questions we must address in order to support brainstorming activities, primarily creative problem solving and active inquiry, for learning. We then describe VIBRANT, a technological infrastructure for supporting this work. We evaluate VIBRANT's ability to offer coherent feedback to students in response to their brainstorming behavior in a CPS task. We conclude by discussing how we will use the VIBRANT system to address our research questions.

THEORETICAL FRAMEWORK

In this section, we identify three types of brainstorming activities connected with learning, which together form a theoretical foundation for guiding the design of VIBRANT and our planned future empirical studies. Table 1 (see Appendix) offers an overview of the types of empirical investigations we either draw from or plan to conduct. We build on previously published theories, interventions, key concepts such as the distinction between performance and learning outcomes, and methodologies for operationalization and measurement.

A conceptual continuum identifying *performance-oriented* and *learning-oriented* activities is delineated in Table 1. The activity of idea-generation, as frequently characterized in the literature of social psychology and organizational studies, focuses more on *performance* rather than *learning*. By contrast, the activity of inquiry learning, which emphasizes students' active knowledge acquisition/construction, is clearly at the other side of the performance-learning continuum. In our view, CPS combines concerns from both ends of this spectrum. On the one hand, when working on CPS tasks, it is the quality/quantity of solutions (i.e., performance) that the organization or individual tends to optimize, and which is a more obvious measure of success. On the other hand, in learning environments, a more open-ended creative problem solving task may encourage problem solvers to explore the domain, to establish links among learned concepts actively, to interact with peers with argumentative discourses, and to detect what knowledge components are to be further acquired. The act of idea generation may lead a student to a new area of the exploration space that offers a valuable opportunity for learning. Thus, one may argue that it is not the exploration per se that is important, but the opportunities for learning that may result from that exploration.

Brainstorming for Idea Generation

The cognitive task of idea generation in brainstorming groups has been extensively studied in social psychology through controlled experiments (Diehl & Stroebe, 1987). In the study of group brainstorming for idea-generation, empirical work has repeatedly revealed phenomena related to process losses, in which a group with mutually-interacting members may not always perform better than a collection of non-interacting individuals whose contributions are pooled statistically (i.e., nominal group) (Hill, 1982; Diehl & Storebe, 1987). Theoretical explanations for process losses have been proposed and tested empirically, including social pressure (e.g., evaluation apprehension), social loafing (e.g., "free riding"), and production blocking (Connolly, 1993; Diehl & Storebe, 1987 & Kraut, 2003). Early electronic brainstorming systems (EBS) designed to avoid process losses resulting from these causes have shown great promise (Connolly, 1993). Thus, we believe it is within the realm of what state-of-the-art technology can support to address concerns in the realm of learning in CPS tasks.

It is considered that idea generation activities can be used as the foundation for stimulating students' active thinking for performing other higher level activities in the taxonomy, including creative problem solving and inquiry learning, although it appears that only doing idea-generation may *not* help students acquire domain knowledge effectively. A series of interventions is proposed in Table 1 in order to test several social psychological hypotheses (intervention ig1-ig4), and to further verify the relation between students' idea generation and domain learning. For example, in a CPS task, whether more and better ideas generated would lead to better problem solving and learning.

Brainstorming for Creative Problem Solving

Based on the literature of creativity research and science education, the Creative Problem Solving (CPS) model generally describes problem solving as a process consisting of two qualitatively different phases: divergent and convergent thinking (Basadur, 1995; Chang & Weng, 2002; Osborn, 1963). The literature proposes that a problem solving process can be decomposed into several stages, typically including fact-finding, problem-solving, idea-finding, and solution-finding. The stages of fact-finding and idea-finding are more divergent thinking oriented, while the stages of problem-solving and solution-finding are considered more convergent thinking oriented. In each stage, some work proposed that two sub-phases of ideation and evaluation can be further identified (Basadur, 1995).

We foresee opportunities for students to *learn*, hypothetically, partly subject matter and partly meta-cognitive skills, during CPS tasks. Obviously, idea-generation appears to play an evident role in the part of divergent thinking in CPS tasks, which may help stimulate students' active thinking and engagement with other members within the group who may help them see the problem from a different perspective. Furthermore, in convergent thinking sub-phases, students are required to evaluate, explain and negotiate ideas they have generated, which appears to offer valuable *knowledge events* similar to the activity of self-explanation that can be employed in shaping effective tutoring (Aleven & Koedinger, 2002). From a different angle, by tracing and coding students' argumentative process in CPS, the analytic framework of argumentative knowledge construction may be applied towards providing process-oriented instructional support (Gweon et al., 2006; Weingerger & Fischer, 2006). In a broader sense, CPS-based tutoring appears to be situated at the intersection of ITS, Computer Supported Cooperative Work (CSCW) and Computer Supported Collaborative Learning (CSCL), or correspondingly, cognitive psychology, social psychology and educational psychology. While engaging in cooperative work in the context of CPS, it is reasonable to expect that valuable opportunities for tutoring and collaborative learning are interwoven within the processes of work, in which workers may need to recall or even acquire skills not used before, and peer workers' opinions may provoke cognitive conflict that may lead to active learning and conceptual changes.

Brainstorming for Inquiry Learning

Inquiry as an approach to learning typically consists of processes of exploring the targeted problems or phenomena, asking questions, and making discoveries, achieving new understanding and fulfilling personal curiosity (NSF, 2002). The underlying educational ideology states that students will be able to develop durable and transferable science process skills and to continue pursuing lifelong learning in a self-directed means based on the meta-cognitive skills of inquiry and discovery that they derived from inquiry learning.

However, as recent review articles and empirical studies have argued, the approach of *pure* discovery learning that strictly prohibits intervention from teachers did not show any evidence of performing better than other approaches, e.g., direct instruction, in terms of subject matter learning, transfer to new situations, and the acquisition of basic scientific skills (Anderson et al., 1998; Klahr & Nigam, 2004; Mayer, 2004). Nevertheless, this is not to say just filling the curriculum with all lecture-style instructions will be fine. Mayer (2004) proposed that guidance, structure, and focused objectives should be incorporated into activities of inquiry and discovery. Instead of making an oversimplified dichotomous choice between "to discover" or "not to discover", what is more important should be an attempt to balance the two extremes of teachers' instructions and students' constructions with the goal of actually promoting cognitive activities (Rosé et al., 2005). The questions researchers are confronted with now are: how much guidance, in what form, and under what condition can serve as scaffolding and make inquiry learning an effective approach (de Jong, 2005).

Virtually every inquiry activity begins with "asking questions", and then students may be motivated to move on to "finding answers", and subsequently, "asking *better* questions" that incrementally leads students to really learn from inquiry and discovery. Idea-generation and CPS activities are potentially useful to be cast as inquiry activities. Instead of asking students to conceive ideas or solutions that can be used in problem solving, if designed properly, we may ask students to conceive some questions at a meta-level against a CPS task, which may lead to better observation of the task structure for performing better idea-generation at a later round. The "question-generation" activity appears to introduce new opportunities and challenges to the design of system behavior and instruction, but should be a worthy goal for our project to pursue in the long term.

RESEARCH QUESTIONS AND HYPOTHESES

Group versus Individual

One set of research questions we must address relate to what aspects of creative problem solving are best addressed with individual learners and which with groups. We expect that the ultimate answer will involve a

balance of these two. Dillenbourg (1999) indicated that a variety of collaborative activities may contribute to learning, which could be course assignment sharing, joint problem solving, or even performing cooperative work. A driving force behind learning in groups could be that group members will typically bring unique resources, perspectives and backgrounds into collaborative activities (Kraut, 2003) However, as argued above, the phenomena of process losses cast doubts on the oversimplified assertion that group must be better than individual in terms of work performance, but also revealed the needs of sophisticated analyses on the micro features and interaction dynamics of a group.

Performance versus Learning

In the taxonomy presented in Table 1, the distinction between performance-oriented and learning-oriented brainstorming activities has been made. However, this is not to say that the activity of idea-generation has nothing to do with learning. As emphasized previously, we consider idea-generation as the basis for other higher-order brainstorming activities that would lead to valuable instructional and learning opportunities. An issue of interest to educational practitioners and organizations is about how do we enable the transfer of what students learn to improve their work performance in solving new problems? It appears that brainstorming activities and the taxonomy may fit well to the investigation of relations among learning process, transfer, and work performance.

The distinction between work performance and learning may be confounded with the comparison of group versus individual performance. It should be clear that the evaluation criteria or desirable results for work performance and learning are rather different. For example, for the debris-flow problem solving task introduced in the Introduction, if the purpose is to evaluate the work performance, the number of good ideas generated or viability of the solutions, either individually or collaboratively, for preventing the damage of debris-flows may serve as the criteria. However, if the purpose is to evaluate how well do *individual* students learn or acquire knowledge, measures addressing individuals' knowledge status or meta-cognitive skills are required. It is unlikely to be valid in measuring learning at the group level. For the measure of learning, beyond traditional educational achievement testing, it is considered that the technology of user profiling may serve the later purpose well. It is argued that different measures should be constituted and applied in research questions regarding performance and learning respectively.

Learning and Transfer

Through the lens of cognitive science, it is essential to ask questions of what do we expect students to learn and transfer by doing these brainstorming activities, either within a group or in a solitary manner, either supported by other agents (human or computer) or not. As prior work in this area have posited, one of the central issues in education is about *transfer*, which refers to how knowledge acquired in one situation or task can be applied in another, perhaps novel and unfamiliar, situation (Anderson, 1993, Anderson et al., 1995).

Two targets of transfer are possible for the research to pursue. *First*, more conventionally, it is intended to enable the learning and transfer of domain knowledge through brainstorming activities. For example, for the debris-flow question, although students can be taught declaratively about the physics, geology, and ecology related to the phenomena of debris flows, an open ended CPS task would give students the opportunity to actively summarize and interconnect what they have acquired via natural language, detect what they have not mastered, and seek instruction accordingly. What is different from what is typical of traditional ITSs is the absence of a prescribed ideal solution path as well as the presence of peer collaboration. Therefore we do allow more exploration and variability of students' behaviors. Empirical studies in comparing CPS-tutoring and typical ITS approach as planned in Table 1 may verify the value of exploration and collaboration in terms of domain learning. *Second*, perhaps debatably, we may go on to investigate the possibility of improving students' meta-cognitive skills such as information foraging, self-explanation, or more generally, creative problem solving ability. In other words, the question is, after students practicing one CPS task, can we expect students to perform better or equally well in another previously unfamiliar CPS task? Will students become more capable or strategic in interacting with peers and seeking for information? Our current focus is not on this aspect, but it could be interesting, also challenging, to touch these questions in the future.

System Tuning and Optimization

In order to optimize VIBRANT's instructional effects, we may incorporate research evidences resulting from intervention studies as proposed in Table 1. Thus, our plan is to use VIBRANT as a research platform in which to address our research questions and then to fold our findings back into the design of new versions of VIBRANT in order to offer students higher quality educational opportunities.

For example, in the brainstorming activity of idea-generation, we are interested in knowing what would be better, either the system instantiates a single agent or a group of agents to interact with students, where agents

may be computer agents, human participants, or some combination of the two. We may go on to ask if a group of agents should be better, whether the group members should be homogeneous or heterogeneous in opinions, and also, whether the brainstorming feedbacks provided by agents should sound evaluative (i.e., more critical) or supportive.

An experimental paradigm used in our prior work (Gweon et al., 2006), which we may adopt is to employ “confederate peer agents”, or humans behaving in a highly prescriptive manner, in order to investigate effects of group dynamics on individuals. We believe that the confederate peer agent experimental paradigm provides an appropriate level of experimental control while allowing us to evaluate the impact of agent capabilities beyond the current state-of-the-art for handling open-ended inputs. However, for some studies we may take a different approach. For example, when the intention is to study the effects of these interventions in groups with larger size, say 50, it may be more practical to program VIBRANT properly to enable such an experiment. It is foreseen that VIBRANT will play dual roles in our future work, at the one hand, it is a tutoring system that can be used in educational activities, while at the other hand, it serves as a research tool to enable simulations of group dynamics for shaping future computer supported collaborative learning.

SUPPORTIVE TECHNOLOGIES

Based on the taxonomy of brainstorming activities, at an abstract level, VIBRANT provides two types of support for each specific activity, which are (1) cognitively oriented support providing brainstorming feedback and (2) socially oriented support for discussion group formation. As an abbreviation, the former is called as the *intelligent support*, and the later is called the *social support*. Certainly, the actual system behavior for each specific brainstorming activity will need further adaptation and tuning, such as a dialogue script capable of differentiating evaluative tone from supportive tone in providing intelligent support, and a matching function for recommending either the most similar peer or the most different peer in providing social support. The design decisions of system behavior can be made in an evidence-based manner by making references to interventions studies proposed in our research framework. VIBRANT, as a web-based system, is designed to be versatile and flexible to mimic particular behaviors suggested by our studies easily. Figure 1 depicts the architecture of VIBRANT, which mainly consists of three functional modules, including the brainstorming agent, the user modeling (UM) agent, and the user interface (UI) at the client side. In this section, we describe the system characteristics at an abstract level that aims to generalize across various brainstorming activities.

User Modeling

In order to provide adaptive support in response to students’ exploration of the task when performing idea-generation, CPS, and active inquiry, a learning environment would require corresponding user-modeling technologies integrated for tracing students’ knowledge status and performing instructional decisions accordingly.

Based on our prior work (Wang et al., 2005a, 2005b), knowledge of solving a CPS task is modeled as a bipartite graph-based formal user profile (fUP), in which the connections between a student’s ideation (i.e., section of divergent thinking) and explanation (i.e., section of convergent thinking) are explicitly represented. For a particular CPS task, student U ’s ideation in solving the problem is represented as a set of ideas $A_U = \{a_{U1}, a_{U2}, \dots, a_{Un}\}$, and explanation is represented as a set of reasons $B_U = \{b_{U1}, b_{U2}, \dots, b_{Um}\}$. A fUP is denoted as an undirected bipartite graph $G_U = (V_U, E_U)$ where $V_U = A_U \cup B_U$ and $A_U \cap B_U = \emptyset$. The mapping between A_U and B_U , modeled as $E_U = \{e_{ij}\}$ representing a linkage between an idea a_i and a reason b_j , is of a many-to-many nature, in which one idea may have several reasons, and several ideas may connect to an identical reason. If the focus of study is only on the activity of idea-generation, then a simple fUP that contains only A_U would suffice.

A prescriptive Domain fUP as shown in Figure 1 is created by the domain expert by using appropriate authoring tools. The domain fUP is employed by the system in building user profiles. A collection of historic fUPs is managed by the system and can be retrieved later for fulfilling specific decision making as well as offline collaboration among peers.

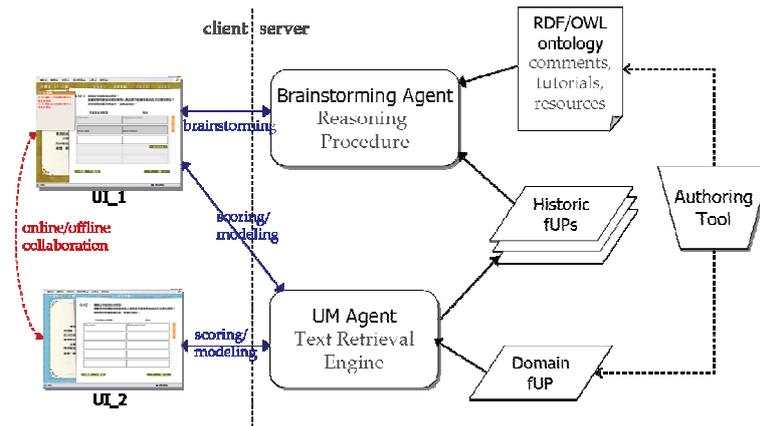


Figure 1. System Architecture for VIBRANT

Intelligent and Social Support

A formal ontology is constructed serving as the core device for organizing experts' CPS ideas and learning resources, including feedback texts and learning materials. The ontology consists of an *is-a* hierarchy organizing experts' ideas into several levels of abstraction. Prescriptive feedback messages are attached to specific idea nodes at lower levels and categorical nodes at higher levels in the ontology.

Finite state machines (FSMs) are designed to retrieve learning resources such as feedback. In FSMs, the finite set of *states*, Q , represents the range of the system's functional behavior, including actions such as *check_coverage* or *move_upward_in_hierarchy*, while the finite *alphabet*, Σ , represents the set of events that can trigger a transition from one state to another, such as *all_sub_nodes_covered* which in this case triggers a transition to a state called *get_new_category*. Transition functions $\delta: Q \times \Sigma \rightarrow Q$ represent designers' instructional decisions of what behavior should be triggered when particular events are observed, which are particularly useful to adapt system behaviors to accommodate specific brainstorming activities.

The *feedback* prepared by the system consists of two parts, a *comment* and a *tutorial*. Two separate FSMs are designed for the generation of each part. *Comments* are evaluative texts responding directly to the most recent idea submitted by the student to the system, while the *tutorial* is the instruction that directs the student to the next logical focus node, which may either be an idea node or a categorical node, selected by the system adaptively based on its model of the student. In the current design, a *comment* is a function of the current idea entry, while a *tutorial* is a function of the current idea entry and the student's fUP built incrementally during the brainstorming process. In other words, a context of students' previous responses in the same session is incorporated into the mechanism for retrieving *tutorials*.

The use of the *is-a* hierarchy is considered beneficial for the FSM-based feedback generation. *First*, the hierarchy of topics provides a basis for supporting a more organized and coherent brainstorming process. The system may select a next focus for tutorial to maximize the students' local coverage of categorical nodes that have been partially addressed by the students' idea entries. The system may then proceed to provide feedback that may lead students to cover all related categorical nodes at a higher level. *Second*, *comments* can be fetched strategically at a more generalized level in the ontological hierarchy when a particular idea proposed by the student is semantically ambiguous and thus results in low similarity scores as computed using vector-based information retrieval methods. The strategy may help remedy the insufficiency of IR-based methods for computing semantic similarity and to improve the relevance of system-prepared comments against students' ideas. Along with a student's producing more ideas, her/his profile, fUP, also evolves incrementally. A later instructional decision, including feedback generation or social recommendation, will further make use of these fUPs.

Given a collection of historic fUPs done by previously visited students, we may re-model the system of fUPs as a tripartite graph with hyperedges. The condition of student p having an idea q and explaining it by reason r can be represented as a hyperedge e_{pqr} , which results in a tripartite graph $H=(V, E)$ where $V=S \cup A \cup B$ and $E=\{e_{pqr}\}$. A variety of analyses can then be computed over the tripartite graph for social structure discovery. *First*, one may conceive local heuristics to extract particular (hyper-)edges as cues for social recommendation. *Second*, we may apply Social Network Analysis (SNA) methods (Wasserman & Faust, 1994) such as co-occurrence analysis against the tripartite graph for clustering users and then forming discussion groups

according to the structural and global information. With this tripartite graph, the same techniques can also be used to recommend a next brainstorming focus of an idea in a data-driven manner.

PRELIMINARY EVALUATION

We conducted an evaluation of VIBRANT feedback generation using a corpus containing 163 entries of self-generated ideas from 25 Taiwanese high school students on the debris-flow question as inputs for simulation. For each input entry, two comment/tutorial generation methods, *with-hierarchy* (H) or *without-hierarchy* (NH), were invoked to generate two versions of simulated comments and tutorials. The H method generates feedback using the aforementioned approach, while the NH method does not make use of a category-based brainstorming plan, so that the brainstorming focus motivating the tutorial is selected randomly from the pool of non-covered concept nodes in the domain model. In the NH method, the comment offered is the comment attached to the most similar node in the domain model, and *no* remedial device was used against potential low relevance of comments retrieved at the instance-level. Note that two types of message were evaluated separately, comment or tutorial texts, and therefore, results of two parts of message generation, comment generation (CG) and tutorial generation (TG), are reported.

We recruited two independent judges to rate the quality of the comment and tutorial offered by the two different methods. The two judges were asked to evaluate the quality of each comment/tutorial text against the idea entry it targets. Each coder assigned a binary score of *acceptance* (Acceptable/Not_Acceptable) to the comment/tutorial message generated by the two methods against the same idea entry. The judges then assigned a nominal score of *subjective preference* (H: With-hierarchy/ NH: Without-hierarchy/ S: Same/ N: Neither) to indicate which version among the pair of messages they preferred or considered better against this idea entry. The coders were blind to the method used in generating the messages. Although the two versions of message were presented in pair for rating, the order of which version presented in the left side and which one in the right was randomized. After data cleaning which excluded pairs containing empty messages generated by either method, totally 138 pairs of comments (85% comment pairs) and 153 pairs of tutorials (94% tutorial pairs) were used in data analyses.

The dependent variable of *acceptance* indicating the proportions of messages accepted for the H and NH methods was first analyzed. Independent-samples chi-square tests were conducted to examine the difference between H and NH. For the CG part, no statistically significant differences in comparing H vs. NH were found in either coder's ratings. For the TG part, tutorials generated by both methods also appear to be equally well with no significance, no matter evaluated by which coder. We did find significantly different distribution on the variable of *subjective preference* for TG over CG. Statistically significance were detected by using chi-square analyses, specifically for coder 1's preference votes: $\chi^2(3, N=291) = 26.621, p < .001$, and for coder 2's: $\chi^2(3, N=291) = 34.843, p < .001$. In a post-hoc inspection of the data, we found that H was preferred more in the TG condition rather than in the CG condition.

In summary, we found that in the TG part, the with-hierarchy method was rated higher than the without-hierarchy method, which is very different from the trend revealed in CG. The results may imply that the with-hierarchy method can help produce better tutorial texts, but not comments, as supports in response to students' ideas generated. This makes sense that the selection of brainstorming foci in TG for guiding students' exploration would benefit from the context (i.e., the ontological hierarchy and students' fUPs), while local information is sufficient for CG. The preliminary evaluation shows evidences of usefulness of our current design on the core component in VIBRANT. Note that what was evaluated here was the quality of feedback but not its instructional effects against real students. Evaluations on instructional effects are planned in the near future. In addition, in this evaluation, the inter-rater reliability between the two coders was found to be low. We also aim to refine the operationalization of "acceptance" and "preference" to enhance the inter-rater reliability in replicating a similar evaluation design.

CURRENT DIRECTIONS

In this paper, we describe our taxonomy and technologies for creative group brainstorming. Three kinds of brainstorming activities identified in the context of science education are idea-generation, creative problem solving, and inquiry learning. The user modeling-based learning environment, VIBRANT, is proposed to incorporate feedback generation and social recommendation technologies to support various brainstorming activities by properly authoring the system behavior. Interventions and corresponding measures are proposed to better understand group dynamics in brainstorming groups in various tasks, which may inform the design of future intelligent tutoring systems for ill-defined domains.

REFERENCES

- Aleven, V., & Koedinger, K. R. (2002). An effective metacognitive strategy: learning by doing and explaining with a computer-based Cognitive tutor. *Cognitive Science*, Vol. 26, pp. 147-179.
- Anderson, J. R. (1993). *Rules of the Mind*. NJ: Lawrence Erlbaum Associates.
- Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995). Cognitive Tutors: Lessons Learned. *Journal of Learning Sciences*, 4(2), 167-207.
- Anderson, J. R., Reder, L. M. & Simon, H. (1998). Radical constructivism and cognitive psychology. In D. Ravitch (Ed.) *Brookings papers on education policy 1998*. Washington, DC: Brookings Institute Press.
- Basadur, M. (1995). Optimal Ideation-Evaluation Ratios. *Creativity Research Journal*, Vol. 8, No. 1, pp.63-75.
- Chang, C-Y., Weng, Y-H. (2002). An Exploratory Study on Students' Problem-Solving Ability in Earth Sciences. *International Journal of Science Education*, 24(5), pp. 441-451.
- Connolly, Terry. (1993). Behavioral decision theory and group support systems. (pp.270-280). In L. Jessup & J. Valacich (Eds.). *Group support systems*. NY: Macmillan.
- Diehl, M., & Storebe, W. (1987). Productivity loss in brainstorming groups: toward the solution of a riddle. *Journal of Personality and Social Psychology*, 53(3), 497-509.
- Dillenbourg P. (1999). What do you mean by collaborative learning?. In P. Dillenbourg (Ed) *Collaborative-learning: Cognitive and Computational Approaches* (pp.1-19). Oxford: Elsevier
- de Jong, Ton. (2005). Scaffolding inquiry learning: How much intelligence is needed and by whom? Invited speech in 12th International Conference on Artificial Intelligence in Education, *Frontiers in Artificial Intelligence and Applications*, 125, 4-4.
- Gweon, G., Rose, C. P., Carey, R., & Zaiss, Z. S. (2006). Providing support for adaptive scripting in an on-line collaborative learning environment. *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI 2006)*.
- Hill, G. W. (1982). Group versus individual performance: are $N+1$ heads better than one? *Psychological Bulletin*, 91(3), 517-539.
- Klahr, D., & Nigam, M. (2004). The equivalence of learning paths in early science instruction: effects of direct instruction and discovery learning. *Psychological Science*, 15(10), 661-667.
- Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education*, 8(1), 30-43.
- Kraut, R. E. (2003). Applying Social Psychological Theory to the Problems of Group Work. In J. Carroll (Ed.), *HCI Models, Theories, and Frameworks* (pp. 325-356). NY: Morgan-Kaufmann Publishers.
- Kumar, R., Rose, C., Aleven, V., Iglesias, A., & Robinson, A. (2006) Evaluating the Effectiveness of Tutorial Dialogue Instruction in an Exploratory Learning Context. *Proceedings of International Conference on Intelligent Tutoring Systems (ITS 2006)*.
- Mayer, R. E. (2004). Should there be a three-strikes rule against pure discovery learning? the case for guided methods of instruction, *American Psychologist*, 59(1), 14-19.
- National Science Foundation (2002). Inquiry: thoughts, views and strategies for the K-5 classroom, *Foundation*, Vol. 2.
- Osborn, A. (1963) *Applied Imagination: Principles and Procedures of Creative Problem Solving*. New York: Charles Scribner's Sons.
- Rose, C. P., Aleven, V., Carey, R., & Robinson, A. (2005). A first evaluation of the instructional value of negotiable problem solving goals on the exploratory learning continuum. *Proceedings of 12th International Conference on Artificial Intelligence in Education (AIED 2005)*, 563-570.
- Wang, H-C., Chang, C-Y., & Li, T-Y. (2005a). Automated scoring for creative problem-solving ability with ideation-explanation modeling. *Proceedings of 13th International Conference on Computers in Education (ICCE 2005)*, 522-529.
- Wang, H-C., Li, T-Y., & Chang, C-Y. (2005b). A user modeling framework for exploring creative problem-solving ability. *Proceedings of 12th International Conference on Artificial Intelligence in Education (AIED 2005)*, 941-943.

- Wang, H-C., Li, T-Y., Rose, C. P., Huang, C-C., & Chang, C-Y. (2006). VIBRANT: A Virtual Brainstorming Agent for Computer Supported Creative Problem Solving. *Proceedings of 8th International Conference on Intelligent Tutoring Systems (ITS 2006)*.
- Wasserman, S., & Faust, K. (1994). *Social Network Analysis: Methods and Applications*. NY: Cambridge University Press.
- Weinberger, A., & Fischer, F. (2006). A framework to analyze argumentative knowledge construction in computer supported collaborative learning. *Computers & Education*, 46, 71-95.
- VanLehn, K., Lynch, C., Schulze, K., Shapiro, J. A., Shelby, R., Taylor, L., Treacy, D., Weinstein, A., & Wintersgill, M. (2005). The Andes physics tutoring system: five years of evaluation. *Proceedings of 12th International Conference on Artificial Intelligence in Education (AIED 2005)*, 678-685.

APPENDIX

Table 1. A research framework for guiding future works

<i>Theory</i>	<i>Idea generation</i>	<i>Creative problem solving</i>	<i>Inquiry learning</i>
	social psychology	creativity research, science education	science education, learning sciences
<i>Continuum (Evaluation)</i> <i>What to be learned?</i>	<i>Performance (ideas/solutions produced)</i>		
<i>Intervention</i>	<i>Subject matters</i>		
	<i>N/A</i>		<i>Meta-cognitive skills</i>
	ig1. group sizes (including single vs. group)* ig2. homogeneity of group* members ig3. anonymity ig4. types of feedback (e.g., evaluative vs. supportive feedback)	ig1-ig4 cps1. different problem solving tasks for the same topic (e.g., less vs. more open ended problem solving tasks) cps2. orientation of tutorial dialogue (e.g., divergent thinking-oriented prompts vs. convergent thinking-oriented prompts)	ig1-ig4 cps1. degree of guidance available (e.g., unstructured vs. structured information sources) iqr2. amount of error information
<i>Measurement</i>	◆ number/quality of generated ideas ◆ member attitude (e.g., satisfaction)	◆ number/quality of generated ideas ◆ completeness/novelty of CPS entries ◆ likelihood of having the problem solved ◆ domain achievement test ◆ member attitude	◆ number/originality of proposed questions ◆ (retentive) domain achievement test ◆ member attitude

* group members can be either real students or simulated agents, so a group can be a mix-up of human peers and simulated agents