

Positive and Negative Learning Impacts from Technological Social Agents

Jonathan S. HERBERG^{a,b,*}, Daniel T. LEVIN^c & Martin SAERBECK^a

^a*Institute of High Performance Computing, A*STAR, Singapore*

^b*Department of Psychology, National University of Singapore, Singapore*

^c*Department of Psychology and Human Development, Vanderbilt University, United States*

*herbergjs@ihpc.a-star.edu.sg

Abstract: Computerized learning environments often include implementations of simulated social agents, incorporating the reasonable design assumption that learning interactions that are more social are more effective. However, recent evidence suggests that computerized social agents can in some circumstances fail to promote or even hinder learning. The current paper outlines evidence both supporting and arguing against the utility of computerized social agents for learning. We propose a framework reconciling this evidence by delineating impacts from a social agent that impinge upon different points of learning, from shallow to deep phases. Shallow social impacts when ineffectual carry little to no potential to actively impede learning, but any potential positive impacts on shallow learning phases are relatively limited in the absence of positive deep impacts. Deep social impacts, on the other hand, carry the potential to strongly drive deep learning, but when ineffectual carry the risk to impede it. The paper concludes with a proposal for future research based on this framework.

Keywords: Social agents, pedagogical agents, social cognition, theory of mind

Introduction

The notion of the beneficial impacts of a social interaction context has been incorporated into various socio-cognitive frameworks for learning. On this basis, computerized learning environments often include simulated social agents, which can be used either to fill the role of the tutor [1] or, to generate learning-by-teaching, to fill the role of the agent that the student is supposed to teach [2].

However, a number of tests suggest limitations to the effectiveness of social agents for learning, in particular of those that fail to provide adaptive explanatory feedback [3]. This suggests there may be added layers of extraneous processing associated with a simulated social agent, which in the absence of cues to channel the behaviors for effective learning can sum to a net cost rather than benefit [4]. Related to this, situations in which people must use their representation of a social agent partner's differing knowledge states typically require extra cognitive effort [e.g., 5, 6]. Thus, in the context of computer based learning environments, it is important to consider whether the cognitions and behaviors that can be elicited by social agents sometimes result in net learning costs when they function as diversions from, rather than reinforcements of the learning context.

A recent study has suggested that a social interaction context may in some situations generate a greater cost than benefit to learning even when the social agent seemingly elicits more active learning behaviors [7]. In this study, participants' transfer of their knowledge of how to solve the 3-ring Tower of Hanoi problem in order to solve the 4-ring problem was less robust after they demonstrated their 3-ring solution to a social agent, a person, as compared to those who demonstrated to a non-anthropomorphic computer. Those

demonstrating to the social relative to non-social agent engaged in more behaviors such as highlighting solution steps for the agent while demonstrating. Therefore certain social interaction behaviors may in some contexts actually impede learning. In the current paper, we develop a framework that classifies different social agent impacts as beneficial versus costly to sustained, deep self-regulated learning. This framework provides general parameters on when social interactions should help and when they should hurt deep learning resulting from knowledge construction, suggesting fruitful avenues of future research.

1. Learning Effects from Social Situations and Computerized Social Agents

1.1 Learning Behaviors and Depth of Learning

We begin by examining behaviors that may stem from an interaction with a social agent that may enhance learning. A framework by Chi separates out active, constructive, and interactive learning behaviors [8]. Active learning behaviors are comprised of physical movements and engagement in learning. Constructive behaviors are seen in the learner's verbal reasoning, logical elaborations and links, and explaining of solutions. Cognitively, the learner now recruits knowledge to infer and construct new knowledge. Finally, interactive behaviors can be observed when two learning partners listen to each other, respond to each other, and argue or collaborate with each other in their problem solutions.

From this sketch of learning behaviors at different depths it is possible to classify specific mechanisms that can drive beneficial learning effects from a social agent. Social agents may generate learning benefits by, on the one hand, eliciting from the learner more attention and engagement in the task and by increasing the activation of already-existing task-relevant knowledge. On top of such active learning behaviors, social agents may shape a path of constructive learning by pushing the learner to engage in deep reasoning, question seeking and answering, and inference of new knowledge from prior knowledge. Social agents may induce this form of deep learning by prompting the learner to explain, reason with, and build conceptual links between different pieces existing knowledge, and to make inferences of new knowledge on the basis of existing knowledge. Social agents may, furthermore, push the learner to reflect on what he or she knows and what needs to be learned. Such metacognitive behaviors, crucial for deep abstract learning, occur most vigorously when the learning interaction is rich and sustained [8, 9]. At the deepest level of interaction between two learning agents, the two will build on each other's knowledge and make joint inferences of and elaborations on new knowledge [10]. Some examples of the most effective computer-based learning environments in generating robust learning outcomes use simulated social agents for dialogic interactions, whether in the form of a kind of Socratic tutor who scaffolds learning [1], or in the form of a peer-agent the student builds concept maps for which the agent in turn indicates understanding on the basis of [2].

1.2 Self-Regulated Learning and Depths of Learning

Ultimately, the goal of computer or robotic implementations with simulated social agents should be toward deep knowledge-building through self-regulated learning (SRL). SRL embodies a set of purposeful, self-directed learning behaviors and cognitions, leading to deeper understanding and better repair of misconceptions than externally-regulated learning. A successful self-regulated learner continuously monitors what he or she knows and how much he or she is learning on a given learning task, and on that basis sets and updates goals and strategies to learn what needs to be learned [1, 11]. Thus, a learner engaging in effective SRL is constantly devoting significant cognitive resources and effort

toward metacognition, that is, monitoring and thinking about his or her own knowledge and learning. Successful learning environments that employ social agents should therefore motivate and guide SRL.

It is possible to identify four phases of SRL as follows [11]. 1) Task definition: The initial phase where the learner, for effective direction of resources, sets in mind the conditions of the overall learning task including the context and materials. 2) Goal-setting: This phase involves setting specific learning goals, including setting standards of adequate versus insufficient understanding, and constructing plans for reaching these goals. 3) Studying-tactics: The learner enacts plans to guide his or her learning. 4) Metacognitive adaptations: The effective self-regulated learner self-evaluates his or her learning by comparing products of learning to the initial standards and on this basis alters the task definition, standards of learning, learning goals, and plans for achieving goals. This last phase is where metacognitive processes are most relevant, which highlights their role in directing the setting of goals for correcting misconceptions.

1.3 Separate Kinds of Social Learning Impacts at Separate Depths

Based on this learning-depth distinction, we may further distinguish between the effects a social agent may have on more shallow learning behaviors, tied to basic engagement in a task, versus deeper learning impacts from a social agent, tied to metacognition and misconception correction. That is, tweaking for the current context the general idea from Social Impact Theory of various distinct measurable social impacts that can impinge on task performance [12], we propose that separable effects arise from shallow versus deep social impacts. We propose that social agents have the potential to generate a basic facilitation of learning from automatic or simple mechanisms such as boosted neural arousal [cf. 13], selective attention to task dimensions [14], and potentially more emotion to increase motivation [15]. When such impacts on earlier learning phases fail to generate robust outcomes, however, they do not actively *interfere* with learning. On the other hand, we propose that deeper social impacts have the potential to either *enhance* metacognitive and related deep learning processes, or *divert* cognitive resources away from such processes.

2. Potential Negative Impacts on Deep Learning from a Social Agent

With our developed framework of separable social impacts we are now in a position to better understand the results of the recent study broached in the introduction, which suggested a net negative learning impact from a social versus non-social agent [7]. Recall that in this study participants first learned how to solve the 3-ring Tower of Hanoi problem, then demonstrated the solution to either a social or non-social agent, and then as a measure of their learning solved for themselves the more difficult 4-ring problem. Participants' less optimal solutions to the 4-ring problem after demonstrating to the social agent presents a mystery, in that the social and the non-social agent were both represented by a simple picture, and subjects' demonstrations for the social agent were more active (i.e., there were more looks to the picture and more social highlighting, i.e., pointing behaviors, for the social agent). In particular, this presents a challenge for the Chi framework, which conceptualizes active behaviors as always being potentially better, never worse for learning, than passive engagement [8]. However, in the context of the framework we have developed for shallow and deep social impacts, the limitations of shallow social effects on learning become easier to see, as does the potential traps of deep social impacts that are not guided by the interaction to enhance learning, thereby carrying the capacity to function as an *opposing* rather than merely ineffectual force. This is made clear by a multiple regression analysis that

suggested that pointing in the 3-ring demonstration adversely affected performance in the 4-ring task, while looks at the agent were actually linked to better performance. That is, the increased looking behaviors represent effects confined to shallow levels of learning, that while improving 4-ring performance, are swamped by the negative impacts that can directly interfere with deep learning cognitions. Pointing behaviors on the other hand seem tied to deep but negative social impacts. That is, with increased metacognition for the imagined social agent's insufficient task knowledge and the resulting increase in action highlighting behaviors, resources are diverted away from building a richer representation of the solution that is more readily transferable to solving the subsequent more difficult problem. The main contrast making the metacognition here detrimental to learning-by-teaching versus the strong learning benefits found in computerized teachable-agent implementations is that in the latter case the social interaction guides the student's metacognitive behaviors to cohere with correcting and building the student's own knowledge, rather than interfering with this end. Thus such implementations channel the student's metacognition for the teachable agent toward the student's own knowledge construction [2, 16]. However, with no guidance from the appropriate social interaction, ironically the deep social impact arising in the context of demonstrating to an imagined social agent results in *shallow* learning behaviors [7].

3. Conclusions and Future Research Directions

These considerations suggest paths of future research for testing the relative strength of different social effects, positive and negative, on both shallow and deep learning, and for thereby generating insights for maximizing the learning impacts of simulated social agents. In our sketch of learning program implementations we traced out a number of different kinds of social impacts that a social agent can carry. As we noted there are simple effects of increased arousal and engagement that come more or less automatically from a social interaction. Future work can be carried out to investigate whether distinct kinds of effects at this level can be differentiated, and whether the *social* character of such cues are necessary for such effects. For instance, in tasks which benefit from cues that elicit attention at specific timepoints in learning, is there a difference in the cue's effectiveness if it is a *social* cue as opposed to a non-social signal (e.g., a disembodied beep)? Furthermore, is the generalized learning impact of a watchful social agent dissociable from specific cues for directing attention? Turning to deeper social impacts, our overview highlighted two distinct metacognitive effects positive for learning. Social agents may dialogically guide a learner to notice misconceptions, by pushing the learner via prompts and hints which provide just enough information for the learner to figure out misconceptions on his or her own [1]. As a distinct pro-metacognitive and learning-enhancing impact, social agents can lift a natural impediment for acknowledging and correcting misconceptions, when the learner engages in these essential deep learning behaviors for the agent rather than his or herself [4, 16]. On the flip side, however, limited social interactions may push metacognitive behaviors that are *detrimental* to learning, resulting ironically in shallow-learning behaviors, with resources for constructing a deep understanding of the problem solution being diverted instead toward actions like highlighting solution steps for the social agent [7].

A related line of follow-up research suggested by our framework might, given an empirically determined set of both positive and negative social impacts on learning, test the relative *strength* of each impact by keeping the social learning situation a constant except for a manipulation of one impact across conditions at a time. For instance, experiments may be designed to test the strength of social agent behaviors which push learning-relevant metacognition in their feedback, versus the strength of the general metacognitive impact that comes from the creation of knowledge representations for a teachable agent. The

strength of these effects can in turn be compared to a social interaction that entails a *negative* learning impact from the social agent. It may also be possible to determine if each of these sorts of impacts, whether positive or negative, affect a learner's performance in a pattern similar to that well established for a wide range of social impact factors, such as the impact of a group's opinion's on conformity [12]. For instance, if after learning how to solve a 3-ring Tower of Hanoi problem the task were to demonstrate to, across conditions, one, two, or three social agents, would the number of participants' looks and points fit with the pattern suggested by the psychosocial law, of a diminishing increase in these behaviors for each subsequent agent? Would the net (negative) impact on subsequent performance for the 4-ring task in terms of time and number of solution steps also fit this pattern? Would other principles of social impact theory be applicable in studying the effects of social agents on learning and learning-behaviors? These considerations highlight the potential of developing from the current framework methods of pinpointing the various components of positive and negative learning effects that may arise from a social agent and a social interaction context.

Acknowledgement

This work was supported in part by an NSF grant (#0826701) to the second author.

References

- [1] Graesser, A. C. & McNamara, D. (2010). Self-regulated learning in learning environments with pedagogical agents that interact in natural language. *Educational Psychologist*, 45(4), 234-244.
- [2] Leelawong, K., & Biswas, G. (2008). Designing learning by teaching agents: The Betty's Brain system. *International Journal of Artificial Intelligence in Education*, 18(3), 181-208.
- [3] Moreno, R. (2004). Animated pedagogical agents in educational technology. *Educational Technology*, 44(6), 23-30.
- [4] Moreno, R., & Mayer, R. (2007). Interactive multimodal learning environments. *Educational Psychological Review*, 19, 309-326.
- [5] Keysar, B., Shuhong, L., & Barr, D.J. (2003). Limits of theory of mind use in adults. *Cognition*, 89:1, 25-41.
- [6] Epley, N., Boven, L. V., Keysar, B., & Gilovich, T. (2004). Perspective taking as egocentric anchoring and adjustment. *Journal of Personality and Social Psychology*, 87(3), 327-339.
- [7] Herberg, J. S., Levin, D. T., & Saylor, M. M. (2011). Social audiences can disrupt learning by teaching. *Journal of Experimental Social Psychology*, 48(1), 213-219.
- [8] Chi, M. (2009). Active-constructive-interactive: A conceptual framework for differentiating learning activities. *Topics in Cognitive Science*, 1, 73-105.
- [9] De Jaegher, H., Di Paolo, E., & Gallagher, S. (2010). Can social interaction constitute social cognition? *Trends in Cognitive Science*, 14(10), 441-447.
- [10] Chi, M. T. H., Roy, M., & Hausmann, R. G. M. (2008). Observing tutorial dialogues collaboratively: Insights about human tutoring effectiveness from vicarious learning. *Cognitive Science*, 32(2), 301-341.
- [11] Azevedo, R., Moos, D. C., Johnson, A. M., & Chauncey, A. D. (2010). Measuring cognitive and metacognitive regulatory processes during hypermedia learning: Issues and challenges. *Educational Psychologist*, 45(4), 210-223.
- [12] Latané, B. (1981). The psychology of social impact. *American Psychologist*, 36(4), 343-356.
- [13] Zajonc, R. B. (1965). Social facilitation. *Science*, 149, 269-274.
- [14] Huguet, P., Galvaing, M. P., Monteil, J. M., & Dumas, F. (1999). Social presence effects in the stroop task: Further evidence for an attentional view of social facilitation. *Journal of Personality and Social Psychology*, 77(5), 1011-1025.
- [15] Locke, E. A. & Latham, G. P. (2002). Building a practically useful theory of goal setting and task motivation. *American Psychologist*, 57(9), 705-717.
- [16] Chase, C. C., Chin, D. B., Opezzo, M. A., & Schwartz, D. L. (2009). Teachable agents and the protégé effect: Increasing the effort towards learning. *Journal of Science Education and Technology*, 18(4), 334-352.