

Optimal Action-based or User Prediction-based Haptic Guidance: Can You Do Even Better?

Hee-Seung Moon
Yonsei University
Incheon, Korea
hs.moon@yonsei.ac.kr

Jiwon Seo*
Yonsei University
Incheon, Korea
jiwon.seo@yonsei.ac.kr

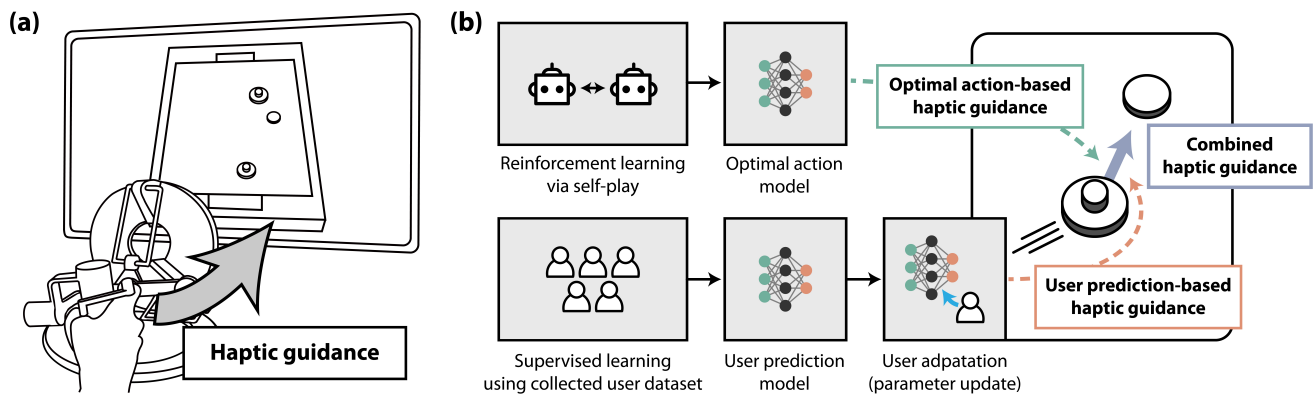


Figure 1: Overview of this research. (a) We built a virtual air hockey task environment controlled by a haptic device. Haptic guidance that assists a user’s task performance is delivered through this device. **(b)** We implemented optimal action-based and user prediction-based haptic guidance using deep learning-based approaches, proposed a combined haptic guidance for better performance, and experimentally compared them.

ABSTRACT

The recently advanced robotics technology enables robots to assist users in their daily lives. Haptic guidance (HG) improves users’ task performance through physical interaction between robots and users. It can be classified into optimal action-based HG (OAHG), which assists users with an optimal action, and user prediction-based HG (UPHG), which assists users with their next predicted action. This study aims to understand the difference between OAHG and UPHG and propose a combined HG (CombHG) that achieves optimal performance by complementing each HG type, which has important implications for HG design. We propose implementation methods for each HG type using deep learning-based approaches. A user study ($n=20$) in a haptic task environment indicated that UPHG induces better subjective evaluations, such as naturalness and comfort, than OAHG. In addition, the CombHG that we proposed further decreases the disagreement between the user intention and HG, without reducing the objective and subjective scores.

*Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '21, May 8–13, 2021, Yokohama, Japan

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8096-6/21/05...\$15.00

<https://doi.org/10.1145/3411764.3445115>

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI)**; *Haptic devices*; *User models*; • **Computing methodologies** → *Machine learning*.

KEYWORDS

Haptic guidance, Physical human–robot interaction (pHRI), User adaptation, Deep learning

ACM Reference Format:

Hee-Seung Moon and Jiwon Seo. 2021. Optimal Action-based or User Prediction-based Haptic Guidance: Can You Do Even Better?. In *CHI Conference on Human Factors in Computing Systems (CHI '21)*, May 8–13, 2021, Yokohama, Japan. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3411764.3445115>

1 INTRODUCTION

With the recent advancements in artificial intelligence and robotics technology, it is becoming increasingly common for users to be assisted by robots or computing machineries. In the context of physical human–robot interaction (pHRI), haptic modality has a high use potential; haptic feedback can directly pass through the neuromuscular system without going through a high-level recognition process [1]; therefore, a user can respond faster than when other modalities (e.g., visual feedback) are used. Recently, the haptic guidance (HG) system, also called the haptic shared control system, has been accepted as a promising approach for human–machine interfaces or pHRI situations [2]. The HG system is defined as a method in which the control input determined through physical

interaction between the force exerted by a human operator and the guiding force of a robot is applied to the target system [1]. Compared to the previous support systems using haptic cueing [12, 29] or input-mixing shared control [13], an advantage of the HG system is that the user can not only recognize the intention of the robot (e.g., direction and strength), but also choose to what extent this intention to be reflected. For example, a user usually follows the robot’s guidance, but when his/her choice is necessary, he/she can apply a force that exceeds the HG to perform the desired control.

HG technologies presented in earlier studies can be classified into two categories according to their design method: *optimal action-based haptic guidance* (OAHG) and *user prediction-based haptic guidance* (UPHG). OAHG is designed to convey the optimal movement for performing a task in the current state (e.g., guiding the user toward the center of the road in a steering task [15, 35, 53] or through the movement of a skilled expert in a peg-in-hole task [38]). Meanwhile, UPHG is designed to provide proactive guidance in the direction the user intends to move based on their behavior prediction. For example, in a steering task, it guides the users to their individually preferred courses rather than the center of the road [10]. Both types of HG systems have been proven to have several positive effects in terms of task performance improvement [5, 10, 15, 19, 27, 35, 55], user workload reduction [19, 27, 35], and user subjective satisfaction [15, 19, 27] in various recent studies.

While most of the previous studies deal with the design methods and effects of OAHG and UPHG, very few studies have clearly compared OAHG and UPHG. The two HGs are expected to have different characteristics because each goal behavior is different. OAHG informs the user of the most optimal action for the current task. However, if the guidance is in conflict with the user’s intention (i.e., a disagreement occurs), it will lead to an undesired physical interaction between the user and the robot, which can induce discomfort and frustration for the user [13, 19]. On the other hand, UPHG supports comfortable movements of the users by reducing trajectory mismatch with the robot [10], but has a limitation in that it cannot present more optimal movements, although they may exist. In this context, this study aims to answer the following questions, which have important implications for pHRI design, “What is the difference between OAHG and UPHG in terms of user acceptance?” and “Is it possible to design a better HG by combining OAHG and UPHG?”

To impartially compare OAHG and UPHG, we need to implement each type of HG to achieve its best performance for a given application. Consequently, we present the following implementation methods to improve the performance of each HG based on deep learning approaches, which are attracting attention in recent HG studies [7, 46, 52]. First, a deep reinforcement learning algorithm is used to train an optimal action model (for OAHG) to learn the optimal policy (i.e., the optimal way of behaving from a specific state during a task) through self-play between AI agents [3] (the upper flow in Figure 1(b)). Second, we train a user prediction model (for UPHG) in a supervised manner with multiple user’s behavior data. To deal with individual differences between human operators, we apply a meta-learning approach to enable to adapt the model parameters according to the current user (the lower flow in Figure 1(b)). Third, both the optimal action model and the user prediction model are designed to infer the model uncertainty from

their outputs. By making HGs consider their model uncertainty at every timestep, we aim to prevent the decrease in HG performance because of an inaccurate model.

Another essential purpose of our study is to explore the possibility of complementing each HG type. If OAHG and UPHG are in a trade-off relationship with each other, it can be expected to produce HG that achieves optimal performance in terms of objective and subjective metrics by properly harmonizing the two HG types. In particular, we focus on the fact that the disagreement between the guiding force and the user intention reduces the OAHG performance [13, 19]. Therefore, we devised a method to implement the combined haptic guidance (CombHG) utilizing the similarity of guiding forces generated by the two HG types every timestep to minimize the disagreement. With this similarity-based method, the combined guiding force is adjusted according to the difference between the two guiding force directions.

We conducted a user experiment with 20 participants to verify how each type of HG has different effects on objective and subjective evaluations. We built a task environment for confronting an AI agent in an *air hockey* game that can be operated and assisted through a haptic device, as shown in Figure 1(a). Each participant played the air hockey game with the HGs we implemented, and subjective evaluations including user interviews were conducted after each task. Our user experiment results indicate the following important points. First, all types of HG—OAHG, UPHG, and CombHG—led to a significant improvement in objective user performance compared to the case where the user did not receive any HG, but there was no significant difference in objective performance between the three HG types. Second, UPHG and CombHG elicited a significantly higher score in subjective metrics, such as perceived naturalness and comfort, than OAHG. Finally, CombHG significantly lowered the disagreement between the user intention and HG compared to the cases of OAHG and UPHG, without reducing the objective and subjective scores.

Overall, this paper has following three key contributions:

- We present deep learning-based approach to implement OAHG and UPHG to achieve their best performance, applying a self-play-based reinforcement learning framework for OAHG and a meta-learning framework for UPHG. In particular, we propose and verify two novel implementation methods—uncertainty-based thresholding and user adaptation—to improve the HG performance.
- We propose and verify a combined approach (CombHG) of OAHG and UPHG that can complement each HG type, utilizing our similarity-based combination method. To the authors’ knowledge, this is the first attempt to combine OAHG and UPHG to achieve better performance.
- We experimentally compare OAHG, UPHG, and CombHG in terms of objective and subjective metrics through a user study and interview (n=20).

2 RELATED WORK

2.1 Haptic Guidance

Early HG studies were based on control assistance methods in virtual or teleoperation environments, such as a virtual fixture [42] and an artificial force field [54], which help users accurately move

toward goals and prevent access to dangerous areas. Robotic devices enabled the delivery of haptic feedback generated by the virtual fixture or artificial force field techniques to users, allowing them to perform the tasks with improved stability [4, 28]. Over the past few decades, the implementation of HG has made significant progress, and it has been embedded in a variety of forms, including haptic devices [5, 10, 38, 55], sleeve devices [8, 16], pen devices [25], and steering wheels [9, 21, 35, 50, 53]. Accordingly, the range of HG applications has also expanded, such as in surgical assistance [20, 41, 55], driving assistance [21, 35, 50, 53], teleoperation of robots [38, 47] and UAVs [28, 49], and desktop computer interfaces [11, 26].

A majority of the previous studies have reported positive effects of HG, such as improvement in task performance and user comfort, and reduction of user workload. Nevertheless, some elements of HG that hamper usability remain. As the most representative example, a conflict between the user and the HG can lead to a temporary increase in the force exerted by the user [6, 36], or even a decrease in performance [37]. In addition, if the interference of HG is excessive for a user, then the user requires more physical effort, which leads to a deteriorated user evaluation (e.g., low comfort and controllability) [30]. A detailed analysis of the factors affecting HG will bring useful implications in designing a user-friendly HG. In this study, we aim to determine what users expect (and not expect) from HG through a comparison between OAHG and UPHG, which was not sufficiently covered in previous studies. Furthermore, we attempt, for the first time, to optimally combine OAHG and UPHG to achieve better performance based on the understanding of the factors that influence HG performance.

2.2 Implementing Optimal Action-based HG

The most straightforward way of implementing OAHG is to set a reference trajectory for performing a task and deliver a continuous guiding force so that the user does not leave the trajectory. In the context of a steering task such as when driving, a number of studies have implemented OAHG in the form of a guiding force directed toward the center of the path [9, 15, 21, 35, 53]. As another example, in a backward parking situation, Tada *et al.* [50] implemented OAHG by utilizing a Bezier curve between a start point and a target parking point as a reference trajectory. Meanwhile, when the reference trajectory could not be clearly defined in advance, a demonstration by skilled experts also served as the reference trajectory, for example, in the case of a handwriting task [5, 51] and a peg-in-hole task [38].

Recently, attempts have been made to apply reinforcement learning techniques in HG implementation to train an optimal policy for a task. Scobee *et al.* [46] attempted to determine the underlying value function of each observation state of a steering task from the movement of an expert operator through the inverse reinforcement learning method, and developed OAHG based on the trained value function. Meanwhile, the deep Q-network (DQN) was applied in [52] to train HG that minimizes the magnitude of the steering wheel angle during a steering task. These previous reinforcement learning-based HGs showed promising results, but had a limitation in that they cannot provide a perfectly optimal action because they were trained by data from very few (one or two) human operators. In this study, we trained an optimal action model for implementing

OAHG by using a latest reinforcement learning framework based on self-play between AI agents only [3].

2.3 Implementing User Prediction-based HG

Research on UPHG emerged later than research on OAHG and has received less attention. De Jonge *et al.* [10] presented UPHG according to a reference trajectory adapted to each user based on his/her previous trial in a steering task, whereas OAHG utilized a fixed reference trajectory (i.e., center of the path). Their personalized UPHG demonstrated the positive effect of reducing the conflict between the human operator and HG. Meanwhile, data-driven approaches have also been used to model a user's movement. Hernández *et al.* [19] stochastically modeled each user's individual movement based on a hidden Markov model (HMM) in a task of moving a virtual object to a target position avoiding obstacles, and presented UPHG to assist the user's movement based on the user model. Previous studies such as [41] and [55] also presented HG based on a user movement model using HMMs in a surgical task environment.

In this study, we implement a user model for UPHG based on deep neural networks. Such an approach has recently been highlighted for user modeling in HCI and HRI fields [32, 34, 39, 43]. To apply UPHG tailored for an individual user, it is necessary to learn on each individual's data. However, collecting enough data to train the neural network-based model from scratch for every new user is very inefficient, particularly, when targeting numerous users. To solve this challenging problem, we apply a meta-learning approach, enabling the trained model to adapt the model parameters to a new user based on his/her short trial data, which is the first attempt in UPHG implementation.

3 HAPTIC GUIDANCE DESIGN

3.1 Target Task Environment

As a target task for the experimental investigation of various types of HG, we developed a virtual *air hockey* game environment that can be controlled using a haptic device (Figure 2(a)). In our air hockey game environment, a player competes with an opponent AI on a slippery surface by moving his/her *paddle*, which is used to hit a *puck* (Figure 2(b)). The player's successful task execution is judged by how many times he/she wins against the opponent AI. The player wins a round if he/she succeeds in putting a puck into the opposing goalpost, whereas he/she loses the round if the opponent puts a puck in the player's goalpost. Each paddle is set to move only within its own side (i.e., above or below the half-line). The player's paddle is controlled by a two-dimensional action vector, which corresponds to the desired paddle location on the xy coordinates with the midpoint of the player's side as the origin (Figure 2(b)). The action vector has one-to-one correspondence with the position of the end effector of the haptic device. For intuitive control of the player, the end effector is also set to move only on the 2D plane, and the moving directions of the end effector and the paddle are matched (Figure 2(a)).

A video game environment is suitable for comparing OAHG and UPHG in that each player can adopt various strategies according to his/her preference [22, 40] while having a common winning formula. A player's basic strategy to win the air hockey game is to accurately *smash* the puck toward the opponent's goal and *defend*

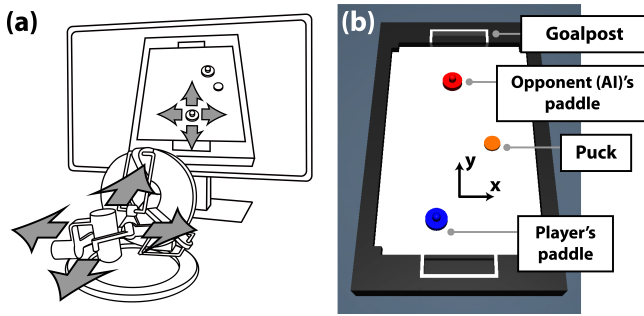


Figure 2: The air hockey game environment we developed. (a) Example of control based on a haptic device. (b) In-game screen and composition of our air hockey game.

against the puck from the opponent heading inside the player’s goal. However, in a detailed process, players can have several choices. A player can smash along a path that goes directly toward the goal, or he/she can choose a path through one or two reflections using the wall, avoiding the opponent’s paddle. In addition, in a defensive situation, that is, when the opponent has the puck, a player can press the opponent near the half-line to narrow the angle of attack or wait for the opponent’s attack right in front of the goal. We expect UPHG to predict the player’s choices in advance and assist him/her with the corresponding action, while OAHG informs him/her of the most optimal action he/she can choose.

3.2 Optimal Action Model and User Prediction Model

We implemented the following two models based on deep neural networks: 1) an optimal action model that outputs the most optimal action that a user can take and 2) a user prediction model that outputs the expected action to be selected by the current user. Because both models use the same type of input (i.e., current game state) and output (i.e., target action vector), we designed them to have the same structure but with different hyperparameters (i.e., depth and size of the hidden layer) that are fine-tuned for each model’s training. Each model was trained through different learning approaches based on reinforcement learning and meta-learning, respectively.

3.2.1 Model Architecture. Figure 3 shows the structure of our model. The air hockey game state—the two-dimensional position and velocity vectors of two paddles and a puck (total, size of 12)—is mapped to the two-dimensional distribution data, which are the mean and standard deviation (STD) of the distribution of each xy coordinate of the action vector (total, size of 4). The output distribution represents the distribution of the target action vector used to generate the guiding force (i.e., optimal action for OAHG or predicted action for UPHG). Notably, we designed the model to output the mean and STD of the distribution rather than a single action value. In our HG implementation, we utilize the mean of the distribution as a reference action to guide the users, and also consider the model uncertainty which can be inferred from the STD of the distribution. Fully connected (FC) hidden layers with rectified linear unit (ReLU) activation functions are present between the

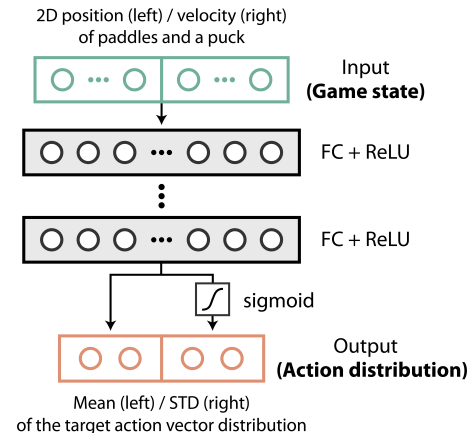


Figure 3: Overview of the model structure applied to both the optimal action model and the user prediction model.

input and output nodes. We used grid search to determine the optimal hyperparameters that exhibit the best learning performance in each training method; therefore, the depth and size of the hidden layers are different for the optimal action model and user prediction model: 2×64 and 4×80 , respectively. The output STD goes through an additional sigmoid activation function.

3.2.2 Optimal Action Model Training. Reinforcement learning techniques have succeeded in acquiring an optimal policy that surpasses a human player in simulated video game environments [31, 48]. To train the optimal action model for our air hockey task, which is based on the competition between two players, we applied a latest learning framework based on self-play between two AI agents [3]. In this framework, a training agent of the first generation initially confronts an opponent with random movements and updates the action model for increasing the expected reward. As the learning progresses, the training agent of the next generation competes with the agents of the previous generations, which also grow gradually, therefore the optimal action model can be progressively developed without human operator intervention. Meanwhile, the training agent updates the model using the sampled action from the action distribution output of the model. Therefore, when the specific action is confidently more optimal than other actions, the model will train to output a lower STD value, that is, the STD reflects the model uncertainty at current state.

For implementation details, trust region policy optimization (TRPO) [44] was used as the learning algorithm that updates the policy of the training agent, because it internally had the best learning performance for our air hockey task among the latest algorithms such as proximal policy optimization (PPO) [45] and soft actor-critic (SAC) [18]. Self-play-based procedural learning was performed over 100 generations. Each generation was trained with a simulation of 200K timesteps, and batch updates were applied every 5K timesteps.

3.2.3 User Prediction Model Training. In utilizing the user prediction model, we mainly focus on how to train the model to adapt to different users. Inspired by the fact that the existing meta-learning frameworks train models to quickly learn to perform new unseen tasks using only a few datapoints of each new task, we applied a meta-learning approach to train the user prediction model for

HG to quickly adapt to different users. We demonstrated in our prior study [33] that model-agnostic meta-learning (MAML) [14], a recent mainstream meta-learning algorithm, can be used to train a model that mimics user behavior for dynamic difficulty adjustment and effectively adapts to a new user with his/her minimal demonstration data. In the current study, we present a *user adaptation* (UA) method that trains the user prediction model based on the MAML algorithm and updates the model parameters for each user. The training process of our user prediction model is as follows. The training data consists of a set of task execution data D^U of different users U , and D^U is divided into D_{demo}^U for model adaptation, and D_{valid}^U for meta-update. First, the algorithm updates the model parameter θ by a single gradient descent that reduces $L(\theta, D_{demo}^U)$, which is the model loss function for D_{demo}^U . Therefore,

$$\theta_{adapt}^U = \theta - \alpha \nabla_{\theta} L(\theta, D_{demo}^U), \quad (1)$$

where α denotes the inner learning rate. Subsequently, a meta-update is performed to reduce model loss for D_{valid}^U using the updated parameters. Overall, the objective of our meta-learning is as follows:

$$\min \sum_U L(\theta_{adapt}^U, D_{valid}^U). \quad (2)$$

Through this two-fold backpropagation, the user prediction model eventually learns to rapidly adapt to unseen users. Meanwhile, we aim to train the model to output the STD of the predicted action distribution, which is impossible with the mean-squared error loss function commonly used in general regression learning. Inspired by a previous work [24] that modeled uncertainty in deep learning models for computer vision tasks, we use the following loss function to train our user prediction model:

$$L(\theta, D) = \frac{1}{N} \sum_{(x_i, y_i) \in D} \frac{1}{2\|\hat{\sigma}_i\|^2} \|y_i - \hat{y}_i\|^2 + \frac{1}{2} \log \|\hat{\sigma}_i\|^2, \quad (3)$$

where N denotes the size of the training data D (i.e., D_{demo}^U or D_{valid}^U), subscript i indicates the i -th sample from D , y_i is an actual action taken by a player at a game state x_i , and \hat{y}_i and $\hat{\sigma}_i$ are the predicted mean and STD output by feeding x_i into the model with parameter θ . With this loss function, the model can be trained to output the STD that implies the uncertainty of the predicted action; that is, the lower the STD values, the more confident the model is about the prediction.

For the user prediction model training, we collected data from nine participants aged between 22–29 (mean=25.33, STD=2.16) while performing the air hockey task without haptic guidance. Each participant was provided sufficient practice time to be familiar with the task prior to data collection. A total of 360K timesteps of recorded behavioral data were employed to train the user prediction model. During our meta-learning process, the adaptation was conducted with an inner learning rate of 0.1, and the meta-update was conducted using an Adam optimizer with a learning rate of 0.001. The batch size of D_{demo}^U and D_{valid}^U was 1K timesteps, and the entire training was conducted for 200 epochs.

To verify the practical effectiveness of the UA method, in a subsequent user study, we experimentally compared the performance of two UPHGs based on models trained with and without UA. For the UPHG implementation without using the UA method, we trained

another user prediction model in a typical supervised manner, excluding the meta-learning approach, employing the same dataset used in the training with UA. In this case, the model trained without UA will predict generalized user movement because it is trained to use fixed network parameters for various user behaviors. An Adam optimizer with a learning rate of 0.001 was used and the batch size was 1K timesteps. The entire training was conducted for 100 epochs.

3.3 Generating HGs Based on Trained Models

The trained models output the target action for assisting the user's task. A straightforward and reliable approach to generate HG is to implement a spring force toward the target action so that users can match their actions accordingly [10, 15, 35, 46]. Meanwhile, one of the important implications from the previous studies on HG implementation is that a human operator cannot respond to HG immediately. Accordingly, a lookahead method [15] where HG should be proactively applied based on a slight future game state, considering the reaction time of a human, has been applied as a common technique. Following the method described in [15], we used the anticipated game state $x_{lookahead}$, which is obtained by virtually moving the puck and paddles with the current velocity for a short time $T_{lookahead}$, instead of the current game state x as the model input. Therefore, the guiding force of our OAHG and UPHG is generated by the following equation:

$$F_{HG} = clip(-K(\mathbf{u} - \hat{\mathbf{y}})), \quad (4)$$

where K is the stiffness gain, \mathbf{u} is the user's current action vector, and $\hat{\mathbf{y}}$ is the mean of the target action vector distribution obtained by feeding $x_{lookahead}$ into the model. A clipping function was applied to make the guiding force bounded, because momentary excessive force may induce safety problems to the user.

Additionally, we propose an *uncertainty-based thresholding* (UT) method that adjusts the magnitude of HG by considering model uncertainty. The importance of model uncertainty has been recognized in several previous pHRI studies [17, 19, 23], as the results from inaccurate models of robot or human movements can lead to human discomfort or safety issues during human-robot collaborations. We infer the model uncertainty at a specific state from the STD of the target action; that is, if the output STD is large, the model at current state is judged to be highly uncertain. Using this, we defined the uncertainty weight according to the squared magnitude of the STD as follows:

$$W_{unc} = \begin{cases} 1, & \text{if } \|\hat{\sigma}\|^2 < T_{low} \\ \frac{T_{high} - \|\hat{\sigma}\|^2}{T_{high} - T_{low}}, & \text{if } T_{low} \leq \|\hat{\sigma}\|^2 < T_{high} \\ 0, & \text{otherwise,} \end{cases} \quad (5)$$

where $\hat{\sigma}$ is the output STD, and T_{low} and T_{high} are threshold values. With the UT method, calculated F_{OAHG} or F_{UPHG} from (4) is multiplied by the uncertainty weight, which ranges from 0 to 1, and then provided to users. In other words, only when the STD is less than a certain level (i.e., T_{high}), we determine the HG to be confident in its purpose and can assist the users. Figure 4(a) summarizes the process of generating OAHG and UPHG, including the UT method.

Table 1: Proposed methods applied to our HG implementation.

Abbrev	Full Form	OAHG	UPHG	CombHG
UT	Uncertainty-based thresholding	✓	✓	✓
UA	User adaptation		✓	✓
SC	Similarity-based combination			✓

✓A checkmark indicates whether it was used for the corresponding HG type.

A straightforward approach to implement CombHG is to use the average of two guiding force vectors from OAHG and UPHG, therefore make the user to be assisted by both HG types simultaneously. This approach is intuitive in that the combined HG guides the user to the midpoint between the two target actions to be guided by OAHG and UPHG. Accordingly, we implement the simplest CombHG as follows:

$$F_{CombHG} = (F_{OAHG} + F_{UPHG})/2. \quad (6)$$

However, simply taking the average of two guiding forces may worsen the results, guiding the user to an unintended third direction if the two HGs have different directions. To solve this problem, we propose a conservative combination method, namely *similarity-based combination* (SC), that considers the similarity of two HGs based on the angle between the two guiding forces and reduces the magnitude of HG when the similarity is low. The main purpose of the SC method is to provide an appropriate guiding force only when the target action is optimal and meets the user's intention,

that is, when the directions of the two HGs match. To do this, we define the similarity weight as follows:

$$W_{sim} = \cos^2(\phi/2), \quad (7)$$

where ϕ is the angle between the two guiding forces according to OAHG and UPHG. The similarity weight ranges from 0 to 1, which corresponds to the angle ϕ from π (opposite direction) to 0 (matched direction). With the SC method, F_{CombHG} from (6) is provided to the user after multiplying by the similarity weight. Figure 4(b) shows the implementation of CombHG using OAHG and UPHG, including the SC method.

4 USER STUDY DESIGN

4.1 Research Questions

The objective of this research is divided into two main parts: a presentation of the implementation methods for three types of HG, that is, OAHG, UPHG, and CombHG, and an experimental investigation into how each type of HG differs in the context of user acceptance. We implemented the HGs by generating a guiding force based on the real-time output of the optimal action model and the user prediction model we trained. To achieve better performance of HG, we applied the three proposed methods (i.e., UT, UA, and SC), according to the HG type specified in Table 1. To aid future HG studies, this user study first aims to experimentally confirm whether the methods actually lead to a HG performance improvement. Subsequently, by using HGs to which all the applicable methods are employed, we determine whether each HG assists users when compared to a no HG (NHG) condition. Further, we investigate how the three types of HG assist users differently in terms of objective and subjective evaluations, which is our primary research objective. Therefore, we formulate the following four research questions:

- **RQ1:** Do the UT, UA, and SC methods that we propose and apply to the HG implementation contribute to an improvement in HG performance?
- **RQ2:** Do OAHG, UPHG, and CombHG improve users' task performance when compared to NHG?
- **RQ3:** What differences do OAHG and UPHG have in users' objective and subjective evaluations?
- **RQ4:** Can CombHG, which integrates OAHG and UPHG, complement each HG or provide better effects in users' objective and subjective evaluations?

4.2 Experimental Method

We conducted an indoor laboratory experiment to measure the actual assisting performance of the implemented HGs. We recruited 20 participants (4 females and 16 males) aged between 21–30 (mean=




Figure 4: The process of implementing (a) optimal action-based haptic guidance (OAHG) and user prediction-based haptic guidance (UPHG), and (b) combined haptic guidance (CombHG). The three proposed methods (i.e., UT, UA, and SC), whose effectiveness is investigated through a subsequent user experiment, are highlighted in gray shades.



Figure 5: Experimental setup of our user study. A participant performs a virtual air hockey task with a haptic device.

25.58, STD=2.09) for this user study. All participants were right-handed, and none of them reported perception defects in their vision and touch.

Figure 5 presents the experimental setup. The air hockey task environment was composed of a haptic interface device (Omega.7, Force Dimension) and a 24-inch monitor, which were connected to a PC. Each participant was instructed to sit down and perform the air hockey task by holding the end effector of the haptic device with his/her dominant hand.

We selected HGs under the following nine conditions for the user experiment for verifying the effectiveness of the HG implementation methods (RQ1) and comparing the effect of each HG type (RQ2–RQ4):

- (a) **NHG**: no haptic guidance.
- (b) **OAHG (vanilla)**: OAHG from Equation (4) with no additional method.
- (c) **OAHG (UT)**: OAHG applying the uncertainty-based thresholding in addition to OAHG (vanilla).
- (d) **UPHG (vanilla)**: UPHG from Equation (4) using a model not trained with the meta-learning algorithm, and therefore without the user adaptation.
- (e) **UPHG (UA)**: UPHG from Equation (4) using a model trained with the meta-learning algorithm (i.e., UA applied).
- (f) **UPHG (UT+UA)**: UPHG applying the uncertainty-based thresholding in addition to UPHG (UA).
- (g) **CombHG (UA)**: CombHG from Equation (6) based on OAHG (vanilla) and UPHG (UA). Note that because the basic concept of UPHG is assisting a user according to personalized prediction, we use UPHG (UA) instead of UPHG (vanilla) to implement the underlying CombHG.
- (h) **CombHG (UA+SC)**: CombHG applying the similarity-based combination on OAHG (vanilla) and UPHG (UA).
- (i) **CombHG (UT+UA+SC)**: CombHG applying the similarity-based combination on OAHG (UT) and UPHG (UT+UA).

Our experimental procedure is as follows. First, all participants were informed that they would experience nine different HGs, and interviews would be conducted about their impressions received

from each HG, immediately after experiencing each HG. Second, each participant was provided with as much practice time as he/she wanted to be familiar with the task environment. Third, for the purpose of updating the user prediction model parameters (i.e., UA), each participant performed one game in the NHG condition, which was not reflected in the evaluation. One game basically consists of seven rounds (one round ends when either a player or an AI scores a goal), but if it ends earlier than the minimum play time that we set (two minutes per task), up to three additional rounds proceed to secure more data from the player. After completing UA with the recorded data, that is, updating the parameters of the user prediction model through Equation (1), the participant sequentially performed all nine HG conditions in a random order with counterbalancing. All participants performed one game for each HG condition, and before the start of each game, they were allowed 30 seconds to adapt to the given HG, which was not reflected in the evaluation. Following the end of each game, participants assessed the subjective scores for the HG they had just been assisted with, and a short interview with a supervising researcher was conducted. All participants were able to rest as much as they wanted between each game. The total duration of the experiment was between 1–1.5 hours, depending on the participants.

4.3 Measured Variables and Metrics

We measured the assisting performance of each HG condition in terms of objective metrics automatically measured by the system and subjective metrics based on participants' evaluations. Four objective metrics were used: win rate, mean smash speed, and defense rate, which indicate the degree of high task performance; and mean disagreement, which indicates the degree of high conflict between a user and HG. The metrics are defined as follows:

- **Win rate**: the ratio of the participant winning the opponent in one game (7–10 rounds).
- **Mean smash speed**: the average speed of the puck hit by the participant over the opponent's side (i.e., smashed). Note that, the speed was measured in relative figures because the air hockey task was built in a virtual environment without specific units.
- **Defense rate**: the proportion of the pucks blocked by the participant among the pucks headed into the participant's goal.
- **Mean disagreement**: the average of the disagreement between the participant and HG, which is proposed in [19] and defined as follows ($\Delta \mathbf{u}$ denotes the user's action change during a timestep after receiving HG).

$$disagreement = \begin{cases} -\frac{\mathbf{F}_{HG}^T \cdot \Delta \mathbf{u}}{\|\Delta \mathbf{u}\|}, & \text{if } \mathbf{F}_{HG}^T \cdot \Delta \mathbf{u} < 0 \\ 0, & \text{otherwise.} \end{cases}$$

For the subjective evaluation, we selected the following four items and asked the participants to rate each HG condition in terms of the items on a 7-point Likert score:

- **Helpfulness**: how much the participants felt the HG *helped* them perform the task.
- **Naturalness**: how *natural* the participants felt the assistance of the HG.

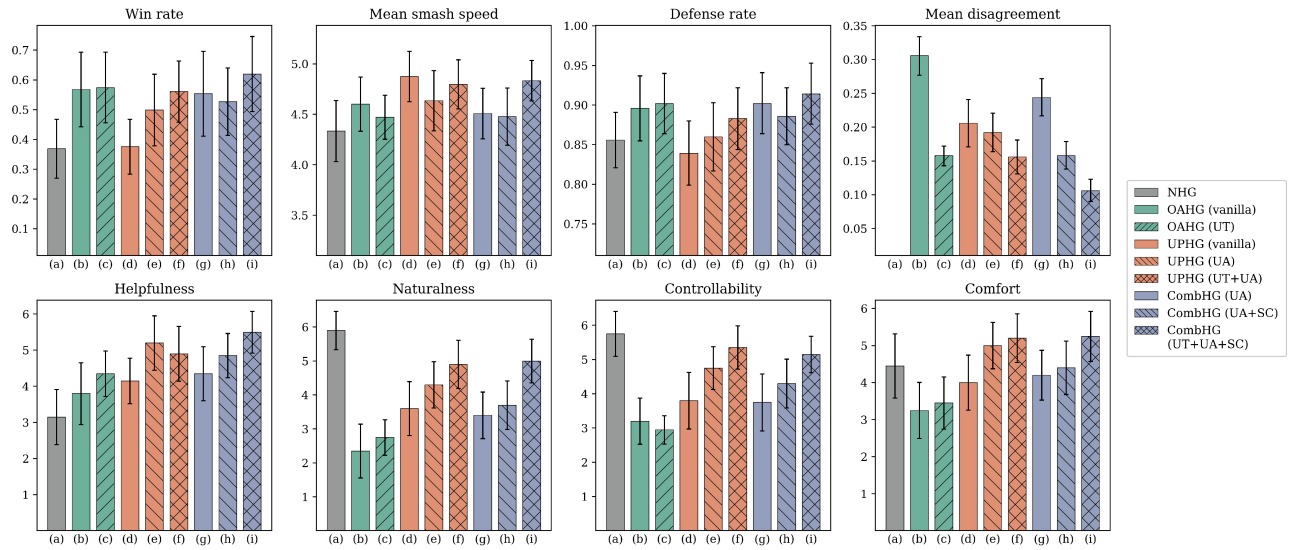


Figure 6: Experimental results of the objective (top) and subjective (bottom) metrics for the nine HG conditions (a)–(i). The HG conditions of the same HG type are indicated by bars of the same color. Error bars represent 95% confidence intervals.

- **Controllability:** how well the participants felt they were *able to control* the paddle under the HG.
- **Comfort:** how *comfortable* the participants were with the assistance of the HG.

5 RESULTS

Figure 6 shows the objective and subjective evaluation results for each HG condition. To answer RQ1, we first investigated how assisting performance varies according to implementation methods within each HG type. Next, we examined how each HG type provides different assistance for the user (RQ2–RQ4), using the results of the HG conditions to which all applicable implementation methods were employed, that is, OAHG (UT), UPHG (UT+UA), and CombHG (UT+UA+SC).

5.1 Comparison Within HG Types

We first compared the evaluated results of the two conditions within the OAHG type according to the implementation methods (i.e., vanilla vs. UT). A paired t-test revealed that the use of the UT method in the OAHG type significantly reduces the participants' mean disagreement ($t = 12.143$, $p < .001$). On the other hand, no significant difference was found in the remaining seven metrics between the OAHG (vanilla) and OAHG (UT) conditions.

To compare the three conditions within the UPHG type (i.e., vanilla, UA, UT+UA), a repeated measures ANOVA with a Greenhouse-Geisser correction was used. The analysis revealed that there were significant effects of the UPHG implementation methods on the following metrics: win rate ($F_{2,38} = 6.430$, $p = .004$, $\eta^2 = .253$), mean disagreement ($F_{2,38} = 19.856$, $p < .001$, $\eta^2 = .511$), helpfulness ($F_{2,38} = 3.422$, $p = .044$, $\eta^2 = .153$), naturalness ($F_{2,38} = 5.156$, $p = .013$, $\eta^2 = .213$), controllability ($F_{2,38} = 10.188$, $p = .001$, $\eta^2 = .349$), and comfort ($F_{2,38} = 7.012$, $p = .006$, $\eta^2 = .270$).

However, there was no significant effect on the mean smash speed and defense rate. For the metrics with significant effects, we conducted post hoc tests with a Bonferroni correction to investigate a significant mean difference between the conditions. Table 2 summarizes the comparisons of the conditions in which significant mean differences exist (i.e., $p < .05$). According to the analysis, the use of methods such as UA and UT contributed individually or together to improve objective evaluations (e.g., an increase in win rate and a decrease in mean disagreement) and subjective evaluations (e.g., an increase in helpfulness, naturalness, controllability and comfort) from participants.

We also conducted a repeated measures ANOVA with a Greenhouse-Geisser correction to compare the results of the conditions representing the three implementation methods within the CombHG type (i.e., UA, UA+SC, and UT+UA+SC). There were significant effects of the implementation methods on the following metrics: mean smash speed ($F_{2,38} = 5.733$, $p = .011$, $\eta^2 = .232$), mean disagreement ($F_{2,38} = 79.038$, $p < .001$, $\eta^2 = .806$), helpfulness ($F_{2,38} = 3.295$, $p = .049$, $\eta^2 = .148$), naturalness ($F_{2,38} = 7.883$, $p = .002$, $\eta^2 = .293$), controllability ($F_{2,38} = 5.607$, $p = .008$, $\eta^2 = .228$), and comfort ($F_{2,38} = 4.057$, $p = .027$, $\eta^2 = .176$). However, there was no significant effect on win rate and defense rate. Post hoc tests with a Bonferroni correction were also conducted, and comparisons between conditions with significant mean differences are summarized in Table 2. Similar to the analysis of the UPHG type, the post hoc analysis indicated that the methods we present for CombHG (i.e., UT and SC) contribute to improving the objective and subjective aspects of assisting performance for users.

5.2 Comparison of HG Types

To compare the assisting performance between HG types, we used the OAHG (UT), UPHG (UT+UA), and CombHG (UT+UA+SC) conditions, which showed the best assisting performance for each

Table 2: Comparisons showing statistically significant differences between implementation methods within HG type.

Type	Metrics	Comparison	p	Type	Metrics	Comparison	p
OAHG	Mean disagreement	(b) > (c)	.000	CombHG	Mean smash speed	(i) > (g)	.003
UPHG	Win rate	(f) > (d)	.009		(i) > (h)	.028	
	Mean disagreement	(d) > (f)	.000		(g) > (h)	.000	
		(e) > (f)	.000		(g) > (i)	.000	
	Helpfulness	(e) > (d)	.041		(h) > (i)	.000	
	Naturalness	(f) > (d)	.032		Helpfulness	(i) > (g)	.023
	Controllability	(e) > (d)	.028		Naturalness	(i) > (g)	.001
		(f) > (d)	.004		(i) > (h)	.043	
	Comfort	(f) > (d)	.001		Controllability	(i) > (g)	.009
					Comfort	(i) > (g)	.023

HG type in Section 5.1. Hereafter, the above three conditions will represent each HG type. A repeated measures ANOVA with a Greenhouse-Geisser correction was conducted to determine the effect of HG types (i.e., NHG, OAHG, UPHG, and CombHG) on each metric. As an exception, since mean disagreement cannot be calculated under the NHG condition, the analysis for mean disagreement was conducted only between the other three HG types. The ANOVA analysis revealed that the HG type has a significant effect on all metrics and statistical values of each metric are as follows: win rate ($F_{3,57} = 8.112, p < .001, \eta^2 = .299$), mean smash speed ($F_{3,57} = 5.740, p = .002, \eta^2 = .232$), defense rate ($F_{3,57} = 3.014, p = .041, \eta^2 = .137$), mean disagreement ($F_{2,38} = 13.586, p < .001, \eta^2 = .417$), helpfulness ($F_{3,57} = 8.614, p < .001, \eta^2 = .312$), naturalness ($F_{3,57} = 30.107, p < .001, \eta^2 = .613$), controllability ($F_{3,57} = 30.083, p < .001, \eta^2 = .613$), and comfort ($F_{3,57} = 7.108, p = .001, \eta^2 = .272$). Post hoc tests with a Bonferroni correction were conducted to figure out whether HG type pairs had a statistically significant difference, and the results are summarized in Table 3.

Based on the post hoc analysis, we can conclude the following: All three HG types, that is, OAHG, UPHG, and CombHG, led to significantly higher win rates of participants than NHG, which demonstrates the objective effectiveness of the HGs presented in this study. Specifically, CombHG induced a significantly higher mean smash speed than NHG, while OAHG induced a significantly higher defense rate than NHG. UPHG also induced a marginally higher smash speed than NHG, but it was not statistically significant ($p = .073$). Meanwhile, CombHG showed a lower mean disagreement than both UPHG and OAHG. In terms of the subjective metrics, UPHG and CombHG scored significantly higher in helpfulness than NHG, whereas OAHG did not. UPHG and CombHG received similar levels of subjective evaluation from participants, scoring significantly better than OAHG for the remaining three metrics, that is, naturalness, controllability, and comfort. NHG received the highest scores in naturalness and controllability, but this can be interpreted as an inevitable result of not exerting any artificial force on the users.

6 DISCUSSION

We summarize the findings of this paper as answers to our research questions based on the analysis results and comments from the

user interviews. In addition, we discuss the implications of our HG design regarding its generalization and utilization outside of the air hockey environment.

6.1 Answers to Research Questions

RQ1: *Do the UT, UA, and SC methods that we propose and apply to the HG implementation contribute to an improvement in HG performance?*

The experimental results showed that the three methods proposed in this study contributed to the enhancement of the assisting performance in several metrics. First, the application of UT led to a reduction in mean disagreement for all HG types. This can be explained by the fact that the conflict between HG and a user because of an inaccurate model output, that is, with a high uncertainty, was effectively prevented by the uncertainty weight. The user interviews provided more details. Five participants (P6, P7, P11, P12, and P17) commented on OAHG (vanilla) that “*The interference of the guidance was excessive*,” but only two participants (P8, P11) made such comments on OAHG (UT). A similar tendency was observed for UPHG. When comparing cases of non-applied and applied UT (i.e., UPHG (UA) vs. UPHG (UT+UA)), the number of participants commenting that “*The interference frequency of the guidance was appropriate*” increased from zero to five (P7, P10, P16, P17, and P20).

On the other hand, the application of UA improved the users’ evaluation of helpfulness and controllability. Moreover, when UA worked with UT on UPHG, there were enhancements in win rate, naturalness, and comfort when compared to UPHG (vanilla). This enhancement can be explained by UA providing the effect of adjusting HG to suit an individual user’s playstyle, as we intended. In the interviews, three participants (P7, P8, and P13) reported on UPHG (UA) that, “*The guidance understood my offensive and defensive intentions*,” whereas no one reported this on UPHG (vanilla).

The application of SC reduced the mean disagreement, which is similar to the effect of UT. This can be explained by the fact that the HG was fully delivered to the user only when OAHG and UPHG have a matched direction, thereby reducing the degree of HG interference. Note that there was no deterioration in other metrics, even though the HG interference was controlled this way. Rather, when SC worked with UT on CombHG (i.e., CombHG (UT+UA+SC)), the mean smash speed and all subjective evaluations were enhanced when compared to CombHG (UA). This suggests that users may

Table 3: Comparisons showing statistically significant differences between HG types.

Metrics	Comparison	<i>p</i>	Metrics	Comparison	<i>p</i>	
Win rate	OAHG > NHG	.014	Naturalness	NHG > OAHG	.000	
	UPHG > NHG	.002		NHG > UPHG	.025	
	CombHG > NHG	.001		UPHG > OAHG	.001	
Mean smash speed	CombHG > NHG	.023		CombHG > OAHG	.000	
	CombHG > OAHG	.045	Controllability	NHG > OAHG	.000	
Defense rate	OAHG > NHG	.038			UPHG > OAHG	.000
Mean disagreement	OAHG > CombHG	.000		CombHG > OAHG	.000	
	UPHG > CombHG	.002	Comfort	UPHG > OAHG	.003	
Helpfulness	UPHG > NHG	.009			CombHG > OAHG	.001
	CombHG > NHG	.001				

prefer to reduce such unnecessary HG interference. A remark from P3 to CombHG (SC+UA), “*It felt positive that the guidance assisted me only when I needed it,*” can supplement this.

Furthermore, we analyzed the distribution of W_{unc} and W_{sim} used in the UT and SC methods to determine the degree to which each method adjusted the guiding force to improve the HG performance. First, we calculated the mean and STD of the distribution of each weight within the same game. After averaging the mean and STD values over all games executed, the distribution (mean \pm STD) of W_{unc} for OAHG and UPHG were 0.600 ± 0.390 and 0.751 ± 0.306 , respectively. Because a lower weight implies a stronger decrease in guiding force, the result indicated that the UT method more aggressively controlled the guiding force for OAHG than for UPHG. Meanwhile, the distribution of W_{sim} was measured as 0.515 ± 0.343 .

RQ2: *Do OAHG, UPHG, and CombHG improve users’ task performance when compared to NHG?*

In the user experiment, all three types of HG led to a significantly higher win rate for users than NHG, which clearly indicates that the user’s haptic task performance was improved by the guiding force. The specific differences in assisting performance between the HG types are discussed in more detail in RQ3 and RQ4.

RQ3: *What differences do OAHG and UPHG have in users’ objective and subjective evaluations?*

There was no significant difference in participants’ win rates between OAHG and UPHG. However, differences between the two HG types were revealed in other objective metrics. OAHG induced a significant increase in defense rate, whereas UPHG induced a marginal increase in mean smash speed when compared to NHG. This tendency was also observed in the user interviews. Six participants (P1, P2, P3, P5, P6, and P19) commented on OAHG that “*The guidance helped to defend the puck toward our goal,*” which is the highest number among the other HG types. P6 additionally mentioned that “*It pinpointed the spot to be blocked for important defense,*” which indicates that OAHG reinforced the user’s insufficient defensive ability. On the other hand, the most frequent comment participants (P1, P3, P4, P7, P10, P13, P15, and P20) mentioned about UPHG was that “*The guidance assisted me in the direction I am moving,*” which may induce an increase in the smash speed. This can be explained as the trained user prediction model could easily expect the user to continue to move in a direction when he/she starts moving, thereby

providing an assisting guidance in that direction. Furthermore, P3 remarked that, “*It informed me where to stop smashing during the smashing process, so I was able to smash stable,*” which indicates that UPHG assisted the user’s smashing motion rather than simply accelerating in the moving direction.

In subjective metrics, user evaluation of OAHG and UPHG exhibited a clear difference. UPHG was evaluated to be significantly higher in most metrics, that is, in naturalness, controllability, and comfort, and marginally higher in helpfulness, when compared to OAHG. We looked for the cause of this one-sided subjective evaluation in the user interviews. Participants mentioned several negative comments on OAHG as follows: “*In an offensive situation, the guidance over-asserted its intention*” (P5, P7, P9, P11, and P17); “*It followed the puck too hard*” (P3, P10, P11, P13, and P15); “*It did not fit my intention sometimes*” (P13, P14, P19, and P20). P3 remarked in detail, “*When the puck was on the opponent’s side, I wanted to wait, but the guidance preferred to move from side to side along the puck.*” These comments indicate that, even if the behavior suggested by OAHG is optimal, it could inconvenience the users when it does not match their intentions. On the other hand, it can be seen that UPHG received higher subjective evaluations in that it did not harm the intention of the participants. P7 remarked on UPHG, “*When I wanted to stay still, I could stay still, and when I tried to move toward the puck, I was assisted,*” and five participants (P7, P10, P16, P17, and P20) also positively commented that “*The interference frequency of the guidance was appropriate.*” Meanwhile, there was also a skeptical view on UPHG. P13 mentioned that “*I tried to stop in front of the puck, but as I was assisted in the direction of movement, it moved further and touched the puck,*” indicating that an occasional inaccurate guidance could induce user’s inconvenience.

RQ4: *Can CombHG, which integrates OAHG and UPHG, complement each HG or provide better effects in users’ objective and subjective evaluations?*

CombHG significantly lowered the mean disagreement than OAHG and UPHG, without reducing other objective and subjective metrics. This implies that CombHG succeeded in assisting users with less interference by effectively combining OAHG and UPHG. In detail, CombHG received a positive comment from six participants (P1, P3, P6, P13, P16, and P20), the highest number along with OAHG, that “*The guidance helped to defend the puck toward our goal.*” Additionally, it also scored high subjective evaluations along

with UPHG. Several user comments can summarize the advantages of CombHG. P3 remarked on CombHG that “*The guidance was actively helping me in the defensive situation and allowing me attack freely in the offensive situation, which felt ideal for me to play.*” P19 also mentioned, “*I felt like the guidance was chasing the puck but giving me a choice to attack.*” Through these interviews, we judged that CombHG satisfied participants by providing adequate guidance in situations where it is necessary (e.g., defense against the fast-approaching puck) and ensuring participants’ autonomy in situations where various strategies are possible.

6.2 Generalization

The proposed OAHG and UPHG implementation methods (i.e., UT and UA) are applicable to other pHRI tasks if the optimal action and user behavior for the given task can be modeled as any neural network structure that receives the task state as input and outputs the action distribution (e.g., Figure 3). Because the UT method utilizes W_{unc} calculated using Equation (5) based on the output STD from the model, it is applicable to any neural network structure that outputs the STD of the action distribution. Because the UA method is based on the parameter update of the MAML with the user demonstration data (Equation (1)), it is applicable to any neural network structure that can be trained by the MAML algorithm. Therefore, our framework is not limited to the video game, but it can be easily extended to other pHRI tasks because data-driven modeling of the optimal action or user behavior has been successfully demonstrated in various HG scenarios (e.g., robot-assisted surgery [41, 55] and steering task [46]). To model the optimal action, we can either apply reinforcement learning in simulated environments, which is not limited to the self-play-based learning, or use the movement of skillful experts as in [5, 38, 51]. To model the user behavior, we can apply the MAML algorithm utilizing sufficient motion data from multiple users.

In addition, our CombHG implementation method (i.e., SC) is further general because it is not limited to the specific implementation methods of OAHG and UPHG. Any OAHG and UPHG implementations (e.g., the traditional OAHG method using a fixed reference path [15]) can be combined by the SC method because it requires only the similarity of guiding force vectors from OAHG and UPHG.

7 CONCLUSION

In this paper, we proposed deep learning-based novel implementation methods for OAHG and UPHG, applying a self-play-based reinforcement learning framework for OAHG and a meta-learning framework for UPHG to achieve their best performance. Further, we proposed CombHG that aimed to complement each HG type and provide better performance than OAHG and UPHG. In detail, the three proposed implementation methods (i.e., UT, UA, and SC) were applied to the given problem and demonstrated clear performance enhancement. Through the user study, we validated the assisting performance of each HG for users conducting a haptic task and investigated the difference in the user’s subjective evaluation for each HG. The user study results indicated that UPHG and CombHG elicited significantly better subjective scores than OAHG. In addition, CombHG exhibited a further decrease in user disagreement compared to OAHG and UPHG, without reducing any objective and

subjective scores. The comparison of each HG type based on our experimental analyses and user interviews can suggest the criteria for general HG design based on the aspects of HG that positively or negatively affect users. Considering that the generalization of the proposed HG implementation methods for other HG applications is straightforward, our findings are expected to contribute to the design of other HG-based pHRI applications beyond the video game environment considered in this study.

ACKNOWLEDGMENTS

This work was supported by the Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (NRF-2018R1D1A1B07043580).

REFERENCES

- [1] David A Abbink and Mark Mulder. 2010. Neuromuscular analysis as a guideline in designing shared control. *Advances in Haptics* (2010), 499–516.
- [2] David A Abbink, Mark Mulder, and Erwin R Boer. 2012. Haptic shared control: smoothly shifting control authority? *Cognition, Technology & Work* 14, 1 (2012), 19–28.
- [3] Trapit Bansal, Jakub Pachocki, Szymon Sidor, Ilya Sutskever, and Igor Mordatch. 2018. Emergent complexity via multi-agent competition. In *Proceedings of the Sixth International Conference on Learning Representations (ICLR)*.
- [4] Alessandro Bettini, Samuel Lang, Allison Okamura, and Gregory Hager. 2002. Vision assisted control for manipulation using virtual fixtures: Experiments at macro and micro scales. In *2002 IEEE International Conference on Robotics and Automation (ICRA)*. 3354–3361.
- [5] Jérémy Bluteau, Sabine Coquillart, Yohan Payan, and Edouard Gentaz. 2008. Haptic guidance improves the visuo-manual tracking of trajectories. *PLoS One* 3, 3 (2008), e1775.
- [6] Rolf Boink, Marinus M Van Paassen, Mark Mulder, and David A Abbink. 2014. Understanding and reducing conflicts between driver and haptic shared control. In *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. 1510–1515.
- [7] Alexander Broad, Ian Abraham, Todd Murphey, and Brenna Argall. 2020. Data-driven Koopman operators for model-based shared control of human-machine systems. *The International Journal of Robotics Research* 39, 9 (2020), 1178–1195.
- [8] Chia-Yu Chen, Yen-Yu Chen, Yi-Ju Chung, and Neng-Hao Yu. 2016. Motion guidance sleeve: Guiding the forearm rotation through external artificial muscles. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. 3272–3276.
- [9] Laura Marchal Crespo and David J Reinkensmeyer. 2008. Haptic guidance can enhance motor learning of a steering task. *Journal of Motor Behavior* 40, 6 (2008), 545–557.
- [10] Arnold W de Jonge, Jeroen GW Wildenbeest, Henri Boessenkool, and David A Abbink. 2015. The effect of trial-by-trial adaptation on conflicts in haptic shared control for free-air teleoperation tasks. *IEEE Transactions on Haptics* 9, 1 (2015), 111–120.
- [11] Jack Tigh Dennerlein, David B Martin, and Christopher Hasser. 2000. Force-feedback improves performance for steering and combined steering-targeting tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '00)*. 423–429.
- [12] Patrizia Di Campli San Vito, Gözel Shakeri, Stephen Brewster, Frank Pollick, Edward Brown, Lee Skrypchuk, and Alexandros Mouzakitis. 2019. Haptic navigation cues on the steering wheel. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. 1–11.
- [13] Anca D Dragan and Siddhartha S Srinivasa. 2013. A policy-blending formalism for shared control. *The International Journal of Robotics Research* 32, 7 (2013), 790–805.
- [14] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*. 1126–1135.
- [15] Benjamin AC Forsyth and Karon E MacLean. 2005. Predictive haptic guidance: Intelligent user assistance for the control of dynamic tasks. *IEEE Transactions on Visualization and Computer Graphics* 12, 1 (2005), 103–113.
- [16] Takashi Goto, Swagata Das, Yuichi Kurita, and Kai Kunze. 2018. Artificial motion guidance: an intuitive device based on Pneumatic Gel Muscle (PGM). In *The 31st Annual ACM Symposium on User Interface Software and Technology Adjunct Proceedings (UIST '18 Adjunct)*. 182–184.
- [17] Jordi Grau-Moya, Eduard Hez, Giovanni Pezzulo, and Daniel A Braun. 2013. The effect of model uncertainty on cooperation in sensorimotor interactions. *Journal of The Royal Society Interface* 10, 87 (2013), 20130554.

- [18] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*. 1861–1870.
- [19] JR Medina Hernández, Tamara Lorenz, and Sandra Hirche. 2015. Synthesizing anticipatory haptic assistance considering human behavior uncertainty. *IEEE Transactions on Robotics* 31, 1 (2015), 180–190.
- [20] Minsik Hong and Jerzy W Rozenblit. 2016. A haptic guidance system for computer-assisted surgical training using virtual fixtures. In *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. 2230–2235.
- [21] Amin Hosseini, Florian Richthammer, and Markus Lienkamp. 2016. Predictive haptic feedback for safe lateral control of teleoperated road vehicles in urban areas. In *2016 IEEE 83rd Vehicular Technology Conference (VTC Spring)*. 1–7.
- [22] Ioanna Iacovides, Anna L Cox, and Thomas Knoll. 2014. Learning the game: breakdowns, breakthroughs and player strategies. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems (CHI EA '14)*. 2215–2220.
- [23] Akira Kanazawa, Jun Kinugawa, and Kazuhiro Kosuge. 2019. Adaptive motion planning for a collaborative robot based on prediction uncertainty to enhance human safety and work efficiency. *IEEE Transactions on Robotics* 35, 4 (2019), 817–832.
- [24] Alex Kendall and Yarin Gal. 2017. What uncertainties do we need in bayesian deep learning for computer vision?. In *Proceedings of the 31st Conference on Neural Information Processing Systems (NIPS)*. 5574–5584.
- [25] Soheil Kianzad, Yuxiang Huang, Robert Xiao, and Karon E MacLean. 2020. Phasking on paper: Accessing a continuum of physically assisted sketching. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. 1–12.
- [26] Ravi Kuber, Wai Yu, and Graham McAllister. 2007. Towards developing assistive haptic feedback for visually impaired internet users. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. 1525–1534.
- [27] Roel J Kuiper, Dennis JF Heck, Irene A Kuling, and David A Abbink. 2016. Evaluation of haptic and visual cues for repulsive or attractive guidance in nonholonomic steering tasks. *IEEE Transactions on Human-Machine Systems* 46, 5 (2016), 672–683.
- [28] Thanh Mung Lam, Harmen Wigert Boschloo, Max Mulder, and Marinus M Van Paassen. 2009. Artificial force field for haptic feedback in UAV teleoperation. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* 39, 6 (2009), 1316–1330.
- [29] Robert W Lindeman, John L Sibert, Erick Mendez-Mendez, Sachin Patil, and Daniel Phifer. 2005. Effectiveness of directional vibrotactile cuing on a building-clearing task. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '05)*. 271–280.
- [30] Franck Mars, Mathieu Deroo, and Jean-Michel Hoc. 2014. Analysis of human-machine cooperation when driving with different degrees of haptic shared control. *IEEE Transactions on Haptics* 7, 3 (2014), 324–333.
- [31] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, and Georg Ostrovski. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529.
- [32] Hee-Seung Moon and Jiwon Seo. 2019. Prediction of human trajectory following a haptic robotic guide using recurrent neural networks. In *2019 IEEE World Haptics Conference (WHC)*. 157–162.
- [33] Hee-Seung Moon and Jiwon Seo. 2020. Dynamic difficulty adjustment via fast user adaptation. In *Adjunct Publication of the 33rd Annual ACM Symposium on User Interface Software and Technology (UIST '20 Adjunct)*. 13–15.
- [34] Hee-Seung Moon and Jiwon Seo. 2020. Sample-efficient training of robotic guide using human path prediction network. (2020). arXiv:2008.05054
- [35] Mark Mulder, David A Abbink, and Erwin R Boer. 2008. The effect of haptic guidance on curve negotiation behavior of young, experienced drivers. In *2008 IEEE International Conference on Systems, Man and Cybernetics (SMC)*. 804–809.
- [36] Mark Mulder, David A Abbink, and Erwin R Boer. 2012. Sharing control with haptics: Seamless driver support from manual to automatic control. *Human Factors* 54, 5 (2012), 786–798.
- [37] Carolina Passenberg, Antonia Glaser, and Angelika Peer. 2013. Exploring the design space of haptic assistants: The assistance policy module. *IEEE Transactions on Haptics* 6, 4 (2013), 440–452.
- [38] Carlos Jesús Pérez-del Pulgar, Jan Smisek, Victor F Munoz, and André Schiele. 2016. Using learning from demonstration to generate real-time guidance for haptic shared control. In *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. 3205–3210.
- [39] Johannes Pfau, Jan David Smeddinck, Ioannis Bikas, and Rainer Malaka. 2020. Bot or not? User perceptions of player substitution with deep player behavior models. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. 1–10.
- [40] Susanne Poeller, Max V Birk, Nicola Baumann, and Regan L Mandryk. 2018. Let me be implicit: Using motive disposition theory to predict and explain behaviour in digital games. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. 1–15.
- [41] Maura Power, Hedyeh Rafii-Tari, Christos Bergeles, Valentina Vitiello, and Guang-Zhong Yang. 2015. A cooperative control framework for haptic guidance of bimanual surgical tasks based on learning from demonstration. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*. 5330–5337.
- [42] Louis B Rosenberg. 1993. Virtual fixtures: Perceptual tools for telerobotic manipulation. In *Proceedings of IEEE Virtual Reality Annual International Symposium*. 76–82.
- [43] Edward Schmerling, Karen Leung, Wolf Vollprecht, and Marco Pavone. 2018. Multimodal probabilistic model-based planning for human-robot interaction. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. 1–9.
- [44] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. 2015. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*. 1889–1897.
- [45] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. (2017). arXiv:1707.06347
- [46] Dexter RR Scobee, Vicene Rubies Royo, Claire J Tomlin, and S Shankar Sastry. 2018. Haptic assistance via inverse reinforcement learning. In *2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. 1510–1517.
- [47] Mario Selvaggio, P Robuffo Giordano, F Ficuciello, and Bruno Siciliano. 2019. Passive task-prioritized shared-control teleoperation with haptic guidance. In *2019 International Conference on Robotics and Automation (ICRA)*. 430–436.
- [48] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, and Adrian Bolton. 2017. Mastering the game of go without human knowledge. *Nature* 550, 7676 (2017), 354–359.
- [49] Jan Smisek, Emmanuel Sunil, Marinus M van Paassen, David A Abbink, and Max Mulder. 2016. Neuromuscular-system-based tuning of a haptic shared control interface for UAV teleoperation. *IEEE Transactions on Human-Machine Systems* 47, 4 (2016), 449–461.
- [50] Shintaro Tada, Kohei Sonoda, and Takahiro Wada. 2016. Simultaneous achievement of workload reduction and skill enhancement in backward parking by haptic guidance. *IEEE Transactions on Intelligent Vehicles* 1, 4 (2016), 292–301.
- [51] Akiko Teranishi, Timothy Mulumba, Georgios Karafotias, Jihad Mohamad Alja'Am, and Mohamad Eid. 2017. Effects of full/partial haptic guidance on handwriting skills development. In *2017 IEEE World Haptics Conference (WHC)*. 113–118.
- [52] Zheng Wang, Zhanhong Yan, and Kimihiko Nakano. 2019. Comfort-oriented haptic guidance steering via deep reinforcement learning for individualized lane keeping assist. In *2019 IEEE International Conference on Systems, Man and Cybernetics (SMC)*. 4283–4289.
- [53] Zheng Wang, Rencheng Zheng, Tsutomu Kaizuka, Keisuke Shimono, and Kimihiko Nakano. 2017. The effect of a haptic guidance steering system on fatigue-related driver behavior. *IEEE Transactions on Human-Machine Systems* 47, 5 (2017), 741–748.
- [54] Dongbo Xiao and Roger Hubbard. 1998. Navigation guided by artificial force fields. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '98)*. 179–186.
- [55] Ehsan Zahedi, Javad Dargahi, Michael Kia, and Mehrdad Zadeh. 2017. Gesture-based adaptive haptic guidance: A comparison of discriminative and generative modeling approaches. *IEEE Robotics and Automation Letters* 2, 2 (2017), 1015–1022.