Supporting Collaborative Idea Generation: A Closer Look Using Statistical Process Analysis Techniques

Hao-Chuan WANG and Carolyn P. ROSÉ
Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213, USA
{haochuan, cprose}@cs.cmu.edu

Abstract. We are developing a conversational agent called VIBRANT to provide adaptive support for brainstorming in pairs in a scientific inquiry context. Our previous experimental study indicated that although adaptive support was effective for learning, it did not mitigate the negative influence of social interaction on idea generation. In this paper, we present a process analysis that makes the effect of the adaptive support on the idea generation process more apparent and suggests directions for future work.

Keywords. Collaborative Learning, Process Analysis, Inquiry Learning

Introduction

In this paper we explore the use of statistical process analysis techniques to examine data from an experimental study in a scientific inquiry learning setting in order to gain insights into how better to design support for productive collaborative learning interactions. Idea generation tasks related to scientific inquiry learning, simulation based learning, or project based learning are often done in groups despite overwhelming evidence of process losses during group brainstorming [2]. Learning may result from the brainstorming process itself, as it provides the impetus to engage in constructing bridging inferences similar to the process of self-explanation [1]. A significant correlation between idea generation productivity and learning in our data supports the view that learning from brainstorming may come from this constructive process [5]. Thus, an important question for collaborative learning for idea generation tasks is how to support productive idea generation and learning using adaptive collaborative learning support technology. We demonstrate how the methodology we propose can be used to gain insights such as these from corpus data.

We illustrate by example a methodology whereby a corpus of conversational interactions is first annotated with a categorical coding scheme, and how this process can be facilitated with a publicly available toolset called TagHelper tools (http://www.cs.cmu.edu/~cprose/TagHelper.html). A statistical process analysis is then conducted on top of this representation. We conclude with a discussion of lessons learned.

1. Motivation

Previously we have presented the results of an experimental study investigating the connection between brainstorming productivity and domain learning during a scientific inquiry learning task, namely the Debris Flow Hazard Task (DFH) [5]. The DFH task involves two related brainstorming tasks, one focused on *problem finding*, and the other related to *problem solving*. In the remainder of the paper we refer to these two brainstorming tasks as task 1 and task 2. The experimental study that provides a running example for this paper is a 2X2 factorial design with two between subjects factors. The first independent variable we manipulated in our experimental study was the Pair factor, i.e., whether students worked individually or in pairs. The second independent factor is referred to as Support. For this factor we manipulated whether or not students received adaptive support from the VIBRANT agent [5], which provides stimulus in the form of suggested categories of ideas [4]. A total of 42 Taiwanese high school students participated in the study. In the two conditions where students worked individually, 7 students were randomly assigned to each condition, while in the two pairs conditions, 14 students were randomly selected to form 7 pairs. The experimental manipulation took place during the first brainstorming task, which is the task related to problem finding. In addition to evaluating the success of idea generation productivity during the main brainstorming task, we also measured pre-/post-test learning gains as well as productivity on the second brainstorming task building on the earlier task, which was performed individually as a practical assessment.

The results of the experiment demonstrated a strong connection between idea generation productivity and learning. Students in the pairs condition were significantly less productive and learned significantly less during the initial brainstorming task than students in the individual condition. On the other hand, the students who brainstormed in pairs during the first task performed better on the second brainstorming task following up on the problem finding task with a related problem solving task. Furthermore, although Support had a positive effect on learning both in the individual and pairs conditions, it did not have a significant positive effect overall on productivity during the initial brainstorming task, although there was a trend for students in the supported conditions to have a higher productivity in their idea generation. We were intrigued to find that students in the individuals with support condition, who performed best on the first brainstorming task and gain most between pre and post-test, performed worst on the second brainstorming task. The best preparation for the second brainstorming task was brainstorming in pairs with support during the first task. Thus, we were left with the mystery of how to achieve success consistently across all of our dependent measures.

The evidence suggests that working in pairs allowed students to approach the first brainstorming task more broadly, seeing it as problem finding in preparation for problem solving rather than problem finding more narrowly. There was a significant main effect of Pairs on the conditional probability that a solution was offered in the second brainstorming task given that a related problem was identified during the initial brainstorming task $F(1,38) = 4.19$, $p < .05$, d=.6. We also have evidence that this conditional probability, operationalizing the extent to which students conceptualized the first brainstorming task as preparation for problem solving, mediated the effect of condition on task 2 success in that there was a significant correlation between this conditional probability and task 2 success, $R^2=.49$, $p<.0001$, N=42. Because working in pairs with support lead to positive effects on the second brainstorming task, in this paper we take a closer look at the impact of the support provided by the VIBRANT agent on the idea generation process to better understand how it can effectively be supported with the ultimate goal of achieving high productivity on both brainstorming tasks as well as high learning gains as measured by the test.

2. Process Analysis

With our process analysis, we evaluate the effect of our experimental manipulation on the process of idea generation in connection with two important characteristics of idea generation, namely coherence [4] and the temporal characteristics of the idea generation process [2]. Specifically, by coherence we are referring to the finding that idea generation is more efficient when topic transitions between related ideas are more frequent than transitions between unrelated ideas. As for the temporal characteristics, it is conjectured that productivity losses occur mainly at the earliest stage of group idea generation. The logical consequence of this hypothesis is that the discrepancy between individual and group idea generation may vary over time. In our analysis, we make use of the discourse from our idea generation study, which has previously been coded with a set of 19 specific ideas that fit into 5 general idea topics, with a reliability of .81 Kappa [5]. This topic based analysis can be replicated automatically with high reliability (Kappa .7) using the TagHelper tools package http://www.cs.cmu.edu/~cprose/TagHelper.html.

In the first analysis, we modeled students using probability distributions that characterized their trajectory of shifting between the five general topics in the DFH domain during idea generation. Given a set of potential brainstorming ideas W, for any student s in the set of participating students S, the list of generated ideas is represented as $C(s) = \langle c_1, \ldots, c_{n-1}, c_n \rangle$ where c_n denotes the general topic of the idea w that s expressed at time stamp t_n . We may then quantify the pattern of topic transitions for a later comparison across students by estimating the conditional probability of an idea being generated within a given general topic given the student and the general topic of the last idea uttered by that student. There were five general topics of "valuable ideas" in this domain, and we also considered non-valuable ideas as a separate category of ideas (i.e., ideas out of the domain model). Altogether, probabilities for 36 transitions were thus computed for each student.

We first explored the effect of our experimental manipulation on the coherence of idea generation. In order to do this analysis, we classified each topic transition as within topic or between topics. If the average of within topic transitions for condition X are higher than in condition Y, we can say topic transitions within condition X are more coherent than in condition Y. We did an analysis of variance (ANOVA) with conditional probability associated with a topic transition as a repeated dependent measure and the two independent variables from our experimental manipulation as well as the within versus between topic factor (WVBT) and Student nested within Pair and Support as independent variables. There was a significant main effect of WVBT such that transitions within topic were more frequent overall than between topics $F(1,1466) = 47.0$, $p < .0001$. More importantly, there was a significant three way interaction between Pair, WVBT, and Support $F(1,1466)= 6.3$, $p < .01$. A Tukey posthoc analysis shows that supported brainstorming in pairs is significantly more coherent than unsupported brainstorming in pairs. Brainstorming for individuals with or without support is not statistically different from the other conditions. Thus, we conclude that our brainstorming support had a positive effect on coherence, at least in the pairs condition. Infrequently students mentioned ideas during the first task that counted as solutions for the second task. These occurred mainly in the pairs with support condition. On the whole this evidence suggests that students who worked with support were encouraged to engage with more vigor in the problem as they construed it. As discussed in the previous section, students who worked individually construed the problem more narrowly, and achieved greater success on that narrowly defined task. Considering the earlier analysis in connection with this coherence analysis, we see that students who worked in pairs construed the problem more broadly and stayed on the same topic longer, which afforded them the opportunity to pursue the problem more fully.

One might be concerned that the effect of support on coherence might mean that we must sacrifice variety in the exploration of the idea space. However, an additional statistical analysis alleviated these fears. By modeling students as probability distributions, we can compare students to one another based on Kullback-Leibler divergence [3]. For two students with probability vectors PV_1 and PV_2, a symmetric distance measure of how similar their behaviour is can be derived by averaging two values $D_{KL}(PV_1 \| PV_2)$ and $D_{KL}(PV_2 \| PV_1)$. For every student, we computed the average distance between her/him and other two types of students—either in the *same* experimental condition or in all of the *different* ones. The two measures in effect captured the information of within-condition and between-condition divergence in terms of topic transitions. A greater average distance measure in a specific condition can be interpreted to mean that condition encourages more variety in idea generation patterns. A repeated measures ANOVA analysis was performed to compare the homogeneity of topic transition patterns across experimental conditions and types of measures (i.e., "Same" or "Different"). Our independent variables included Pair, Support, Same versus Different (SVD), and Student nested within Pair and Support. Most importantly, there was a significant main effect of Support, such that adaptive support enhanced variety $F(1,38) = 177.7$, $p < .0001$.

Now we extend our process analysis in order to look at how idea generation productivity changes over time. To achieve a fair comparison between individual idea generation and group idea generation, it is standard to form "nominal pairs" by randomly selecting people from experimental conditions in which participants worked individually, and then pool their ideas collectively for a comparison with real pairs [2]. We examined the effect of our experimental manipulation at different time points during the 30-minute long brainstorming session. The first five minutes of brainstorming were the most intense, and roughly half of the ideas contributed were contributed during that span of time. Up through the first 5 minute time point, a significant evidence of process loss was found ($F(1,24)=12.22$, $p<.005$), with an effect size of 1 standard deviation. During this time we do not see a positive effect of Support, and in fact see evidence of the opposite, that Support from the VIBRANT agent may have interfered with productivity during that first segment of brainstorming. However, if we then look just at brainstorming behavior after the five minute mark, we still see significant process losses due to working in pairs, but the effect size is smaller, specifically .61 ($F(1,24)=4.61$, $p<.05$). Furthermore, we find a significant main effect in favor of the Support condition, $F(1,24)= 16.43$, $p<.0005$, effect size = 1.37 s.d. Moreover, by comparing nominal pairs who received no support and real pairs who received support after the first 5 minutes, not only is there no significant disadvantage to working in pairs, we see that there is a trend in favor of supported group brainstorming. Thus, we find that support is indeed effective for mitigating process losses, just not within the first 5 minutes.

3. Conclusions

A summative evaluation of our experimental study left us with a mystery to solve, namely the question of why the condition in which students achieved the greatest success in connection with task 1 productivity and with learning was the condition where they performed worst on task 2. A statistical process analysis building upon a categorical coding of the conversational behavior revealed insights that offer us suggestions for moving ahead and gaining success consistently on all of the dependent measures we have discussed. In subsequent studies we plan to require students in all conditions to brainstorm individually for the first five minutes with no support and only then to work in pairs with support for the duration of the first task.

This work was funded in part by National Science Foundation grant number SBE0354420.

References

[1] Chi, M. H., De Leeuw, N., Chiu, M., Lavancher, C. (1994). Eliciting self-explanations improves understanding. Cognitive Science, 18, 439-488.

[2] Diehl, M. & Stroebe, W. (1991). Productivity loss in idea-generating groups: tracking down the blocking effect. Journal of Personality and Social Psychology, 61(3), 392-403.

[3] Kullback, S., and Leibler, R. A. (1951). On information and sufficiency, Annals of Mathematical Statistics 22: 79-86.

[4] Nijstad, B. A., & Stroebe, W. (2006). How the group affects the mind: a cognitive model of idea generation in groups. Personality and Social Psychology Review, 10(3), 186-213.

[5] Wang, H., Rosé, C. P., Cui, Y. Chang, C., Huang, C., & Li, T. (to appear). Thinking Hard Together: The Long and Short of Collaborative Idea Generation, Computer Supported Collaborative Learning Confernece 2007 (CSCL 2007).