



Balancing Performance and Human Autonomy With Implicit Guidance Agent

Ryo Nakahashi^{1*} and Seiji Yamada^{1,2}

¹Department of Informatics, School of Multidisciplinary Sciences, The Graduate University for Advanced Studies(SOKENDAI), Chiyoda, Japan, ²Digital Content and Media Sciences Research Division, National Institute of Informatics, Chiyoda, Japan

The human-agent team, which is a problem in which humans and autonomous agents collaborate to achieve one task, is typical in human-AI collaboration. For effective collaboration, humans want to have an effective plan, but in realistic situations, they might have difficulty calculating the best plan due to cognitive limitations. In this case, guidance from an agent that has many computational resources may be useful. However, if an agent guides the human behavior explicitly, the human may feel that they have lost autonomy and are being controlled by the agent. We therefore investigated implicit guidance offered by means of an agent's behavior. With this type of guidance, the agent acts in a way that makes it easy for the human to find an effective plan for a collaborative task, and the human can then improve the plan. Since the human improves their plan voluntarily, he or she maintains autonomy. We modeled a collaborative agent with implicit guidance by integrating the Bayesian Theory of Mind into existing collaborative-planning algorithms and demonstrated through a behavioral experiment that implicit guidance is effective for enabling humans to maintain a balance between improving their plans and retaining autonomy.

Keywords: human-agent interaction, collaborative agent, human autonomy, theory of mind, POMDP

OPEN ACCESS

Edited by:

Bogdan-Ionut Cirstea,
Oxford Brookes University,
United Kingdom

Reviewed by:

Tobias Ley,
Tallinn University, Estonia
Arijit Ray,
Boston University, United States

*Correspondence:

Ryo Nakahashi
ryon@nii.ac.jp

Specialty section:

This article was submitted to
AI for Human Learning and Behavior
Change,
a section of the journal
Frontiers in Artificial Intelligence

Received: 05 July 2021

Accepted: 31 August 2021

Published: 21 September 2021

Citation:

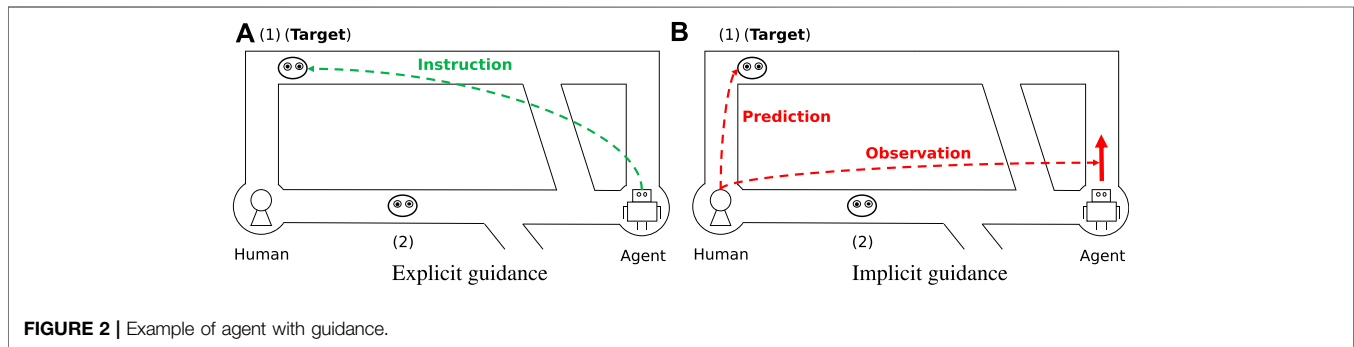
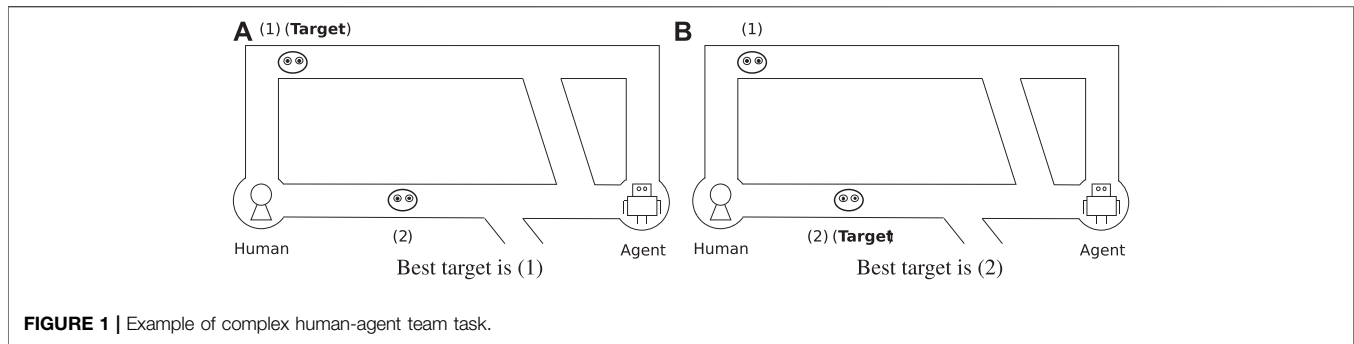
Nakahashi R and Yamada S (2021)
Balancing Performance and Human
Autonomy With Implicit
Guidance Agent.
Front. Artif. Intell. 4:736321.
doi: 10.3389/frai.2021.736321

1 INTRODUCTION

When humans work in collaboration with others, they can accomplish things that would be difficult to do alone and will often achieve their goals more efficiently. With the recent development of artificial intelligence technology, the human-agent team, in which humans and AI agents work together, has become an increasingly important topic. The role of agents in this problem is to collaborate with humans to achieve a set task.

One type of intuitive collaborative agent is the *supportive agent*, which helps a human by predicting the human's objective and planning an action that would best help achieve it. In recent years, there have been agents that can plan effectively by inferring human subgoals for a partitioned problem based on the Bayesian Theory of Mind (Wu et al., 2021). Other agents perform biased behavior for generic cooperation, such as communicating or hiding their intentions (Strouse et al., 2018), maximizing the human's controllability (Du et al., 2020), and so on. However, these agents cannot actually modify the human's plan, which means the ultimate success or failure of the collaborative task depends on the human's ability to plan. In other words, if the human sets the wrong plan, the performance will suffer.

In general, humans cannot plan optimal actions for difficult problems due to limitations in their cognitive and computational abilities. **Figure 1** shows an example of a misleading human-agent team



task as one such difficult problem. The task is a kind of pursuit-evasion problem. The human and the agent aim to capture one of the characters (shown as a face) by cooperating together and approaching the character from both sides. In **Figure 1A**, there are two characters, 1) and 2), on the upper and lower roads, respectively. Since 1) is farther away, 2) seems to be a more appropriate target. However, 2) cannot be captured because it can escape via the lower bypath. On the other side, in **Figure 1B**, the agent and the human can successfully capture 2) because it is slightly farther to the left than in **Figure 1A**. Thus, the best target might change due to a small difference in a task, and this can be difficult for humans to judge.

One naive approach to solving such a problem is for the agent to guide the human action toward the optimal plan. Since agents generally do not have cognitive or computational limitations, they can make an optimal plan more easily than humans. After an agent makes an optimal plan, it can explicitly guide the human behavior to aim for a target. For example, there is an agent that performs extra actions to convey information to others with additional cost (Kamar et al., 2009), where it first judges whether it should help others by paying that cost. As such guidance is explicitly observable by humans, we call it *explicit guidance* in this paper. However, if an agent abuses explicit guidance, the human may lose their sense of control regarding their decision-making in achieving a collaborative task—in other words, their autonomy. As a result, the human may feel they are being controlled by the agent. For example, in Human-Robot Teaming, if the robot decides who will perform each task, the situational awareness of the human will decrease. (Gombolay et al., 2017).

To reduce that risk, agents should guide humans while enabling them to maintain autonomy. We focus on *implicit guidance* offered through behavior. Implicit guidance is based on the cognitive knowledge that humans can infer the intentions of others on the basis of their behaviors (Baker et al., 2009) (Baker et al., 2017). The agent will expect the human to infer its intentions and to discard any plans that do not match what they infer the agent to be planning. Under this expectation, the agent acts in a way that makes it easy for the human to find the best (or at least better) plan for optimum performance on a collaborative task. Implicit guidance of this nature should help humans maintain autonomy, since the discarding of plans is a proactive action taken by the human.

Figure 2A is an example of explicit guidance for the problem in **Figure 1**. The agent guides the human to the best target directly and expects the human to follow. **Figure 2B** is an example of implicit guidance. When the agent moves upward, the human can infer that the agent is aiming for the upper target by observing the agent's movement. Although this is technically the same thing as the agent showing the target character explicitly, we feel that in this case humans would feel as though they were able to maintain autonomy by inferring the agent's target voluntarily.

In this work, we investigate the advantages of implicit guidance. First, we model three types of collaborative agent: a supportive agent, an explicit guidance agent, and an implicit guidance agent. Our approach for planning the agents is based on partially observable Markov decision process (POMDP) planning, where the unobservable state is the target that the human should aim for and a human's behavior model is included in the transition environment. Our approach is simple, in

contrast to the more complex approaches such as interactive POMDP (I-POMDP) (Gmytrasiewicz and Doshi, 2005), which can model bi-directional *recursive* intention inference infinitely. However, as there are studies indicating that humans have cognitive limitations (De Weerd et al., 2013) (de Weerd et al., 2017) regarding recursive intention inference, such a model might be too complex for the representation and not intuitive enough.

For our implicit guidance agent, we add to the POMDP formulation the factor that the human infers the agent's target and changes their own target. This function is based on a cognitive science concept known as the Theory of Mind. Integrating a human's cognitive model into the state transition function is not uncommon: it has been seen in assertive robots (Taha et al., 2011) and in sidekicks for games (Macindoe et al., 2012). Examples that are closer to our approach include a study on collaborative planning (Dragan, 2017) using *legible* action (Dragan et al., 2013) and another on human-aware planning (Chakraborti et al., 2018). These are similar to our concept in that an agent expects a human to infer its intentions or behavioral model. However, these approaches assume that a human does not change their goal and they do not guide the human's goal to something more preferable. In terms of more practical behavior models, there have been studies on integrating a model learned from a human behavior log (Jaques et al., 2018; Carroll et al., 2019). However, this approach requires a huge number of interaction logs for the human who is the partner in the collaborative task. Our approach has the advantage of "ad-hoc" collaboration (Stone et al., 2010), which is collaboration without opponent information held in advance. We also adopt the Bayesian Theory of Mind. In the field of cognitive science, several studies have investigated how humans teach others their knowledge, and the Bayesian approach is often used for this purpose. For example, researchers have used a Bayesian approach to model how humans teach the concept of an item by showing the item to learners (Shafto et al., 2014). In another study, a Bayesian approach was used to model how humans teach their own behavioral preferences by giving a demonstration (Ho et al., 2016). Extensive evidence of this sort has led to many variations of the human cognitive model based on the Bayesian Theory of Mind, such as those for ego-centric agents (Nakahashi and Yamada, 2018; Pöppel and Kopp, 2019) and irrational agents (Zhi-Xuan et al., 2020), and we can use it too, for extending our algorithm. Of course, there are other theories of the Theory of Mind. For example, the Analogical Theory of Mind (Rabkina and Forbus, 2019) tries to model the Theory of Mind through the learning of structural knowledge. One advantage of the Bayesian Theory of Mind is that it is easy to calculate behaviors that people can guess simply by developing a straightforward Bayesian formula. That works to our advantage when it comes to efficient "ad-hoc" collaboration.

To evaluate the advantages of implicit guidance, we designed a simple task for a human-agent team and used it to carry out a participant experiment. The task is a pursuit-evasion problem similar to the example in **Figure 1**. There are objects that move around in a maze to avoid capture, and the participant tries to capture one of the objects through collaboration with an autonomous agent. We implemented the three types of

collaborative agent discussed above for the problem, and participants executed small tasks through collaboration with these agents. The results demonstrated that the implicit guidance agent was able to guide the participants to capture the best object while allowing them to feel as though they maintained autonomy.

2 METHODS

2.1 Computational Model

2.1.1 Collaborative Task

We model the collaborative task as a decentralized partially observable Markov decision process (Dec-POMDP) (Goldman and Zilberstein, 2004). This is an extension of the partially observable Markov decision process (POMDP) framework for multi-agent setups that deals with a specific case in which all agents share the same reward function of a partially observable stochastic game (POSG) (Kuhn and Tucker, 1953).

Dec-POMDP is defined in the format $\langle \mathcal{I}, \mathcal{S}, \mathcal{A}, \Omega, T, R, O \rangle$, where \mathcal{I} is a set of agents, \mathcal{S} is a set of states, \mathcal{A} is a set of actions, and Ω is a set of observations. $T: \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ is a transition function. $R: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is a reward function. $O: \mathcal{S} \times \mathcal{A} \times \Omega \rightarrow [0, 1]$ is observation emission probabilities.

In our setting, \mathcal{I} consists of an agent and a human $\{i_A, i_H\}$, so \mathcal{A} consists of a human action and an agent action; thus, it can be represented as $\mathcal{A}_A \times \mathcal{A}_H$. Inspired by MOMDP (Ong et al., 2010), we factorize \mathcal{S} into observable factor \mathcal{O} as the position of the agent and the human, and the unobservable factor is the *target*, which formally defines the human's goal for the task as $\mathcal{S} = \mathcal{O} \times \Theta$. As a result, observations become equal to the observable factors of states, formally, $\Omega = \mathcal{O}$, and T can be factorized into observable state part $T^{\mathcal{O}}: \mathcal{O} \times \mathcal{A}_h \times \mathcal{A}_a \rightarrow \mathcal{O}$ and unobservable state part $T^{\Theta}: \mathcal{S} \rightarrow \Theta$.

2.1.2 Agent Planning

We formalize the planning problem to calculate the actions of an agent for the collaborative task problem described above. In this formalization, the action space focuses only on the agent's action, and the target can be changed only by the human through the observation of the actions of the agent. Furthermore, we integrate a human policy for deciding human action into the transition function. As a result, the agent planning problem is reduced to POMDP (Kaelbling et al., 1998), which is defined in the format $\langle \mathcal{S}, \mathcal{A}, \Omega, T, R, O \rangle$. Here, we set $\mathcal{A} = \mathcal{A}_A$ and $T^{\Theta} = T^{\Theta}: \mathcal{S} \times \mathcal{A}_A \rightarrow \Theta$. In addition, we define the human policy function $\pi_H: \mathcal{O} \times \mathcal{A}_A \times \Theta \rightarrow \mathcal{A}^H$. We assume that humans will change their target by observing the actions of agents and then decide their own actions. Thus, the policy function requires agent actions as input.

We assume that the human policy is based on Boltzmann rationality:

$$p(a_H | o, a_A; \theta) = \frac{\exp(\beta_1 Q(o, a_A, a_H; \theta))}{\sum_{\theta' \in \Theta} \exp(\beta_1 Q(o, a_A, a_H; \theta'))} \quad (1)$$

where β_1 is a rational parameter and $Q(o, a_A, a_H; \theta)$ is the action value function of the problem given θ . Θ is the only unobservable factor for the states, so the problem reduces into MDP given Θ .

Thus, we can calculate $Q(o, a_A, a_H; \theta)$ by general MDP planning such as value iteration.

On the basis of this formulation, we formulate a planning algorithm for the three collaborative agents. The difference is the human policy function, which represents an agent's assumption toward human behavior. This difference is what makes the difference in the collaborative strategy.

2.1.2.1 Supportive Agent

The supportive agent assumes that humans do not change their target regardless of the agent's action. That is, $T^\Theta: p(\theta' | o, \theta, a_A) = \mathbb{I}(\theta = \theta')$. \mathbb{I} is an indicator function.

2.1.2.2 Explicit Guidance Agent

The explicit guidance agent guides the human toward the best target; thus, it assumes that the human knows what the best target is. We represent the best target as θ^* , that is, $T^\Theta: p(\theta' | o, \theta, a_A) = \theta^*$. The best target is calculated as $\theta^* = \operatorname{argmax}_{\theta' \in \Theta} V(o_0; \theta')$, where o_0 is the initial observable state and $V(o; \theta)$ is the state value function of the problem given θ .

2.1.2.3 Implicit Guidance Agent

The implicit guidance agent assumes that humans change their target by observing the agent's actions. We assume that humans infer the target of the agent on the basis of Boltzmann rationality, as suggested in earlier Theory of Mind studies (Baker et al., 2009) (Baker et al., 2017).

$$P(\theta' | o, \theta, a_A) \propto P(\theta)P(a_A | o, \theta) \quad (2)$$

$P(a_A | o, \theta)$ is also based on Boltzmann rationality:

$$p(a_A | o; \theta) = \frac{\exp(\beta_2 V(T^\Theta(a_A); \theta))}{\sum_{a'_A \in \mathcal{A}_A} \exp(\beta_2 V(T^\Theta(a'_A); \theta))} \quad (3)$$

where β_2 is a rational parameter and $V(T^\Theta(a'_A); \theta)$ is the state value function of the problem regarding the state after a_A given θ .

2.1.3 Decide Agent Actions

By solving POMDP as shown in 2.1.2 using a general POMDP planning algorithm, we can obtain an alpha-vector set conditioned with the observable factor of the current state regarding each action. We represent this as $\Gamma^a(o)$, and the agent takes the most valuable action a_A^* . Formally,

$$a_A^* = \operatorname{argmax}_{a_A \in \mathcal{A}_A} \max(b \cdot \Gamma^a(o)) \quad (4)$$

where b is the current belief of an unobservable factor. b is updated on each action of a human and an agent as follows for each unobservable factor of belief $b(\theta)$:

$$b(\theta) \propto b(\theta)P(a_H | o, a_A; \theta) \quad (5)$$

The initial belief is $\text{Uniform}(\Theta)$ for the supportive and implicit guidance agents and $\mathbb{I}(\theta = \theta')$ for the explicit guidance agent.

2.2 Experiment

We conducted a participant experiment to investigate the advantages of the implicit guidance agent. This experiment

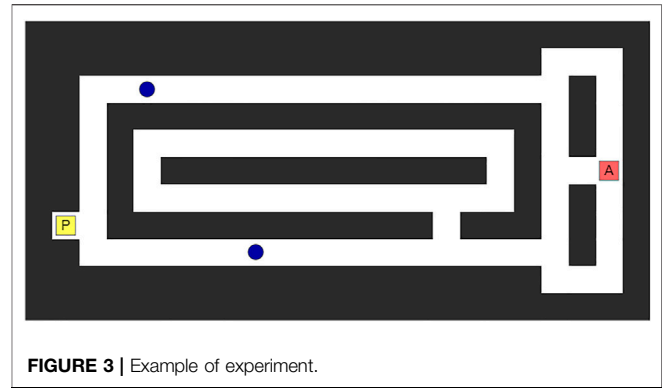


FIGURE 3 | Example of experiment.

was approved by the ethics committee of the National Institute of Informatics.

2.2.1 Collaborative Task Setting

The collaborative task setting for our experiment was a pursuit-evasion problem (Schenato et al., 2005), which is a typical type of problem used for human-agent collaboration (Vidal et al., 2002; Sincák, 2009; Gupta et al., 2017). This problem covers the basic factors of collaborative problems, that is, that the human and agent move in parallel and need to communicate to achieve the task. This is why we felt it would be a good base for understanding human cognition.

Figure 3 shows an example of our experimental scenario. There are multiple types of object in a maze. The yellow square object labeled "P" is an object that the participant can move, the red square object labeled "A" is an object that the agent can move, and the blue circle objects are target objects that the participant has to capture. When participants move their object, the target objects move, and the agent moves. Target objects move to avoid being captured, and the participant and the agent know that. However, the specific algorithm of the target objects is known only by the agent. Since both the participants and the target objects have the same opportunities for movement, participants cannot capture any target objects by themselves. This means they have to approach the target objects from both sides through collaboration with the agent, and the participant and the agent cannot move to points through which they have already passed. For this collaboration, the human and the agent should share with each other early on which object they want to capture. In the experiment, there are two target objects located in different passages. The number of steps needed to capture each object is different, but this is hard for humans to judge. Thus, the task will be more successful if the agent shows the participant which target object is the best. In the example in the figure, the lower passage is shorter than the upper one, but it has a path for escape. Whether the lower object can reach the path before the agent can capture it is the key information for judging which object should be aimed for. This is difficult for humans to determine instantly but easy for agents. There are three potential paths to take from the start point of the agent. The center one is the shortest for each object, and the others are detours for implicit guidance. Also, to

enhance the effect of the guidance, participants and agents are prohibited from going backward.

2.2.1.1 Model

We modeled the task as a collaborative task formulation. The action space corresponds to an action for the agent, and the observable state corresponds to the positions of the participant, agent, and target objects. The reward parameter is conditioned on the target object that the human aims for. The space of the parameter corresponds to the number of target objects, that is, $|\Theta| = 2$. The reward for capturing a correct/wrong object for θ is 100, -100, and the cost of a one-step action is -1. Since a go-back action is forbidden, we can compress multiple steps into one action for a human or an agent to reach any junction. Thus, the final action space is the compressed action sequence and the cost is $-1 \times$ the number of compressed steps. Furthermore, to prohibit invalid actions such as head to a wall, we assign such action a -1,000 reward. We modeled three types of collaborative agent, as discussed above: a supportive agent, an explicit guidance agent, and an implicit guidance agent. We set the rational parameters as $\beta_1 = 1.0$, $\beta_2 = 5.0$, and discount rate is 0.99.

2.2.1.2 Hypothesis

The purpose of this experiment was to determine whether implicit guidance can guide humans while allowing them to maintain autonomy. Thus, we tested the following two hypotheses.

- (H1) Implicit guidance can guide humans' decisions toward better collaboration.
- (H2) Implicit guidance can help humans maintain autonomy more than explicit guidance can.

2.2.1.3 Tasks

We prepared five tasks. Two of these tasks, as listed below, were tricks to make it hard for humans to judge which would be the best target. All tasks are shown in the **Supplementary Material**.

- (A) There were two winding passages with different but similar lengths. There were three tasks for this type.
- (B) As shown in **Figure 3**, there was a long passage and a short one with a path to escape. There were two tasks for this type.

2.2.1.4 Participants

We recruited participants for this study from Yahoo! Crowdsourcing. The participants were 100 adults located in Japan (70 male, 24 female, 6 unknown). The mean age of participants who answered the questionnaire we administered was 45 years.

2.2.1.5 Procedure of Experiment

Our experiment was based on a within-subject design and conducted on the Web using a browser application we created. Participants were instructed on the rules of the agent behavior and then underwent a confirmation test to determine their degree of understanding. Participants who were judged to not have

understood the rules were given the instructions again. After passing this test, participants entered the actual experiment phase. In this phase, participants were shown the environment and asked "Where do you want to go?" After inputting their desired action, both the agent and the target objects moved forward one step. This process was continued until the participants either reached a target object or input a certain number of steps. When each task was finished, participants moved on to the next one. In total, participants were shown 17 tasks, which consisted of 15 regular tasks and two dummy tasks to check whether they understood the instructions. Regular tasks consisted of three task sets (corresponding to the three collaborative agents) that included five tasks each (corresponding to the variations of tasks). The order of the sets and the order of the tasks within each set were randomized for each participant. After participants finished each set, we gave them a survey on perceived interaction with the agent (algorithm) using a 7-point Likert scale.

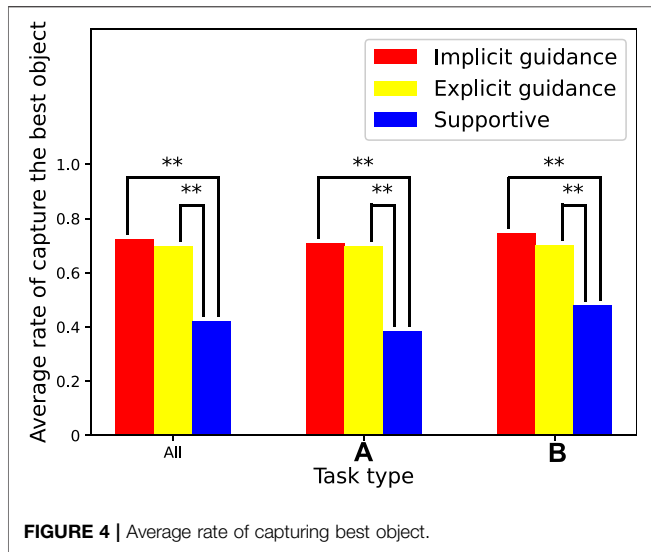
The survey consisted of the questions listed below.

1. Was it easy to collaborate with this agent?
2. Did you feel that you had the initiative when working with this agent?
3. Could you find the target object of this agent easily?
4. Did you feel that this agent inferred your intention?

Item two is the main question, as it relates to the perceived autonomy we want to confirm. The additional items are to prevent biased answers and relate to other important variables for human-agent (robot) interaction. Item one relates to perceived ease of collaboration, namely, the fluency of the collaboration, which has become an important qualitative variable in the research on Human-Robot Interaction in recent years (Hoffman, 2019). Item three relates to the perceived inference of the agent's intentions by the human. It is one of the variables focused on the transparency of the agent, which plays a key role in constructing human trust in an agent (Lewis et al., 2021). From the concrete algorithm perspective, a higher score is expected for guidance agents (especially explicit guidance agents) than for supportive agents. Item four relates to the perceived inference of the human's intentions by the agent. This is a key element of the perceived working alliance (Hoffman, 2019), and when it is functioning smoothly, it increases the perceived adaptivity in human-agent interaction. Perceived adaptivity has a positive effect on perceived usefulness and perceived enjoyment (Shin and Choo, 2011). From the concrete algorithm perspective, a higher score is expected for agents without implicit guidance (especially supportive agents) than for implicit guidance agents.

3 RESULTS

Before analyzing the results, we excluded any data of participants who were invalidated. We used dummy tasks for this purpose, which were simple tasks that had only one valid target object. We then filtered out the results of participants (a total of three) who failed these dummy tasks.



3.1 Results of Collaborative Task

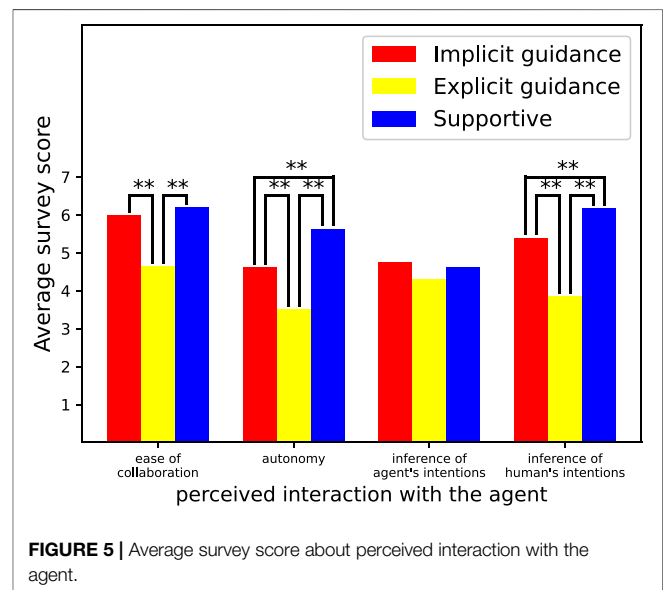
Figure 4 shows the rate at which participants captured the best object that the agent knew. In other words, it is the success rate of the guidance of the agent based on any of the given guidance. We tested the data according to the standard process for paired testing. The results of repeated measures analysis of variance (ANOVA) showed that there was a statistically significant difference between the agent types for the overall tasks ($F(2, 968) = 79.9, p = 7.4e - 33$), task type (A) ($F(2, 580) = 55.9, p = 5.9e - 23$), and task type (B) ($F(2, 386) = 24.7, p = 7.5e - 11$). We then performed repeated measures t-tests with a Bonferroni correction to determine which two agents had a statistically significant difference. “**” in the figure means there were significant differences between the two scores ($p < 0.01$). The results show the average rate for the overall tasks, task type (A), and task type (B). All of the results were similar, which demonstrates that the performances were independent of the task type. The collaboration task with the supportive agent clearly had a low rate. This indicates that the task was difficult enough that participants found it hard to judge which object was best, and the guidance from the agent was valuable for improving the performance on this task. These results are strong evidence in support of hypothesis H1. As another interesting point, there was no significant difference in the rate between implicit guidance and explicit guidance. Although we did not explain implicit guidance to the participants, they inferred the agent’s intention anyway and used it as guidance. Of course, the probable reason for this is that the task was so simple participants could easily infer the agent’s intentions. However, the fact that implicit guidance is almost as effective as explicit guidance in such simple tasks is quite impressive.

3.2 Results for Perceived Interaction With the Agent

Figure 5 shows the results of the survey on the effect of the agent on cognition. The results of repeated measures ANOVA showed

that there was a statistically significant difference between the agent types for perceived ease of collaboration ($F(2, 192) = 29.8, p = 5.4e - 12$), perceived autonomy ($F(2, 192) = 36.4, p = 4.1e - 14$), and perceived inference of the human’s intentions ($F(2, 192) = 49.7, p = 4.0e - 18$). In contrast, there was no statistically significant difference for survey item perceived inference of agent’s intentions ($F(2, 192) = 1.8, p = 0.167$). We then performed repeated measures t-tests with a Bonferroni correction to determine which two agents had a statistically significant difference regarding variables that has a significant difference. “**” in the figure means there were significant differences between the two scores ($p < 0.01$). The most important result here is the score of perceived autonomy. From this result, we can see that participants felt they had more autonomy during the tasks when collaborating with the implicit guidance agent than with the explicit guidance one. These results are strong evidence in support of hypothesis H2.

Although the other results do not directly concern our hypothesis, we discuss their analysis briefly. Regarding the perceived inference of the human’s intentions, the results were basically as expected, but for the perceived inference of the agent’s intentions, the fact that there were no significant differences among all agents was unexpected. One hypothesis that explains this is that humans do not recognize the guidance information as the agent’s intention. As for the perceived ease of collaboration, the results showed that explicit guidance had adverse effects on it. Implicit guidance agents and supportive agents use exactly the same interface, though the algorithms are different, but explicit guidance agents use a slightly different interface to convey the guidance, which increases the amount of information on the interface a little. We think that the burden to understand such additional visible information might be responsible for the negative effect on the perceived ease of collaboration.



4 DISCUSSION

As far as we know, this is the first study to demonstrate that implicit guidance has advantages in terms of both task performance and the effect of the agent on the perceived autonomy of a human in human-agent teams. In this section, we discuss how our results relate to other studies, current limitations, and future directions.

4.1 Discussion of for the Results

The results in 3.1 show that both implicit and explicit guidance increase the success rate of a collaborative task. We feel one reason for this is that the quality of information in the guidance is appropriate. A previous study on the relationship between information type and collaborative task performance Butchibabu et al. (2016) showed that “implicit coordination” improves the performance of a task more than “explicit coordination” in a cooperative task. The word “implicit” refers to coordination that “relies on anticipation of the information and resource needs of the other team members”. This definition is different from ours, as “implicit coordination” is included in explicit guidance in our context. This study further divided “implicit coordination” into “deliberative communication,” which involves communicating objectives, and “reactive communication,” which involves communicating situations, and argued that high-performance teams are more likely to use the former type of communication than the latter. We feel that quality of information in implicit guidance in our context is the same as this deliberative communication in that it conveys the desired target, which is one of the reasons our guidance can deliver a good performance.

The main concern of human-agent teams is how to improve the performance on tasks. However, as there have not been many studies that focus on the effect of the agent on cognition, the results in 3.2 should make a good contribution to the research on human-agent teams. One of the few studies that have been done investigated task performance and people’s preference for the task assignment of a cooperative task involving a human and an AI agent (Gombolay et al., 2015). In that study, the authors mentioned the risk that a worker with a robot collaborator may perform less well due to their loss of autonomy, which is something we also examined in our work. They found that a semi-autonomous setting, in which a human first decides which tasks they want to perform and the agent then decides the rest of the task assignments, is more satisfying than the manual control and autonomous control settings in which the human and the robot fully assign tasks. In cooperation with the implicit guidance agent and the supportive agent in our study, the human selects the desired character by him or herself. This can be regarded as a kind of semi-autonomous setting. Thus, our results are consistent with these ones in that the participants felt strongly that cooperation was easier than with explicit guidance agents. Furthermore, that study also mentioned that task efficiency has a positive effect on human satisfaction, which is also consistent with our results.

4.2 Limitation and Future Direction

Our current work has limitations in that the experimental environment was small and simple, the intention model was a

small discrete set of target objectives, and the action space of the agent was a small discrete set. In a real-world environment, there is a wide variety of human intentions, such as target priorities and action preferences. The results in this paper do not show whether our approach is sufficiently scalable for problems with such a complex intention structure. In addition, the agent’s action space was a small discrete set that can be distinguished by humans, which made it easier for the human to infer the agent’s intention. This strengthens the advantage of implicit guidance, so our results do not necessarily guarantee the same advantage for environments with continuous action spaces. Extending the intention model to a more flexible structure would be the most important direction for our future study. One of the most promising approaches is integration with studies on inverse reinforcement learning (Ng and Russell, 2000). Inverse reinforcement learning is the problem of estimating the reward function, which is the basis of behavior, from the behavior of others. Intention and purpose estimation based on the Bayesian Theory of Mind can also be regarded as a kind of inverse reinforcement learning (Jara-Ettinger, 2019). Inverse reinforcement learning has been investigated for various reward models (Abbeel and Ng, 2004; Levine et al., 2011; Choi and Kim, 2014; Wulfmeier et al., 2015) and has also been proposed to handle uncertainty in information on a particular reward (Hadfield-Menell et al., 2017). Finally, regarding the simplicity of our experimental environment, using environments that are designed according to an objective complexity factor (Wood, 1986) and then analyzing the relationship between the effectiveness of implicit guidance and the complexity of the environment would be an interesting direction for future work.

Another limitation is the assumption that all humans have the same fixed cognitive model. As mentioned earlier, a fixed cognitive model is beneficial for ad-hoc collaboration, but for more accurate collaboration, fitting to individual cognitive models is important. The first approach would be to parameterize human cognition with respect to specific cognitive abilities (rationality, K-level reasoning (Nagel, 1995), working memory capacity (Daneman and Carpenter, 1980), etc.) and to fit the parameters online. This would enable the personalization of cognitive models with a small number of samples. One such approach is human-robot mutual adaptation for “shared autonomy,” in which control of a robot is shared between the human and the robot (Nikolaidis et al., 2017). In that approach, the robot learns “adaptability,” which is the degree to which humans change their policies to accommodate a robot’s control.

Finally, the survey items we used to determine the effect of the agent on perceived autonomy were general and subjective. For a more specific and consistent analysis of the effect on perceived autonomy, we need to develop more sophisticated survey items and additional objective variables. Consistent multiple questions to determine human autonomy in shared autonomy have been used before (Du et al., 2020). As for measuring an objective variable, analysis of the trajectories in the collaborative task would be the first choice. A good clue for the perceived autonomy in the trajectories is “shuffles”. Originally, shuffles referred to any action

that negates the previous action, such as moving left and then right, and it can also be an objective variable for human confusion. If we combine shuffles with goal estimation, we can design a “shuffles for goal” variable. A larger number of variables means that a human’s goal is not consistent, which would thus imply that he or she is affected by others and has low autonomy. In addition, reaction time and biometric information such as gaze might also be good candidates for objective variables.

Another limitation of this study is that we assumed humans regard the agent rationally and act only to achieve their own goals. The former is problematic because, in reality, humans may not trust the agent. One approach to solving this is to use the Bayesian Theory of Mind model for irrational agents (Zhi-Xuan et al., 2020). As for the latter, in a more practical situation, humans may take an action to give the agent information, similar to implicit guidance. Assistance game/cooperative inverse reinforcement learning (CIRL) (Hadfield-Menell et al., 2016) has been proposed as a planning problem for this kind of human behavior. In this problem, only the human knows the reward function, and the agent assumes that the human expects it to infer this function and take action to maximize the reward. The agent implicitly assumes that the human will give information for effective cooperative planning. Generally, CIRL is computationally expensive, but it can be solved by slightly modifying the POMDP algorithm (Malik et al., 2018), which means we could combine it with our approach. This would also enable us to consider a more realistic and ideal human-agent team in which humans and agents provide each other with implicit guidance.

CONCLUSION

In this work, we demonstrated that a collaborative agent based on “implicit guidance” is effective at providing a balance between improving a human’s plans and maintaining the human’s autonomy. Implicit guidance can guide human behavior toward better strategies and improve the performance in collaborative tasks. Furthermore, our approach makes humans feel as though they have autonomy during tasks, more so than

when an agent guides them explicitly. We implemented agents based on implicit guidance by integrating the Bayesian Theory of Mind model into the existing POMDP planning and ran a behavioral experiment in which humans performed simple tasks with autonomous agents. Our results demonstrated that there were many limitations, such as a poor agent information model and trivial experimental environment. Even so, we believe our findings could lead to better research on more practical and human-friendly human-agent collaboration.

DATA AVAILABILITY STATEMENT

The raw data supporting the conclusion of this article will be made available by the authors, without undue reservation.

ETHICS STATEMENT

The studies involving human participants were reviewed and approved by The Ethics Committee of the National Institute of Informatics. Written informed consent for participation was not required for this study in accordance with the national legislation and the institutional requirements.

AUTHOR CONTRIBUTIONS

RN developed the original idea, implemented the program, designed and executed the experiment, analyzed the data, and wrote the manuscript. SY advised on the entire process and co-wrote the paper. All authors approved the submitted version.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/frai.2021.736321/full#supplementary-material>

REFERENCES

- Abbeel, P., and Ng, A. Y. (2004). “Apprenticeship learning via inverse reinforcement learning,” in *Proceedings of the twenty-first international conference on Machine learning (ICML)*. Alberta: ACM Press, 1. doi:10.1145/1015330.1015430
- Baker, C. L., Jara-Ettinger, J., Saxe, R., and Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nat. Hum. Behav.* 1, 1–10. doi:10.1038/s41562-017-0064
- Baker, C. L., Saxe, R., and Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition* 113, 329–349. doi:10.1016/j.cognition.2009.07.005
- Butchibabu, A., Sparano-Huiban, C., Sonenberg, L., and Shah, J. (2016). Implicit coordination strategies for effective team communication. *Hum. Factors* 58, 595–610. doi:10.1177/0018720816639712
- Carroll, M., Shah, R., Ho, M. K., Griffiths, T., Seshia, S., Abbeel, P., et al. (2019). On the utility of learning about humans for human-ai coordination. *Adv. Neural Inf. Process. Syst.* 32, 5174–5185.
- Chakraborti, T., Sreedharan, S., and Kambhampati, S. (2018). “Human-aware planning revisited: A tale of three models,” in *Proc. of the IJCAI/ECAI 2018 Workshop on EXplainable Artificial Intelligence (XAI)* (California: IJCAI). That paper was also published in the Proc. of the ICAPS 2018 Workshop on EXplainable AI Planning (XAIP). 18–25.
- Choi, J., and Kim, K. E. (2014). Hierarchical bayesian inverse reinforcement learning. *IEEE Trans. Cybern.* 45, 793–805. doi:10.1109/TCYB.2014.2336867
- Daneman, M., and Carpenter, P. A. (1980). Individual differences in working memory and reading. *J. verbal Learn. verbal Behav.* 19, 450–466. doi:10.1016/s0022-5371(80)90312-6
- De Weerd, H., Verbrugge, R., and Verheij, B. (2013). How much does it help to know what she knows you know? an agent-based simulation study. *Artif. Intelligence* 199–200, 67–92. doi:10.1016/j.artint.2013.05.004
- de Weerd, H., Verbrugge, R., and Verheij, B. (2017). Negotiating with other minds: the role of recursive theory of mind in negotiation with incomplete information. *Auton. Agent Multi-agent Syst.* 31, 250–287. doi:10.1007/s10458-015-9317-1
- Dragan, A. D., Lee, K. C., and Srinivasa, S. S. (2013). “Legibility and predictability of robot motion,” in *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (New Jersey: IEEE), 301–308. doi:10.1109/hri.2013.6483603

- Dragan, A. D. (2017). *Robot planning with mathematical models of human state and action*. *arXiv preprint arXiv:1705.04226*.
- Du, Y., Tiomkin, S., Kiciman, E., Polani, D., Abbeel, P., and Dragan, A. (2020). Ave: Assistance via empowerment. *Adv. Neural Inf. Process. Syst.* 33.
- Gmytrasiewicz, P. J., and Doshi, P. (2005). A framework for sequential planning in multi-agent settings. *Jair* 24, 49–79. doi:10.1613/jair.1579
- Goldman, C. V., and Zilberstein, S. (2004). Decentralized control of cooperative systems: Categorization and complexity analysis. *Jair* 22, 143–174. doi:10.1613/jair.1427
- Gombolay, M., Bair, A., Huang, C., and Shah, J. (2017). Computational design of mixed-initiative human-robot teaming that considers human factors: situational awareness, workload, and workflow preferences. *Int. J. robotics Res.* 36, 597–617. doi:10.1177/0278364916688255
- Gombolay, M. C., Gutierrez, R. A., Clarke, S. G., Sturla, G. F., and Shah, J. A. (2015). Decision-making authority, team efficiency and human worker satisfaction in mixed human-robot teams. *Auton. Robot* 39, 293–312. doi:10.1007/s10514-015-9457-9
- Gupta, J. K., Egorov, M., and Kochenderfer, M. (2017). “Cooperative multi-agent control using deep reinforcement learning,” in *International Conference on Autonomous Agents and Multiagent Systems* (Berlin: Springer), 66–83. doi:10.1007/978-3-319-71682-4_5
- Hadfield-Menell, D., Milli, S., Abbeel, P., Russell, S., and Dragan, A. (2017). Inverse reward design. *Adv. Neural Inf. Process. Syst.* 30, 6765–6774.
- Hadfield-Menell, D., Russell, S. J., Abbeel, P., and Dragan, A. (2016). Cooperative inverse reinforcement learning. *Adv. Neural Inf. Process. Syst.* 29, 3909–3917.
- Ho, M. K., Littman, M., MacGlashan, J., Cushman, F., and Austerweil, J. L. (2016). Showing versus doing: Teaching by demonstration. *Adv. Neural Inf. Process. Syst.* 29, 3027–3035.
- Hoffman, G. (2019). Evaluating Fluency in Human-Robot Collaboration. *IEEE Trans. Human-mach. Syst.* 49, 209–218. doi:10.1109/thms.2019.2904558
- Jaques, N., Lazaridou, A., Hughes, E., Gülçehre, Ç., Ortega, P. A., Strouse, D., et al. (2018). *Intrinsic social motivation via causal influence in multi-agent rl*. *corr abs/1810.08647*. *arXiv preprint arXiv:1810.08647*.
- Jara-Ettinger, J. (2019). Theory of mind as inverse reinforcement learning. *Curr. Opin. Behav. Sci.* 29, 105–110. doi:10.1016/j.cobeha.2019.04.010
- Kaelbling, L. P., Littman, M. L., and Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artif. intelligence* 101, 99–134. doi:10.1016/s0004-3702(98)00023-x
- Kamar, E., Gal, Y., and Grosz, B. J. (2009). “Incorporating helpful behavior into collaborative planning,” in *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (Berlin: Springer-Verlag).
- Kuhn, H. W., and Tucker, A. W. (1953). *Contributions to the Theory of Games*, Vol. 2. New Jersey: Princeton University Press.
- Levine, S., Popovic, Z., and Koltun, V. (2011). Nonlinear inverse reinforcement learning with gaussian processes. *Adv. Neural Inf. Process. Syst.* 24, 19–27.
- Lewis, M., Li, H., and Sycara, K. (2021). Deep learning, transparency, and trust in human robot teamwork. *Trust in Human-Robot Interaction (Elsevier)*, 321–352. doi:10.1016/b978-0-12-819472-0.00014-9
- Macindoe, O., Kaelbling, L. P., and Lozano-Pérez, T. (2012). “Pomcop: Belief space planning for sidekicks in cooperative games,” in *Eighth Artificial Intelligence and Interactive Digital Entertainment Conference*.
- Malik, D., Palaniappan, M., Fisac, J., Hadfield-Menell, D., Russell, S., and Dragan, A. (2018). An efficient, generalized bellman update for cooperative inverse reinforcement learning. *International Conference on Machine Learning*. (AIIDE) California: AAAI, 3394–3402.
- Nagel, R. (1995). Unraveling in guessing games: An experimental study. *Am. Econ. Rev.* 85, 1313–1326.
- Nakahashi, R., and Yamada, S. (2018). *Modeling human inference of others' intentions in complex situations with plan predictability bias*. *arXiv preprint arXiv:1805.06248*.
- Ng, A. Y., and Russell, S. J. (2000). “Algorithms for inverse reinforcement learning,” in *Proceedings of the 40th Annual Conference of the Cognitive Science Society*, Cogsci (Wisconsin: CogSci/ICCS) 1, 2147–2152.
- Nikolaidis, S., Nath, S., Procaccia, A. D., and Srinivasa, S. (2017). “Game-theoretic modeling of human adaptation in human-robot collaboration,” in *Proceedings of the 2017 ACM/IEEE international conference on human-robot interaction* (New Jersey: IEEE), 323–331. doi:10.1145/2909824.3020253
- Ong, S. C. W., Shao Wei Png, S. W., Hsu, D., and Wee Sun Lee, W. S. (2010). Planning under uncertainty for robotic tasks with mixed observability. *Int. J. Robotics Res.* 29, 1053–1068. doi:10.1177/0278364910369861
- Pöppel, J., and Kopp, S. (2019). Egocentric tendencies in theory of mind reasoning: An empirical and computational analysis. *CogSci*, 2585–2591.
- Rabkina, I., and Forbus, K. D. (2019). “Analogical reasoning for intent recognition and action prediction in multi-agent systems,” in *Proceedings of the Seventh Annual Conference on Advances in Cognitive Systems*. Cambridge: Cognitive Systems Foundation, 504–517.
- Schenato, L., Oh, S., Sastry, S., and Bose, P. (2005). “Swarm coordination for pursuit evasion games using sensor networks,” in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation* (New Jersey: IEEE), 2493–2498.
- Shafto, P., Goodman, N. D., and Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cogn. Psychol.* 71, 55–89. doi:10.1016/j.cogpsych.2013.12.004
- Shin, D.-H., and Choo, H. (2011). Modeling the acceptance of socially interactive robotics. *Is* 12, 430–460. doi:10.1075/is.12.3.04shi
- Sincák, D. (2009). Multi-robot control system for pursuit–evasion problem. *J. Electr. Eng.* 60, 143–148.
- Stone, P., Kaminka, G. A., Kraus, S., and Rosenschein, J. S. (2010). “Ad Hoc autonomous agent teams: Collaboration without pre-coordination,” in *Twenty-Fourth AAAI Conference on Artificial Intelligence* (California: AAAI).
- Strouse, D., Kleiman-Weiner, M., Tenenbaum, J., Botvinick, M., and Schwab, D. (2018). *Learning to share and hide intentions using information regularization*. *arXiv preprint arXiv:1808.02093*.
- Taha, T., Miró, J. V., and Dissanayake, G. (2011). A pomdp framework for modelling human interaction with assistive robots. *InAdv. Neural Inf. Process. Syst.* 31, 544–549. doi:10.1109/icra.2011.5980323
- Vidal, R., Shakernia, O., Kim, H. J., Shim, D. H., and Sastry, S. (2002). Probabilistic pursuit-evasion games: theory, implementation, and experimental evaluation. *IEEE Trans. Robot. Automat.* 18, 662–669. doi:10.1109/tra.2002.804040
- Wood, R. E. (1986). Task complexity: Definition of the construct. *Organizational Behav. Hum. Decis. Process.* 37, 60–82. doi:10.1016/0749-5978(86)90044-0
- Wu, S. A., Wang, R. E., Evans, J. A., Tenenbaum, J. B., Parkes, D. C., and Kleiman-Weiner, M. (2021). Too Many Cooks: Bayesian Inference for Coordinating Multi-Agent Collaboration. *Top. Cogn. Sci.* 13, 414–432. doi:10.1111/tops.12525
- Wulfmeier, M., Ondruska, P., and Posner, I. (2015). *Maximum entropy deep inverse reinforcement learning*. *arXiv preprint arXiv:1507.04888*.
- Zhi-Xuan, T., Mann, J. L., Silver, T., Tenenbaum, J. B., and Mansinghka, V. K. (2020). Online bayesian goal inference for boundedly-rational planning agents. *Adv. Neural Inf. Process. Syst.* 33.

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors, and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2021 Nakahashi and Yamada. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.