

Examining the Effects of Anticipatory Robot Assistance on Human Decision Making

Benjamin A. Newman*, Abhijat Biswas*,
Sarthak Ahuja, Siddharth Girdhar, Kris K. Kitani, and Henny Admoni

Robotics Institute, Carnegie Mellon University, Pittsburgh PA, USA
{newmanba, abhijat, sarthaka, sgirdhar, kkitani, henny}@cmu.edu

Abstract. Collaborative robots that provide *anticipatory assistance* are able to help people complete tasks more quickly. As anticipatory assistance is provided before help is explicitly requested, there is a chance that this action itself will influence the person’s future decisions in the task. In this work, we investigate whether a robot’s anticipatory assistance can drive people to make choices different from those they would otherwise make. Such a study requires measuring intent, which itself could modify intent, resulting in an observer paradox. To combat this, we carefully designed an experiment to avoid this effect. We considered several mitigations such as the careful choice of which human behavioral signals we use to measure intent and designing unobtrusive ways to obtain these signals. We conducted a user study ($N=99$) in which participants completed a collaborative object retrieval task: users selected an object and a robot arm retrieved it for them. The robot predicted the user’s object selection from eye gaze in advance of their explicit selection, and then provided either collaborative anticipation (moving toward the predicted object), adversarial anticipation (moving away from the predicted object), or no anticipation (no movement, control condition). We found trends and participant comments suggesting people’s decision making changes in the presence of a robot anticipatory motion and this change differs depending on the robot’s anticipation strategy.

Keywords: Anticipatory Motion-Eye Gaze-Human Robot Interaction-Human Decision Making.

1 Introduction

Anticipatory assistance is the ability to continuously forecast a person’s actions and take steps to assist in executing those predicted actions. Anticipation is an important capability for collaborative human-robot systems to facilitate seamless interactions.

Consider a customer ordering a smoothie from a robot barista in a cafe. The customer aims to choose ingredients that align with their tastes. The robot barista aims to maximize customer satisfaction by making the smoothie accurately to the customer’s choices while maximizing efficiency by acting in anticipation of their choice. Prior research [11, 12] has shown that in such situations the robot can capitalize on human behaviors exhibited during the interaction (*e.g.*, the customer’s eye gaze) to forecast customer actions (*e.g.* smoothie ingredient selections) and use these forecasts to inform its own action selection. This process increases the fluidity of the overall interaction

* The first two authors contributed equally to this work




Fig. 1: *Can anticipatory robot motions affect a user's choice?*

and reduces the overall smoothie making time. These desirable interaction qualities are achieved due to the anticipatory assistance delivered by the robot.

Now consider a situation in which the robot has incorrectly predicted the customer's desired choice of ingredient. For example, the robot might anticipate the customer wants to add artificial banana flavoring, while the user is actually deciding to add a real banana to the smoothie. This creates a decision point for the customer: should they change their original intent (and thereby their initially proposed action) or hold true to their original plan? Regardless of the customer's eventual decision, their action is now clearly dependent on the robot's decision to act in anticipation. While prior research has shown that user actions can influence robot actions in order to engender anticipatory assistance, **we hypothesize that this influence is actually bidirectional, and that the robot's actions influence user actions, as well.** Additionally, while this example highlights an explicit decision point for clarity, we hypothesize that this process can happen implicitly, as well.

To study this hypothesis we conducted a user study with a robot arm that implements a collaborative 1-among-3 selection task. We designed this study to explore our main research question: **do anticipatory robot actions affect user decision making?** Participants in our study were presented with three bins, each filled with a unique variety of candy. The robot monitored the participants' eye gaze and used an online model to map eye gaze to preferred candy bins. The robot then took no action (*i.e.* control) or acted in anticipation of its expectation of the participants' action, as determined through their eye gaze. This action was performed in either a collaborative (towards a user's predicted choice) or adversarial manner (away from a user's predicted choice, *i.e.*, a mistake), depending on the condition. We explore the idea that collaborative anticipation and adversarial anticipation may lead to differences in the robot's influence on the user. Following this, the participant selected their preferred bin by pressing a button prompting the robot to hand them the corresponding bin. This setup is shown in Figure 1.

Our study was designed to limit measurements that would influence user intent outside of the robot's anticipatory action. Such measurements included survey questions requiring the user to explicitly articulate their intent and questions or prompts revealing that the user's choices would be closely monitored. This restriction prohibited us

from obtaining ground truth measurements for initial user intent. We thus designed, implemented, and tested a method to measure user intent before its explicit expression. In line with previous research, we use a model of the user’s eye gaze in order to measure this unarticulated intent [11, 12, 14].

With this study we address the following research questions:

Robot action-based intent re-shaping: Can a robot making anticipatory movements (movement towards/away from a participant’s intended goal) affect a participant’s eventual decision?

Unarticulated intent prediction through eye gaze: Does participant eye gaze preceding an explicit decision accurately predict that decision?

We hypothesize that a user’s decision making process can be affected by a robot’s anticipatory actions. We then present a user study that continuously measures the development of user intent using eye gaze naturally exhibited during a selection task and investigate if user intent can be altered by displaying anticipatory robot motions. We find quantitative trends and qualitative results suggesting the bidirectional nature of human robot interactions in anticipatory systems as well as differing responses to different types of anticipatory systems. Finally, we discuss challenges we encountered during our study and propose strategies to mitigate these challenges in future studies.

2 Related Work

Anticipatory robot behavior in human-robot interactions Robots utilizing anticipatory actions have been shown to improve user engagement during human-robot interactions. Previous work suggests that participants working with an anticipatory robot attribute more human qualities to the robot [10] than robots that do not anticipate a user’s actions. Modelling human behavior explicitly can even lead to beneficial outcomes as evidenced by robots that explore more safely [3] and learn more effective policies [6, 19] as compared to robots that do not model their human counterparts. Especially relevant to our own study is work that used eye gaze as a signal of a user’s intent, and thus as the input into an anticipatory robot system. This work showed that interactions such as our smoothie making example can be accelerated when compared to reactive systems [12]. Our inquiry, to understand whether anticipatory robot motions have an effect on user choice, was not considered in this previous work.

Anticipation is not the only factor at play in these types of interactions, however, as the quality of the expression of the robot’s anticipation can shape the success of the collaboration. For example, an anticipatory action that is not correctly perceived (*i.e.*, is not legible) by the user can prove to be counter-productive [4]. One conclusion of this work is that a robot has the potential to disrupt a person’s intentions by deliberately choosing legible or illegible actions at opportune moments during the interaction. Even when this expression is not intentional, research on human teams show that team members adapt to each other during collaborative tasks, and that this mutual adaptation leads to improved team performance [15]. This effect has even been shown to translate to human-robot teams [16], where the robot models a user’s adaptability and changes its actions based upon this estimation.

Furthermore prior work in human intent reshaping has shown that intentional robot movement can cause a person to change their initial course of action when acting independently in a room [5]. The authors infer human intent by using a hidden Markov




Fig. 2: Typical trial timeline. In this case, the participant saw the *collaborative* condition and accepted the robot’s suggestion. Video here.

model based on the human’s location and posture. They observe that robots were able to reshape the intention of 68% of the 15 study participants. This work also reports that the robot’s potential to reshape a user’s intention decreases as the user interacts and becomes more familiar with the system.

Collaborative and adversarial actions have been studied in the context of how obedient and rebellious robots affect a participant’s cooperation during task completion [20]. This research reports that working with robotic systems that make collaborative actions can lead to enhanced human robot cooperation when compared to rebellious, or adversarial, robotic systems. In this prior work, the user had a stated, explicit goal. Our work explores the effects of anticipatory actions in tasks where the user has no specific objective.

Predicting user intent from human behavior Robots can infer intent by sensing a range of verbal and non-verbal signals exhibited by people. For example, existing work has shown that a person’s desired target object can be inferred from the history of user control inputs on a joystick [13] during a human-robot joint teleoperation task. Similarly, the handover intention of a participant can be inferred from a combination of wearable sensors and natural language commands [21].

Other work has shown that by leveraging these types of patterns in a user’s past actions, robots can further improve collaborations by anticipating future user actions [19]. Additional research in human-robot teaming shows that if a robot disregards a human’s preference it can lose the human’s trust, leading to a deterioration in team performance. This shows that anticipating and adapting to a user’s evolving preferences throughout a task is crucial to effective human-robot collaboration [8].

Most relevant to our own study, natural eye gaze is highly predictive of user goals, future actions, and mental states - particularly in unaccompanied human manipulation tasks [9, 14, 17]. Eye gaze has been used as a supplementary modality to infer human intent primarily alongside head-hand tracking [7, 22, 23]. Recent research shows great promise of integrating eye gaze with joystick control for shared manipulation tasks [1, 2]. For simple collaboration tasks such as our smoothie making example, eye gaze has also been shown to be a strong stand-alone modality for intent prediction using a support-vector machine on handcrafted features [11]. In our work, we use spatial binning on a moving exponential average over a temporal sequence of natural eye gaze as a heuristic measure for anticipating human intent as depicted in Figure 1a.

3 Methods

To answer our research questions, we conducted a user study in which participants interacted with a robot arm in a handover game. Participants were asked to select a

target object from among three options located on a table in front of them. The robot used participant eye gaze behavior to model their preferences among the objects, then asked for their selection and finally handed them their selected object. We examined whether participants’ object selections, decision-making process and perceptions of the robot would be affected if the robot took anticipatory actions. These actions were either collaborative or adversarial (toward the most or least likely object respectively). A typical trial is illustrated in Figure 2 and in a video [here](#).

3.1 Study design

Anticipation model Human signals used to estimate intent prior to its explicit expression must be natural and measurable. Eye gaze has been shown to be correlated with intent in similar settings [11]. Due to this correlation, we map user eye gaze to predicted user choice in the following way.

We model the relationship between a user’s eye gaze and their eventual choice as an exponentially decaying moving average of the temporal history of their eye gaze. Our model takes as input a sequence of planar eye gaze locations and classifies each eye gaze measurement into one of the four segments shown in Figure 1a. The model then filters out the eye gaze in the top segment, which is considered robot-viewing gaze, and applies a temporal, exponentially decaying moving average over counts in each of the remaining three segments. This creates temporally evolving probability distribution that maps to the expected probability of a user’s choice.

Experimental variables Our study uses a between subjects design with one independent variable, “robot suggestion,” and three conditions. In each condition, participant preference is determined by the aforementioned anticipation model: *collaborative*: robot moves toward the **most preferred** bin; *adversarial*: robot moves toward the **least preferred** bin; *no-movement*: control, in which the robot takes **no action**.

All participants first saw two trials of *no-movement*, followed by a single experimental trial in which one of the three conditions above was randomly applied. This design allows participants to become familiar with the study procedure and robot operation in the first trial. This trial is treated as practice, and therefore these data were not included in analysis. In total, each participant experienced 3 trials in the following manner: *practice*, *control*, and *experimental*.

3.2 Procedure

Participants were seated across the table from the robot with three bins, filled with unique candy types (Figure 1a). Bins were initially hidden from participant view. Participants were told that once the candy was revealed, they would be given some time to decide between the candies and choose one, which they could keep. In the practice and control trials, after the candy was revealed, participants were given 5 seconds to decide, after which the buttons lit up prompting participants to make a selection (Figure 1a). In the experimental trials, the anticipatory motion began at 5 seconds, preceding the illumination of the selection buttons. Once participants pressed the button indicating their selection, the robot retrieved the corresponding bin, offered it to them, and allowed them to take a candy (Figure 2).

All participants saw the same candies, in the same locations, in the first two trials. For the experimental trial, we discarded the two types of candy previously selected and

randomly replaced from the four remaining types and randomized all candy locations. This randomization and replacement ensured that participants had not previously expressed preference for the types of candy present in the experimental trial. We recorded the **Predicted Bin** (bin our anticipation model predicts the user will eventually choose), **Selected Bin** (bin selected by the user), **Suggested Bin** (target bin of the robot’s anticipatory motion), as well as survey data following each trial (Sec. 4.3).

Materials and Participants We used a Kinova MICO 6-DOF manipulator (Figure 1) robot with pre-computed trajectories for all actions. Participant eye gaze was measured using a Tobii 4C eye tracker bar, typically used for 2D planar eye tracking. We adapt it for use in this 3D setting by segmenting the user viewing plane into sections corresponding to three candy segments and one robot segment (Figure 1a). The stimuli used were 6 flavors of an international candy, each with a distinct label. These stimuli were chosen to be unfamiliar to the majority of the participant pool, consistent in shape and size, have little correlation between wrapper patterns and flavor, and have intricate labels. Thus, the stimuli required attentive viewing before being chosen and would not be chosen due to prior familiarity or preference with the candies’ brand or flavor.

One hundred and eighteen participants (72 Female/44 Male/2 Other) were recruited through an online subject pool, mailing lists, and word-of-mouth. Twenty one participants were excluded due to eye gaze sensor drop out or participant non-compliance. Our analyses were performed on $N = 99$ (33 for each experimental condition). The research was approved by an institutional review board. Participants were compensated for their time with 5 USD, for an average of 15 minutes.

Design considerations Investigating whether or not anticipatory robot motions affect human decision making is challenging due to the observer’s paradox. Attempts to measure intent in a human decision making task can result in altered behavior, especially if participants are aware of the observation [18].

First, we take care to capture the user’s eye gaze in a subtle manner. Accurate eye gaze data is typically obtained by instrumenting participants with wearable eye-trackers [11, 12]. However, an implied social presence via a worn eye-tracker has been shown to change looking behavior [18]. We minimized this interference by avoiding the instrumentation of participants, instead adapting a discreet, screen-based eye tracker (Section 3.2) for use in this study (Figure 1). Second, we diverted the focus of the study away from the anticipatory motions by advertising our study as a “handover game” indicating to potential participants that we were interested in the robot’s handover mechanics. We also ran a between subjects experiment with only one experimental trial, so that anticipatory motions are always novel when participants encounter them. Further, we incentivized participants to select their true preferences by instructing them that could keep the candies they selected. This differentiates our work from previous studies using collaborative or adversarial robots [20] in which participants received virtual rewards. Finally, we avoided priming participants during the study by not asking them for explanations of their thought processes before the termination of the experiment. Asking for this information prior to the last trial would potentially prime them to be conscious of their decisions in subsequent trials. Asking for it after they view the anticipatory motion would encourage *post hoc* rationalization. Instead of using these explicit measures, we provide participants with an opportunity to provide free-form feedback at the end of the experiment, which some participants used to share their internal decision making processes with us unprompted (Sec 4.3).




Fig. 3: Top 1 and 2 intent predictor accuracy in control and no-movement, experimental trials. Observed accuracy is in blue, chance is in green. * indicates $p < 0.05$.




Fig. 4: How likely people are to accept the robot’s suggestion in *adversarial* versus *collaborative*. * indicates $p < 0.05$.

4 Results & Discussion

4.1 Intent prediction

The accuracy of the intent prediction algorithm is the number of times our algorithm predicted the user’s selected the bin out of all users. The algorithm uses the eye gaze from the beginning of the trial until the buttons light up in order to make this prediction. We report this accuracy against uniform random chance, the expected baseline value. Figure 3 reports these results for the control ($n=99$) and experimental ($n=33$, participants in the *no-movement* condition) trials. A binomial test (2-sided) showed that our algorithm performed significantly better than chance at predicting the correct target (ours: 44% versus chance: 33%, $p=0.025$), though the overall accuracy of our top 1 model shows the difficulty of forecasting user choice from historical eye gaze.

From Figure 4 we see that users shown *collaborative* motion are significantly more likely to choose the bin suggested by the robot than those who see the *adversarial* condition, as determined by a χ^2 -squared test, $\chi^2=0.242, p=0.028$, showing that our manipulation is valid.

Figure 5 shows the continuous (anytime) accuracy of the intent predictor across all participants during the ‘control’ trial. We see qualitative evidence for a choice hierarchy in that the predicted first choice bin has higher accuracy than the predicted second choice bin which has a higher accuracy than the predicted third choice bin. Additionally, we can see that people cycled between their top two choices after the 3s mark. Combining these results with those from Figures 3 and 4, we show that we can generate predictions about a user’s choice hierarchy before they have explicitly made a decision.

Our results also suggest that our anticipation model can predict a user’s least preferred choice. Figure 3 shows our model predicted top 2 accuracy significantly above chance, with a respectable accuracy (ours: 80% versus chance: 66%, $p=0.005$), indicating historical eye gaze might be a better indicator of those items a user does not show a preference toward. It is further supported by qualitative evidence in the evolution of participant intent, Figure 5. As time progresses, the top 2 accuracy rises consistently, implying that we are increasingly sure of which bin the participant does not want. The “cycling” between 1st and 2nd choices suggests people quickly (around 3s) eliminate




Fig. 5: Anytime intent prediction accuracy across all participants in the control trial, our model




Fig. 6: Anytime intent prediction accuracy across all participants in the control trial, using the SVM model (Sec. 4.1)

one option and spend their remaining time making their decision amongst the other two options. These trends also held in the experimental trial.

Our model’s ability to predict user decisions disappears in the *no-movement* group of the experimental trial, Figure 3. This indicates user’s in this condition base their decisions on external factors. One possibility is that these users believe this to be a simple handover task, and are testing to the robot’s handover capabilities, (*e.g.*, whether it is able to reach all three bins). In this scenario, the user’s choice is made through a process of elimination, thus disentangling their gaze from their choice, as suggested by the following quotes:

“MICO had an easier time reaching and holding onto the right and left positions than the center one. I purposely chose a candy in each location to test that.” • “..watched the robot go through it’s full range of motion before selecting a button. I selected the item on my left because I wanted to see the robot pick up each box.”

SVM model To check if the necessarily simplistic nature of our zero-data anticipation model was the reason for limited intent prediction accuracy, we explored a data-driven model. After collecting the data by doing this study, we trained a modified version of the SVM intent predictor from [11], which uses four features for each bin: number of glances, duration of first glance, total duration of all glances, if the bin was the most recently glanced at. We add a fifth feature, a unique identifier for each bin, and then use the same testing and training paradigm as [11]. This SVM model gives us a validation accuracy of 54% on the control trial and 45% on the no-movement condition in the experimental trial. While this is better than our heuristic based model (44% and 33% respectively), the SVM is only able to be trained after collecting eye gaze data on a large population.

We explored remapping our conditions in the experimental trial to match the SVM model’s classification by comparing the robot’s bin suggestion to a post-hoc SVM prediction. If a participant originally in the *adversarial* condition saw the robot move toward their 1st choice bin, as predicted by the heuristic model, but the SVM predicted this bin to be the user’s last choice, then we reassign this user to the *collaborative* condition.

Under this remapping, we obtained 20 Anticipatory and 26 Adversarial trials. Additionally, 20 participants moved to a new condition, *Adversarial⁺*, where the robot moved to the user’s 2nd choice as predicted by the SVM. Most participants remained in their original condition, but about a third from both *collaborative* and *adversarial* conditions remapped to the new condition. Under the new remapping, we did not see a significant change in the distribution of user choice vs model-ranked preferences (Table 2), and consequently in the statistical results throughout this paper.

Discriminating unarticulated intent using eye gaze Our results show that gaze most accurately predicts user intent close to a user’s explicit selection. Figure 5 shows anytime intent prediction accuracy through the entire trial. This accuracy increases

	Ant.	Adv.	Cont.	Adv. ⁺
Ant.	17	7	0	9
Adv.	3	19	0	11
Cont.	0	0	33	0

Table 1: Confusion matrix of remapping shifts between conditions in the experimental trial. Old labels (Rows) vs New labels (Columns)

	Original	Remapped
1 st pref	0.424	0.515
2 nd pref	0.364	0.242
3 rd pref	0.212	0.242

Table 2: Proportion of users that selected their model-ranked N^{th} choice bin in the *Anticipatory* condition in the experimental trial




Fig. 7: Anytime intent prediction accuracy across participants in the ‘experimental’ trial, during the time from buttons lighting up to when the decision was finally made.

monotonically as time progresses towards the user’s explicit decision, after the buttons light up and a decision is prompter. Additionally, following the evaluation method in [11], we see the modified SVM method behaving similarly (Fig. 6), achieving 90.4% accuracy, on average about 1.48 seconds before explicit choices are made by users (in [11], accuracy is 76%, 1.8s prior). Users in our study made their choice 3.93 seconds after the buttons lit up on average. Any gaze after the buttons lit up was not available to the predictor for anticipation actions.

Both models yielding these results indicates not enough information is present to make accurate predictions as early as 4 seconds before the users made their decision. This was hinted at by [11], but the curves in Figure 5 and Figure 6 provide more evidence that eye gaze reveals intent accurately under 2 seconds before intent is expressed. In our study, we did not want to enforce a strict time limit so as to get the most natural eye gaze and decision making behavior, so we allowed users to take as much time as they needed after the buttons lit up to make their selection.

4.2 Intent reshaping

Exploring aggregated anytime prediction accuracy can inform our understanding of the robot’s ability to influence user decisions (Figures 7a & 7b). Since user preferences are accurately predicted towards the end of the trial (as we saw in Sec. 4.1), we see how user intent becomes clearer as the trial draws to a close. (Note that in the *collaborative/adversarial* condition, the buttons light up only after the robot completes its anticipatory motion, so the gaze is not associated with tracking robot movement)

In Fig. 7a, the accuracy of the user’s (predicted) 1st choice reduces dramatically after robot motion occurs before finally recovering towards the end. This indicates robot motion induces users to reconsider their decision, suggesting a restart of the decision




Fig. 8: Comparing percentage of users selecting the suggested bin in *collaborative* and our prediction of their top choice in the *no-movement* trial.




Fig. 9: Comparing the percentage of users selecting the suggested bin in *adversarial* and our prediction of their last choice in the *no-movement* trial.

making process, which is supported by user comments (Sec. 4.3). Contrast this to the *no-movement* condition, Fig. 7b, where the estimate of user 1st choice becomes increasingly accurate as the time to decision decreases. Further, we see that user intent becomes apparent earlier, during trials in the *no-movement* condition than in the *adversarial* condition.

To further study intent reshaping we can measure the percentage of time users chose the robot’s suggested bin in the *collaborative* and *adversarial* conditions, Figures 8 and 9. We see that participants did not choose the robot’s suggested bin significantly more frequently in the *collaborative* condition ($\chi^2=1.47, p=0.612$) or the *adversarial* conditions ($\chi^2=0.663, p=0.751$).

4.3 Qualitative analysis

In a post-experiment survey, we asked users about differences between the first two and the final trials, to determine if they were aware of the anticipatory motion. Eight out of 66 participants who saw motion reported either not noticing or seeing only minor differences (e.g., “the candies were different”), with the rest identifying robot movement as the difference. We conjecture that this is due to these participants not ascribing intent to the anticipatory actions. An optional free-response section revealed further themes related to our hypotheses:

Several users reported negative feelings associated with the robot or its movement in the *adversarial* condition. No such comments were present in the *no-movement* and *collaborative* conditions:

“After I chose the container on the left, the robot arm went up and there was a pause before it reached for the container I chose. Like it was annoyed with my choice.” • “It moved a lot unnecessarily.” • “Somehow robot was waiting to pick up the left side box before I pressed the button, actually it was the most different side of my desire (I wanted right side candy)”

Several users reported positive feelings associated with the robot or its movement in the *collaborative* condition:

“The robot prepared to select the candy which I was looking at. The robot was more intelligent in the last trial.” • “As soon as I made the decision in my head, I noticed the robot move directly in front of the candy I wanted.” • “The robot seemed guess my choice well in the last trial. The more I have interfaced with the robot the more it has understood me”

Additionally, several users revealed their thought processes, which indicated that anticipatory motion can influence participant decision making in both the *collaborative* and *adversarial* conditions, but not *no-movement*:

“I felt like the robot ‘wanted’ me to choose the middle bin after moving there the third trial. I picked a different bin after the one I originally wanted due to this.” • “The robot moved and kinda guess the candy I wanted ... decision that the robot made definitely influenced my decision however I was going to choose the candy either way. I was surprised.” • “While I noticed the robot move its arm in front of the candy I wanted, I thought about picking a different candy for a split second. I did not end up doing it because I figured it would make its job easier even though it’s a robot.” • “Robot moved to be close to one box, I guess if I wanted the robot to work less I would have chosen that box but the wrapper color wasn’t appealing to me.”

From these responses, we see that the robot action can create an inflection point during the interaction causing some users to reconsider their original intention. This reconsideration results in a variety of outcomes including users changing their original decision to match the robot’s suggestion, users explicitly disavowing the robot’s choice even if it aligned with their own, as well as users acting without regard to the robot’s motion. The specific factors that contribute to these various outcomes is out of scope for our work, but should be studied in the future.

5 Conclusion and Future Work

In this paper, we explored whether anticipatory robot motion displayed during a collaborative handover interaction can influence the choices people make during that interaction. Our study found several small quantitative effects suggesting distinct types of anticipatory motions (collaborative and adversarial) can cause users to reconsider their decisions. Qualitative data show further support for this claim. Future work can consider more targeted studies focusing on one particular aspect of the effect of anticipatory robot behavior on human robot interactions such as how specific task context may amplify or diminish this effect.

We showed evidence supporting previous work that eye gaze is indeed indicative of user preferences in selection tasks. We provided detailed analysis of this correlation over the duration of such tasks, indicating that eye gaze is most effective in predicting user choice only a few seconds before it is made explicit, rather than a percentage of overall task time.

Taken together, our findings suggest that robot paradigms designed to anticipate user choice should model the effects of this anticipation during task execution or else anticipatory models may end up changing the very behaviors they are intended to anticipate. Future efforts should explore how this effect changes during complex, sequential tasks, long term interactions with a single robot in a particular task, with a change in the subtlety or timing of the robot’s motions, or the perceived expertise or authority of the robot partner. In summary, our work shows initial evidence that human-robot interactions are bidirectional, meaning that researchers need to consider how robot motions designed to anticipate user actions may in turn affect those actions.

Acknowledgements

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE1252522 and DGE1745016,

as well as NSF award IIS-1755823, the Tang Foundation, and ONR award N00014-18-1-2503. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

1. Admoni, H., Srinivasa, S.: Predicting user intent through eye gaze for shared autonomy. In: 2016 AAAI Fall Symposium Series (2016)
2. Aronson, R.M., Santini, T., Kubler, T.C., Kasneci, E., Srinivasa, S., Admoni, H.: Eye-hand behavior in human-robot shared manipulation. In: ACM/IEEE International Conference on Human-Robot Interaction. ACM, New York, NY, USA (March 2018)
3. Chen, M., Hsu, D., Lee, W.S.: Guided exploration of human intentions for human-robot interaction. In: International Workshop on the Algorithmic Foundations of Robotics. pp. 921–938. Springer (2018)
4. Dragan, A.D., Bauman, S., Forlizzi, J., Srinivasa, S.S.: Effects of robot motion on human-robot collaboration. In: ACM/IEEE International Conference on Human-Robot Interaction. pp. 51–58. ACM (2015)
5. Durdu, A., Erkmén, I., ERKMEN, A.M.: Estimating and reshaping human intention via human–robot interaction. *Turkish Journal of Electrical Engineering & Computer Sciences* **24**(1), 88–104 (2016)
6. Fisac, J.F., Bronstein, E., Stefansson, E., Sadigh, D., Sastry, S.S., Dragan, A.D.: Hierarchical game-theoretic planning for autonomous vehicles. *arXiv preprint arXiv:1810.05766* (2018)
7. Grigore, E.C., Eder, K., Pipe, A.G., Melhuish, C., Leonards, U.: Joint action understanding improves robot-to-human object handover. In: 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 4622–4629. IEEE (2013)
8. Hancock, P.A., Billings, D.R., Schaefer, K.E., Chen, J.Y., De Visser, E.J., Parasuraman, R.: A meta-analysis of factors affecting trust in human-robot interaction. *Human factors* **53**(5), 517–527 (2011)
9. Hayhoe, M., Ballard, D.: Eye movements in natural behavior. *Trends in cognitive sciences* **9**(4), 188–194 (2005)
10. Hoffman, G.: Anticipation in human-robot interaction. In: AAAI Spring Symposium Series (2010)
11. Huang, C.M., Andrist, S., Saupé, A., Mutlu, B.: Using gaze patterns to predict task intent in collaboration. *Frontiers in psychology* **6**, 1049 (2015)
12. Huang, C.M., Mutlu, B.: Anticipatory robot control for efficient human-robot collaboration. In: ACM/IEEE International Conference on Human-Robot Interaction. pp. 83–90 (2016)
13. Javdani, S., Srinivasa, S.S., Bagnell, J.A.: Shared autonomy via hindsight optimization. *Robotics science and systems: online proceedings* **2015** (2015)
14. Johansson, R.S., Westling, G., Bäckström, A., Flanagan, J.R.: Eye–hand coordination in object manipulation. *Journal of Neuroscience* **21**(17), 6917–6932 (2001)
15. Mathieu, J.E., Heffner, T.S., Goodwin, G.F., Salas, E., Cannon-Bowers, J.A.: The influence of shared mental models on team process and performance. *Journal of applied psychology* **85**(2), 273 (2000)
16. Nikolaidis, S., Hsu, D., Srinivasa, S.: Human-robot mutual adaptation in collaborative tasks: Models and experiments. *The International Journal of Robotics Research* (2017)
17. Pelz, J., Hayhoe, M., Loeber, R.: The coordination of eye, head, and hand movements in a natural task. *Experimental brain research* **139**(3), 266–277 (2001)
18. Risko, E.F., Kingstone, A.: Eyes wide shut: implied social presence, eye tracking and attention. *Attention, Perception, & Psychophysics* **73**(2), 291–296 (2011)

19. Sadigh, D., Landolfi, N., Sastry, S.S., Seshia, S.A., Dragan, A.D.: Planning for cars that coordinate with people: leveraging effects on human actions for planning and active information gathering over human internal state. *Autonomous Robots* **42**(7), 1405–1426 (Oct 2018)
20. Stolzenwald, J., Mayol-Cuevas, W.W.: Rebellion and obedience: The effects of intention prediction in cooperative handheld robots. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2019
21. Wang, W., Li, R., Chen, Y., Jia, Y.: Human intention prediction in human-robot collaborative tasks. In: Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction. pp. 279–280. ACM (2018)
22. Yi, W., Ballard, D.H.: Recognizing behavior in hand-eye coordination patterns. *International journal of HR : humanoid robotics* **6** **3**, 337–359 (2009)
23. Yu, C., Ballard, D.H.: Learning spoken words from multisensory input. 6th International Conference on Signal Processing, 2002. **2**, 998–1001 vol.2 (2002)