Artificial Intelligence and Social Computing, Vol. 122, 2024, 62–74
https://doi.org/10.54941/ahfe1004642

Positive Interactions With Intelligent Technology Through Psychological Ownership: A Human-in-the-Loop Approach

Bianca Dalangin[1], Stephen Gordon[1], and Heather Roy[2]

[1]DCS Corp., Alexandria, VA, USA
[2]DEVCOM Army Research Laboratory, Aberdeen Proving Ground, MD, USA

ABSTRACT

While human-agent interaction is intended to ease daily and critical burdens on human operators, issues such as trust, lack of transparency, and system performance often negatively impacts the process to yield sub-optimal outcomes. Here, we propose a human-in-the-loop approach, in which users train an AI, as a potential avenue to remedy this complex problem. We use Tetris® as a use case and require participants to provide trial-by-trial inputs to train the AI model. Improvements in trust correlated with increased satisfaction levels during the training process but not final AI performance. Users' preference for their trained AI, compared to a pre-trained AI, demonstrated increased improvements in trust. Personality and AI literacy did not affect these relationships. Results suggest positive perceptions towards AI systems can be elicited through psychological ownership pathways. We discuss how users' involvement in constructing the system may influence ownership giving rise to positive human-agent interactions.

Keywords: Human-agent interaction, Human-in-the-loop, Psychological ownership, Trust towards systems, Individual differences

INTRODUCTION

Research in human-agent interaction (HAI) has rapidly grown over recent decades. However, there remain several fundamental problems related to the integration of these two systems. Some users may be slow to trust or to adopt new technology (Efendić et al., 2020) while others may become too reliant and use technology when it is not optimal to do so (Vasconcelos et al., 2022). At the same time, AI systems that aim to produce optimal outcomes and purely follow data-driven rules can overwhelm or intimidate the user, leaving the user distrustful towards the system and dissuaded to use it. Users may not always understand what goes on "under the hood" of an AI system (Miller, 2023). Balancing the cost of verifying an AI's decision with task complexity may be important for overcoming overreliance (Vasconcelos et al., 2022). However, to push the burden of adaptation on either human or AI alone risks overwhelming the human, producing AI models that favor transparency over optimality and destabilize system performance (e.g., constantly adapting precisely tuned complex systems that were designed to operate within specified performance ranges). HAI solutions must engender trust and end-user acceptance, reduce AI system overreliance (or misuse), *but not sacrifice performance or stability of the AI system or push the burden of adaptation onto any single component*.

Here, we investigate a novel approach to address these issues by enabling the end-user to play a role in the development and construction of the AI system, with the caveat that this construction process has been scaffolded to maintain AI performance within pre-specified boundaries. In other words, the AI system is largely "pre-cooked" such that many parameters within the system have been optimized for the task at hand while a few key parameters are left blank for the end-user to fill. It is important to note that final system performance is not completely invariant to the user and training is not an illusion. Neither is the system simply learning the preferences of the user (e.g., recommender system). Rather, analogous to a pre-cooked meal kit, users who follow laid-out instructions will have a largely predetermined outcome (e.g., like a tasty meal); users who deviate largely from it (overcook or leave out the ingredients) could still alter system performance.

Humans' role in interacting with their intelligent systems can positively shape their post-adoption behaviors and attitudes towards those systems (Delgosha & Hajiheydari, 2021; LaCroix et al., 2023). This is based on the theory of Psychological Ownership (PO), in which individuals perceive an object as "MINE", eventually leading to satisfaction, commitment, and post-adoption behaviors (Pierce et al., 1991). PO is a construct that describes people and their perceptions of ownership to a material or immaterial item (Pierce et al., 1991). Studied initially in organizational and consumer psychology, PO develops through experienced control, intimate knowledge of the object, and self-investment (Pierce et al., 1991). Increased PO has been found to predict commitment (Han et al., 2010), user self-esteem (Lee & Suh, 2015), and satisfaction (Lee & Suh, 2015). For example, in a study that assessed users' interactions with consumer robots, Delgosha and Hajiheydan (2021) found that users' perceived control with the robot and self-investment in the robot predicted PO, suggesting that people who dedicate their time and energy into a system that they perceive they can control shapes their perceptions of ownership. Furthermore, users' measures of PO and trustworthiness in their robots predicted users' willingness to explore new technologies and pay more attention to those systems (Delgosha & Hajiheydari, 2021). As people are found to have an innate desire to possess, perceptions of ownership emulate a symbolic extension to the self (Pierce et al., 1991). This in turn engenders the endowment effect, or the idea that people place more psychological and economic value on the item they perceive ownership with (Thaler, 1980).

In a technology context, playing a part in constructing and carrying out the actions of a technological system reduces uncertainty and improves predictability for the user, leading to positive post-adoption outcomes, such as trust towards the system (Delgosha & Hajiheydari, 2021) and reuse of the system (Lin et al., 2021). Prior work has shown that designing a robot, through customization, has an effect on intimate knowledge, self-investment, and perceived control, leading to higher PO and higher levels of affective trust (Lacroix et al., 2023). Other studies exploring human-technology interactions have measured factors of perceived control through questionnaires (Delgosha & Hajiheydari, 2021), levels of self-customizations in their robots (Lacroix et al., 2023), and automated vs. manual manipulations of the system (Jörling et al., 2019).

Individual traits have also been found to influence the relationship between PO and post-adoption behaviors. Essig and Soparnot (2021) found that certain personality traits (i.e., extraversion) and their facets (e.g., those who show altruism) from the Big Five Inventory (John et al., 1991) are related to PO towards their company. A study that explored trust in AI and robots found that openness and exposure to robots positively predicted trust, while conscientiousness negatively predicted trust towards a robot in a trust game (Oksanen et al., 2020). In a separate study, conscientiousness was positively associated with perceived and actual technology use, while neuroticism had a negative association with technology use (Barnett et al., 2015). Notably, these relationships were not mediated by expressed intentions to use the system (Barnett et al., 2015).

To our current knowledge, there are virtually no studies on the relationship between PO and AI usage or acceptance. Given the relationship between PO and post-adoption behaviors in other domains, we believe that the development of PO from training an AI system can provide an alternative path for user satisfaction, trust, and re-use. Due to the paucity of research in this area, we also believe such work can address a critical scientific gap as AI tools become an increasingly pervasive presence in our everyday lives. Here, we elicit the antecedent factors of PO by allowing users to train their own AI. By giving users a more active role in shaping the system's decisions, we predict that users will be more satisfied with the system and will prefer using their system rather than a pre-trained system. We also investigate how individual differences such as AI literacy (Pinski & Benlian, 2023) and constructs of the Big Five personality traits (John et al., 1991) relate to these relationships.

1) How does being in the loop affect end-users' attitudes post-adoption behaviors towards pre-cooked systems?
2) Do individual differences (e.g., AI Literacy and Big 5) moderate this relationship?

METHODS

Participants

This study (ARL 23-054) was approved by and conducted in compliance with the U.S. Army Research Laboratory (ARL) accredited Institutional Review Board (IRB) and Human Research Protection Program (HRPP). All participants were 18 years of age or older, fluent in English, had corrected to normal vision, no color blindness, and no significant brain trauma/injuries in the last three months. Furthermore, because this was an online study, participants were required to meet certain system requirements to participate in the study (e.g., did not have a Mac). Participants were recruited via Prolific (Prolific, 2023) and were compensated \$10.00 for completing the pre-screener and an additional \$30.00 for completing the main study. In both the pre-screener and main study, participants were given information sheets that described the study with any associated risks or benefits.

Data from up to 370 participants in total was collected for the Horizontal AI and Vertical AI training conditions, examined in this paper. Considering the Horizontal AI condition, data from 355 participants, including 99 females, with mean (M) and ± standard deviation (SD) of age in years (M = 37.06, SD = 10.05), and 256 males (M = 36.62, SD = 8.73) was included in analyses. Data from the Vertical AI condition consisted of data from 349 participants, including 98 females (M = 37.27, SD = 10.01) and 251 males (M = 36.64, SD = 8.74). Data was excluded from participants with incomplete data sets for these conditions.

Procedure

Potential participants first completed a prescreener that determined their eligibility and contained the following questionnaires: AI Literacy Scale (Pinski & Benlian, 2023), Big Five Inventory (BFI) (John et al., 1991), Flexible Thinking in Learning Scale (Barak & Levenberg, 2016), Mental Rotation Test (Vandenberg & Kuse, 1978), and a short gaming inventory. Then, participants played a short trial of classic Tetris®. Eligible participants with verified data were invited to volunteer for the main study.

For the main study, participants answered demographic and experience questionnaires related to AI and technology use. Then, participants completed six conditions of the game Tetris® and answered some questions regarding their/the AI's performance and strategies. After completing all game conditions, participants completed the Intrinsic Motivation Inventory (Ryan, 1982) and the Depression, Anxiety, and Stress Scale (Lovibond & Lovibond, 2011). For the purposes of this paper, analyses explored four of the conditions and four questionnaires, described below.

AI Training Conditions

Participants completed the Horizontal AI and Vertical AI conditions. Prior to these conditions, participants completed Horizontal and Vertical Baseline conditions, respectively, to familiarize themselves with the game without the AI (see Figure 1). Each condition took 10 minutes. For the Horizontal Baseline and Horizontal AI condition, the game consisted of classic Tetris® where 10 continuous blocks must be stacked horizontally to clear a row and earn points. For the Vertical Baseline and Vertical AI conditions, this relationship was flipped and 10 continuous blocks had to be stacked vertically to clear that portion of a column and earn points. This was a novel and more difficult way to play Tetris® which required participants to adapt to the new parameters of the game quickly to maximize performance. In each condition, participants' goal was to maximize performance by clearing as many rows (or columns) as possible.

In both AI training conditions, participants had a maximum of 20 minutes to train each AI. Participants were able to preview their AI's performance as desired throughout training by speeding up the game and letting the AI play on its own. This capability could be toggled on or off by the participant with a single key press. At the end of the 20 minutes or when the participant indicated that they were ready, participants were able to submit their trained AI.

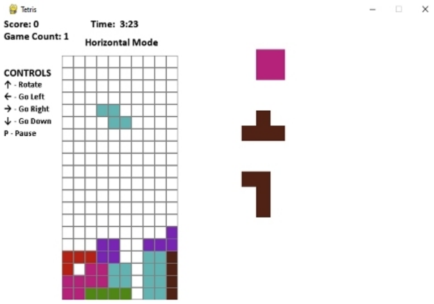

Figure 1: A participant's view of the horizontal baseline condition of Tetris®.

TAMER AI Assessment

The AI algorithm used for this work was 'Training an Agent Manually via Evaluative Reinforcement (TAMER)' originally designed by Knox and Stone (2008). TAMER is a framework for enabling humans to rapidly adapt an untrained AI system; however, in order for the "untrained" system to adapt within a limited number of human trials, a significant amount of initial knowledge must be imparted in the system. In the original work, which utilized Tetris® as the testbed, this involved a highly optimized, and pre-computed, feature vector for the task (e.g., Horizontal Tetris®) (Knox & Stone, 2008). This feature engineering significantly reduces the amount of "free parameters" in the system that must be adapted by the end-user, thus reducing the amount of learning required and, at the same time, reducing the performance variability of the final trained system. Here, we use a variant of the original feature representation used by Knox and Stone (2008) and introduced by Bertsekas and Tsitsiklis (1996). In total, there were 47 free parameters of the system, which were all initialized to 0 prior to training. We argue this approach creates a "pre-cooked" system in which the AI was, essentially, optimized for the task at hand but required user interaction to finalize its weight vector.

To train the AI, the users observed each move, or piece placement, performed by the system. At the end of each move, while the next move was being performed, the users had the option to ENCOURAGE or DISCOURAGE the system by pressing an appropriate key. This feedback was applied to the previous move and used to perform a single-trial gradient update of the AI weights (Knox & Stone, 2008). Notably, the precomputed feature vector used for this approach was optimized for horizontal Tetris® and, thus, suboptimal for the vertical version.

Once the AI was trained, the participants submitted the final weights for evaluation. These weights were used offline in post-processing to evaluate the performance of each participant's trained AI. This evaluation process included allowing each trained AI instance to play the game (horizontal or vertical) for 1000 randomly generated pieces. The cumulative performance of each AI instance was used to provide a numerical estimate of the quality of each AI solution. Naturally, there are many ways to evaluate an individual's (or AI's) play at Tetris®. For the current work we used points per piece, or the average number of points per piece across all games played. During the baseline conditions, if a game ended before time was up, a new game automatically started. During the AI evaluation, if a game ended before the limit of 1000 pieces was met, a new game started.

Questionnaires

Here, questionnaires that are used for analyses will be described:

Initial Experience Questions. These questions were designed to understand the participant's familiarity and experience with Tetris overall, their experience with training AI, how comfortable they are with technology, and how trustworthy they are with technology and AI.

AI Literacy Scale. This 16-item scale is used to measure subjective AI literacy (Pinski & Benlian, 2023). Responses are provided along a 7-point Likert scale (strongly disagree to strongly agree). The subscale Overall AI Literacy was calculated as the average of the last three items.

Big Five Inventory. This 44-item questionnaire was designed to measure the Big Five (BFI) personality dimensions (openness, conscientiousness, extraversion, agreeableness, and neuroticism) along a five-point Likert scale. Appropriate items were reverse scored, and a final total derived for each subscale (John et al., 1991).

Post Condition Questions. After each AI training condition, participants were asked to respond to questions relating their experience to training their AI, their AI's performance, their trust in their AI, and whether they would choose their AI over a pre-trained AI. This last question acted as our proxy for PO.

RESULTS

From 370 participants collected, 355 participants were pulled from the larger pool of data for the Horizontal AI condition (36.74 ± 9.10, Female = 99) and 349 for the Vertical AI condition (36.81 ± 9.10, Female = 98) to measure improvements in trust, satisfaction, preference, and score. Data analyses were done in R using the stats package (R Core Team, 2023). For parametric tests, homogeneity of variance was measured using Levene's test. Normality is robust against those with sample sizes of 30 or more (Ghasemi & Zahediasl, 2012). Outliers, defined as three standard deviations from the mean, were removed before analyses. Three data points were identified as outliers from the Horizontal AI condition for this analysis.

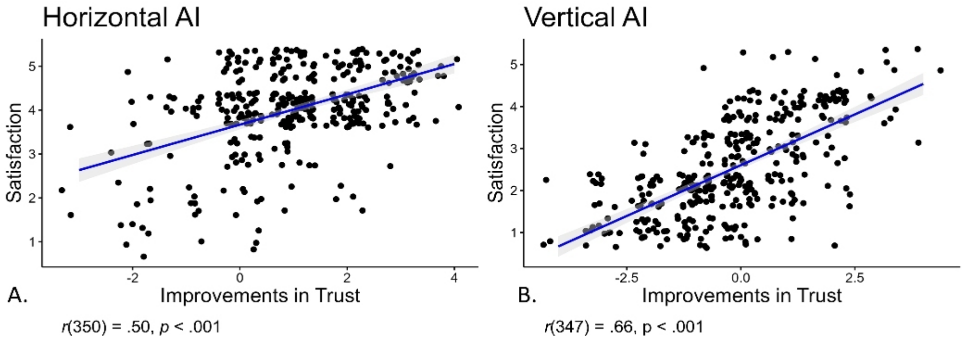

Figure 2: Satisfaction as a function of improvements in trust in the A. Horizontal AI condition and B. Vertical AI condition.

Improvements in Trust and User Satisfaction

To measure improvements in trust, the difference between participants' initial trust in AI ("*How trustworthy do you find AI*?") and the trust in the AI system they trained, which was asked after each condition (Horizontal AI and Vertical AI) pertaining to the AI they trained ("*How trustworthy is your AI?*"), were calculated. Both questions were on a Likert scale from 1 to 5, with a higher difference between the responses indicating increased trust in the system. Similarly, satisfaction in the trained AI was measured on a Likert scale from 1 to 5 and was asked after each condition ("*How satisfied are you with the final performance of your Tetris AI?*") in which higher responses indicated higher levels of satisfaction. Pearson's correlation was run to evaluate the correlation between improvements in trust and self-reported satisfaction in both conditions (Horizontal AI and Vertical AI). Satisfaction in the system increased as people increasingly trusted the system they trained in both the Horizontal AI condition ($r(350) = .50$, $p < .001$) and the Vertical AI condition ($r(347) = .66$, $p < .001$) (see Figure 2).

User Preference

To evaluate users' preference of AI, we had asked, "*Given the choice, would you choose to use your AI or a pre-trained AI?*", and participants selected either "My AI" or "Pre-trained AI". This question acted as our proxy for PO. The Welch's t-test was run to compare the mean improvement in trust between the preference types in each condition. Across both conditions, those who preferred their own AI after training had a higher average improvement in trust (Horizontal AI: Mean = 1.24, SD = 1.30; Vertical AI: Mean = 0.70, SD = 1.57) as compared to those who preferred the pre-trained system (Horizontal AI: Mean = 0.53, SD = 1.38; Vertical AI: Mean = −0.41, SD = 1.42) in both the Horizontal AI condition ($t(272) = 4.81$, $p < .001$) and Vertical AI condition ($t(120.07) = -5.68$, $p < .001$) (see Figure 3). Notably, improvements of trust on average *decreased* for those who preferred the pre-trained system in the Vertical AI condition. Supplementary analyses reveal that those who preferred their own AI and those who preferred the pre-trained AI did not have differences in the number of responses, average score, response rate, or response time. This suggests that both groups had participated in the training process similarly.

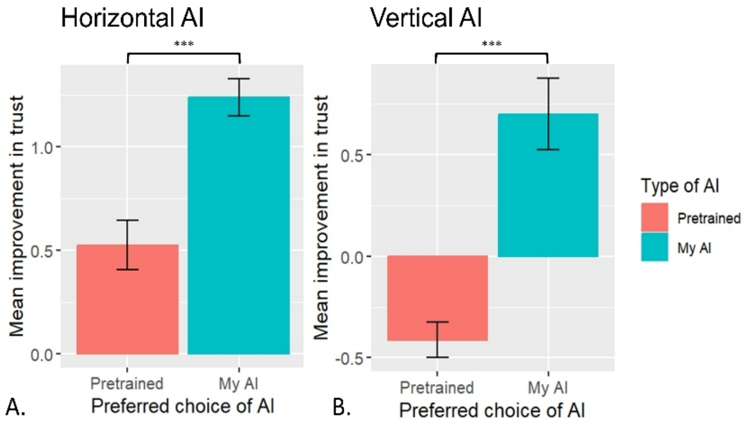

Figure 3: Bar plots in the A. Horizontal AI condition and B. Vertical AI condition showing improvements in trust between those who preferred the pre-trained system and those who preferred their own AI system.

Improvements in Trust and Score

A Pearson's correlation was used to run improvements in trust and score, defined as the average number of points earned per piece across all games. When looking at improvements in trust and score, we find that score does not change as a function of improved trust (see Figure 4). This suggests that even if people increased their trust in the system, it does not actually correlate with how well the system performed. We also found no relationship between an individual's ability to play Tetris® measured by their score during the Baseline task and the score of their final AI (e.g., for Horizontal Tetris® $r(340) = .08$, $p > .14$ and for Vertical Tetris® $r(339) = .02$, $p > .72$).

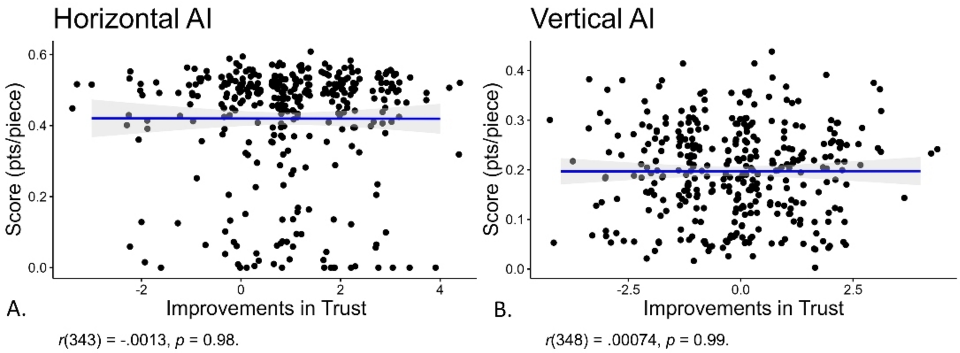

Figure 4: Scatter plots in the A. Horizontal AI condition and B. Vertical AI condition showing a non-significant relationship between improvements in trust and score.

Individual Differences

327 participants from the Horizontal AI condition with mean (M) and $\pm$ standard deviation (SD) in years (M = 36.49, SD = 9.18, Female = 95) and 321 from the Vertical AI condition, (M = 36.51, SD = 9.19, Female = 94) were pulled to measure the effects of AI literacy. 329 participants from the Horizontal AI condition (M = 36.47, SD = 9.15, Female = 95) and 323 participants from the Vertical AI condition, including 94 females with mean (M) and $\pm$ standard deviation (SD) in years (M = 36.53, SD = 9.17) were pulled to measure effects of the BFI. Before running analyses, we first looked at possible covariates of age, gender, and Tetris® experience ("*How many times have you approximately played Tetris in your lifetime?*") in improvements of trust, satisfaction, and score. We did not find any statistically significant patterns or findings. When looking at the effects of the BFI and AI Literacy on satisfaction, preference, and performance, we did not find any linear relationships or statistically significant patterns.

DISCUSSION

The aim of our study was to look at end-users' preferences, satisfaction, and performance as it relates to trust when they are "in the loop" (i.e., have an invested role in training the AI system) and explore individual differences that may influence these relationships. Our results show that with improvements in trust, people were more likely to report higher levels of satisfaction with their trained AI. Additionally, those who preferred their own AI over the pre-trained system showed increased improvements in trust. These results suggest that the human-in-the-loop approach, through perceived control of the user, can positively change their perceptions of their system.

Both conditions used in our study were built to be robust against any modifications to the system aside from extreme user actions, however, the AI was optimized for the Horizontal AI condition and not as optimized for the Vertical AI condition. Patterns with preference, satisfaction, and system performance in relation to improvements in trust were consistent across both conditions. In the Vertical condition, though, those who preferred the pre-trained AI had a decreased improvement in trust on average. We did not find any differences in response rate, response time, or average score between the two groups. Given that both groups participated in the construction of the system in similar ways, we predict that the pre-trained group may possess individual differences that may potentially prevent facilitation of user-system interactions. Future work should be dedicated to identifying these individual differences.

In our study, improvements in trust did not predict system performance. This result suggests that how well the system performed did not necessarily improve trust. This speaks to the design of our testbed, which was intentionally developed to be relatively invariant to performance outcomes aside from extreme disruptive actions by the user. As previously stated, as the number of free parameters in a system increases, so too does the time (or trial) investment needed to train the system increase and the range of potential variants of the system increase (e.g., a high-dimensional, untrained AI system could learn any one of a number of underlying functions). Conversely, as the number of free parameters decreases, so does the required time investment for training. As a result, the performance of the system becomes, successively, constrained. To ensure that training could converge within the amount of time a human would reasonably devote to the task, thus, required a system with a high amount of initial knowledge and a low number of free parameters. We argue, therefore, that allowing a human in the loop for any such system can potentially help the human (via investment and PO) without significantly altering system performance.

In the current work, this pre-cooked approach did not compromise either the user or the system. In other words, the user did not have to change in order to work with the system; similarly, the system did not have to bend its performance or approach to the user (e.g., transparent or explainable models). Pre-cooking in this way enabled the final system to largely perform within tight bounds across users, but still enabled users to experience PO in the process. Naturally, though, these results are limited to the current testbed and warrant further research in more immersive situations where the strength of these psychological bonds can be formally tested.

We were not able to find effects of AI literacy level and personality traits from the Big 5 to satisfaction, preference, or average score. This suggests that relationships between trust and perception towards technology through constructing an AI is relatively stable, at least across levels of AI literacy and personality traits—even those who may not know much about AI systems can still use the Tetris® system equally well and feel satisfied after training it. This is in contrast to other work that have found links to personality dimensions and PO (Essig & Soparnot, 2021). We believe these differences are due to not having direct measurements of PO. Future work aims to look at other individual differences that have been found to relate to PO such as regulatory focus and intrinsic motivation (Delle et al., 2022; Dai et al., 2021), materialism (Jami et al., 2021), emotional intelligence (Kaur et al., 2013), and the extent to which people associate what they own as themselves (i.e., "mine-me" sensitivity) (Jami et al., 2021). Additionally, some traits may be situational and based on the target of ownership or the situation at hand.

Overall, findings indicated users' trust in the AI increased with their reported satisfaction in the system. Increased trust was further linked to preferring their own AI to a pre-trained AI. It is important to note that increased trust did not significantly correlate with system performance. Thus, this link to trust is stronger when considering their personal preference and investment in the system they trained, suggesting an underlying PO mechanism. This is further illustrated when considering those who preferred the pre-trained system over their own. Improvements in trust on average decreased for those who preferred the pre-trained system in the Vertical AI condition. This suggests that lacking attachment or underlying PO may be a detriment to trust across system interactions. Furthermore, supplementary analyses revealed that those who preferred their own AI and those who preferred the pre-trained AI did not have differences in the number of responses, average score, response rate, or response time. This suggests that both groups had participated in the training process similarly and therefore their preference for system may drive trust. This is further underscored by the non-significant relationship to AI literacy and BFI, indicating that we did not find any individual factors driving this relationship at this time beyond their system preference.

Based on our results, we assert that training a system is one way that users can be a part of the system's execution and increase trust and future investment in that system. As this elicits psychological bonds such as satisfaction with the system and the preference to use the system they designed, our results suggest a form of ownership between the user and system at play. These results have broader design considerations. Designers of technology can provide useful features that allow for partial user control, such as personalized and customizable interfaces or guiding users to understand and modify the inner workings of the system. We also champion the use of pre-cooked systems, which offer a potential solution of enabling users to help train an AI system while still allowing the system to perform within predetermined margins.

There are limitations to this study. First, because we do not have a condition where participants rely completely on a pre-trained AI without interaction, we cannot claim conclusively that these emergent relationships are a result of training an AI system. Future work should compare these metrics to a condition in which participants worked with an AI that they did not train. Future work should also include other measures that are known to reliably measure PO such as questionnaires and its link to technology. This work offers a potential solution on how the integration of humans and AI facilitates effective and positive interactions. By bringing together data optimized designs and intentional user execution, we can promote trust, satisfaction, and continued usage through self-investment in the system.

ACKNOWLEDGMENT

We would like to acknowledge and thank Ashley Rabin for facilitating data collection.

REFERENCES

Barak, M., & Levenberg, A. (2016). Flexible thinking in learning: An individual differences measure for learning in technology-enhanced environments. *Computers & Education*, *99*, 39–52. https://doi.org/10.1016/j.compedu.2016.04.003

Barnett, T., Pearson, A. W., Pearson, R., & Kellermanns, F. W. (2015). Five-factor model personality traits as predictors of perceived and actual usage of technology. *European Journal of Information Systems*, *24*(4), 374–390. https://doi.org/10.1057/ejis.2014.10

Bertsekas, D. and Tsitsiklis (1996). *Neuro-dynamic programming*. Athena Scientific.

Dai, Y. D., Altinay, L., Zhuang, W. L., & Chen, K. T. (2021). Work engagement and job burnout? Roles of regulatory foci, supervisors' organizational embodiment and psychological ownership. *Journal of Hospitality and Tourism Management*, *46*, 114–122. https://doi.org/10.1016/j.jhtm.2020.12.001

Delgosha, M. S., & Hajiheydari, N. (2021). How human users engage with consumer robots? A dual model of psychological ownership and trust to explain post-adoption behaviours. *Computers in Human Behavior*, *117*, 106660. https://doi.org/10.1016/j.chb.2020.106660

Delle, E., Kumasey, A. S., Agyabeng, A. N., Alfa, A. A., & Hossain, F. (2023). Work engagement: The role of psychological ownership and self-regulatory focus behavior. *Journal of Organizational Effectiveness: People and Performance*, *10*(1), 77–93. https://doi.org/10.1108/JOEPP-12-2021-0344

Efendić, E., Van De Calseyde, P. P. F. M., & Evans, A. M. (2020). Slow response times undermine trust in algorithmic (But not human) predictions. *Organizational Behavior and Human Decision Processes*, *157*, 103–114. https://doi.org/10.1016/j.obhdp.2020.01.008

Essig, E., & Soparnot, R. (2021). Psychological ownership meets the big five. *Revue de gestion des ressources humaines*, 120(2), 22–38. https://doi.org/10.3917/grhu.120.0022

Ghasemi, A., & Zahediasl, S. (2012). Normality tests for statistical analysis: A guide for non-statisticians. *International Journal of Endocrinology and Metabolism*, *10*(2), 486–489. https://doi.org/10.5812/ijem.3505

Han, T.-S., Chiang, H.-H., & Chang, A. (2010). Employee participation in decision making, psychological ownership and knowledge sharing: Mediating role of organizational commitment in Taiwanese high-tech organizations. *The International Journal of Human Resource Management, 21*(2), 2218–2233. https://doi.org/10.1080/09585192.2010.509625

Jami, A., Kouchaki, M., & Gino, F. (2021). I own, so I help out: How psychological ownership increases prosocial behavior. *Journal of Consumer Research*, *47*(5), 698–715. https://doi.org/10.1093/jcr/ucaa040

John, O. P., Donahue, E. M., & Kentle, R. L. (2012). *Big five inventory* [dataset]. https://doi.org/10.1037/t07550-000

Jörling, M., Böhm, R., & Paluch, S. (2019). Service robots: Drivers of perceived responsibility for service outcomes. *Journal of Service Research*, *22*(4), 404–420. https://doi.org/10.1177/1094670519842334

Kaur, D., Sambasivan, M., & Kumar, N. (2013). Effect of spiritual intelligence, emotional intelligence, psychological ownership and burnout on caring behaviour of nurses: A cross-sectional study. *Journal of Clinical Nursing*, *22*(21–22), 3192–3202. https://doi.org/10.1111/jocn.12386

Knox, W. B., & Stone, P. (n.d.). Tamer: Training an agent manually via evaluative reinforcement. In *2008 7th IEEE International Conference on Development and Learning*. https://doi.org/10.1109/DEVLRN.2008.4640845

Lacroix, D., Wullenkord, R., & Eyssel, F. (2022). I designed it, so i trust it: The influence of customization on psychological ownership and trust toward robots. In F. Cavallo, et al., *Social Robotics* (Vol. 13818, pp. 601–614). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-24670-8_53

Lee, J., & Suh, A. (2015). How do virtual community members develop psychological ownership and what are the effects of psychological ownership in virtual communities? *Computers in Human Behavior, 45*, 382–391. https://doi.org/10.1016/j.chb.2014.12.002

Lin, Y., Zhou, X., & Fan, W. (2021). How do customers respond to robotic service? A scenario-based study from the perspective of uncertainty reduction theory. *Hawaii International Conference on System Sciences 2021 (HICSS-54)*. https://aisel.aisnet.org/hicss-54/cl/human-robot_interactions/4

Lovibond, S. H., & Lovibond, P. F. (2011). *Depression anxiety stress scales* [dataset]. https://doi.org/10.1037/t01004-000

Miller, T. (2023). *Explainable ai is dead, long live explainable ai! Hypothesis-driven decision support*. https://doi.org/10.48550/ARXIV.2302.12389

Oksanen, A., Savela, N., Latikka, R., & Koivula, A. (2020). Trust toward robots and artificial intelligence: An experimental approach to human–technology interactions online. *Frontiers in Psychology*, *11*, 568256. https://doi.org/10.3389/fpsyg.2020.568256

Pierce, J. L., Rubenfeld, S. A., & Morgan, S. (1991). Employee ownership: A conceptual model of process and effects. *The Academy of Management Review, 16*(1), 121. https://doi.org/10.2307/258609

Pinski, M., & Benlian, A. (2023). *Ai literacy—Towards measuring human competency in artificial intelligence*. https://hdl.handle.net/10125/102649

Prolific (2024). Retrieved from https://www.prolific.com/

R Core Team (2023). R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.

Ryan, R. M. (1982). Control and information in the intrapersonal sphere: An extension of cognitive evaluation theory. *Journal of Personality and Social Psychology*, *43*(3), 450–461. https://doi.org/10.1037/0022-3514.43.3.450

Thaler, R. (1980). Toward a positive theory of consumer choice. *Journal of Economic Behavior & Organization*, *1*(1), 39–60. *ScienceDirect*, https://doi.org/10.1016/0167-2681(80)90051-7.

Vandenberg, S. G., & Kuse, A. R. (1978). Mental rotations, a group test of three-dimensional spatial visualization. *Perceptual and Motor Skills*, *47*(2), 599–604. https://doi.org/10.2466/pms.1978.47.2.599

Vasconcelos, H., Jörke, M., Grunde-McLaughlin, M., Gerstenberg, T., Bernstein, M., & Krishna, R. (2022). *Explanations can reduce overreliance on ai systems during decision-making*. https://doi.org/10.48550/ARXIV.2212.06823