UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title
Modelling Competitive Human Action using Dynamical Motor Primitives for the Development of Human-Like Artificial Agents

Permalink
https://escholarship.org/uc/item/0mt3k725

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors
Ekdawi, Sarah
Patil, Gaurav
Kallen, Rachel W.
et al.

Publication Date
2022

Copyright Information

Peer reviewed

Modelling Competitive Human Action using Dynamical Motor Primitives for the Development of Human-Like Artificial Agents

Sarah Ekdawi (sarah.ekdawi@students.mq.edu.au)
School of Psychological Sciences, Macquarie University, Sydney, NSW 2109 Australia

Dr. Gaurav Patil (gaurav.patil@mq.edu.au)
School of Psychological Sciences, Centre for Elite, Performance, Expertise and Training, Macquarie University, Sydney, NSW 2109 Australia

Rachel. W. Kallen (rachel.kallen@mq.edu.au)
School of Psychological Sciences, Centre for Elite, Performance, Expertise and Training, Macquarie University, Sydney, NSW 2109 Australia

Professor Michael. J. Richardson (michael.j.richardson@mq.edu.au)
School of Psychological Sciences, Centre for Elite, Performance, Expertise and Training, Macquarie University, Sydney, NSW 2109 Australia

Abstract

With artificial intelligence technologies becoming commonplace today, enhancing the efficiency of human-artificial agent (AA) interactions has become increasingly important. A growing body of research has revealed how dynamic motor primitives (DMPs) of human perceptual-motor behavior can be used to create 'human-like' AAs, primarily focusing on cooperative tasks. Using *air hockey* as a representative task, the current experiment is the first part of a large study aimed at determining the utility of DMP-based models for developing 'human-like' competitive AAs. Participants played against a preliminary DMP model and the differences in their behaviors were analyzed. Based on these observed differences, a revised model is proposed, with preliminary results revealing that the new model exhibits behaviors more consistent with those of humans. A major implication of this work is that it presents a framework for creating 'human-like' AAs that capture the essential human decision and movement dynamics without requiring large human gameplay datasets.

Keywords: dynamic motor primitives; human-machine interaction; human behavior modelling; artificial intelligence

Introduction

Recent advances in machine learning methodologies and a speedy adoption of artificial intelligence (AI) systems have resulted in rapid and innovative uses of these systems in science, medicine, and industry (Dobbe et al., 2021; Jamilly et al., 2018). Although the capabilities of AI systems vary, contingent on underlying algorithms, training procedures, and training data (Shank et al., 2019), AI systems are now capable of performing a wide array of tasks, including speech recognition and translation, decision-making, image processing, visual perception, and financial forecasting (Norris, 2020). Human-AI systems are also becoming commonplace, with AI agents mediating social networks and experiences (Papadimitriou, 2016), how video games are played (Spronck et al., 2006), training and performing workplace activities (Kovacevic, & Radenkovic, 2020; Mihailidis et al., 2016), and seeking medical and clinical advice (Miller et al., 2020).

In many instances, the key to effective human-AI interaction is the ability of AI agents to exhibit 'human-like' behavior that can be predicted and easily understood by human co-actors (Nalepka et al., 2019). This is particularly true for physical or perceptual-motor human-AI interactions, where the behavioral actions of an AI agent can significantly influence the stability and effectiveness of human behavior and learning.

Indeed, a lack of understanding of how to best develop 'human-like' AI agents remains a significant barrier to the full adoption of AI systems within organizations requiring physical human-AI interaction (Lorica & Paco, 2019; Washburn et al., 2019). Importantly, recent research has suggested that one way to develop 'human-like' AI agents is to define the actions of such *artificial agents* (AA) using the same *dynamical motor primitives* (DMPs) that generatively approximate human actions (Patil et al., 2021; Nalepka et al., 2019). See (Patil et al., 2021) for an overview of using DMPs for modelling human action dynamics. Despite the robustness of the latter research, the effectiveness of the DMP approach for creating 'human-like' AAs has only been demonstrated using cooperative human-AA perceptual motor tasks (e.g., object pick and place tasks, multiagent herding, and collection tasks) (Carroll et al., 2019; Lamb et al., 2017, 2019; Nalepka et al., 2019; Patil et al., 2021). Furthermore, the general framework followed by previous research uses data collected from human co-actors to create and parameterize the DMP models, which are then validated with more human participants.

The aim of the current study was to explore the potential use of DMP controlled AAs in competitive human-AA contexts. The novelty of this methodology is that it does not rely upon human-human interactions to model AA behaviors, but it still uses the primitive models that capture the essential

human decision making and movement behaviors as a baseline to reveal the complexities of human behaviors. More specifically, we compared the behavior of human players to those of a preliminary DMP controlled AA to identify the differences/similarities between human and AA air hockey gameplay. This further allowed us to ascertain how a DMP model should be functionally and parametrically realized to best capture 'human-like' gameplay. Accordingly, we explored the general effectiveness of human and AA gameplay (i.e., wins versus losses) and the fundamental aspects of offensive and defensive gameplay (Chang et al., 2020; Chowdhury et al., 2018).

Method

Participants

Fifteen students (7 male, 8 female) with ages ranging from 18 to 35 years ($M = 20.73$, $SD = 4.82$) from an *Anonymous* University participated in this study for course credit. Two of the 15 participants were left-handed (13.33%), with the remaining 13 being right-handed (86.67%). All participants used their dominant hand to play and disclosed that their air hockey experience was minimal at best, consisting only of experience in an arcade environment.

Task Environment

Participants stood on the short side of a 2.2 m x 1.2 m table with a height of 0.81 m in a laboratory room of size 2.8 m x 4.1 m and completed this task in a virtual environment. The virtual environment consisted of a room resembling the laboratory room with a virtual table that was the same size as the real table. The physical table provided a solid surface on which participants could glide the controller (that tracks the position of the mallet) against. This is synonymous with the way a mallet would glide along an actual air hockey table (as seen in Figure 1). The virtual environment was created using Unity (Unity Technologies, CA, USA) and was presented to participants using a HTC Vive Pro virtual reality headset. The headset and the controller were tracked by 2 HTC Vive base stations which had a positional accuracy of 0.5 cm, latency of 11.11ms and communicated wirelessly with the computer running the air hockey environment.

Figure 1: Laboratory table that mirrored the dimensions of the air hockey table in the virtual reality environment

The VR environment, as can be seen in Figure 2, depicted circular mallets with knobs (blue for participants and red for the AA), a circular puck (green), the air hockey table and the scoreboard (presented above the table). The size of both the opponent's and participant's goals were 30 cm wide and lay in the center of the short end of the table on their respective sides.
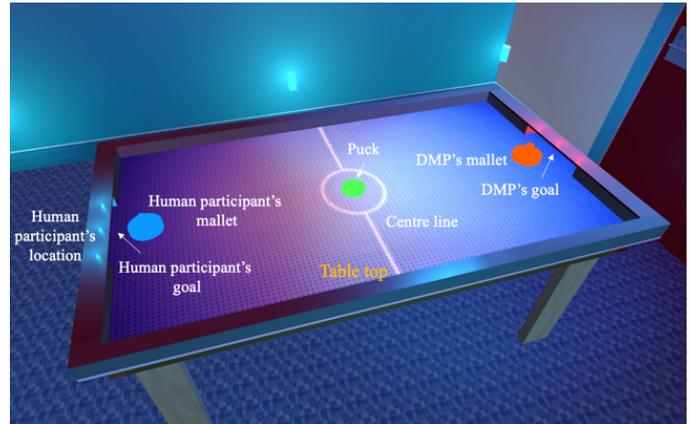

Figure 2: An illustration of the virtual reality air hockey environment.

Within the virtual environment, the mallet moved in alignment with the movements of the tracker controlled by the participant. The mallet had a diameter of 13 cm and height of 1.5 cm whilst the puck had a diameter of 10 cm and height of 1 cm. These dimensions were chosen such that they resembled the size of the mallet and puck in an ordinary game of air hockey. Participants' mallets could be used to hit the puck, and in situations where the puck hit the extremities of the table (table walls), the puck would bounce off it such that it resembled the table walls in a real air hockey environment. Additionally, the mallet could not travel past the walls in VR and the laboratory table had an adhesive strip which provided haptic feedback, resembling the wall hits by the virtual mallet. Players scored by hitting the puck into their opponent's goal.

Artificial Agent

The AA was modelled using the fundamental DMP equations used to model dynamic reaching behavior (Nalepka et al., 2017, 2019; Patil et al., 2020; Saltzman & Kelso, 1987). This DMP model employed discrete mass spring equations to determine the x- and z-position of the AA's mallet (i.e., one for x or forward/backward and one for z or left/right) of the respective form

$$\ddot{x} = -b\dot{x} - k(x - T_x) \qquad (1a)$$

$$\ddot{z} = -b\dot{z} - k(z - T_z) \qquad (1b)$$

where b and k represent the damping and the stiffness parameters and (T_x, T_z) is the target location the AA's mallet is attracted to (moves towards). Essentially, b (friction)

"resists" motion (for $b > 0$) and k is a restoring ("spring") force that induces motion (for $k > 0$) when the system is not at equilibrium (i.e., when the mallet is not at the target position). The parameter values for $b = 10$ and $k = 50$ were derived from previous research on comparable dynamic reaching/target selection tasks (Nalepka et al., 2019; Patil et al., 2020) which previously identified that this ratio of b/k generates general movement velocities (re-normalized to the specific dimensions of the air hockey table) like human arm/hand end effector movements (Lamb et al., 2019).

The target position (T_x, T_z) that the AA's mallet was attracted to was determined based on a critical x position (i.e., $x_{critical} = 0$; middle of the table) of the puck. When the puck was on the AA's side of the table, the AA's mallet was attracted to the puck's location, whereby

$$T_x = x_{puck} \text{ and } T_z = z_{puck}.$$

However, if the puck was on the human participant's side of the table, the AA's mallet was attracted to

$$T_x = -0.85 \text{ and } T_z = 0.25 \times z_{puck}.$$

This resulted in the AA moving towards the puck (offensive gameplay) when the puck crossed the center line into the AA's side of the table and the AA guarding the goal (defensive gameplay) when the puck was in the participant's playing area.

Procedure

All participants stood on the same side of the table chosen primarily out of convenience. Participants were given approximately two minutes to practice hitting the puck without their opponent moving from its starting position to feel comfortable in the VR environment and to familiarize themselves with the game mechanics. After they completed the practice block and indicated that they understood the objective of the task, participants played air hockey for a total of 15 minutes or until 75 goals were scored in total between the AA and participant. At the start of each game, the puck appeared on the side of the player who lost the previous round (randomly assigned for the first game).

Mallet Position Measures

Positions of participants and AA's mallets and the puck were recorded during all trials at 50Hz. In addition, every time the puck came into contact with any object (e.g., mallet, walls, or goals) the name of the object was recorded. The x (forward/backward) and z (left/right) coordinates of players' mallets, at the time of each hit, were extracted and split by whether the player or their opponent hit the puck (whether they were in offensive or defensive positions). This positional data was used to calculate the mean mallet positions in both directions during offensive and defensive positions in order to discern the general strategies used by the players. The positional data was additionally grouped into bins of 10cm each in the x and z directions to further identify the areas of the table frequented by both players while attacking and defending.

Mallet Velocity Measures

In addition to the positions, the velocities of players' mallets were recorded throughout gameplay. Similar to the positions, the velocity of their movements were split by whether they were in defensive or offensive positions. Average velocity in each direction was calculated to identify the differences in movement behaviors between human participants and the AA. In addition, the distribution of the velocities in both directions were calculated to get an overall idea of how players modulated their velocities throughout gameplay.

Puck Hit Angle Measures

At the instant when the puck was hit by the mallet, the direction of the puck bouncing off was determined as the angle at which players hit the puck. These angles were analyzed to investigate if players preferred hitting the puck in a certain direction when they were in a particular area of the table. The angle was 0 if they hit the puck straight ahead, -90 if they hit the puck straight towards the right edge of the table, 90 if they hit the puck straight towards the left edge, and 180 if they hit it straight back (as can be seen in Figure 4). The distribution of the puck angles were calculated for all the hits that happened in the left side ($-.5 < z < -.3$), right side ($.3 < z < .5$), and middle ($-.1 < z < .1$) of the air hockey table in the z direction and high ($x < -.5$) and low ($x >= -.5$) areas in the x direction (refer to Figure 3).
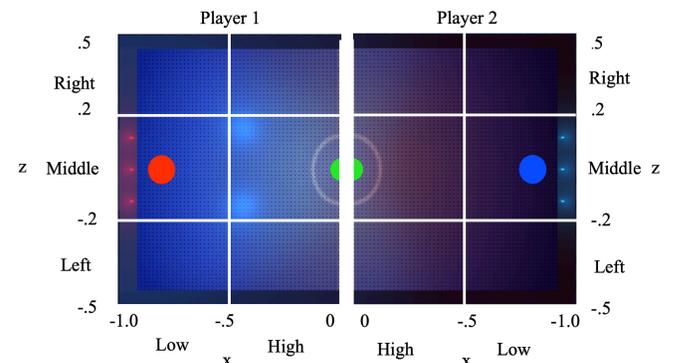

Figure 3: (*top*) Division of air hockey table for angular analysis. Puck angle hits were split by positional data, such that they were separated according to where hits occurred on the air hockey table. Correspondingly, these puck angle hits are displayed separately for puck hits when the player/agent positions were x < -.5 (low) and x >= -.5 (high), where x = -1 is the player/agents end of the table and when player/agent positions were z < -.2 (left side hits), z >=.2 (right side hits), and -.2 < z < .2 (middle of table hits). Data for player 2 was remapped from 0 < x < 1 to -1 < x < 0 before analysis to remain consistent with human data and allow for further examination of the differences in gameplay between these players. (bottom).

Results

As the behavior of the AA was not independent of participant (i.e., the AA behavior was coupled to each participant, and therefore participant specific), paired sample t-tests and Wilcoxon's Signed-Ranks tests (the non-parametric equivalent) were the primary methods of analysis, the latter employed when the assumption of normality was violated. Although numerous variables were assessed, no adjustment was made to the alpha level (set at $p = .05$) due to the exploratory nature of the research.

Game performance

Fourteen participants (93.33%) beat the AA, with only one participant failing to win most rounds. The average proportion of round wins for human participants ($M = .58$, $SD = .08$) was significantly more than chance, at .5 ($t(14) = 4.10$, $p = .001$, $d = 2.12$), indicating that human participants ($M = 37.20$, $SD = 9.95$) won significantly more rounds of air hockey than their AA competitor ($M = 27.07$, $SD = 7.41$).
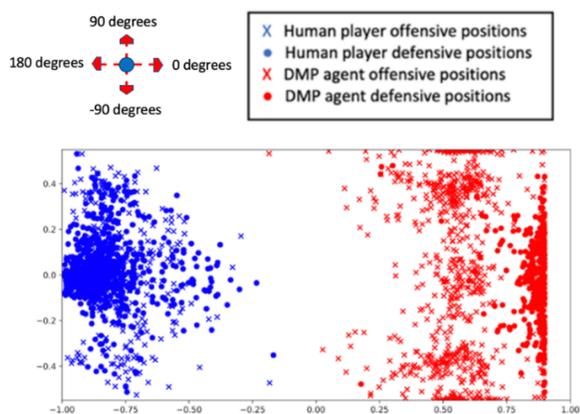

Figure 4: Exemplar scatter plot of a human player and the AA's offensive (puck hit) and defensive mallet positions when the puck was hit (see Figure 6 for other representative human examples).

Player Position

A representative example of the offensive and defensive positions of human players and the AA, and the constituents of offensive play (i.e., players' mallet position, velocity and angle the puck moved in when players hit the puck) are displayed in Figures 4 and 5 respectively.

Offensive Play As can be seen from an inspection of the positional distributions in Figure 5, the x-coordinate for the average position of the AA's mallet ($M = -.56$, $SD = .01$) was significantly different to the position of the human participant's mallet ($M = -.79$, $SD = .04$) when hitting the puck ($t(14) = -24.99$, $p < .001$, $d = 7.89$). In short, human players kept their mallet close to their goal rather than toward the center of the table while the AA tended to spend most of their time further from their goal. This suggests that human

players were more comfortable attacking from a position closer to their goal.

As seen in Figure 5, a bimodal distribution existed for the AA and a normal distribution for human players. Hence, for analysis of the positional distributions in the z-axis, the axis was split into three separate bins, such that the distribution when player/agent positions were $-.5 < z < -.3$ (left side hits), $.3 < z < .5$ (right side hits), and $-.1 < z < .1$ (middle of table hits) were analyzed separately. The proportion of hits within each third of the air hockey table, representative of the positional distribution of player's mallets, were measured as the amount of hits that occurred in each third divided by the amount of hits overall.

A significant difference between the proportion of hits by human players ($M = .34$, $SD = .08$) and the AA ($M = .12$, $SD = .02$) was found in the middle third of the air hockey table ($-.1 < z < .1$), $t(14) = 10.31$, $p < .001$, $d = 3.77$. The odds of the human participant's mallet occupying space in the middle of the air hockey table was 2.76 times that of the AA's mallet. There was also significant difference between the proportion of human player hits ($M = .17$, $SD = .05$) and AA hits ($M = .33$, $SD = .03$) within the left third of the air hockey table ($-.5 < z < -.3$; $t(14) = -11.96$, $p < .001$, $d = -3.88$), and the proportion of human player hits ($M = .13$, $SD = .06$) and AA hits ($M = .32$, $SD = .04$) within the right third of the air hockey table ($.3 < z < .5$; $t(14) = -12.33$, $p < .001$, $d = -3.73$). Indeed, the odds of the AA's hitting the puck on the left or right of the air hockey table was 1.9 times that of the human participant's, with human players tending to remain close to their goal location, whereas the AA spent time on either side of the air hockey table.

Defensive Play No significant differences in the x- or z-positions of the human and AA were found, with both agents remaining close to their goal location during defensive period of play and the results of this analysis are not presented here is favor of conserving space.

Mallet Velocity

Offensive Play The AA's mallet movements were significantly faster, in the x-axis (forward and backward), when the AA hit the puck ($M = 5.56$, $SD = .34$) compared to the human participant's ($M = 1.24$, $SD = .39$), $z = -3.41$, $p < .001$, $d = 11.81$. The AA's mallet also moved significantly faster in the z-axis (left and right), when the AA hit the puck ($M = 3.34$, $SD = .29$) compared to the human participant's ($M = .21$, $SD = .03$), $z = -3.41$, $p = .001$, $d = 15.18$. Furthermore, it can be observed from the distribution of velocities while hitting the puck (see Figure 5) that participants used a wider range of velocities as compared to the AA.

Defensive play The AA's mallet moved significantly faster, in the x-axis (forward and backward), when their opponent hit the puck ($M = .69$, $SD = .23$) compared to the human participant ($M = .29$, $SD = .11$), $t(14) = 5.07$, $p < .001$, $d = 2.22$. The AA's mallet also moved significantly faster, in the z-axis (left and right), when their opponent hit the puck ($M = $
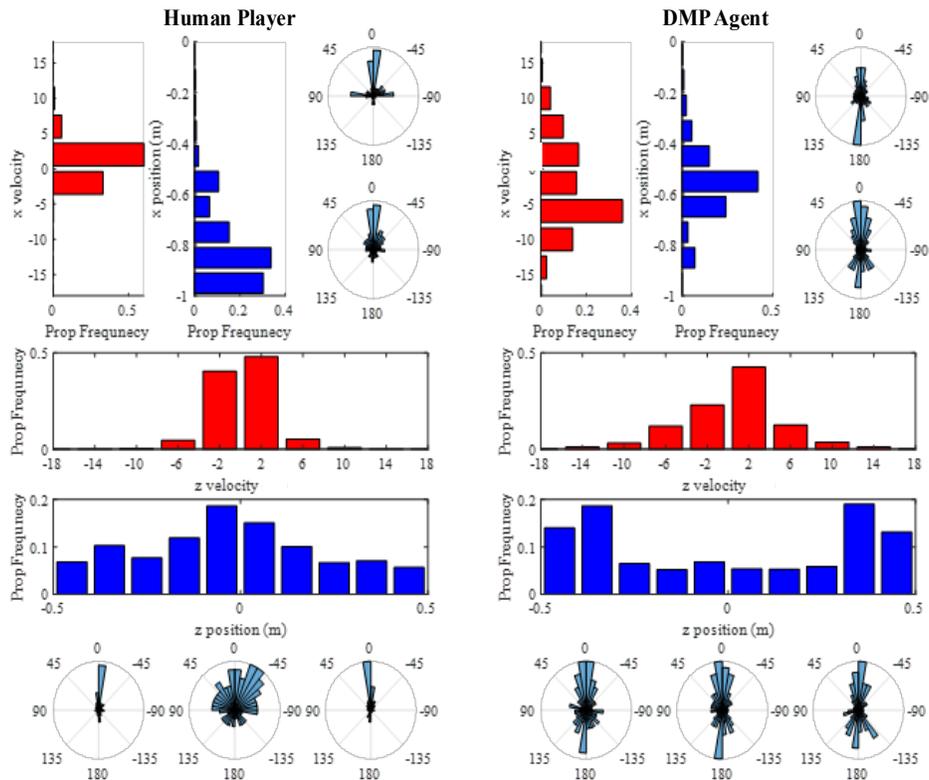
Figure 5: The average proportional frequencies of (*left panel*) human player and (*right panel*) AA (*x, z*) positions (blue), hit velocities (red) and angles when hitting puck (i.e., offensive play). *x* corresponds to the forward-backward table direction. *z* corresponds to the left-right table direction. (*top-right of each panel*) the puck hit angles are displayed separately for puck hits when the player/agent positions was x < -.5 and x >= -.5, where x = -1 is the player/agents end of the table. (*bottom of each panel*) the puck hit angles are displayed separately for puck hits when the player/agent positions was z < -.2 (left side), z >= .2 (right side), and -.2 < z < .2 (middle of table).

1.53, *SD* = .21) compared to human players (*M* = .14, *SD* = .04), *t* (14) = 28.55, *p* < .001, *d* = 9.20.

Puck Hit Angles

Strategies employed by players within an air hockey environment can be differentiated by the angle at which the puck is hit. As displayed in Figure 5, these puck angle hits were split by positional data such that they were separated according to where hits occurred on the air hockey table.

Overall, there were very few differences in the puck hit angles between the human players and the AA. The only significant difference was that puck hit angles were more variable for the human player compared to the AA when the puck was hit in the -.5 > x < .5 area of the air hockey table *(z = -3.07, p = .001)*. Note also that the AA did hit the puck back toward their own goal more than human players (see Figure 5), although this difference was not significant.

Discussion

The main aim of this study was to assess the 'human-like' nature of an DMP controlled AA by examining motor movements employed when playing against a human opponent in a competitive air hockey environment. Overall, the analysis of revealed that human players were better than the AA and that the offense gameplay of the AA differed

from that of human participants. However, there were no significant differences in the positional play of human players and the AA, indicating that the current DMP model could particle replicate "human-like" gameplay (i.e., defending near the goal location).

Key Behavioral Differences

The analysis of the human player's and AA' positions and velocities revealed three key differences in gameplay.

First, that human players kept their mallet closer to their own goal location while hitting the puck as compared to the AA. Note that the DMP model was configured to initiated offense gameplay as soon as the puck crossed the center line (see equations 1a and 1b).

Second, the positions of players while hitting the puck in the z-direction was also significantly different, in that, the human players spent the majority of their time in the middle of the table and the AA spent their time on both the left and right sides of the table.

Third, the AA's mallet movements were significantly faster than that of the human players in both the x- and z-directions.

The discrepancies in the offensive positions and puck hits of the human players and the AA indicated that the threshold used for switching the DMP model between offensive and defensive play needed to be refined. Instead of using a
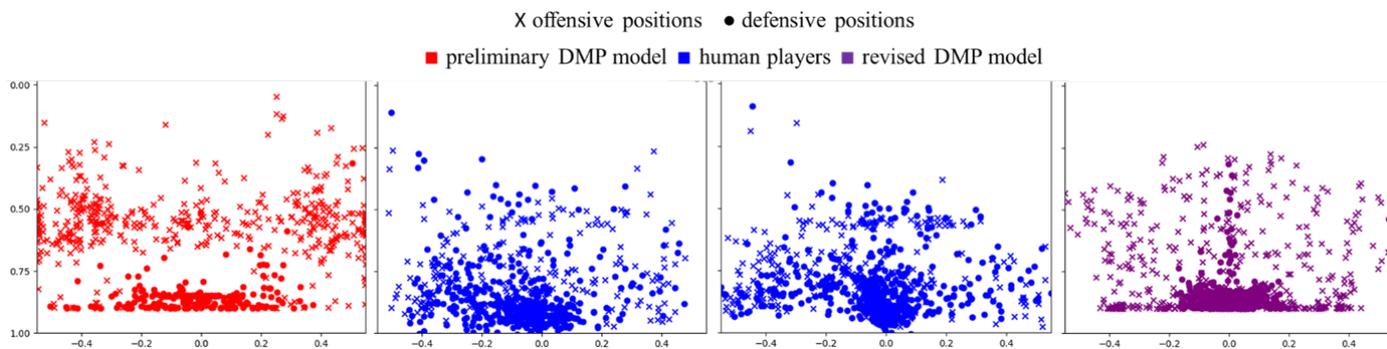
Figure 6: Exemplary scatter plots comparing offensive and defensive positions of the preliminary DMP model, human players, and revised DMP model

constant threshold (the center line in this case), a more adaptive threshold is proposed that not only relies on the puck's x position but also on its x velocity and z position. More specifically, we propose that the critical threshold should be determined at any given time point by

$$x_{critical} = (d_x - \gamma(\dot{x}_{puck}^2 + z_{puck}^2))\theta$$

where d_x is the baseline $x_{critical}$ value, which based on the participants data collected here is expected to be between $|.5|$ and $|.7|$, $\theta = \pm 1$ depending on whether the player is on the positive or negative side of the table, and γ is the fixed parameter that determines the strength of the DMP model's defensive/offensive tendency. In short, an AA controlled by this new DMP model is more likely to venture away from the goal to hit the puck when the puck is travelling slower and is not coming directly towards the goal.

Furthermore, to address the differences in velocities, the ratio of b/k should be reduced by either increasing damping (i.e., b = 20 instead of 10) and/or reducing the value of the stiffness (i.e., k = 35 instead of 50). This would result in slower DMP movement trajectories and may also result in less puck over shooting, which appeared to be why the AA tended to hit the puck back towards its own end of the table more often than the human participant.

Finally, the analysis of the puck hit angles also revealed significant differences in the variance of the angles with which players hit the puck between the AA and human players. Specifically, the variance of angle hits by humans was significantly larger than that of the AA. This can be addressed by adding stochasticity to the DMP models movements and future work will investigate the appropriate amount and type of noise that leads to more 'human-like' behaviors.

Although the initial plan was to test the modified model with human participants, lockdowns in 2021 due to the ongoing COVID-19 pandemic have delayed this data collection to 2022. However, scatter plot of positions when the puck was hit (similar to Figure 4) from a pilot trial with the revised DMP model are displayed in Figure 6 alongside the preliminary DMP model and exemplary participant data from playing against the preliminary DMP model. It can be observed that the offensive and defensive hits by the revised DMP model have a larger overlap as compared to the preliminary model and the distribution of the hits better captures the behaviors observed in participants' data. Furthermore, in a competitive task context like air hockey, human behavior can significantly change due to the skill level and the behavior of their opponent and further testing is required to unravel these effects using the revised DMP model. Additionally, future research may consider the height, wingspan, and reaching length at the table for human players, such that examining these human features can allow for the determination of any invariant structure between these physiological constraints and the distribution of player's mallet positions (Babajanyan et al., 2022).

Conclusion

Using a competitive air hockey environment, the current study was able to examine differences in human and a preliminary DMP controlled AA. A strength of this study was that significant differences in gameplay, between human and DMP competitive behaviors, were identified and quantified, and hence, a re-parametrized DMP model was proposed. Furthermore, no data from human-human interactions in the current task context was utilized to construct or parameterize the DMP model and the baseline model was formulated only based on the preconceived notions of human actions and decision making identified from other tasks. This can be advantageous in scenarios where human-human data is difficult to record as compared to human-AI interactions. Furthermore, having a model whose behavior is closer to that of human players opens avenues for using hybrid deep reinforcement learning-DMP approaches to model the AA behavior (Patil et al., 2021). The advantage of the latter methodology is that it can create AAs with expert human level or even better performance while being scaffolded by the essential human action behaviors. These 'super-human' AAs can further be used as a tool for skill learning by novice or expert human players.

References

Babajanyan, D., Patil, G., Lamb, M., Kallen, R, W., Richardson, M, J. "I Know Your Next Move: Action Decisions in Dyadic Pick and Place Tasks." Proceedings of the Annual Meeting of the Cognitive Science Society, 2022. [Accepted, In press]

Chang, C. L., Chen, S. T., Chang, C. Y., & Jhou, Y. C. (2020). Application of machine learning in air hockey interactive control system. *Sensors (Basel, Switzerland)*, *20*(24), 7233. https://doi.org/10.3390/s20247233

Chowdhury, S. S., Mashbu, R., Shaan, A. A., Al Ashfaq, K., Niloy, F. M., & Rahman, R. M. (2018). Perceiving attributes of game AI using fuzzy logic. *Modern Approaches for Intelligent Information and Database Systems*, 161–172. Springer International Publishing. https://doi.org/10.1007/978-3-319-76081-0_14

Dobbe, R., Krendl Gilbert, T., & Mintz, Y. (2021). Hard choices in artificial intelligence. *Artificial Intelligence*, 103555. https://doi.org/https://doi.org/10.1016/j.artint.2021.103555

Jamilly, M., Nagle, F., & Ross, C. (2018). Limitations of AI. *ITNow*, *60*(1), 16–17. https://doi.org/10.1093/itnow/bwy006

Kovacevic, A., & Radenkovic, S. D. (2020). SAWIT-Security awareness improvement tool in the workplace. *Applied Sciences*, *10*(9), 3065. https://doi.org/10.3390/app10093065

Lamb, M., Kallen, R. W., Harrison, S. J., Di Bernardo, M., Minai, A., & Richardson, M. J. (2017). To pass or not to pass: Modeling the movement and affordance dynamics of a pick and place task. *Frontiers in Psychology*, *8*, 1061. https://doi.org/10.3389/fpsyg.2017.01061

Lamb, M., Nalepka, P., Kallen, R. W., Lorenz, T., Harrison, S. J., Minai, A. A., & Richardson, M. J. (2019). A hierarchical behavioral dynamic approach for naturally adaptive human-agent pick-and-place interactions. *Complexity (New York, N.Y.)*, *2019*, 1–16. https://doi.org/10.1155/2019/5964632

Lorica, B., & Paco, N. (2019). *AI adoption in the enterprise: how companies are planning and prioritizing AI projects in practice* (First edition). O'Reilly Media.

Mihailidis, A., Melonis, M., Keyfitz, R., Lanning, M., Van Vuuren, S., & Bodine, C. (2016). A nonlinear contextually aware prompting system (N-CAPS) to assist workers with intellectual and developmental disabilities to perform factory assembly tasks: System overview and pilot testing. *Disability and Rehabilitation: Assistive Technology*, *11*(7), 604–612. https://doi.org/10.3109/17483107.2015.1063713

Miller, S., Gilbert, S., Virani, V., & Wicks, P. (2020). Patients' utilization and perception of an artificial intelligence–based symptom assessment and advice technology in a british primary care waiting room: Exploratory pilot study. *JMIR Human Factors*, *7*(3), e19713–e19713. https://doi.org/10.2196/19713

Nalepka, P., Kallen, R. W., Chemero, A., Saltzman, E., & Richardson, M. J. (2017). Herd those sheep: Emergent multiagent coordination and behavioral-mode switching. *Psychological Science*, *28*(5), 630–650. https://doi.org/10.1177/0956797617692107

Nalepka, P., Lamb, M., Kallen, R. W., Shockley, K., Chemero, A., Saltzman, E., & Richardson, M. J. (2019). Human social motor solutions for human-machine interaction in dynamical task contexts. *Proceedings of the National Academy of Sciences - PNAS*, *116*(4), 1437–1446. https://doi.org/10.1073/pnas.1813164116

Norris, M. (2020). *The value of AI-powered business intelligence* (First edition.). O'Reilly Media.

Papadimitriou, A. (2016). The future of communication: Artificial intelligence and social networks.

Patil, G., Nalepka, P., Kallen, R. W., & Richardson, M. J. (2020). Hopf bifurcations in complex multiagent activity: The signature of discrete to rhythmic behavioral transitions. *Brain Sciences*, *10*(8), 536. https://doi.org/10.3390/brainsci10080536

Patil, G., Nalepka, P., Rigoli, L., Kallen, R. W., & Richardson, M. J. (2021). Dynamical perceptual-motor primitives for better deep reinforcement learning agents. *Advances in Practical Applications of Agents, Multi-Agent Systems, and Social Good. The PAAMS Collection*, 176–187. https://doi.org/10.1007/978-3-030-85739-4_15

Saltzman, E., & Kelso, J. A. S. (1987). Skilled actions: A task-dynamic approach. *Psychological Review*, *94*(1), 84–106. https://doi.org/10.1037/0033-295X.94.1.84

Shank, D. B., Graves, C., Gott, A., Gamez, P., & Rodriguez, S. (2019). Feeling our way to machine minds: People's emotions when perceiving mind in artificial intelligence. *Computers in Human Behavior*, 98, 256–266. https://doi.org/10.1016/j.chb.2019.04.001

Spronck, P., Ponsen, M. J. V, Sprinkhuizen-Kuyper, I. G., & Postma, E. O. (2006). Adaptive game AI with dynamic scripting. *Machine Learning*, *63*(3), 217–248. https://doi.org/10.1007/s10994-006-6205-6

Washburn, A., Kallen, R. W., Lamb, M., Stepp, N., Shockley, K., & Richardson, M. J. (2019). Feedback delays can enhance anticipatory synchronization in human-machine interaction. *PloS One*, *14*(8), e0221275–e0221275. https://doi.org/10.1371/journal.pone.0221275