

Decoding Partner Type in Human-Agent Negotiation using functional MRI

Eunkyung Kim (eunkyung@usc.edu)¹, Jared Gilbert (jaredgil@usc.edu)¹,
Charlotte Horowitz (shchar@umich.edu)², Jonathan Gratch (gratch@ict.usc.edu)³,
Jonas T. Kaplan (jtkaplan@usc.edu)¹ and Morteza Dehghani (mdehghan@usc.edu)¹

¹Brain and Creativity Institute, University of Southern California, Los Angeles, CA 90089

²Department of Psychology, University of Michigan, Ann Arbor, MI 48109

³Institute for Creative Technologies, University of Southern California, Playa Vista, CA 90094

Abstract

People interact differently with humans than they do with computers, but there is minimal research on what brings about these differences. Using agents labeled as either “another participant” or a “computer program”, we investigated the differences in people’s behavior and brain activity during the course of a negotiation paradigm. Our results indicate that people perceive human-labeled agents more human-like than computer-labeled agents, and the level of concession in the negotiations is dependent on agent type. We have also found that these differences can be captured in brain activation by showing that parts of the Theory of Mind neural correlates are activated in human-labeled agent conditions, but not in computer-labeled agent conditions. We further demonstrate that brain activity can predict whether the negotiation agent was introduced as a competing human player or a computer program. Overall, our study suggests that labeling an interaction partner as either another human or a computer program leads to significant impacts on one’s decision making.

Keywords: Human-Agent Interaction; Negotiations; fMRI

Introduction

Can computer agents act as substitutes for human beings during the course of an interaction? This has been a popular topic in sci-fi movies for decades. In fact, some computer agents that were thought to exist only in movies a few decades ago, are now widely used in daily life. Smartphones, for example, are commonly used to execute voice commands through programs like Siri, and movie theaters now have more ticket vending machines than guest personnel.

However, human-agent interactions are often quite different from human-human interactions (Gray, Gray, & Wegner, 2007; Melo, Marsella, & Gratch, 2016), and many factors contribute to these differences. Researchers have continuously tried to identify what these disparities are and why they occur, with the hopes to bridge the gap between human-human and human-robot/agent encounters.

A robot’s appearance has been found to be paramount in the interaction style of the human subjects. For example, when people interface with robots that have mechanical, nonhuman like features, even when the robot performs human-like actions, they are often unable to overlook these traits (Hegel, Krach, Kircher, Wrede, & Sagerer, 2008). Thus, robots designed to have eyes similar to humans (Banh, Rea, Young, & Sharlin, 2015), or baby face-like heads (Powers & Kiesler, 2006), were found to be more effective in evoking a more human-like interaction.

These differences have also been extensively studied using brain imaging techniques, especially regions associated

with Theory of Mind (ToM), due to their importance in social interactions. ToM refers to the ability of one person to reason about another person’s mental states, including their intentions and beliefs (Premack & Woodruff, 1978). Previous fMRI studies have demonstrated that cortical activity in the neural structures related to ToM tend to be more active when participants were told they were facing a human partner compared to a computer program (Kircher et al., 2009). Research also demonstrated that activity in these same regions scaled according to the human-likeness of their interaction partner when using computer-animated characters or nonhuman agents (Chaminade, Hodgins, & Kawato, 2007; Krach et al., 2008).

Many previous behavioral and neuroimaging studies have used negotiation platforms to examine human-agent communication, because negotiations involve complex cognitive effort and established social interaction techniques. For example, studies show that when participants play the Ultimatum Game against computer partners, they are more likely to accept unfair offers compared to when they play against human partners, where they tend to be less willing to accept offers of unequal value (Sanfey, Rilling, Aronson, Nystrom, & Cohen, 2003). Studies have also shown that people have stronger emotional reactions to unfair offers made by other humans (Vant Wout, Kahn, Sanfey, & Aleman, 2006).

Using a multi-round, multi-object negotiation platform for our research, we explored whether a computer agent introduced as another human was perceived more anthropomorphically than one that was introduced as a computer program. We then investigated whether agent type produced behavioral differences, and whether one type of agent resulted in more concessions compared to the other. In a follow up experiment, we compared brain activity during interactions with human-labeled and computer-labeled agents to determine whether these perceptual differences were also observable in brain patterns. Following collection of fMRI data, we investigated whether classifiers could be trained to determine whether the participant was playing against a human-labeled or computer-labeled agent.

We hypothesized that participant behavior and brain activity would be different during interactions with human-labeled agents, compared to interactions with computer-labeled agents, even though both agents used exactly the same strategies and emotions. Our initial experiment consisted of an online negotiation task intended to explore perceptual and

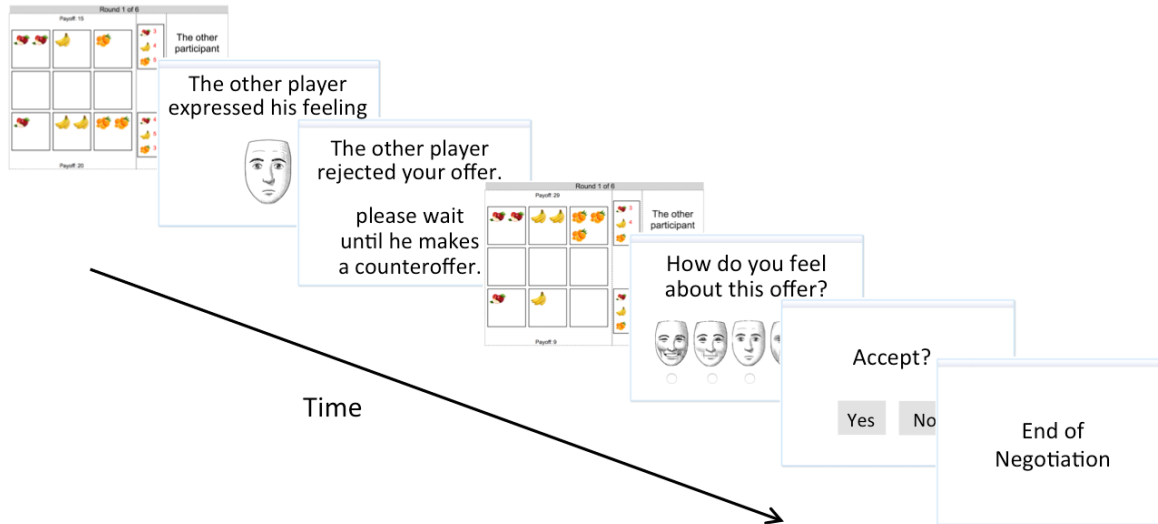


Figure 1: Objects Negotiation Task Timeline

behavioral differences pertaining to anthropomorphic characteristics in human and computer-labeled agents. Next, we adapted our negotiation framework into an fMRI experiment, attempting to find neural differences for the two distinct partner conditions. In addition to these studies, we also ran a prediction algorithm and multi-voxel pattern analysis (Norman, Polyn, Detre, & Haxby, 2006) based on the fMRI data.

This work is distinct from previous studies due to the use of an identical computer agent, regardless of what partner type was specified. The majority of previous studies employed computer-animated characters or robots that had differing levels of anthropomorphism. We demonstrate that even though the same computer agent is used, perceptual differences were captured in behavioral and brain data. To the best of our knowledge, this is one of the first lines of research that uses a multi-round negotiation platform to investigate perceptions of anthropomorphism. We believe that natural interactions take place over multiple rounds/sessions, and it is therefore important to investigate perception differences through multi-round negotiations.

This paper is structured as follows. First, we introduce and explain the Object Negotiation Task, the platform used for both experiments described. Next, we outline an exploratory, behavioral experiment performed to examine the differences between interactions with a human-labeled agent and a computer-labeled agent. After which we discuss our fMRI experiment, using the same paradigm as the first, trying to identify differences in brain activity. Lastly, we discuss the implications of our results and future work.

Objects Negotiation Task

The Objects Negotiation Task is a web-based multi-round negotiation task where a participant and a computer agent can take turns distributing objects (Dehghani, Carnevale, & Gratch, 2014). The original version was designed for behav-

ioral data collection only, so in this paper, a modified version of this task was used for collection of both behavioral and brain data. Figure 1 shows the timeline of the modified version of this task. Some of the modifications made included an emotion-reporting phase and offer-review phase, which were added to separate collection of brain data between differing phases. In addition, a partner introduction phase was added, allowing participants to receive a notification specifying whether their partner type was another participant or the computer program before the negotiation began.

The sequence of the modified Objects Negotiation Task is as follows. When the task begins, the negotiation partner type is displayed. Specifically, in the human-labeled agent condition, the message shown to the participant is ‘In this task, you will be negotiating with the other participant.’ In a computer-labeled agent condition, the same message was shown but ‘the other participant’ was changed to ‘a computer program’. Next, a ‘connection establishment’ message for the human condition and a ‘program setup’ message for the computer condition appear on screen, to persuade participants of their partner setting. Throughout the negotiation, the partner type is constantly included on screen so the participant clearly recognizes his/her partner type. The partner is labeled as ‘the other participant’ or ‘the computer program.’

In the first negotiation round, items are positioned in the middle row, indicating that those items belong to neither player. The participant is asked to propose an initial offer by moving items into his/her own set of boxes (bottom row) or their partner’s set of boxes (top row). Once the initial offer is made, the partner (agent) chooses an emotion pertaining to the offer which is then displayed to the participant. Available emotions include: happy, content, neutral, angry and sad. The partner only shows the predefined emotion for each round. After the emotion is displayed, the partner decides whether to accept or reject the offer. This decision is based on a pre-

defined offer value; when the payoffs of the predefined offer are more than the participant's current offer, the partner rejects, when the payoff is less, the partner accepts.

If the participant's offer is accepted, the items are distributed as proposed and the participant is notified. If the participant's offer is rejected, the partner then proposes a counteroffer. When the counteroffer is received, the participant has 5-seconds for review. During this time, the participant can only observe; no items can be transferred. The review time was specifically introduced for optimal brain activation, as we wanted to record an active decision-making process. For the same reason, our analysis was focused on data collected during this review phase. After the review phase, the participant reports his/her emotion about the proposed offer by choosing from the following descriptive options: happy, content, neutral, angry, or sad. The participant also decides whether to accept or reject the offer. If the participant rejects the offer, a new round begins, and all phases are repeated. The negotiation can last for a maximum of six rounds. If no agreement is made in six rounds, neither party receives anything.

Study 1: Online Experiment

We designed an online experiment to determine whether people perceive human-labeled agents differently than computer-labeled agents during interactions in our negotiation game, as well as to find behavioral differences between agent type in concession-making.

Negotiation Partners Two sets of strategies and two types of emotions were used for the partner agents. Agent strategies included tough and soft. A tough strategy starts with a greedy offer, and a soft strategy starts with a relatively generous offer. Figure 2 shows payoffs for the agent and the participant when the agent uses tough or soft strategies.

Agent emotions included anger and neutral (no emotion). Anger was chosen because it was found to be the most effective emotion in yielding concessions during negotiation tasks (Van Kleef, De Dreu, & Manstead, 2004). For the anger condition, the agent displayed an angry face in rounds 2, 4, and 6, and a neutral face in rounds 1, 3, and 5. For the neutral condition, the agent reported a neutral face in every round.

Procedure 420 subjects (237 male and 183 female; mean age = 33.5) living in the United States were recruited via Amazon Mechanical Turk (MTurk). Each participant was asked to read a hypothetical scenario in which they acted as a restaurant owner, and negotiated for fruit with another restaurant representative due to a fruit shortage as a result of a recent fire in a local market. Each subject was then told to negotiate with either a computer program or another (hypothetical) MTurk player. Regardless of type label, the negotiation partner was always a pre-programmed computer agent. After completing all negotiations, subjects were asked to fill out an anthropomorphism questionnaire (Bartneck, Kulić, Croft, & Zoghbi, 2009) about their partner, as well as a demographic




Figure 2: Payoffs for agents and participants across both agent strategies.

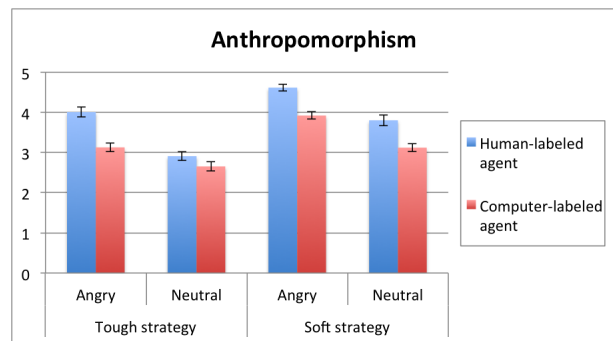


Figure 3: Anthropomorphism Scores for human-labeled agents and computer-labeled agents. Higher score means the agent is perceived as more human-like. The error bar shows standard errors.

questionnaire. In the anthropomorphism questionnaire, participants rated their impression of their partner using a scale from 1 to 7, where 7 means human-like and 1 means machine-like. Subjects were also given a simple attention-check question, implemented to make sure the participants were paying attention; it merely asked what type of partner they were assigned during the task. Each participant was compensated \$1.

Data Analysis We excluded subjects who had participated in our previous negotiation studies or failed to give the correct answer to the attention-check question. After exclusion, we had data from 329 subjects. Scores from each condition were calculated for the anthropomorphism questionnaire to verify whether participants perceived human-likeness differently between agents. In addition, we calculated concessions across partner type in each condition to analyze behavioral differences. Concession was calculated by subtracting payoffs of agreed offers from payoffs of initial offers. A three-way between-subjects analysis of variance (ANOVA) was used to find the interaction between partner type, partner strategy, and partner emotion during concession.

Results The anthropomorphism scores of the human-labeled agents and the computer-labeled agents are shown in Figure 3. 1-way ANOVA results show that people consis-




Figure 4: Concessions to human-labeled agents and computer-labeled agents. The error bars show standard errors.

tently thought their partner to be more human-like when told their partner was a human player, no matter what the negotiation strategy ($F(1, 327) = 14.09, p < 0.001$).

Concessions to human-labeled agents and computer-labeled agents during negotiations are shown in figure 4. ANOVA results show that there is a $2 \times 2 \times 2$ interaction between agent type (human/computer) \times agent emotion (angry/neutral) \times agent strategy (tough/soft) for concession ($F(1, 321) = 3.387, p = 0.066$). We also ran a 2×2 ANOVA after dropping each strategy. A two-way interaction was found for tough strategy ($F(1, 164) = 4.699, p = 0.031$).

Discussion Our findings from anthropomorphism scores suggest that there are perception differences in the interactions between human-labeled agents and computer-labeled agents. Also, our ANOVA results suggest that there is an interaction between agent type \times agent emotion \times agent strategy for concession. This indicates that not only are people’s perceptions of the two agents distinct, but their behaviors also vary depending on agent type. To study whether these behavioral differences have neural correlates, we designed the following fMRI experiment. Because the largest concession differences were found in the tough conditions, implying the tough strategy was best suited to observe those differences in behavior, we mainly employed tough agents in the following experiment.

Study 2: fMRI Experiment

We hypothesized that perceptual and behavioral differences could be captured in brain activity, especially in ToM related brain regions, as they were found to be correlated with human-likeness of physically existing human-like robots (Chaminade et al., 2007; Krach et al., 2008). Each subject performed the negotiation task with both types of agent in order to compare brain activity from human-labeled vs. computer-labeled agent interaction. To make interactions with human-labeled agents more realistic, we introduced a confederate into the study, so participants believed they would be competing against another human player.

Participants 20 healthy American subjects (10 male and 10 female), recruited via the University of Southern California online bulletin board, took part in this study. Subjects were 21.4 years old on average ($SD = 2.58$). All participants were right-handed and had no history of neurological or psychiatric disorders.

Negotiation Partners Although we were only interested in tough agent type, we used two types of soft agents on top of four types of tough agents (human/computer-labeled \times angry/neutral agents). This modification was implemented because subjects participated in a series of consecutive negotiations, unlike our online experiment where each subject only played in a single negotiation. Including soft agents ensured that participants did not play with the same agent over and over again. Each subject negotiated with six types of agents. While every subject negotiated with four types of tough agents, 10 subjects (5 male and 5 female) negotiated with two types of emotion-neutral soft agents, while the remaining 10 subjects negotiated with two types of emotion-angry soft agents. Agent order was randomized.

Procedure Each participant was greeted by an experimenter and introduced to the confederate as the competing player. The participant and the confederate were guided to a preparation room where they filled out an informed consent form, incidental findings form, and safety screening form. After forms were completed, the confederate was guided to a separate MRI room for “setting up”. The participant was given the instructions and rules regarding the negotiation task, and played a trial negotiation against a computer program before starting the experiment. During the trial, a trackball mouse similar to one used in the scanner environment was provided, so that the participant became familiarized with its operation. The participant was then guided to the actual MRI room and was told that while in the scanner he/she would run through three negotiation tasks with the participant in the other MRI room, and three negotiation tasks with the computer program. The task was back projected on a screen, seen through a mirror attached to the head coil, and operated a trackball mouse to navigate negotiations. In each task, a different set of negotiation items were used, and payoffs for these items varied with position, in order to give an impression that each negotiation was unique. Participants answered a shortened version of the anthropomorphism questionnaire at the completion of each round. After a maximum of six negotiation rounds, participants filled out a handedness and demographic questionnaire. Before leaving, subjects were debriefed and compensated \$30.

fMRI Data Acquisition fMRI scans were performed at the USC Dana & David Dornsife Cognitive Neuroscience Imaging center. Images were acquired using a 3-Tesla Siemens PRISMA MRI scanner with a 20-channel matrix head coil. Two sets of high-resolution anatomical images were acquired for registration purposes. Six sets of echo-planar images (EPI), one set for each negotiation, were acquired continu-




Figure 5: Frontal medial cortex from Harvard-Oxford atlas (left) and overlaid accuracy map for all participants (right). A part of frontal medial cortex was included in the accuracy map (white dotted box).

ously with the following parameters: TR = 2,000ms, TE = 25ms, flip angle = 90° , 64×64 matrix, one shot per repetition, in-plane resolution $3 \times 3 \text{ mm}^2$, 41 transverse slices, each 3mm thick, covering the whole brain. Total scan time for each participant was approximately 50 minutes.

Data Analysis We conducted a general linear model (GLM) analysis and then used the results as input for multivoxel pattern analysis (MVPA). GLM analysis was performed using FMRIB's Software Library (FSL) to locate brain regions activated during proposal review phases. As mentioned earlier, this phase was specifically targeted due to the likelihood of collecting data pertaining to decision making for subsequent negotiation rounds. For GLM analysis, data pre-processing steps included motion correction, brain extraction, spatial smoothing, slice timing correction, and high-pass temporal filtering. After completing data pre-processing, we modeled brain activity during proposal review phases using a double gamma hemodynamic response function. Data collection from all other time points were used as baseline. We then performed MVPA to find brain regions that illustrated different patterns across agent type. In MVPA, neural representations were decoded by applying pattern-classification algorithms on fMRI data (Norman et al., 2006). We used detrended and z-scored GLM analysis results as inputs for MVPA, and trained a linear Support Vector Machine (SVM) classifier using feature selection, introduced to improve classification performance by picking the most relevant features as inputs for the classifier (Guyon & Elisseeff, 2003). Searchlight analysis (Kriegeskorte, Goebel, & Bandettini, 2006) was used as the feature selection method to analyze contents multivariately at each location in the brain. We implemented a leave-one-participant-out cross validation as balance for MVPA. More details on fMRI data analysis can be found in (Kim, Gimbel, Litvinova, Kaplan, & Dehghani, 2016), as the same analysis methods were used.

Results The right side of Figure 5 shows an overlaid accuracy map for all participants. Accuracy maps were acquired from the top 5% of all t-maps, which were generated using t-tests versus chance. Interestingly, medial prefrontal cortex, one of the ToM brain regions, was included in the accuracy map, suggesting that people indeed perceived the

human-labeled agent to have more human-like qualities.

MVPA, with searchlight as a feature selection method, revealed that agent type (human/computer) can be predicted based on brain activity during proposal-review phases. Prediction accuracy for agent type was 58.41%, with a standard error 0.01%, where chance level is 50%. The improvement was found to have statistical significance (Two-tailed t-test: $p < 0.001$).

Discussion The results of Study 2 demonstrates that differences in how we perceive the 'humanness' of an agent can be captured using fMRI. Specifically, our results show that parts of the ToM neural correlates are activated in human-labeled agent conditions, but not in computer-labeled agent conditions. This finding is consistent with previous studies that reported increased brain activity in ToM brain regions corresponding to human-likeness of interacting partners (Chaminade et al., 2007; Krach et al., 2008). Our MVPA analysis further revealed that these differences are great enough that classifiers can be trained that can reliably distinguish brain activity between the two types of agents.

Overall Discussion

Our goal was to investigate differences in behavior and brain activity during human-agent negotiations. Focusing on partner type, we hypothesized that both parameters would be distinct when comparing computer-labeled and human-labeled interactions.

Results from our online experiment indicate that people perceive human-labeled agents more human-like than computer-labeled ones, even though both used parallel strategies and emotions. This suggests that people's attitudes towards computer partners are distinguishable from those towards human partners. Furthermore, a 3-way interaction between agent type \times agent emotion \times agent strategy was found for concession.

Results from our fMRI experiment suggest that brain patterns observed during interactions with human-labeled agents are different from ones with computer-labeled agents. More specifically, the medial prefrontal cortex, part of the ToM-related neural structures, was found to be included in accuracy maps, indicating neural activity during interactions with human-labeled agents are distinct from ones with computer-labeled agents. This is in line with a previous finding, where the medial prefrontal cortex was found to be activated while playing rock-paper-scissors with a human player, but not activated when playing the same game with a known, pre-defined computer algorithm (Gallagher, Jack, Roepstorff, & Frith, 2002). Using a negotiation paradigm, a more complicated task than rock-paper-scissors due to the inclusion of multiple decision-making rounds, we found that the same effect exists with differently labeled agents. Negotiations require a more substantial cognitive effort than a game like rock-paper-scissors; there are a larger amount of possible actions to consider, an increased opportunities for loss, and a greater obligation to compete, or cooperate and come to some sort of

agreement. The negotiation tasks used in the aforementioned experiments are able to advantageously measure interactions that require high levels of cognitive energy, and are therefore more useful when attempting to explore human-robot interactions.

Our findings are also consistent with previous studies that found ToM-related neural structures to be more activated when interacting with agents that had more human-like characteristics, regardless of whether those agents happened to be robots or computer-generated characters (Chaminade et al., 2007; Krach et al., 2008). These findings imply perceptual variance between interactions with human-like and nonhuman-like agents, and leads us to believe that human participants do engage in greater mentalising efforts when faced with human-labeled or human-like robots.

In conclusion, we examined the relationship between negotiation partner type and behavioral and neural measures regarding individual's perceptions of human-agent interactions. Our study suggests that when either by labeling agents as other humans or as computer programs significantly impacts one's perception of the situation; this is ultimately demonstrated through negotiation behavior and brain activity. The results give us further insight into the counter-play between emotional and cognitive processes, leading us to believe that our emotions may have greater impact on decision making than which we are consciously aware. Ultimately, these results inform us that there is a noticeable and consistent difference between the perceptual and emotional reactions that humans have towards other humans when compared to those same reactions with computer agents. Further research needs to be executed to more thoroughly understand why these differences in interaction occur, but this study has illustrated that computers and technology do indeed impact the way humans interact with the world, and each other. This is important to consider as computers continue to be increasingly implemented in everyday life and society.

Acknowledgments

This research is supported by AFOSR FA9550-14-1-0364.

References

- Banh, A., Rea, D. J., Young, J. E., & Sharlin, E. (2015). Inspector baxter: The social aspects of integrating a robot as a quality inspector in an assembly line. In *Proceedings of the 3rd international conference on human-agent interaction* (pp. 19–26).
- Bartneck, C., Kulić, D., Croft, E., & Zoghbi, S. (2009). Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International journal of social robotics*, 71–81.
- Chaminade, T., Hodgins, J., & Kawato, M. (2007). Anthropomorphism influences perception of computer-animated characters' actions. *Social cognitive and affective neuroscience*, 2(3), 206–216.
- Dehghani, M., Carnevale, P. J., & Gratch, J. (2014). Interpersonal effects of expressed anger and sorrow in morally charged negotiation. *Judgment and Decision Making*, 9(2).
- Gallagher, H. L., Jack, A. I., Roepstorff, A., & Frith, C. D. (2002). Imaging the intentional stance in a competitive game. *Neuroimage*, 16(3), 814–821.
- Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science*, 315(5812), 619–619.
- Guyon, I., & Elisseeff, A. (2003). An introduction to variable and feature selection. *The Journal of Machine Learning Research*, 3.
- Hegel, F., Krach, S., Kircher, T., Wrede, B., & Sagerer, G. (2008). Understanding social robots: A user study on anthropomorphism. In *Robot and human interactive communication, 2008. ro-man 2008. the 17th ieee international symposium on* (pp. 574–579).
- Kim, E., Gimbel, S. I., Litvinova, A., Kaplan, J. T., & Dehghani, M. (2016). Predicting decision in human-agent negotiation using functional mri. *Proceedings of the 38th Annual Meeting of the Cognitive Science Society*.
- Kircher, T., Blümel, I., Marjoram, D., Lataster, T., Krabben-dam, L., Weber, J., ... Krach, S. (2009). Online mentalising investigated with functional mri. *Neuroscience letters*, 454(3), 176–181.
- Krach, S., Hegel, F., Wrede, B., Sagerer, G., Binkofski, F., & Kircher, T. (2008). Can machines think? interaction and perspective taking with robots investigated via fmri. *PLoS one*, 3(7), e2597.
- Kriegeskorte, N., Goebel, R., & Bandettini, P. (2006). Information-based functional brain mapping. *Proceedings of the National Academy of Sciences of the United States of America*, 103(10).
- Melo, C. D., Marsella, S., & Gratch, J. (2016). People do not feel guilty about exploiting machines. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 23(2), 8.
- Norman, K. A., Polyn, S. M., Detre, G. J., & Haxby, J. V. (2006). Beyond mind-reading: multi-voxel pattern analysis of fmri data. *Trends in cognitive sciences*, 10(9).
- Powers, A., & Kiesler, S. (2006). The advisor robot: tracing people's mental model from a robot's physical attributes. In *Proceedings of the 1st acm sigchi/sigart conference on human-robot interaction* (pp. 218–225).
- Premack, D., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and brain sciences*, 1(04), 515–526.
- Sanfey, A. G., Rilling, J. K., Aronson, J. A., Nystrom, L. E., & Cohen, J. D. (2003). The neural basis of economic decision-making in the ultimatum game. *Science*, 300(5626).
- Van Kleef, G. A., De Dreu, C. K., & Manstead, A. S. (2004). The interpersonal effects of anger and happiness in negotiations. *Journal of personality and social psychology*, 86(1).
- Vant Wout, M., Kahn, R. S., Sanfey, A. G., & Aleman, A. (2006). Affective state and decision-making in the ultimatum game. *Experimental brain research*, 169(4), 564–568.