Investigating the Role of Metacognition for Joint Decision-Making in Human-Robot Collaboration

Raunak Bhattacharyya,[1] Clara Colombatto,[2] Stephen Fleming,[2] Ingmar Posner,[1] Nick Hawes[1]

[1] Oxford Robotics Institute, University of Oxford

[2] Department of Experimental Psychology, University College London

{raunakbh, ingmar, nickh}@robots.ox.ac.uk, {c.colombatto, stephen.fleming}@ucl.ac.uk

Abstract

Human operators are increasingly working with robots in many safety-critical applications. Operators in such domains have to make decisions under uncertainty due to incomplete information and time pressure. A common type of such decision making involves making a choice between two options. Research in joint decision making in human-human dyads has shown that two heads are better than one: joint decisions can lead to better performance compared to the best individual performing alone. Critically, joint decisions are better only when the individuals are allowed to exchange confidence estimates-suggesting a key role for metacognition in joint decision-making. In this paper, we propose a research plan to investigate the role of metacognition in joint decision-making, and pose broad questions as well as specific hypotheses pertaining to joint decisions in robot tasks.

Introduction

Human operators are increasingly working with robots in domains such as nuclear decommissioning (Nagatani et al. 2013; Budd et al. 2020; Chiou et al. 2022), inspection (Hawes et al. 2017; Chiou, Hawes, and Stolkin 2021; Budd et al. 2023), and search and rescue (Casper and Murphy 2003; Dole et al. 2015). Operators in such domains are often working under time pressure, with incomplete information and communication latencies, such that they often have to make decisions under uncertainty.

An important type of such decision making under uncertainty is choosing between options. For example, a human operator may choose between direct teleoperation or autonomous robot operation (Lee, Mehmood, and Ryu 2016), or may have to select which out of a fleet of robots to assist in the case of robot failure (Ji, Dong, and Driggs-Campbell 2022). Automated agents can also be tasked with making choices. For example, an automated agent may have to decide between operating autonomously or querying the human operator for a demonstration (Rigter, Lacerda, and Hawes 2020), or a decision making agent may select between giving control to the human operator or an automated controller (Costen et al. 2022).

While such implementations often involve decision-making by a single agent (usually the human operator), an open question is whether two decision makers instead of one may yield better decisions than either individual decision maker alone. The two decision makers could be both human (Boschetti et al. 2021; Szczurek et al. 2023), or a human with an automated decision-support system.

Joint decision-making has been studied in human-human dyads (Bahrami et al. 2010; Koriat 2012; Bahrami et al. 2012a; Bang et al. 2014). Specifically, work in joint decision-making in visual-perception tasks has revealed interesting results. Bahrami et al. found that under certain conditions, team performance was better than either individual, provided the participants were allowed to exchange confidence estimates in their decisions. This raises the question whether the exchange of confidence estimates between two decision makers can also elevate team performance beyond either individual decision maker alone, in tasks involving human operators working with robots.

The importance of confidence estimation in joint decision-making suggests a key role for metacognition (Dunlosky and Metcalfe 2008), the ability to assess one's own task performance. Accurate metacognition is thought to optimise learning strategies, for instance, in education by contributing to efficient cognitive offloading, or protecting against cognitive biases (Fleming 2021). Therefore, understanding the role of metacognition is a key goal for decision-making in human-robot collaboration.

We pose the following questions:

1. It has been shown that humans show metacognitive sensitivity when doing visual-perception tasks, i.e., their confidence estimates track task success despite the absence of feedback. Do humans show similar metacognitive sensitivity when working with robots in action-oriented tasks such as selecting between robot controllers?
2. Is human metacognitive sensitivity in visual-perception tasks a good predictor of their metacognitive sensitivity in tasks involving robots?
3. Are two heads better than one even in the robotics domain, i.e., does a pair of human decision-makers perform better than either of the single decision makers, when they share confidence estimates?
4. Does a human decision maker with an automated decision-support system perform better than either of them alone, when they share confidence estimates?

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Background

Cognition refers to the set of processes by which people understand the world (eg. reading a textbook chapter), whereas metacognition refers to judgements that we make about our own cognition (eg. how confident we are in our understanding of the chapter). Metacognitive judgements enable us to direct our cognition (eg. if we are not confident in the chapter, we decide to dedicate more time to understanding it). For example, people are often aware of their mistakes, and report levels of confidence in their choices that correlate with objective performance. While metacognition has been investigated in the context of automated agents in the context of introspective perception (Grimmett et al. 2016), and terrain assessment (Berczi, Posner, and Barfoot 2015), we focus on human metacognition.

The fidelity of metacognition is typically assessed by asking how subjective judgements - such as confidence - track objective performance. Metacognitive "sensitivity" is defined as the trial-by-trial relationship between confidence and performance. In a person with high levels of metacognitive sensitivity, trials with high confidence would be very likely to be correct trials, while trials with low confidence would be at chance levels of performance. This person's confidence would closely track their accuracy, on a trial-by-trial basis. In contrast, a person with very low levels of metacognitive sensitivity would be no more likely to be correct on trials when their confidence judgements were high than those when their confidence judgements were low (Fleming and Lau 2014).

Bahrami et al. showed that metacognition plays an important role in joint decision-making. In their experiments, participants judged which of two briefly-presented stimuli contained an oddball target. Participants worked in dyads. They first made their decision individually, then shared their decisions, and if they disagreed, they discussed the matter until they reached a joint decision. The results led to the conclusion that "for two observers of nearly equal visual sensitivity, two heads were definitely better than one provided they were given the opportunity to communicate freely." In discussing the mechanism for the two-heads-better-than-one (2HBT1) effect, the authors assumed that each individual can monitor the accuracy of their performance and can communicate their confidence accurately to the other member.

Method

We will conduct user studies with human participants recruited from the Prolific Academic platform. Each participant will complete two tasks: visual-perception decision-making, and robot-driving decision-making. Participants will be able to perform the tasks online through a browser. Both tasks will have a series of trials and on each trial, the participants will first make a decision, and then rate their confidence using a Likert scale. We have obtained ethics approval for our study from the University of Oxford's Research Ethics Committee. Next, we describe both tasks.

Perceptual Task

We will use a perceptual decision-making task together with trial-by-trial confidence ratings (Rouault et al. 2018). On each trial, participants will be first presented with a fixation cross for 1000 ms. Two black boxes filled with differing numbers of randomly positioned white dots will then be presented for 300 ms. One box will be always half-filled (313 dots out of 625 positions), while the other box will contain an increment of +1 to +70 dots compared to the standard. After 300 ms, the dots will disappear, leaving the black boxes on the screen until a keyboard button press response is made. Participants will then be asked to judge which box had the highest number of dots. The left/right position of the target box will be pseudo-randomised across all trials. The chosen box will be highlighted for 500 ms. On every trial, subjects will then be asked to report their confidence in their response on a rating scale ranging from 1 ("not confident at all") to 4 ("very confident").

The difference in dots will be determined via a calibration procedure to maintain a constant level of performance during the experiment and across participants. We will implement a two-down one-up staircase procedure with equal step-sizes for steps up and down. The step-size will be calculated in log-space, with a starting point of 4.2 (+70 dots), changing by ± 0.4 for the first 5 trials, ± 0.2 for the next 5 trials and ± 0.1 for the rest of the task.

Robot Driving Task

In this task, the goal will be to drive a Jackal robot through a doorway to a target cone within a time limit of 15 seconds. On each trial, participants will first see the world configuration, i.e., the position and orientation of the robot, and the position of the door. Participants will then be asked to provide a confidence rating (about reaching the target within 15 seconds) after which the confidence box will disappear. At this point, the timer will begin a countdown of 15 seconds and the participant will drive the robot using the keyboard to reach the cone within the time limit. Success (participant reaches the cone within the time limit) or failure (time runs out) will cause the world to reset to a new configuration, and a new trial will begin.

Similar to the perceptual task, we will add a calibration procedure to maintain a constant level of performance during the experiment and across participants. We will achieve this by adding a delay to the control such that key presses made by the participant will take effect on the robot after a delay. We will implement a two-down one-up staircase procedure with equal step-sizes for steps up and down. If the participant fails twice, the delay will reduce by 10 milliseconds and if the participant succeeds twice, delay will increase by 10 milliseconds.

Research Plan

In this section, we describe our research questions and our proposed approach towards investigating them. Our first set of research hypotheses considers the case where the human operator has to choose between two automated controllers. The second set considers the case where the human chooses between themself and a controller. The final set considers the effect of learning.

Selecting Between Two Controllers

The first set of hypotheses consider controller selection. A single operator has to choose between two available controllers that will then autonomously drive the robot in our robot driving task. One of the controllers is very maneuverable but drives slower, and the other has a faster speed but struggles with hard maneuvering. Therefore, the operator is uncertain about their performance as it varies depending on the configuration of the world. We want to investigate whether humans show metacognitive sensitivity in this problem, whether metacognitive sensitivity generalises across tasks, and whether the 2HBT1 effect holds true in this problem. Participants will go through a familiarisation phase at the beginning of the session where they will observe the behavior of the two controllers.

Hypothesis 1: Humans show metacognitive sensitivity in controller selection The participant will be first shown the world configuration based on which they will have to select one of the two controllers. Every trial would have a correct choice, i.e., one of the two controllers will be able to complete the task better than the other, and this correct answer will be known to the experimenters in advance. We hypothesise that trial-by-trial confidence ratings given by participants will track the outcome of the robot driving task. When the confidence is high, the chosen controller will succeed in reaching the goal within the time limit, and when the confidence is low, the chosen controller will fail half the time. We will gather participant confidence ratings and the resulting task outcomes and extract metrics of metacognitive sensitivity such as meta-d' (Fleming and Lau 2014).

Hypothesis 2: Humans metacognitive sensitivity generalises across tasks We hypothesize that the metacognitive sensitivity of participants in the visual-perception task will be correlated to that in the controller-selection task (Mazancieux et al. 2020). To analyse this, we will gather confidence ratings and task outcome data in the visual perception task, and analyse it for correlation against the data gathered in the controller selection task.

Hypothesis 3: Two heads are better than one for controller selection We will have two participants jointly perform the controller-selection task. Both participants will be asked to first individually choose a controller. In case the two participants' choices do not agree, the final choice will be made either by random selection or by picking that choice which was made with higher confidence. Our hypothesis is that the choice resulting from the higher confidence selection will result in better team performance as opposed to random selection. Further, the choice will also result in better performance than either of the participants' individual performance. We will compare the resulting success percentage against both the participants' individual performance.

Selecting Between Self and Controller

Our second set of hypotheses will investigate the condition when the operator has to choose between themself and an automated controller. The operator is uncertain about the performance of themself as well as the controller, depending on the world configuration. The operator will go through a familiarisation phase where they will get to drive the robot themself as well as watch the controller drive the robot. We want to investigate the same hypotheses in the previous section but applied to this setting.

Hypothesis 4: Humans show metacognitive sensitivity when selecting between self and controller We will gather participant task outcomes and confidence ratings, and extract meta-d'. The difference from hypothesis 1 is that rather than choosing between two external controllers, the participant will have to drive the robot themself to establish whether the trial is successful.

Hypothesis 5: Humans metacognitive sensitivity generalises across tasks We will gather confidence ratings and task outcome data in the visual-perception task, and analyse it for correlation against the data from the robot-driving task.

Hypothesis 6: Two heads are better than one for selecting between two human operators Two participants will be asked to jointly decide who will drive the robot. There will be an initial familiarisation phase where both participants will be able to drive the robot, as well as watch the other participant drive the robot. In the experiment phase, both participants will first give their individual choice of who should drive the robot, along with a confidence rating. If their choices are different, the choice with the higher confidence will be selected. The corresponding participant will then drive the robot. The overall success rate will be compared against both the participants' individual success rates.

Learning

The final set of hypotheses is focused on learning. The question here is whether increased experience with the task leads to improved metacognitive sensitivity? And, what is the role of decision outcome feedback on joint decision making?

Hypothesis 7: Human metacognitive sensitivity increases with increased task experience We will compare metacognitive sensitivity between the initial and the final stages across various task settings: selecting between two controllers, selecting between self and controller, and the presence/absence of trial outcome as feedback.

Hypothesis 8: Decision outcome feedback is not needed Studies have shown that team performance is independent of the presence or absence of task outcome as feedback (Bahrami et al. 2012b). We will investigate whether the same observation holds in our robot driving task. In the two participant joint choice experiments, we will provide task outcome as feedback to one group of participants. We hypothesise that the team performance will not be affected by the presence or absence of the feedback about the outcome of the trial.

Conclusion

We considered the problem of joint decision-making involving human operators working with robots, and described our proposed approach to investigate whether two heads are better than one in such tasks. Through a visual-perception and a robot-driving task, we established hypotheses about the presence and transfer of human metacognitive sensitivity in humans. The human operator could choose between autonomous controllers, or between themself and an autonomous controller. Finally, we proposed to investigate the impact of learning through increased task experience. We hypothesised that metacognitive sensitivity in human operators would increase with further exposure to the task. Overall, this work will elucidate mechanisms of decision-making in human-robot interactions.

References

Bahrami, B.; Olsen, K.; Bang, D.; Roepstorff, A.; Rees, G.; and Frith, C. 2012a. Together, slowly but surely: the role of social interaction and feedback on the build-up of benefit in collective decision-making. *Journal of Experimental Psychology: Human Perception and Performance*, 38(1): 3.

Bahrami, B.; Olsen, K.; Bang, D.; Roepstorff, A.; Rees, G.; and Frith, C. 2012b. What failure in collective decision-making tells us about metacognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1594): 1350–1365.

Bahrami, B.; Olsen, K.; Latham, P. E.; Roepstorff, A.; Rees, G.; and Frith, C. D. 2010. Optimally interacting minds. *Science*, 329(5995): 1081–1085.

Bang, D.; Fusaroli, R.; Tylén, K.; Olsen, K.; Latham, P. E.; Lau, J. Y.; Roepstorff, A.; Rees, G.; Frith, C. D.; and Bahrami, B. 2014. Does interaction matter? Testing whether a confidence heuristic can replace interaction in collective decision-making. *Consciousness and cognition*, 26: 13–23.

Berczi, L.-P.; Posner, I.; and Barfoot, T. D. 2015. Learning to assess terrain from human demonstration using an introspective gaussian-process classifier. In *ICRA*.

Boschetti, G.; Bottin, M.; Faccio, M.; and Minto, R. 2021. Multi-robot multi-operator collaborative assembly systems: a performance evaluation model. *Journal of Intelligent Manufacturing*, 32: 1455–1470.

Budd, M.; Duckworth, P.; Hawes, N.; and Lacerda, B. 2023. Bayesian reinforcement learning for single-episode missions in partially unknown environments. In *CoRL*.

Budd, M.; Lacerda, B.; Duckworth, P.; West, A.; Lennox, B.; and Hawes, N. 2020. Markov decision processes with unknown state feature values for safe exploration using gaussian processes. In *IROS*.

Casper, J.; and Murphy, R. R. 2003. Human-robot interactions during the robot-assisted urban search and rescue response at the world trade center. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 33(3): 367–385.

Chiou, M.; Epsimos, G.-T.; Nikolaou, G.; Pappas, P.; Petousakis, G.; Mühl, S.; and Stolkin, R. 2022. Robot-assisted nuclear disaster response: Report and insights from a field exercise. In *IROS*.

Chiou, M.; Hawes, N.; and Stolkin, R. 2021. Mixed-Initiative variable autonomy for remotely operated mobile robots. *ACM Transactions on Human-Robot Interaction*, 10(4): 1–34.

Costen, C.; Rigter, M.; Lacerda, B.; and Hawes, N. 2022. Shared Autonomy Systems with Stochastic Operator Models. In *IJCAI*.

Dole, L. D.; Sirkin, D. M.; Murphy, R. R.; and Nass, C. I. 2015. Robots need humans in the loop to improve the hopefulness of disaster survivors. In *RO-MAN*.

Dunlosky, J.; and Metcalfe, J. 2008. *Metacognition*. Sage Publications.

Fleming, S. M. 2021. *Know thyself: The new science of self-awareness*. Hachette UK.

Fleming, S. M.; and Lau, H. C. 2014. How to measure metacognition. *Frontiers in human neuroscience*, 8: 443.

Grimmett, H.; Triebel, R.; Paul, R.; and Posner, I. 2016. Introspective classification for robot perception. *The International Journal of Robotics Research*, 35(7): 743–762.

Hawes, N.; Burbridge, C.; Jovan, F.; Kunze, L.; Lacerda, B.; Mudrova, L.; Young, J.; Wyatt, J.; Hebesberger, D.; Kortner, T.; et al. 2017. The strands project: Long-term autonomy in everyday environments. *IEEE Robotics & Automation Magazine*, 24(3): 146–156.

Ji, T.; Dong, R.; and Driggs-Campbell, K. 2022. Traversing supervisor problem: An approximately optimal approach to multi-robot assistance. In *RSS*.

Koriat, A. 2012. When are two heads better than one and why? *Science*, 336(6079): 360–362.

Lee, K.-H.; Mehmood, U.; and Ryu, J.-H. 2016. Development of the human interactive autonomy for the shared teleoperation of mobile robots. In *IROS*.

Mazancieux, A.; Fleming, S. M.; Souchay, C.; and Moulin, C. J. 2020. Is there a G factor for metacognition? Correlations in retrospective metacognitive sensitivity across tasks. *Journal of Experimental Psychology: General*, 149(9): 1788.

Nagatani, K.; Kiribayashi, S.; Okada, Y.; Otake, K.; Yoshida, K.; Tadokoro, S.; Nishimura, T.; Yoshida, T.; Koyanagi, E.; Fukushima, M.; et al. 2013. Emergency response to the nuclear accident at the Fukushima Daiichi Nuclear Power Plants using mobile rescue robots. *Journal of Field Robotics*, 30(1): 44–63.

Rigter, M.; Lacerda, B.; and Hawes, N. 2020. A framework for learning from demonstration with minimal human effort. *IEEE Robotics and Automation Letters*, 5(2): 2023–2030.

Rouault, M.; Seow, T.; Gillan, C. M.; and Fleming, S. M. 2018. Psychiatric symptom dimensions are associated with dissociable shifts in metacognition but not task performance. *Biological psychiatry*, 84(6): 443–451.

Szczurek, K. A.; Prades, R. M.; Matheson, E.; Rodriguez-Nogueira, J.; and Di Castro, M. 2023. Multimodal multi-user mixed reality human–robot interface for remote operations in hazardous environments. *IEEE Access*, 11: 17305–17333.