Social audiences can disrupt learning by teaching ☆,☆☆

Jonathan S. Herberg *, Daniel T. Levin, Megan M. Saylor

Vanderbilt University, USA

ABSTRACT

To investigate the effect of a social audience on learning-by-teaching, we examined participants' solutions of the 4-ring Tower of Hanoi problem after they demonstrated the 3-ring problem to a social agent (a person) or a non-social agent (a computer). In Experiments 1 and 2 participants produced less optimal solutions of the 4-ring problem after demonstrating the 3-ring problem to a social agent. An analysis of pointing behavior demonstrated that social highlighting contributed substantially to this effect. Together, these findings indicate that more social highlighting may produce a cost, rather than a benefit, on how deeply the demonstrator encodes the problem solution. Experiment 3 clarified that these results were not simply caused by the disruptions inherent to social highlighting. Taken together, the results suggest that social highlighting does not come for free — producing the highlighting may lead to more shallow encoding of demonstrated actions.

© 2011 Elsevier Inc. All rights reserved.

Introduction

Research on human learning has shown that simple models of learning as a unidirectional transfer of information from teacher to learner are inadequate (Topping, 2005). Rather, in a wide array of learning interactions, deep learning of complex topics and problems is driven not just by actions by the tutor, but also by the tutee, and by the social interactions that support constructive knowledge-building (Chi, Siler, Jeong, Yamauchi, & Hausmann, 2001; Roscoe & Chi, 2007). A particularly important example of the constructive nature of learning is learning-by-teaching, in which teaching another person facilitates the tutor's understanding, sometimes even more than trying to learn the material for oneself (Roscoe & Chi, 2007). Enhanced benefits from learning-by-teaching have been shown in domains ranging from reading skills (e.g., Juel, 1996) to math and science concepts (e.g., Topping, Campbell, Douglas, & Smith, 2003; Topping, Peter, Stephen, & Whale, 2004).

On one view, the benefits of learning by teaching may derive from the default social response to the mere presence of an audience. Not only can an audience produce basic social facilitation that heightens attention and arousal (Bargh, 1994; Guerin, 1986; Zajonc, 1965), but the presence of social actors can activate non-egocentric perspective taking and simulation of actions (Castiello, 2003; Iacoboni, 2009; Levine, Resnick, & Higgins, 1993). Even the mere suggestion that social actors will later be exposed to persuasive statements produced by the subject can increase subjects' attitude change (Nel, Helmreich, & Aronson, 1969; Wicklund & Duval, 1971; Wood, 2000). More to the point, research has previously shown that demonstrating simple tasks to an imagined social audience (represented by a picture of a person) produces a range of social responses, including social highlighting behaviors (such as pointing to task-relevant goals) that do not occur for a nonsocial audience (a computer; Herberg, Saylor, Ratanaswasd, Levin, & Wilkes, 2008).

However, other research and theory suggest that learning-by-teaching benefits require a deeper interaction with the audience that promotes metacognition (e.g., Aleven & Koedinger, 2002; VanLehn, Siler, Murray, Yamauchi, & Baggett, 2003; Wagster, Tan, Wu, Biswas, & Schwartz, 2007) and knowledge-building (Roscoe & Chi, 2007). On this view, benefits from learning by teaching do not occur (or at least are not maximized) by the mere presence of social audience, but rather are a product either of extensive pre-teaching preparation (e.g., Bargh & Schul, 1980; Benware & Deci, 1984), or of subsequent interactions with the audience (e.g., Chi, 2009).

Based on the research reviewed above, it appears that benefits from learning by teaching may stem from the simple presence of a social audience, or they may require extensive interaction and/or explicit consideration of the knowledge of the tutee. Understanding the factors that produce benefits from social audiences is particularly important because social audiences may invoke costs as well, and if these costs are borne in the absence of benefits, it becomes possible that social audiences may actually *interfere* with learning by teaching. Research on theory of mind suggests that many of the basic processes required to determine that another person's knowledge differs from one's own are resource-intensive (Apperly, Riggs, Simpson, Chiavarino, & Samson, 2006; Epley, Boven, Keysar, & Gilovich, 2004; Leslie, Friedman, & German, 2004). Although adults may not always pay these costs because they sometimes treat agents in a shallow

manner (e.g., Barr & Keysar, 2005), there are clearly many situations where adults do engage deeper cognitions, and would therefore be expected to experience these cognitive costs (e.g., Csibra & Gergely, 2007). The resource-intensive nature of theory of mind reasoning, combined with the possibility that this resource expenditure will not improve learning, makes it critical to test what happens when one needs to focus enough attention on an agent's knowledge states to teach the agent something, but there is no opportunity for deeper interaction. If learning-by-teaching can produce benefits in the absence of the explicitly constructive reasoning that comes from deep interactions with the learner, we would expect benefits to be exhibited in such a situation. Alternatively, if social processing costs sometimes overwhelm social benefits of learning by teaching, it is possible that teaching a social other will *lessen* learning.

Experiment 1

Two conditions tested whether the initial cognitive responses induced by an agent improve, or possibly interfere with, learning by teaching. We employed a design partly based on research establishing social effects from the mere presence of a passive audience (Guerin, 1986; Zajonc & Sales, 1966) and similar to that used in Herberg et al. (2008) in which participants taught a human or a computer (represented by a simple picture) to solve the Tower of Hanoi problem. Thus, our goal was to test the degree to which the mere presence of a social audience would affect learning, independent of any actual interaction with that audience. The Tower of Hanoi puzzle involves three poles and three or more rings of different sizes. The rings start on the leftmost pole and the goal is to get them to the rightmost pole. The constraints are that one is allowed to move only one ring from the top of any pole to another pole at a time, and that one is not allowed to place a larger ring on top of a smaller ring. In the current experiment, participants first learned to solve the 3-ring Tower of Hanoi problem, then demonstrated their solution to a computer or human agent. Finally, they solved the more difficult 4-ring Tower of Hanoi problem for themselves, in the absence of the agent. The idea was to measure how deeply they learned from their demonstration of the 3-ring problem by examining how effectively they generalized to the 4-ring problem.

Method

Participants

56 Vanderbilt undergraduates and members of the surrounding community participated for class credit or $10. Six were excluded from analysis for having prior experience solving Tower of Hanoi problems (2 from the Human condition and 4 from the Computer condition), 4 for violating task rules (1 from the Human condition and 3 from the Computer condition), and 2 for failing to solve the 4-ring task within 6 min (both from the Human condition). Therefore, the data for 44 participants (22 in each condition, alternately assigned) were analyzed (Mean age = 26.4 years, 25 females).

Materials

For the Human condition, participants saw a picture of a college-aged male. For the Computer condition, a picture of a computer and monitor with a mounted camera was used. Participants were also shown a picture of a mechanical device capable of grasping objects when the computer audience was described, though this picture was not present when participants demonstrated their solution.

The materials participants used to solve the Tower of Hanoi tasks were three plastic tubes and three or four plastic rings (a small red ring, medium sized yellow ring, and large green ring for the 3-ring problem, plus a larger blue ring for the 4-ring problem). Participants' solutions and demonstrations were videotaped for later coding.

Procedure

Three plastic poles with three rings on the leftmost pole were placed on a table in front of the participant. Participants were told that the goal of the task was to get the three rings from the leftmost pole to the rightmost pole. They were told that they could move one ring at a time from one pole to any other pole, and that they could not move a larger ring on top of a smaller ring.

After explaining the three-ring Tower of Hanoi task, the experimenter asked the participant to complete the task, then went behind a curtain to avoid influencing the participant's solutions and demonstrations. When the participant indicated he or she had finished the task, the experimenter re-entered the room and set up the materials for demonstrating it. A picture of the audience (human or computer) was placed on a stand to the participant's front and right, as well as the three rings and three poles in their original setup. The experimenter described the audience the participant was going to demonstrate to. The human audience was described as being ready to learn the Tower of Hanoi task by watching the participant. The computer audience was described as being able to take in action information through a camera and to carry out actions by moving a mechanical gripping device. In both cases participants were instructed to imagine the computer or human audience in the room while they demonstrated the task in whatever way they felt natural to allow the audience to be able to do the same action. Participants were told not to use language. After describing the audience and explaining the demonstration task, the experimenter again left the participant's section of the room. After the participant indicated he or she had finished the demonstration of the 3-ring Tower of Hanoi problem, the experimenter set up the 4-ring Tower of Hanoi problem with the same three rings as before, plus one larger ring on the bottom. Participants were told that that their task was to solve a more difficult version of the same problem, with the same rules, constraints, and goal, but with four instead of three rings. The audience was removed, and participants were told to simply solve the problem and not worry about the audience anymore. The experimenter then left the participant alone to solve the problem. If the participant failed to solve the problem in 6 min, the experimenter ended the session (excluding the participant's data from analysis).

Coding

Videos of participants' solutions to the Tower of Hanoi problems were coded for solution time and the number of solution steps (movements of a ring from one pole to another). If a step was undone the initial movement was not counted as a step. Both shorter solution times and fewer steps indicate a better understanding of how to solve the task (e.g., Goel & Grafman, 1995). If the participant restarted the task then the number of steps only in the final, complete solution was counted. For all experiments, all statistically significant findings remain when participants who restarted the 4-ring task are excluded from analysis.

Additionally, in participants' demonstrations of the 3-ring solution, the number of looks at the picture of the audience, and the number of gestures highlighting objects and actions (usually points) were counted, to measure participants' social response behaviors. This was done both as a manipulation check and to test whether different kinds of social behaviors may impact participants' learning in different ways. There may be mere-presence effects on learning-by-teaching, which may be captured best by investigating participants' looking frequency at the audience. There may also be different effects from the more effortful reasoning one undertakes in engaging in social highlighting behaviors, which can be ascertained by looking at how frequently participants point in their demonstrations.

A preliminary check was made to ensure there were no mean differences in Tower of Hanoi problem solving abilities in participants in the Computer vs. Human conditions. Participants in the Computer vs. Human conditions did not reliably differ in either variable in the 3-

ring problem, before the computer or human audience was introduced, ts (42) $\leq$ 1.63, *n.s.*

Results

Social cue differences in demonstrations of the three-ring problem

In a replication of Herberg et al. (2008), participants in the Computer condition looked less frequently at the audience picture (M = 4.04, SD = 3.17) than participants in the Human condition (M = 9.00, SD = 3.98), t (42) = 4.54, p<0.01. They also pointed less frequently for the computer audience (M = 1.86, SD = 2.70 versus M = 5.27, SD = 4.07), t (42) = 3.24, p<0.01.

Solution-time and number of steps for the four-ring problem

Participants in the Human condition required more steps to solve the problem than participants in the Computer condition, t (42) = 2.72, p<0.05. There was also an interaction between task (four-ring vs. three-ring) and condition on number of steps, F (1,42) = 4.52, p<0.05. This interaction effect suggests participants solved the 4-ring task in more steps in the Human than Computer condition even when the number of steps in the initial 3-ring task is subtracted out from their 4-ring number of steps. Although as noted above the preliminary analysis did not indicate a statistically significant difference in number of steps or solution time in the initial 3-ring task for participants in the Human vs. Computer conditions, the difference went in the direction of higher values for the Human condition. Therefore we also conducted ANCOVA analyses to test for condition differences in number of steps and in solution time on the 4-ring task with the initial 3-ring task number of steps and solution time entered as covariates. For 4-ring number of steps, the difference was still present, F(1, 40) = 6.56, p<0.05. Participants in the Human condition also solved the 4-ring Tower of Hanoi problem more slowly than participants in the Computer condition, although this difference only approached significance, t (42) = 1.84, p = 0.07, with no interaction between task and condition, F (1, 42) = 1.32, p = 0.26, and no effect when 3-ring solution time and 3-ring number of steps were entered as covariates in the ANCOVA, F (1, 40) = 1.67, p = 0.20. See Figs. 1 and 2.

Discussion

Even though the audience did not actually interact with the participants and was only represented by a picture, participants who demonstrated the solution of the three-ring Tower of Hanoi problem to a human audience solved the four-ring problem less efficiently than those who demonstrated to the computer audience. Participants demonstrating to the social agent appeared to have been paying a cost to do the extra social reasoning, in line with effortful theory of mind use suggested in other research (e.g., Keysar, Shuhong, & Barr, 2003). One possibility is that interactive and constructive behaviors, as highlighted by Chi (2009), may be necessary to offset these costs to generate learning-by-teaching benefits. Interactive behaviors are those in which the tutor and learner build on what each other know to jointly construct new knowledge. Constructive behaviors refer to the inference of new knowledge, as opposed to the mere reiteration of what one knows. These kinds of behaviors based on deeper cognition appear to drive enhanced learning-by-teaching effects, and it is possible that being stuck on the level of merely highlighting what one already knows without any push for these more elaborate behaviors generates a cost to one's capacity for generating a deeper understanding based on the prior knowledge.

One basic question about the nature of this interference is whether participants in the Human condition were less likely to develop any sort of abstract rule to adapt their knowledge of the 3-ring solution to the 4-ring task, relying instead on trial-and-error. This would suggest that their demonstrations led to a generally shallow encoding of the 3-ring solution (cf. Vollmeyer, Burns, & Holyoak, 1996). Alternatively, participants in the Computer condition may have simply hit upon a simpler heuristic for succeeding. While the most explicitly rule-based strategy is to develop a recursive algorithm for solving a Tower of Hanoi problem of any number of rings, it is also possible to develop less precise strategies, such as a strategy to correctly isolate the largest ring on the first step (Simon, 1975). For the four-ring problem, the optimal first step is to move the top ring from the leftmost pole to the middle pole. In the three-ring problem, the optimal first step is to the rightmost pole. More participants in the Computer condition (81.8%) made the optimal first step in the four-ring problem than participants
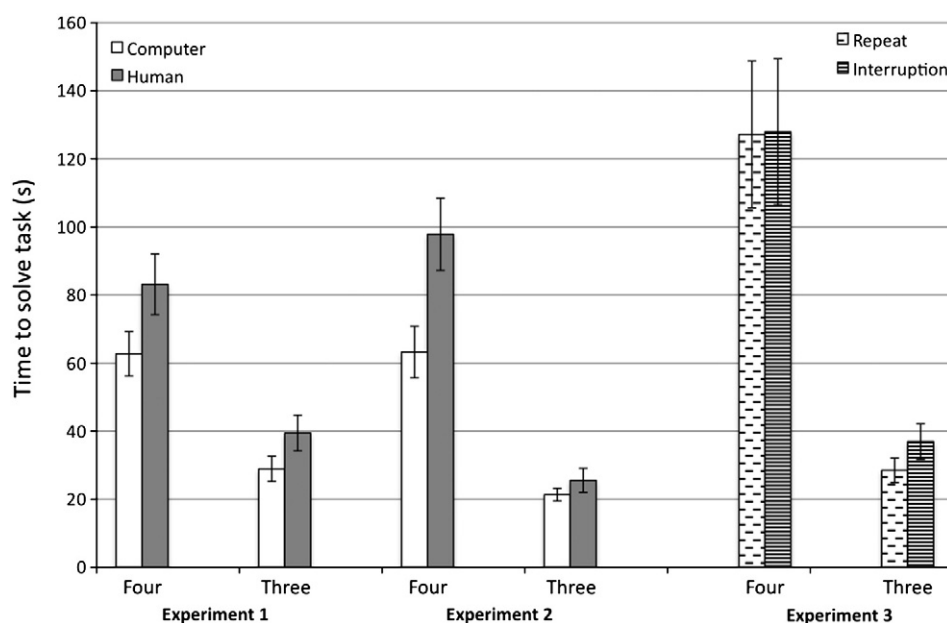

Fig. 1. Mean 4-ring and 3-ring Tower of Hanoi solution times for participants in the Computer and Human conditions in Experiments 1 and 2 and the Repeat and Interruption conditions in Experiment 3.
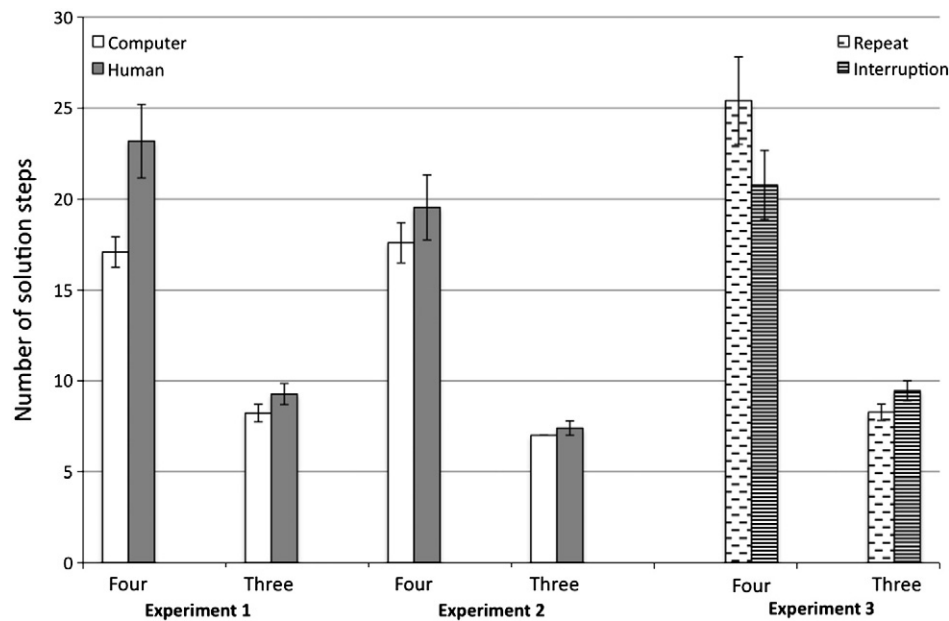
Fig. 2. Mean 4-ring and 3-ring Tower of Hanoi number of steps for participants in the Computer and Human conditions in Experiments 1 and 2 and the Repeat and Interruption conditions in Experiment 3.

in the Human condition (63.8%), $\chi^2(1) = 1.83$; $p = 0.09$. It is therefore important to know whether the impaired 4-ring task performance comes solely from being led down the wrong garden path toward a less optimal solution by taking the wrong first step, or whether participants' learning is impaired by a social agent more generally in other important ways as well. To test whether a suboptimal first step was responsible for relatively poor performance in the Human condition, participants were told the first step for each Tower of Hanoi task. If they still show a decreased ability to solve the four-ring problem in the Human condition, this would indicate that the social interference lessens participants' capacity to form a four-ring strategy in general. Another key purpose for Experiment 2 was to replicate the social interference effect.

Experiment 2

Method

Participants

Out of 59 Vanderbilt undergraduates who participated, 8 (4 from each condition) were excluded for having prior experience solving Tower of Hanoi problems, 7 (3 from the Human condition and 4 from the Computer condition) for violating task rules, and 4 (1 from the Human condition and 3 from the Computer condition) for failing to solve the 4-ring task within 6 min. This left 40 participants (20 in each condition) who completed the experiment in exchange for class credit (Mean age = 19.7 years, 35 females).

Procedure

The materials, design, and procedure were the same as for Experiment 1, with the following changes. After the three-ring Tower of Hanoi problem was described, the participant was told that in order to solve the problem in the minimum number of steps, the first move should be from the leftmost pole to the rightmost pole. Likewise, after the 4-ring problem was set up, the participant was told that the first move should be from the leftmost pole to the middle pole. As in Experiment 1, the question was how the manipulation of audience, and the resulting demonstration behaviors, apart from an explicit minimum number of steps requirement in the problem description, would impact the efficiency of participants' solutions. Therefore participants were

told that they need not necessarily solve either the 3-ring or 4-ring problem in the minimum number of steps possible, but that they should regardless do the optimal first step as shown to them.

A preliminary check on participants' solution time and number of steps for the three-ring problem was again made. Participants in the Computer and Human conditions did not reliably differ, $ts(38) \leq 1.05$, *n.s.*

Results

Social cue differences in demonstrations of the three-ring problem

Participants in the Computer condition again exhibited less pointing behaviors (M = 1.80 SD = 3.14) than participants in the Human condition (M = 6.55, SD = 5.97), $t(38) = 3.08$, $p <. 01$. Participants in the Computer condition (M = 5.65, SD = 4.56) also tended to demonstrate the 3-ring solution with fewer looks to the picture than participants in the Human condition (M = 8.65, SD = 6.00), though this tendency only approached significance, $t(38) = 1.78$, $p =. 08$.

Solution-time and number of steps for the four-ring problem

Participants took less time to solve the 4-ring problem in the Computer condition than in the Human condition, $t(38) = 2.64$, $p < 0.05$, with the task by condition interaction also significant, $F(1,38) = 5.27$, $p < 0.05$. The effect of condition was also significant when initial 3-ring task solution time and number of steps were entered as covariates, $F(1, 36) = 6.10$, $p < 0.05$. Participants in the two conditions did not differ in the number of steps they took to solve the task $t(38) = 0.93$, $p = 0.36$. See Figs. 1 and 2.

Effects of social cues in 3-ring demonstration on 4-ring solution in Experiments 1 and 2

Solution time. To assess the impact of social cue production (looks and points) in the 3-ring demonstrations on participants' solutions of the 4-ring problem, a multiple regression analysis on the Experiments 1 and 2 data combined was conducted. A multiple regression analysis with 4-ring solution time as the dependent variable and with audience-condition, 3-ring solution time, 3-ring number of steps, demonstration solution time, demonstration number of steps, number of looks, and number of points as the independent variables produced an overall model with adjusted $R^2 = 0.30$, $F(7,76) = 6.01$,

$p<0.01$. As shown in Table 1, the Experiments 1 and 2 data suggest that producing more points in their 3-ring demonstrations may have resulted in participants taking longer to solve the 4-ring task. This is the case even though more looks goes in the direction of predicting quicker 4-ring solutions.

Number of steps. A multiple regression analysis with 4-ring number of steps as the dependent variable produced an overall model with adjusted $R^2 = 0.15$, $F(7,76) = 3.06$, $p<0.01$. As Table 2 shows, producing more points in their 3-ring demonstrations is tied to participants developing less optimal strategies for solving the 4-ring task (while more looks predicts more optimal 4-ring solutions).

Discussion

Experiment 2 replicates Experiment 1 by demonstrating that even when participants were given the first step for an optimal solution to the Tower of Hanoi problem, they took longer to solve the 4-ring task after demonstrating the 3-ring task for a person. Demonstrating the 3-ring problem for a person vs. a computer again appears to have impeded learning by teaching. In Experiment 2, this was shown via participants' longer solution times for the 4-ring task in the Human condition, rather than in the solutions to the 4-ring task for the person containing more steps as in Experiment 1. Both measures appear to reflect the depth of participants' understanding of the 3-ring solution, and their capacity to transfer this understanding to solving the 4-ring task, and some measures combine both (e.g., Goel & Grafman, 1995). Similarly, cognitive performance in other tasks, such as visual search efficiency, can be reflected in distinct measures, such as either reaction time or errors, which are also sometimes combined (Smilek, Enns, Eastwood, & Merikle, 2006). Therefore, solution time and number of steps may both measure participants' learning.

On the other hand, while both variables measure learning they may reflect distinct specific ways learning may be impaired, and distinct solution strategies. Experiment 1 suggests that without having the optimal first step "figured out" for them, after demonstrating their 3-ring solution for a social agent, participants tend to settle for non-optimal 4-ring solution paths. Experiment 2 reveals that when given the first step of the 4-ring task participants who first demonstrate their 3-ring solution for a social rather than non-social agent are able to obtain optimal 4-ring task solutions, but take a longer time to do so. In this situation, since the first available path to a less optimal solution, the wrong first step, is blocked off, participants are largely able to find their way to the optimal solution even if they first demonstrate to a social agent. However, it requires more time and effort for them to transfer their 3-ring task understanding in figuring out the 4-ring solution. This suggests less generalized knowledge, and less of an understanding of the logic behind different 3-ring steps after demonstrating to a social vs. non-social agent. The multiple regression analysis further demonstrated that social highlighting behaviors (pointing) in the demonstrations may at least partially drive both of these kinds of social interference effects.

Table 1
Multiple regression for predicting four-ring solution time in Experiments 1 and 2.

Variable	B	SE (B)	ß
Condition	19.43	8.91	0.239[*]
3-ring solution time	1.49	0.312	0.681[**]
3-ring number of steps	−9.19	3.29	−0.486[*]
Demonstration solution time	−0.236	0.279	−0.112
Demonstration number of steps	5.56	3.44	0.223
Looks	−1.74	1.05	−0.209
Points	2.15	1.07	0.240[*]

[*] p<0.05.
[**] p<0.01.

Table 2
Multiple regression for predicting four-ring number of steps in Experiments 1 and 2.

Variable	B	SE (B)	ß
Condition	3.41	1.75	0.236
3-ring solution time	0.132	0.061	0.340[*]
3-ring number of steps	−0.935	0.644	−0.279
Demonstration solution time	−0.062	0.055	−0.164
Demonstration number of steps	1.23	0.674	0.279
Looks	−0.423	0.207	−0.285[*]
Points	0.450	0.210	0.283[*]

[*] p<0.05.

Experiment 3

Before we can conclude that social highlighting of actions interferes with knowledge-building, an alternative explanation to consider is that the basic act of planning and executing a set of pointing behaviors interfered with the planning and execution of the movements in the Tower of Hanoi task. Thus, the presence of a social partner induced highlighting behaviors that in turn interrupted the participant's completion of the task. On this view, the social/communicative nature of the pointing was not necessary to produce interference. Experiment 3 therefore tested whether interrupting one's action sequence by being externally prompted to point in the absence of an audience would raise the cognitive load enough to interfere with learning. In the Repeat condition, participants solved the three-ring problem a second time without any interruptions. In the Interruption condition, however, every time they heard a beep they had to point at the ring they were going to pick up next or at the place they would put a ring they were holding. If the basic disruptiveness of pointing is primarily responsible for the decreased learning in the Human conditions in the previous experiments, participants in the Interruption condition should do worse on the 4-ring problem than participants in the Repeat condition.

Method

Participants

Out of 57 Vanderbilt undergraduates who participated, 5 (4 from the Repeat and 1 from the Interruption condition) were excluded for prior experience solving Tower of Hanoi problems, 3 (1 from the Repeat and 2 from the Interruption condition) for violating the rules of the task, and 5 (3 from the Repeat and 2 from the Interruption condition) for failing to solve the 4-ring task within 6 min. This left 44 participants (22 in each condition) who completed the experiment in exchange for class credit (Mean age = 20.7 years, 35 females).

Procedure

The same procedure as in Experiment 1 was used to introduce participants to the Tower of Hanoi task. After participants solved the 3-ring task for the first time, in the Repeat condition the experimenter simply instructed participants to solve the task a second time, and went behind the screen. In the Interruption condition participants were instructed to solve the task a second time, but while they were doing this, they were told to listen for a beep that would occur periodically. Whenever they heard the beep, the instructions were to point at the ring they were going to next pick up, or the one that they were already holding, and then point to where they were going to move it. If they had not yet decided where to move (and therefore not yet pointed) when the next beep occurred they were told to complete their point and movement before responding to new beeps. A rate of beeping (1 beep per 10 s) was chosen to generate a comparable mean number points for participants in the Interruption condition (M = 5.95) as participants in the Human conditions of Experiments 1 and 2 (Ms = 5.27 and 6.55). After giving these instructions the experimenter pressed a button on the

laptop to start the beep-presenting movie file and then went behind the screen. After participants completed the 3-ring task the second time, the experimenter presented them with the 4-ring task as in Experiment 1.

A preliminary check on participants' solution time and number of steps for the three-ring problem was again made. Participants in the Repeat and Interruption conditions did not reliably differ, ts (42) ≤ 1.65, n.s.

Results and discussion

Solution-time and number of steps for the four-ring problem

Participants in the Repeat and Interruption conditions did not significantly differ in the time they took to solve the 4-ring tower of Hanoi problem, t (42) = 0.03, p = 0.98. Additionally, there was a trend of *fewer* rather than more 4-ring steps for participants in the Interruption vs. Repeat conditions, t (42) = 1.51, p = 0.14, with a marginally significant task by condition interaction, F (1,42) = 3.11, p = 0.08, although no effect when the initial 3-ring task performance measures are entered as covariates, F (1, 40) = 1.41, p = 0.24. See Figs. 1 and 2. This demonstrates that the results from Experiments 1 and 2 were not due to participants in the Human condition having a generally higher cognitive load as a result of having to interrupt their action sequence to point. Rather, the tendency of participants in Experiments 1 and 2 to show less optimal solutions after demonstrating to a human audience versus a computer audience was specific to *social* pointing.

General discussion

We have repeatedly demonstrated that a social audience can interfere with learning-by-teaching, and have further shown that social highlighting behaviors at least partially explain this effect. Some responses elicited by the mere presence of a social agent, such as looking at the agent, may facilitate learning. However, our findings imply that social highlighting behaviors generate interference effects that overwhelm such facilitation effects. Our results suggest the need to consider potential costs of demonstrating a task with active social highlighting. These results are therefore consistent with Chi's (2009) emphasis of the critical importance of deep constructive and interactive behaviors promoted in a learning situation. Such behaviors may be necessary to overcome the costs of social highlighting behaviors in these situations. We have also shown that these costs are specifically tied to the social reasoning aspects of the highlighting rather than simply the increased interruptions caused by pointing.

This burden may be tied to participants having to represent how a social agent's knowledge differs from their own. (cf. Epley et al., 2004; Lin, Keysar, & Epley, 2010; Nickerson, 1999). The increased effort needed to use one's theory of mind appears to siphon off cognitive resources that could otherwise be used to learn the task one is demonstrating. In producing more "knowledge-telling" highlighting behaviors (Roscoe & Chi, 2007), tutors may be inhibiting the full depth of their task knowledge and be more inclined to stick to the surface level of the task. In the Tower of Hanoi demonstration context, this means that social highlighting will lessen the observer's capacity to represent the task using abstract recursive or subtower-based strategies (Simon, 1975). In this context, it is interesting to consider the trend toward fewer 4-ring steps in the Interruption condition of Experiment 3. In the case of pointing without a social agent, there was no social interference with developing a strategy. In fact, this condition may have motivated participants to visualize and plan steps ahead, to better deal with a possible beep should one occur.

These results are novel in demonstrating that social audiences do not always facilitate learning. Although a social partner no doubt can facilitate learning in many situations, we have demonstrated that this is not universally true, especially when the audience induces learners to produce the *wrong kind* of learning behaviors, which can cause more shallow learning. In this light, it is important to consider possible specific reasons behind the link between pointing behaviors and impaired learning, and also why there is no such link, but rather an apparent trend

for facilitated learning, in the case of looking behaviors. As discussed earlier, Chi (2009) offers a framework that seems well-suited to apply to learning-by-teaching situations. More specifically, this framework separates out merely "active" behaviors from the deeper constructive and interactive behaviors we discussed above. The framework specifies that behaviors that merely highlight information, such as gazing and pointing behaviors, are active learning behaviors. Our findings suggest that some active learning behaviors, such as pointing, are more resource-intensive than others, such as looking. We would argue that this cost is tied to social cognitive and theory-of-mind reasoning processes, related to explicitly figuring out what steps in the task the social agent needs to see highlighted, based on what the agent knows and does not know, and producing the manual gestures for the purpose of changing the agent's knowledge about the task. We should point out that while we do not have direct evidence that pointing is associated with theory of mind reasoning, it seems likely because the pointing appears to reflect a resource-intensive cognitive process that is differentially invoked when interacting with a person, and that requires planning predicated upon that person's putative knowledge states. These are very good fits with the tasks typically attributed to theory of mind reasoning (especially the more cognitive and deliberative elements of the system; e.g., Apperly et al., 2006; Leslie et al., 2004) and it is difficult to imagine that some other reasoning system or that a general purpose reasoning system would be relied upon when teaching a person, leaving the theory of mind system unoccupied during this difficult task. Looking behaviors, on the other hand, likely require fewer cognitive resources, and may be triggered in a more automatic fashion by the teaching situation. They may, however, reflect how "actively" engaged the participant is in the demonstration task, and the amount of facilitating arousal being generated from the mere presence of the social agent. Therefore, active behaviors that are burdensome to a tutor may be contrasted with active behaviors that are conducive to, or at least not in opposition to, constructive learning.

It is important to also note some limitations in what can be concluded from this set of studies that future research will need to clarify. One possibility is that having to imagine a social agent may be more cognitively taxing than having one actually present. This possibility can be tested in a similar setup to the experiments here, where one demonstrates to an actually present person (say through a video-link, so that interactive behaviors are similarly limited across conditions) vs. a picture of a person that one imagines. In addition, these experiments raise the possibility that looking behaviors during one's demonstration may facilitate learning while pointing behaviors may generate a cost to one's learning. This possibility should be rigorously tested in an experiment with two conditions, where participants *only* point to highlight the actions they are doing or *only* look at the agent. They might for instance be instructed to point or to look in response to beeps (thus generating similar frequencies of each behavior across conditions). Future research might also look to classify costly versus beneficial active, constructive, and interactive learning behaviors. In these ways, we may gain an enriched understanding of the crucial links between social cognition and learning-by-teaching.

References

Aleven, V. A., & Koedinger, K. R. (2002). An effective metacognitive strategy: learning by doing and explaining with a computer-based cognitive tutor. *Cognitive Science*, 26, 147–179.

Apperly, I. A., Riggs, K. J., Simpson, A., Chiavarino, C., & Samson, D. (2006). Is belief reasoning automatic? *Psychological Science*, 17, 841–846.

Bargh, J. A. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In R. S. WyerJr., & T. K. Srull (Eds.), *Handbook of Social Cognition: Basic Processes* (pp. 1–40). (2nd ed.). New York: Guilford Press.

Bargh, J. A., & Schul, Y. (1980). On the cognitive benefits of teaching. *Journal of Educational Psychology*, 72(5), 593–604.

Barr, D., & Keysar, B. (2005). Mindreading in an exotic case: The normal adult human. In B. F. Malle, & S. D. Hodges (Eds.), *Other minds: How humans bridge the divide between self and others* (pp. 271–283). New York: Guilford Press.

Benware, C. A., & Deci, E. L. (1984). Quality of learning with an active versus passive motivational set. *American Educational Research Journal*, 21(4), 755–765.

Castiello, U. (2003). Understanding other people's actions: Intention and attention. *Journal of Experimental Psychology*, *29*(2), 416–430.

Chi, M. (2009). Active–constructive–interactive: A conceptual framework for differentiating learning activities. *Topics in Cognitive Science*, *1*, 73–105.

Chi, M., Siler, S. A., Jeong, H., Yamauchi, T., & Hausmann, R. G. (2001). Learning from human tutoring. *Cognitive Science*, *25*, 471–533.

Csibra, G., & Gergely, G. (2007). Obsessed with goals: Functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychologica*, *124*, 60–78.

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.

Goel, V., & Grafman, J. (1995). Are the frontal lobes implicated in "planning" functions? Interpreting data from the Tower of Hanoi. *Neuropsychologia*, *33*(5), 623–642.

Guerin, B. (1986). Mere presence of humans: A review. *Journal of Experimental Social Psychology*, *22*(1), 38–77.

Herberg, J. S., Saylor, M. M., Ratanaswasd, P., Levin, D. T., & Wilkes, D. M. (2008). Audience-contingent variation in action demonstrations for humans and computers. *Cognitive Science*, *32*, 1003–1019.

Iacoboni, M. (2009). Imitation, empathy, and mirror neurons. *Annual Review of Psychology*, *60*, 653–670.

Juel, C. (1996). What makes literacy tutoring effective? *Reading Research Quarterly*, *31*(3), 268–289.

Keysar, B., Shuhong, L., & Barr, D. J. (2003). Limits of theory of mind use in adults. *Cognition*, *89*(1), 25–41.

Leslie, A. M., Friedman, O., & German, T. P. (2004). Core mechanisms in "theory of mind". *Trends in Cognitive Sciences*, *8*, 528–533.

Levine, J. M., Resnick, L. B., & Higgins, E. T. (1993). Social foundations of cognition. *Annual Review of Psychology*, *44*, 585–612.

Lin, S., Keysar, B., & Epley, N. (2010). Reflexively mindblind: Using theory of mind to interpret behaviors requires effortful attention. *Journal of Experimental Social Psychology*, *46*, 551–556.

Nel, E., Helmreich, R., & Aronson, E. (1969). Opinion change in the advocate as a function of the persuasibility of his audience: A clarification of the meaning of dissonance. *Journal of Personality and Social Psychology*, *12*(2), 117–124.

Nickerson, R. S. (1999). How we know–and sometimes misjudge–what others know: Imputing one's own knowledge to others. *Psychological Bulletin*, *125*, 737–760.

Roscoe, R., & Chi, M. (2007). Understanding tutor learning: Knowledge-building and knowledge-telling in peer tutors' explanations and questions. *Review of Educational Research*, *77*(4), 534–574.

Simon, H. A. (1975). The functional equivalence of problem solving skills. *Cognitive Psychology*, *7*(2), 268–288.

Smilek, D., Enns, J. T., Eastwood, J. D., & Merikle, P. M. (2006). Relax! Cognitive strategy influences visual search. *Visual Cognition*, *14*(4), 543–564.

Topping, K. (2005). Trends in peer learning. *Educational Psychology*, *25*(6), 631–645.

Topping, K., Campbell, J., Douglas, W., & Smith, A. (2003). Cross-age peer tutoring of in mathematics with seven and 11-year olds: Influence on mathematical vocabulary, strategy, dialogue, and self-concept. *Educational Research*, *45*(3), 287–308.

Topping, K., Peter, C., Stephen, P., & Whale, M. (2004). Cross-age peer tutoring of science in the primary school: Influence on scientific language and thinking. *Educational Psychology*, *24*(1), 57–75.

VanLehn, K., Siler, S., Murray, C., Yamauchi, T., & Baggett, W. B. (2003). Why do only some events cause learning during human tutoring? *Cognition and Instruction*, *21*(3), 209–249.

Vollmeyer, R., Burns, B. D., & Holyoak, K. J. (1996). The impact of goal specificity on strategy use and the acquisition of problem structure. *Cognitive Science*, *20*, 75–100.

Wagster, J., Tan, J., Wu, Y., Biswas, G., & Schwartz, D. (2007). Do learning by teaching environments with metacognitive support help students develop better learning behaviors? *The 29th annual meeting of the cognitive science society* (pp. 695–700).

Wicklund, R. A., & Duval, S. (1971). Opinion change and performance facilitation as a result of objective self-awareness. *Journal of Experimental Social Psychology*, *7*, 319–342.

Wood, W. (2000). Attitude change: Persuasion and social influence. *Annual Review of Psychology*, *51*, 539–570.

Zajonc, R. B. (1965). Social facilitation. *Science*, *149*, 269–274.

Zajonc, R. B., & Sales, S. M. (1966). Social facilitation of dominant and subordinate responses. *Journal of Experimental Social Psychology*, *2*(2), 160–168.