

Cognitive Dissonance as a Measure of Reactions to Human-Robot Interaction

Daniel T. Levin, Julie A. Adams, Caroline Harriott, and Tao Zhang
Vanderbilt University

When people interact with intelligent agents, they likely rely upon a wide range of existing knowledge about machines, minds, and intelligence. This knowledge not only guides these interactions, but it can be challenged and potentially changed by interaction experiences. We hypothesize that a key factor mediating conceptual change in response to human-machine interactions is cognitive conflict, or dissonance. In this experiment, we tested whether interactions with a robot partner during a realistic medical triage scenario would cause increased levels of cognitive dissonance relative to a control condition in which the same task was performed with a human partner. In addition, we tested whether heightened levels of dissonance would affect concepts about agents. We observed increased cognitive dissonance after the human-robot interaction, and found that this dissonance was correlated with a significantly less intentional (e.g. human-like) view of the intelligence inherent to computers.

Keywords: Human-robot interaction, theory of mind, concepts, cognition.

Introduction

When people interact with intelligent agents such as computers and robots, they must rely on knowledge ranging from understandings of specific functions to more general intuitions about the fundamental constraints inherent to machine thinking. In previous research, we have developed measures of the degree of intentionality (e.g. human-like thought) attributed to different agents, and have shown that these attributions not only can effectively predict the quality of human-machine interactions (Hymel et al., 2011), but that they can change with experience in predictable ways (Levin et al., in review). Understanding this malleability is crucial to effectively modeling users knowledge, because it allows principled predictions not only for a range of different users, but also for individual users as their understandings change with experience. We hypothesize that a key factor driving this conceptual change is cognitive dissonance. That is, the feeling of cognitive conflict that users experience when their assumptions are being challenged, or when they detect facts in a given setting that appear to conflict. In this experiment, we test the degree to which a new measure of cognitive dissonance can detect the cognitive conflict that might occur during human-robot interaction, and test the degree to which this conflict is associated with changes in basic concepts about machine thinking.

Recent research has explored people's concepts about how humans and machines "think", and has tested for links between these concepts and people's experiences, while using a variety of artificial agents (for review see Epley, Waytz, & Cacioppo, 2007). In a recent series of studies, we have attempted to go beyond simply asking participants whether machines think like people or can be said to have "goals" or be "intelligent" to asking better-defined questions in which participants predict the behavior of people and

machines in a small number of scenarios that directly test these concepts. The most basic of these scenarios is based on research demonstrating that a precursor of children's early-developing theory of mind is the basic understanding that intentional goal-directed action is usually directed toward objects, and not locations (Woodward, 1998). In this scenario, adult participants are asked to imagine that different agents (a human, a robot, and a computer) have "acted upon" one object (a toy duck), while ignoring another nearby object (a small truck), participants are then asked to predict what the agent will do when the locations of the objects are swapped. Specifically, will the agent reach to the new object in the previously acted-upon location (a response indicating that participants believe that the agent does not have human-like intentional goals), or will it reach to the old object now in the new location (implying that the agent does have human-like intentional goals)? We have previously demonstrated that adults generally predict more intentional responses for humans than for computers and robots (Levin et al., in review; Levin, Killingsworth, & Saylor, 2008; Levin, Saylor, & Lynn, invited revision in review). However, it is important to note that this tendency is far from universal, that it is related to participants' attributions of goal understanding to machines (even when controlling for their attributions of overall intelligence), and that the human-machine contrast is smaller in older adults than in younger adults (Levin et al., in review; Levin, Killingsworth, & Saylor, 2008; Levin, Saylor, & Lynn, invited revision in review).

A key reason to measure these concepts is to track how basic intuitions about machine intelligence change as people experience different kinds of agents. We hypothesize that a key cause of conceptual change is detecting conflicts between situation-relevant facts, or between existing knowledge and new facts. To assess this kind of conflict, we have been developing a brief questionnaire to measure *cognitive dissonance* resulting from human-machine interactions. Our definition of cognitive dissonance is that it is a state of discomfort associated with detection of conflicting concepts, or with concepts that conflict with observations or experiences. This definition is similar to that used in research on social cognition (Eliot and Devine, 1994), except that it lacks a strong role for threat to self-image. Our hypothesis is that cognitive conflict users experience between their existing knowledge and their experiences is a key stimulus to reflection about agents and ultimately conceptual change. In this initial test of this questionnaire, our focus is on testing whether exposure to different artificial agents causes dissonance, and whether this dissonance is related to concepts about agents, as measured by our behavioral prediction scenarios.

Why should exposure to a novel agent cause dissonance? We predict that this agent-induced dissonance should occur because existing concepts about agents represent a collection of intuitions, biases, and explicit knowledge drawn from a wide range of sources and interactions. So, this information not only comes from direct interactions with people and technology, but it also includes a wide variety of secondhand information coming from sources as diverse as news and fiction. However, in contrast to knowledge about human agency, which is constantly tested and subtly refined during everyday social interactions, information about agents such as robots may be available from second hand sources, and inferences based on intuitions about similar agents, but it is rarely tested in direct interactions. Accordingly, direct interactions with specific robots are likely to produce strong conflicts with existing knowledge as broad untested inferences, and intuitions derived from unreliable (perhaps even fictional) sources are observed to be unwarranted.

In this experiment, our primary aim was to test whether a human-robot interaction would produce more cognitive dissonance than a similar human-human interaction, and further to test the degree to which dissonance produced by this situation would affect concepts about agents. Participants completed a realistic medical triage scenario in one of two conditions. In a human-robot condition, the participant partnered with a robot that directed the participant to make a set of assessments on each of several victims, and to report the results of the assessments to the robot. In the human-human condition, the same requests were made of the participant, and they responded with the same information, but in this case the communication occurred over a walkie-talkie (used to emulate a cell phone) to a human partner, and no robot was present. In addition to predicting that the human-robot condition would produce dissonance, we tested whether this increased level of dissonance would produce changes in participants' concepts about agency more generally. Therefore, participants completed the behavioral prediction measure of agency after completing the medical triage exercise. There were two plausible impacts the human-robot interaction might have on these predictions. First, it is possible that participants in the human-robot condition would, overall, differentiate more, or less, strongly between human and machine agents. However, our hypothesis about dissonance is that it is a sign of an individual's specific response to a cognitive challenge. Therefore, in some situations low dissonance individuals will become more certain of their preexisting commitments, while individuals who experience high levels of dissonance change their concepts.

Method

The experiment was a between-subjects manipulation of the presence of a human or robot partner. During the human-human condition, an evaluator played the role of a first responder located outside of the contaminated area. The evaluator provided instructions to the uninjured victim – the participant. The human-robot condition paired the participant with a robot. Both the participant and the robot were located in the contaminated incident area. A human evaluator supervised both the participant and the robot remotely. Participants in the human-human condition completed the experiment prior to participants in the human-robot condition.

Subjects. A total of 34 participants completed the present experiment, 19 in the human-human condition and 15 in the human-robot condition. The average age of all participants was 24.6 and age ranged between 18 and 57 years. The human-human condition mean age was 23.1 years and the human-robot condition mean was 26.1. All participants had at least some college education. The participants rated their level of first aid experience on a Likert scale with 1 representing no experience and 9 representing an expert level of experience. The average level of first aid experience was 3.75, with the human-human condition mean = 3.6 and the human-robot condition mean = 3.9. All participants rated their level of robotics experience on the same scale. The average experience level was 2.7, with the human-human condition mean = 2.8 and the human-robot condition mean = 2.5.

Experimental Environment. The evaluation occurred in a training center within our School of Medicine. During the evaluation, the lights were dimmed and a background noise track, incorporating explosion noises, street noise, people coughing and screaming,

construction noise, and sirens created a more realistic environment. The volume was low enough that participants were able to clearly hear the teammate.




Figure 1. Four of the six mannequins in the experimental environment

Six medical mannequins were distributed in the evaluation room (see Figure 1). All mannequins were dressed as civilians. Four mannequins' breathing rate, pulse and responsiveness levels were predetermined and controlled by the experimenters. The fifth mannequin was an infant and only able to cry, while the sixth mannequin was a child that had no responsive behaviors. The four active mannequins had predetermined breathing and pulse rates. Two mannequins had speakers that emulated speech, permitting responses to questions. Another active mannequin was a toddler that did not have a speaker. The fourth active mannequin's eyes blinked via experimenter control.

The Pioneer 3-DX robot teammate was equipped with a laser range finder and navigated the room autonomously on a pre-planned path. The robot's speech was scripted and controlled by the evaluator. When the robot asked a question, the evaluator logged the response in the robot's script and moved the speech process to the next instruction. If the participant required clarification, the evaluator selected pre-scripted clarifications or repeated the instruction.

Procedure. After completing initial forms and questionnaires, a script was read that introduced the disaster response scenario and informed the participant that he or she would be working with a teammate (either a human or a robot). Each participant viewed a four-minute video intended to set the scene of a mass-casualty incident. The video was comprised of scenes from David Vogler's live footage from the September 11th attacks in New York City (Vogler 2010). After the video, the participant was instructed that his or her role was an uninjured, "contaminated" victim who is unable to leave incident area until responders had set up the decontamination area.

During the briefing, participants assigned to the human-human condition were told that they had called 9-1-1, but could not yet leave the contaminated incident area. Because human responders were not permitted into the incident area, participants were asked if they would be willing to assist a remote human first responder to triage victims.

Participants were told that they would be transferred to a human first responder who would lead them through the triage steps, and would record the participant's responses to questions, along with the GPS location of the victims based on the participant's cell phone GPS signal. The participants identified which victim to treat next. The participants used a walkie-talkie with a headset and microphone (in place of a cell phone) to communicate with the remote human teammate (a remotely located evaluator acting as a first responder).

During the human-robot condition briefing, participants were told that they would be co-located with a robot because human responders were not permitted in the contaminated incident area. After participants indicated they would be willing to work with the robot, it led them to the each of victims in turn. The robot communicated with the participants using a digitally synthesized voice, while leading the participants through the tasks. The participants wore a wireless microphone that transmitted responses to the voice interaction system and the evaluator, and the robot's speech was monitored and advanced by the remote evaluator. The participants were able to ask questions, and in a "Wizard of Oz" manner, the remote evaluator either had the system repeat the robot's statement/question or provided a pre-programmed response.

The victims were positioned in a manner that led the human-human condition participants to visit the victims in the nearly the same order as the human-robot condition during the initial triage. However, it was possible for participants to visit victims in a different order than planned during the human-human condition. If this occurred (usually it was a switch of Victims 3 and 4), the evaluator adjusted the script to assess the alternate order during the first round. During the follow-up triage, the first responder provided instructions to the human-human condition participants that guided them to the proper victim based upon the initial triage results and the GPS location collected from the participant's "cell phone."

The triage instructions provided and questions asked were identical across conditions. The teammate (e.g. the remote evaluator or the robot) guided the participant through the steps to identify a victim's triage level. The participants in both conditions started at the same position in the room and moved from victim to victim during the initial triage (Round 1). After completing the initial triage of all six victims, the participant was led back to the five surviving victims for a second triage check (Round 2). During the second triage for the human-human condition, the next victim was specified by referring to the order in which victims were first visited, for example, "please go to the first victim you triaged." The robot led the participant to the appropriate victim during the human-robot condition. Upon reaching a victim, the teammate provided a summary and led the participant through the triage assessment again. The human-robot condition required the participant to place a color-coded triage card on the victim upon completing the triage. The cards were located on the robot platform and the robot instructed the participant which color card to choose. The human-human condition participants were simply told the victim's triage level. Note that the second triage was ordered by severity of triage level.

After triaging each victim, participants responded to a set of subjective workload questions. These questions, along with other data recorded during the scenario were part of a separate analysis (Harriott, Zhang, & Adams, 2011; Harriott, Zhang & Adams, 2011).

Upon completion of the exercise, participants completed a series of post-experimental questionnaires including the measure of cognitive dissonance, a set of four behavioral prediction scenarios, and a teammate evaluation questionnaire. The scenarios

were prefaced by a brief description and illustration of three agents: a computer (illustrated by a monitor, a keyboard and a mouse), a robot (illustrated by the same pioneer 3DX robot used in the human-robot condition), and a person (illustrated by a head shot of a college-aged adult male). Beneath each agent was a brief description of the agent, and the reminder, “When making your responses for the robot [person/computer], consider what kinds of processes characterize a robot [person/computer], as opposed to any other kind of thing”.

Three of the behavioral prediction scenarios were similar to the one described in the introduction: they pitted object-based responses against location (or spatial pattern-based) responses against each other. The fourth scenario asked participants to predict whether agents would choose to classify a set of objects based on their features or their taxonomic category. According to research and theory in cognitive development, a taxonomic classification is characteristic of an intentional agent (Bloom, 1997).

The cognitive dissonance scale included six self-statements: “Sometimes I was uncomfortable answering these questions”, “At times I worried that some of my answers were inconsistent with my other answers”, “If I were allowed to, I would go back and change some of my responses”, “Some of the answers I gave in this experiment were inconsistent with my previous beliefs about the subject”, “I was always certain about my responses” (reverse scored) and “I never had difficulty putting together all of the facts in this experiment” (reverse scored). Participants responded on a seven-point likert scale with anchor points ranging from “completely disagree” to “completely agree”. The questionnaire was scored such that high scores reflected high levels of cognitive dissonance.

The teammate evaluation questionnaire included eight questions that participants responded to using a 9-level likert scale. Six of the questions were direct evaluations of the teammate and the task setting, one was a self evaluation of the participants effectiveness in completing the tasks, and one was a self report about the amount of stress experienced during the task (the specific questions are listed in Figure 3).

Results

Cognitive Dissonance. Participants reported significantly higher levels of cognitive dissonance in the human-robot condition (mean = 3.90; SD = .877) than in the human-human condition (mean = 3.12; SD = .976; $t[32] = 2.409$, $p = .022$). Because the conditions were run successively, and because there was a small difference in mean age between conditions, the effect of condition on dissonance was further tested in the context of a multiple regression controlling for age and sex, and a range of variables related to experience. The experience variables included ratings of first aid experience, robot-related experience, and overall level of education (ranging from “some college” to a completed doctorate). In the full regression including all of these variables, only condition was a significant predictor of dissonance, (Beta = .364, $t = 2.164$, $p = .041$). None of the other variables approached significance (p 's > .20) except for robot experience (Beta = -.316, $t = -1.799$, $p = .083$).

Behavioral Prediction Scenarios. As in previous research (Levin et. al, in review; Levin, Killingsworth, & Saylor, 2008; Levin, Saylor, & Lynn, in review), participants made many more intentional predictions for the human agent, than for the computer or the robot. Overall, participants gave significantly more intentional response for humans, than

for computers ($t[33] = 10.771, p < .001$) or for robots ($t[33] = 9.268, p < .001$). Also consistent with this previous research, participants did not make more intentional predictions for the robot than for the computer ($t[33] = 1.475, p = .150$). The overall pattern of behavioral predictions was very similar in both conditions (see figure 2). The contrast between the human and the robot was significant in both the human-human condition ($t[18] = 5.771, p < .001$) and the human-robot condition ($t[14] = 7.875, p < .001$), as was the contrast between the human and the computer (human-human: $t[18] = 6.553, p < .001$; human-robot: $t[14] = 9.865, p < .001$). The overall contrast between humans and machines (e.g. the average of computer and robot predictions) was slightly larger in the Human-Robot condition (68%) than in the human-human condition (55%; $t[32] = 1.29, p = .20$).

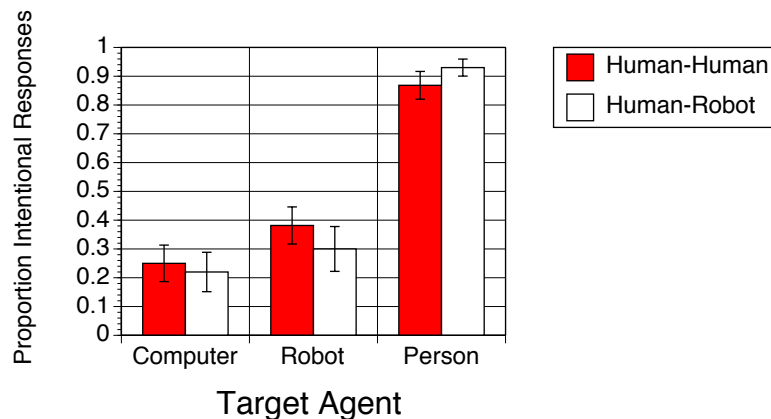



Figure 2. Proportion of intentional behavioral predictions for humans, robots, and computers in the human-human and human-robot conditions. Error bars represent standard errors.

Links between behavioral prediction scenarios and dissonance. We tested the degree to which cognitive dissonance could be correlated with behavioral predictions. There were no significant correlations between cognitive dissonance and behavioral predictions overall (r 's $< .28, p > .10$), but there were strong contrasts in patterns of correlation between conditions. In the human-robot condition, the correlation between dissonance and intentionality of computer predictions was $r = -.57 (p = .023)$, while the same correlation was $+.14$ (ns) in the human-human condition. The predictions for the robots were similar and near zero in both conditions ($r < .15$). The correlations between the human predictions and dissonance were also not significant in both conditions (human-human, $r = -.390, p = .098$; human-robot, $r = .255, p = .359$).

Teammate evaluation questions. An analysis of the teammate evaluation questions demonstrated that participants generally found the robot to be a less satisfactory partner

than the human (Figure 3). Participants indicated that the robot gave less clear instructions ($t[31] = 3.017, p < .01$), and understood them less, ($t[31] = 3.387, p < .01$). Participants also gave lower trust ratings for the robot ($t[31] = 5.654, p < .001$) and felt less comfortable communicating with the robot ($t[31] = 4.406, p < .001$). Finally, participants reported that they did not do as good a job with the robot teammate, ($t[31] = 2.549, p = .016$). Ratings of confusion and stress were not significantly different between conditions.



- Q1. My teammate gave me clear instructions
- Q2. I trusted my teammate
- Q3. I felt comfortable communicating with my teammate
- Q4. My teammate understood what I was trying to communicate
- Q5. I did a good job on the tasks I was assigned
- Q6. I often felt confused about my teammate's instructions
- Q7. I often felt confused as to the purpose of my actions
- Q8. I felt stressed during the scenario

Figure 3. Teammate evaluation ratings by condition. Error bars are standard errors. Between-condition t-tests: * $p < .05$, ** $p < .01$, *** $p < .001$.

To test whether these differences in rated reactions between conditions might explain the differences in cognitive dissonance between conditions, we first ran an exploratory factor analysis on the eight items to reduce the number of dimensions to a usable level. The principle components analysis extraction of the items revealed two components with eigenvalues greater than 1. The first, explaining 50% of variance included questions 1, 2, 3, 4, 6, and 7. The absolute values of loadings for these items were all quite high, ranging from .706 to .837. Based on the content of these items, this factor can be described as the level of trust and clarity inherent to the situation. The second factor explained 19% of total variance and it included items 2 and 8, the stress and success self report, with loadings of .893, and -.694 respectively. Summary scores were created for each of these two factors, and pair of regressions was run testing whether each factor would eliminate the effect of condition in predicting cognitive dissonance.

In the first regression, the trust and understanding items did not significantly predict dissonance (Beta = -.012), while condition was a barely nonsignificant predictor (Beta = .411, $p = .053$). The beta weight for condition was almost identical to the weight for condition without controlling for the trust and understanding items (Beta = .392). In the second regression the stress and success items and the effect of condition were barely not significant (Beta = .312, $p = .069$; Beta = .316, $p = .067$, respectively), and again, the beta weight for condition was only slightly less in the regression controlling for stress and success than in the regression including only condition. These results suggest that the effect of condition on dissonance cannot be reduced to differences in the participant questionnaire responses.

Finally, two regressions tested whether the link between dissonance and computer behavioral predictions would remain for the HR condition, when controlling for the teammate evaluation questionnaire variables. In the first regression ($R = .617$, $p = .072$) trust and understanding was not a significant predictor of computer intentionality (Beta = .169, $t < 1$), while cognitive dissonance remained so, Beta = -.568, $p = .037$. In the second regression ($R = .869$, $p < .001$), both stress and success (Beta = .717, $p = .001$), and cognitive dissonance (Beta = -.926, $p < .001$) were significant predictors of computer intentionality. These regressions suggest that the link between dissonance and computer intentionality cannot be explained by individual differences in responses measured by the teammate evaluation questions.

Discussion

We observed that human-robot interaction produced measurable cognitive dissonance in a controlled experiment characterized by very similar interactions between the human-robot experimental condition and the human-human control. In addition, although the experience of interacting with a robot did not change the mean level of intentionality participants associated with machines, this increased level of dissonance was associated with a link between dissonance and more general concepts about computer intelligence: participants in the human-robot condition who experienced particularly high levels of dissonance were more likely to predict less intentional actions for a computer, while predictions for humans and robots were unaffected. There are three key questions to consider when interpreting these findings. First, why did the human-robot condition lead to increased dissonance? The other two questions concern the secondary finding: why did increased dissonance lead to less intentional behavioral predictions, rather than more intentional predictions, and why was this link present for computers and not robots or humans?

There are a number of plausible reasons why interacting with a robot produced dissonance. One likely alternative is that participants were disappointed in the robot's ability to interact with them. Thus, at least some participants may have experienced conflict between their initial belief that a robot could effectively communicate with and understand them, and the reality that the current state of technology makes this level of interaction difficult. The relatively lower trust and communication ratings reinforce this hypothesis. However, lower trust is unlikely to provide a full explanation for the dissonance effect, because the regression analyses showed very little effect of adding the factor representing these questions to the effect of condition in predicting the dissonance ratings. Accordingly, our current hypothesis is that any effect of trust in impacting dissonance is strongly mediated by the independent likelihood of participants engaging in

additional cognitive processes relevant to their concepts about agency. For example, it is possible that dissonance is mediated by the availability of concepts about agency - participants who invoke background knowledge about agency may experience more dissonance than participants who ignore it.

Another important issue in this experiment is that the correlation with dissonance was stronger with attributed computer intentionality than with robot intentionality despite the fact that the dissonance was caused by an interaction with a robot. It is important to note that this correlation does not appear to be a fluke. In another recent experiment we have observed the same result - observation of an anthropomorphized robot produced a significant link between dissonance and behavioral predictions about computers (Levin et al., in prep). In contrast to the present finding, increased dissonance was associated with increased intentionality for computers. Accordingly we hypothesize that direct interactions such as the result presented in this paper have more potential to lessen attributions of intentionality, while indirect interactions and more simple observations allow the robots to appear more intentional than they are, as participants observe behaviors that they elaborate upon by attributing goals to the agents.

However, the important question remains, why did concepts about computers change after interactions with a robot? This may have occurred because the behavioral prediction questionnaire illustrated the computer with a relatively generic picture of a computer monitor and keyboard, whereas the robot agent was illustrated with a picture of the specific robot used in the experiment. One interesting possibility is that the generic computer allows more flexibility for dissonance-responsive reframing because participants can choose any of a broad array of computer-relevant experiences to reinforce their newly-reframed concepts. For example, participants who had experienced a high level of dissonance may have been surprised the robot's specific pattern of responding, and therefore lessened that dissonance by activating situations where computers had failed to act in a goal-directed manner. In contrast, participants had just experienced the robot much more concretely, and would have fewer opportunities to select from a range of experiences to reinforce any dissonance-induced conceptual change. This line of reasoning is very similar to that employed in research on the above-average effect, which hypothesizes that some personal traits (such as "sophisticated") are very broad, allowing participants the flexibility to select ego-enhancing trait relevant experiences (Dunning, Meyerowitz, & Holtzberg, 1989). This flexibility leads the majority of participants to rate themselves as "above-average". In contrast, other traits (such as "punctuality") are much more constrained, and less susceptible to ego-enhancing reframing. Therefore, fewer participants positively distort their ratings of punctuality.

Conclusion

In summary, this experiment represents the first reported use of a measure of cognitive dissonance in response to human-robot interaction, as it demonstrates a significant increase in dissonance in response to an interaction with a robot. We also observed the dissonance induced by the human-robot interaction, predicted a lessening of intentional predictions for a computer agent. We plan to incorporate these findings into a broader explanation that incorporates concepts of agency into a model that can incorporate conceptual change in response to interactions with a wide range of agents.

References

- Bloom, P. (1997). Intentionality and word learning. *Trends in Cognitive Sciences*, 1, 9-12, 1997.
- Dunning, D., Meyerowitz, J.A., Holtzberg, A.D. (1989). Ambiguity and self-evaluation: The role of idiosyncratic trait definitions in self-serving assessments of ability. *Journal of Personality and Social Psychology*, 67, 1082-1090.
- Eliot, A.J. and Devine, P.G. (1994). On the motivational nature of cognitive dissonance: Dissonance as psychological discomfort. *Journal of Personality and Social Psychology*, 67, pp. 1082-1090, 1994.
- Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On seeing human: A three-factor theory of anthropomorphism. *Psychological Review*, 114, 864-886.
- Harriott, C.E., Zhang, T., & Adams, J.A. (2011). Evaluating the applicability of current models of workload to peer-based human-robot teams. In Proceedings of 6th ACM/IEEE Inter. Conf. on Human-Robot Interaction, 45-52.
- Harriott, C.E., Zhang, T., & Adams, J.A. (2011). Predicting and Validating Workload in Human-Robot Peer-Based Teams. In Proceedings of 21st Conference of Behavioral Representation in Modeling & Simulation, pp. 162-169.
- Hymel, A., Levin, D.T., Barrett, J., Saylor M.M., and Biswas, G. (in press). The interaction of childrens' concepts about agents and their ability to use an agent-based tutoring system.
- Levin, D.T., Killingsworth, S.S., & Saylor, M.M. (2008). Concepts about the capabilities of computers and robots: A test of the scope of adults' theory of mind. In Proceedings of the 3rd Annual IEEE International Workshop on Human and Robot Interaction, 3, 57-64.
- Levin, D.T., Saylor, M.M., Killingsworth, S.S., Gordon, S., and Kawamura, K. (in review). Tests of concepts about different kinds of minds: Predictions about the behavior of computers, robots, and people.
- Levin, D.T., Saylor, M.M., & Lynn, S.D. (invited revision in review). Distinguishing first-line defaults from second-line conceptualization in reasoning about humans, robots, and computers.
- Levin, D.T., Steger, A.P., Lynn, S.D., Adams, J.A., & Saylor, M. (in prep) Conceptual Change in Beliefs About Agency: The Joint Role of Deep Concept Availability and Cognitive Dissonance.
- Vogler, D. (2001). "Raw Video Footage / WTC 9.11.01." 2001. Web. 31 Aug. 2010. <<http://davidvogler.com/911>>.
- Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, 69, 1-34.

Daniel T. Levin, Department of Psychology and Human Development, Vanderbilt University, Nashville, USA. Email: daniel.t.levin@vanderbilt.edu; Julie A. Adams, Department of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, USA. Email: julie.a.adams@vanderbilt.edu; Caroine Harriott, Department of Psychology and Human Development, Vanderbilt University, Nashville, USA. Email: caroline.e.harriott@vanderbilt.edu; Tao Zhang, Department of Electrical Engineering and Computer Science, Vanderbilt University, USA. Email: tao.zhang@vanderbilt.edu.

This research is supported by National Science Foundation Grant No. 0826701, AFOSR award FA9550-09-1-0108, National Science Foundation Grant IIS-0643100, and an Office of Naval Research MURI Program award N000140710749.)