The Impact of Virtual Assistant Advice on Human Trust: An Investigation in a game scenario

Raul Benites Paradeda[1], **Álisson de Oliveira Alves**[2], **Daniel Torres**[1], **Rodrigo Neves**[1], **Althierfson Lima**[1]

[1]State University of Rio Grande do Norte (UERN), Brazil,

[2]Federal Institute of Rio Grande do Norte (IFRN), Brazil,

{raul.paradeda}@uern.br

{althierfsonlima,rodrigoneves}@alu.uern.br

{alissoncc.oliveira}@gmail.com

Abstract. *This study investigated the impact of incorrect advice given by virtual assistants on people's trust. An interactive memory game using a virtual robotic agent was developed, and a study was conducted with 25 participants. The study was divided by conditions; in some, the agent gave wrong clues about cards; in others, it gave the right ones. The results suggest that the players' perceptions of the agent were affected by its behaviour. Also, the person's trust towards the agent was affected first by the task itself and secondly by the agent's behaviour being or not obstructive. The study contributes to a better understanding of how virtual assistants affect human decision-making and reliance and to the development of more engaging and interactive virtual assistants. Future research could use these findings to develop more effective virtual assistants that foster greater user trust and engagement.*

Resumo. *Este estudo investigou o impacto de conselhos incorretos dados por assistentes virtuais na confiança das pessoas. Foi desenvolvido um jogo de memória interativo utilizando um agente robótico virtual e realizado um estudo com 25 participantes. O estudo foi dividido em condições; em algumas o agente deu pistas erradas sobre as cartas; em outras, forneceu pistas certas. Os resultados sugerem que as percepções dos jogadores sobre o agente foram afetadas por seu comportamento. Além disso, a confiança da pessoa no agente foi afetada primeiro pela tarefa em si e, em segundo lugar, pelo fato de o comportamento do agente ser ou não obstrutivo. O estudo contribui para uma melhor compreensão de como os assistentes virtuais afetam a tomada de decisão e a confiança humanas e para o desenvolvimento de assistentes virtuais mais envolventes e interativos. Pesquisas futuras podem usar essas descobertas para desenvolver assistentes virtuais mais eficazes que promovam maior confiança e engajamento do usuário.*

1. Introduction

The use of technological resources to enhance various aspects of daily life has become increasingly prevalent. These resources encompass a wide range of purposes, including communication, education, and physical assistance. However, one of the most popular types of resources is personal assistance technology, which is used to help humans with specific tasks, such as navigation with GPS or controlling smart homes using voice commands. Popular examples of virtual assistants include ALEXA, SIRI, Cortana, and Google Assistant. However, these technologies are imperfect and have been known to fail or provide incorrect advice, which can affect a person's trust in a virtual assistant.

Although researchers are exploring mechanisms for dealing with these faulty behaviours, it has yet to be determined if virtual assistant/agent failure or wrong advice will positively or negatively impact a person's level of trust and the human-agent perception. Therefore, this analysis is essential to create a compelling connection between humans and agents. For [Hancock et al. 2011], trust is strongly influenced by the agent's performance and other attributes, such as transparency. Furthermore, some studies have shown that erroneous behaviours can negatively affect how humans perceive agents [Ragni et al. 2016, Salem et al. 2013].

In this sense, to explore these ideas, the aim of this paper is to investigate the impact of virtual assistant advice on human trust and perception, specifically in the context of a digital casual memory game. Therefore, we developed a digital casual memory game with a virtual agent that gives right or wrong clues about which card to turn. Each participant/player will go through one of the five scenarios (study conditions) designed for the task. Besides, the agent is considered emotive since it simulates emotions (sadness, anger, and happiness), trying to persuade the person to choose a specific card. To measure participants' trust levels, we applied the questionnaire used by [Benbasat and Wang 2005]. Also, we used the well-known Godspeed questionnaire [Bartneck et al. 2009] to measure human-agent perception. Taking advantage of the game scenario, we measured whether the agent's behaviour as an assistant affects the experience felt in the game. The questionnaire available at [Poels et al. 2007] was used for this.

To verify the level of confidence people could have in a virtual agent, the influence of agent behaviour on human-agent perception and the task experience, we formulated the following hypotheses:

- **H1:** A person's level of trust towards an agent is affected if the agent's suggestions are incorrect.
- **H2:** False information the agent provides can affect a person's perception of this agent.
- **H3:** The presence of an agent within the game positively affects the player's experience if it provides accurate advice.

The findings of this study will contribute to a better understanding of how virtual assistants affect human decision-making and reliance, especially in situations where their advice can be erroneous. Furthermore, this study highlights the need to develop more reliable and trustworthy virtual assistants to better meet user needs and expectations. Additionally, the methodology and findings of this study could have implications for other research areas, such as human-robot interaction, human-computer interaction, and cognitive psychology.

In the following sections, we will first review relevant literature on trust, human perception of agents, and the impact of erroneous behaviors. This literature review will provide a theoretical foundation for our study. We will then present our methodology, including the development of the digital memory game and the study conditions. Next, we will describe the measurement instruments used to assess participants' trust levels, human-agent perception, and the impact of the agent's behavior on the game experience. We will also detail the hypotheses formulated to guide our investigation. Subsequently, we will present and discuss the results of our study. Finally, we will draw conclusions based on our findings and discuss their implications for the development of more reliable and engaging virtual assistants.

2. Background and Related Work

In this work, we have developed an agent that utilizes a set of bodily movements and behaviors, referred to as emotions [Darwin and Prodger 1998], to persuade the player. It is crucial that these emotional expressions are well-designed, clearly defined, and easily comprehensible for individuals to understand the conveyed sentiment [Johnston and Thomas 1981]. Additionally, the character must incorporate elements such as body movements or environmental cues that enable the audience to perceive the conveyed emotions. These elements enhance the character's ability to effectively communicate emotions to the person watching or interacting with it. Our research focused on simulating emotions through facial expressions and bodily movements, including sadness, anger, happiness, surprise, thinking, shame, and joy.

Another important concept explored in this research is persuasion, which is a field of psychology that investigates the causes and effects of influence on individuals. Persuasion techniques can range from subtle approaches, such as verbal communication, gaze, pointing behaviors, or displaying pleasant images, to more overt tactics, such as physical coercion, threats, or displaying unpleasant images [Gass and Seiter 2018]. Moreover, it is well-known that not all individuals are affected similarly by the same persuasion technique [Moyer-Gusé 2008]. Therefore, discovering the most effective way to subtly influence a person without their realization is crucial [Paradeda 2020]. In our work, we programmed the agent to persuade participants using a combination of facial expressions and emotions. For instance, if a participant did not follow the agent's advice, the agent would perform a sad bodily movement and facial expression.

In [Hashemian et al. 2019], researchers examined the influence and trust in an agent by programming it to use expressions simulating emotions and sharing a sad story. Results showed that participants' trust in the agent increased when the interaction began with "small talk", and the agent's display of joy or sadness affected participants' trust. The study also noted that participants' behavior differed based on gender when exposed to the agent's emotions. Therefore, our work aims to investigate the effects of simulating emotions in a virtual agent within a casual game scenario to understand its impact on trust, specifically in this context.

In [Chowanda et al. 2016], researchers explored the player's experience in a game where Non-Player Characters (NPCs) could express and perceive emotions. Questionnaires were used to measure the impact of interacting with these agents on player engagement and immersion in the game. The study revealed that participants reported higher levels of engagement and immersion when interacting with agents that expressed emotions. Furthermore, players could perceive specific personality and emotional traits in these agents. Based on these findings, the authors concluded that this agent enhances the user experience.

In [Yang et al. 2017], researchers aimed to explore how individuals perceive emotions in a virtual agent named Zara. The study involved the development of an agent capable of perceiving some emotions of users and simulating emotions through animations. Participants were divided into two groups: one interacted with an agent simulating emotions, while the other group interacted with an agent without emotion simulation. The findings indicated that participants perceived the emotionally expressive agent as more confident in its communication. Participants also reported greater empathy towards the agent that simulated emotions than the agent that did not.

Therefore, these related studies provide valuable insights and context for our research, as they align with the central theme of investigating the effects of simulating emotions in a virtual agent, particularly within the context of a casual game scenario.

3. Virtual Agent Development

3.1. Agent's Characteristics

A crucial decision in achieving the project's goal is defining the characteristics of the virtual agent. Research has demonstrated that an agent's appearance can influence human interaction with it [Türkgeldi et al. 2022, Torre et al. 2019]. Therefore, this project focused on defining the agent's appearance, shape, behavior, and unique traits. We gave the agent robotic features to ensure a neutral and inclusive representation, avoiding any specific gender representation [Shiban et al. 2015]. Additionally, we named the agent "Roboldo" to enhance approachability.

3.2. Emotional Simulation

After finalizing the agent's appearance, we selected a set of emotional states for the agent to simulate, including happiness, anger, sadness, neutrality, concentration, shame, joy, and surprise. Each emotion was associated with specific positions and movements of the arms, eyes, mouth, antenna, and overall body expression. We aimed to create a simulation of easily perceptible emotions by combining facial and body expressions. Figure 1 visually represents the defined emotions.

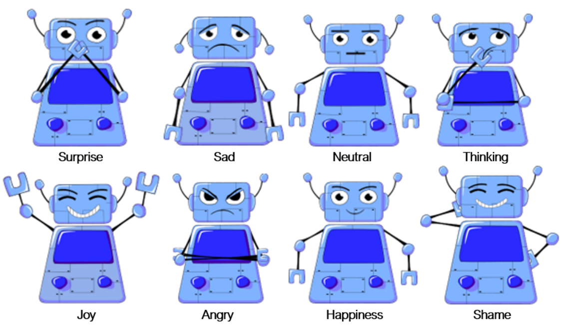

Figure 1. *Roboldo* simulating emotions.

3.3. Animation Development

To enhance the agent's expressiveness, dynamic movements were added through animations. These animations were designed to convey the agent's currently simulated emotion, improving the player's perception of the agent's emotional state. Specifically, animations were created for each expression depicted in Figure 1. All animations began with a neutral expression and smoothly transitioned to the expression representing a specific emotion, thereby simulating the emotional change experienced by the agent.

4. Game Scenario

For this study, a 2D casual game already familiar to most people, the "Memory Game", was developed. In this game, the participant's objective is to flip over two identical cards, and the game ends when the player successfully flips over all of the cards. A game with familiar mechanics was chosen to ensure participants did not waste time becoming familiar with the rules. The game was developed using a 2D and 3D game creation engine called Unity[1].

The memory game was divided into three units (Figure 2): (1) the ''*Main Control Unit*", which manages communication and interaction between the units; (2) the "*Game Unit*"; and responsible for controlling the mechanics; (3) the "*Affective Agent Unit*", responsible for managing the agent's interactions and emotions (details in the following sections).

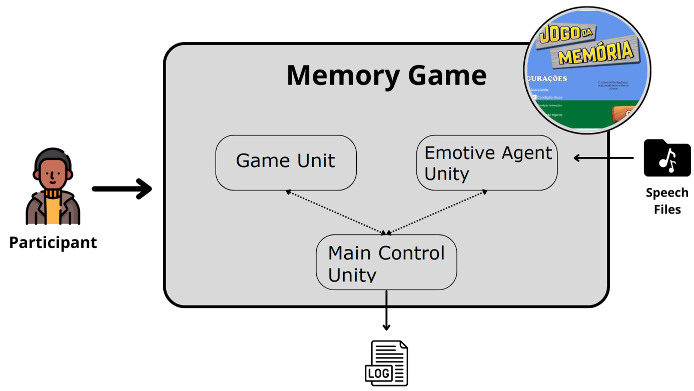

Figure 2. System architecture.

The participant/player can interact with the game by choosing which cards to flip over in an attempt to make a pair. During this interaction, the game provides visual feedback (cards flipping over with each selection) and sound feedback (e.g., a "beep") to confirm the action. A file of "speeches" stores all the information for agent interaction, such as phrases and emotional states that the agent should express. A *log* file is responsible for keeping records of all interactions made in the game, whether made by the agent (speeches and emotions) or the player (selected cards).

4.1. Main Control Unit

This unit controls the interactions between the Game Unit and the Affective Agent Unit. It detects the state of the game and the interactions made. Additionally, this unit is responsi-

[1]Unity Technologies. Available at: `https://store.unity.com/#plans-individual`. Accessed on: 25th September 2020.

ble for activating the Affective Agent Unit and requesting emotions and phrases based on the player's actions. In summary, the Main Control Unit handles all communication. For example, when the participant chooses two cards, it verifies whether the selected cards match the ones indicated by the agent and reports to the Affective Agent Unit whether the player followed the guidance. This unit also manages the log file, recording all interactions between the player and the agent.

4.2. Game Unit

In the game project, two screens are essential to achieve the goal of this research: the settings screen, which is only accessible by the game developer, and the main screen, where the cards and the agent are presented.

On the settings screen (Figure 3A), the necessary information for the game's correct functioning is provided and the scenario is configured according to the desired conditions. This information includes an identification created by the researcher for each participant/player, which is necessary for the log to identify each player's actions and the agent's condition, as it is essential to measure the agent's different behaviors.

On the game screen, the participant interacts with the developed system/game (Figure 3B). The cards are marked with numbers to facilitate player identification and allow the agent to provide straightforward guidance/suggestions. During the game, all interactions made by the participant are recorded and saved in the log file.

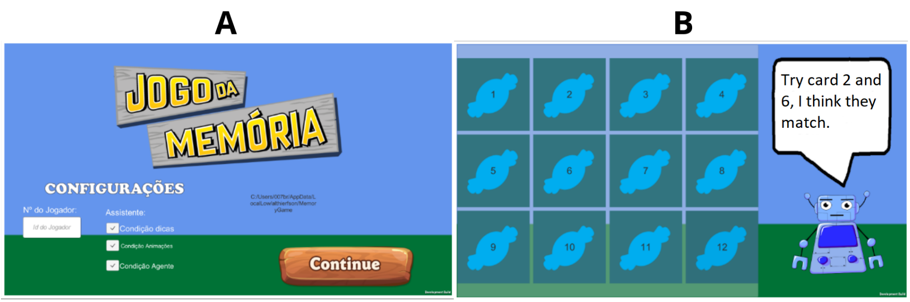

Figure 3. Memory game scenario.

4.3. Emotive Agent Unit

During the game, the agent, positioned on the right side of the cards, assists the player by providing suggestions through animations, text, and speech. The game is designed to partially block when the agent interacts with the player, aiming to prevent the player from ignoring the agent and to encourage them to pay attention to the agent's actions or suggestions.

The Affective Agent Unit is divided into three modules: the Agent Control Module, the Emotions Module, and the Voice Module.

The Agent Control Module controls the agent's response to player interactions, providing true or false instructions. For example, the agent may suggest that the player flip two cards that may or may not match, such as "Try flipping cards 2 and 3". If the player follows the agent's suggestion, the agent reacts positively with an animation expressing joy. Otherwise, an animation of sadness informs the player that the agent is disappointed with their decision. Figure 4 displays the game screen with the agent expressing happiness and a supportive message when the player accepts its suggestion.

The Emotions Module receives information from the Agent Control Module on whether or not the player followed the agent's advice and executes the corresponding animation based on the received data. This module also controls other animations, such as when the agent speaks and passive animations executed while the player is playing or thinking. These latter animations were programmed to give players a sense of immersion and make the agent appear lively rather than a static figure in the game environment.

The Voice Module is responsible for managing the agent's voice. A Text-to-Speech software is used to reproduce the phrases received from the Agent Control Module as sound. This module also controls a dialogue balloon that presents the text of the spoken phrase from the Text-to-Speech. This process is necessary because if the player does not understand what the agent is saying, they can read the text of the speech in the balloon.

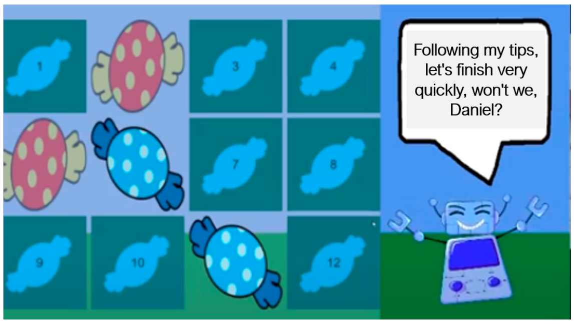

Figure 4. Memory game scenario.

5. Research Methods

5.1. Manipulation

Before starting the investigation, it was determined that there would be no age restrictions or requirements regarding familiarity with games. We aimed to maintain a similar distribution of male and female participants in each study condition. Participants were divided into five (5) conditions, with an equal number of participants in each condition. The conditions are as follows:

- **C1:** Game without the agent's presence.
- **C2:** Game with the agent simulating emotions and providing correct suggestions.
- **C3:** Game with the agent not simulating emotions and suggesting correctly.
- **C4:** Game with the agent simulating emotions and providing false suggestions.
- **C5:** Game with the agent not simulating emotions and providing wrong suggestions.

These study conditions were created to understand whether the agent's presence, affective behavior (expression of emotion), and the accuracy of suggestions could influence trust, human-agent perception, and the participant's experience in the task.

5.2. Procedures and Survey Design

People were invited to participate voluntarily, with the indication that they would play a memory game and interact with a virtual agent (conditions with the presence of the agent). It was explained that the study would consist of three parts: a pre-questionnaire, interaction with the game, and a post-questionnaire. After accepting the invitation, participants were shown the consent form, the study was explained, and their participation was obtained.

The pre-questionnaire collected demographic data from the participants and their level of trust and perception towards the agent. As the interaction did not start at this moment, an image of the virtual agent (Roboldo) was presented on the screen to measure the level of trust and participants' perception. Participants were asked to answer questions based on this image to measure trust and perception.

Next, the researcher left the room and allowed the participants to interact with the agent in three rounds of a memory game. Three rounds were chosen to reduce the game's randomness and give participants enough time to develop a bond of trust or distrust with the agent without becoming bored. After completing the three rounds, the researcher returned and provided the post-questionnaire. This questionnaire measured the participants' level of trust and perception towards the agent, as well as their level of immersion in the game. This strategy allowed investigating whether these levels changed due to the interaction.

Pre-Questionnaire and Post-Questionnaire

The questionnaires played a significant role in obtaining and analyzing data on the interactions. The questionnaires were developed and made available through the Google Forms platform. Participants who interacted with the agent were assigned to one of the following conditions: C2, C3, C4, or C5. Participants who did not interact with the agent were placed in condition C1.

The questionnaires aimed to evaluate the following aspects: (1) Demographic data, (2) participants' experience with games [Poels et al. 2007], (3) participant-agent perception [Rosenthal-von der Pütten et al. 2018], and (4) the level of participant-agent trust [Benbasat and Wang 2005].

For the demographic data, the following characteristics were collected: participants' age and gender. To collect and analyze the participants' experience with games (2), the questionnaire developed in the work of [Poels et al. 2007] was applied in the pre- and post-questionnaires. According to the authors, the Game Experience Questionnaire has a modular structure consisting of the core questionnaire, the Social Presence Module, and the Post-game module. Due to its extensive nature, we focused on measuring the terms Negative and Positive Affect, which probe the fun and enjoyment of gaming. Participants indicated their responses using a 5-point unipolar intensity-based answering scale, ranging from "not at all" (0) to "extremely" (4). Additional questions were added to the pre-questionnaire, such as the number of hours played per week, preferred game style, and most played platforms. The questionnaire can be found in Supplementary Material link folder 1.

For the participant-agent perception (3), the Godspeed questionnaire [Bartneck et al. 2009] was used. This questionnaire has five dimensions measuring specific human-agent perceptions: anthropomorphism, animacy, likability, perceived intelligence, and perceived safety. Participants indicated their responses using a 5-point Likert scale with bipolar options. The questionnaire can be found in Supplementary Material link folder 2.

In short, anthropomorphism refers to attributing human form, characteristics, and behaviours to nonhuman agents. Animacy refers to assigning real properties to an entity by a user. Likeability is defined as the development of an agent's positive impression. Perceived intelligence is defined as the ability of an agent to adapt its behaviour to varying situations. Finally, perceived safety is the user's perception of the level of danger when interacting with an agent and the user's level of comfort during the interaction. We did not measure the last dimension since, in our study, the task performed by the agent did not represent a dangerous situation. The volunteer should point out, using a 5-point Likert scale with bipolar options, their perception regarding questions to measure the Godspeed dimensions.

To measure the level of participant-agent trust (4), the questionnaire presented in the work of [Benbasat and Wang 2005] was used. This questionnaire evaluates trust in an agent across six dimensions: competence, benevolence, integrity, perceived usability, perceived ease of use, and intention to adopt. Participants responded on a 9-point Likert unipolar scale (strongly distrust - strongly believe), indicating their perception of the agent's activity. For this questionnaire, the dimensions that we believe are most related to the activity performed by Roboldo and the questionnaire creators recommend that for trust measurement are the first three dimensions. The questionnaire used can be found in Supplementary Material link folder 3.

5.3. Statistical Analysis

We used IBM SPSS V.28.0 packages and Microsoft Excel for the statistical analysis. The Shapiro-Wilk test [Shapiro and Wilk 1965] was performed in our data set to identify the normality of data. As our data deviated from a normal distribution ($p > 0.05$), we applied the Wilcoxon signed-rank test to understand whether there was a difference in the level of trust and Godspeed perception between the conditions. In those tests, we considered 5% of statistical significance, i.e., p values of $<= 0.05$.

6. Results

This investigation was conducted in an isolated room with a random selection of people. A sample of 25 subjects (avg. 25 years old; $SD = 4.33$) participated and was assigned five per study condition (6 females (avg. 24.33; $SD = 5.12$) and 19 males (avg. 24.89; $SD = 4.03$)).

H1: Firstly, it was verified if there was any statistical difference between the participants' data who interacted with the virtual agent. For this purpose, the Wilcoxon test was applied between the pre- and post-questionnaire data. Significant differences were found for the perceived competence ($Z = -2.092, p = .036$) and benevolence ($Z = -2.175, p = .030$) factors.

Next, the data were divided by conditions created for the study and the Wilcoxon test was applied. Thus, only in condition C4 (game with the agent simulating emotions and providing false suggestions) presented statistical differences in the factors of competence ($Z = -2.032, p = .042$) and benevolence ($Z = -2.023, p = .043$) before and after the interaction. In the others conditions, we did not find statistical differences.

H2: We analysed the participants' answers from the Godspeed questionnaires in the conditions with the agent's presence (C2-C5). Hence, no significant differences were found between pre- and post-questionnaires for each of the evaluated dimensions (anthropomorphism ($Z = -0.825$, $p = .409$), animacy ($Z = -1.585$, $p = .113$), likeability ($Z = -1.121$, $p = .262$), and perceived intelligence ($Z = -0.261$, $p = .794$)).

Next, the data were analysed by study conditions. Our results suggest that participants had a different perception of the dimensions of perceived intelligence in C2 ($Z = -2.041$, $p = .041$) and C3 ($Z = -2.023$, $p = .043$) and anthropomorphism ($Z = -2.000$, $p = .046$) in C4. In C5, we did not find any statistical differences.

H3: Using the game experience questionnaire answers were observed statistical differences for the condition without *Roboldo* (C1) for negative ($Z = -2.023, p = 0.043$) and positive affects ($Z = -2.032, p = 0.042$).

Next, for the same parameters (positive and negative affect), it was checked whether there were statistical differences between the conditions. For condition C1 (without the virtual agent), there were differences according to the Wilcoxon test for negative affect ($p = .043, Z = -2.023$) and positive affect ($p = .042, Z = -2.032$). Regarding the conditions with the virtual agent, the following differences were obtained: C2, only positive affect had statistical differences ($p = .042, Z = -2.032$), whereas no differences were found for negative affect ($p = .225, Z = -1.214$); Similar behaviour was observed for C3, where only positive affect had statistical differences ($p = .043, Z = -2.023$), whereas no differences were found for negative affect ($p = .686, Z = -.405$); C4, only positive affect had statistical differences ($p = .046, Z = -1.997$), whereas no differences were found for negative affect ($p = .498, Z = -.677$); C5, only positive affect had statistical differences ($p = .042, Z = -2.032$), whereas no differences were found for negative affect ($p = .136, Z = -1.490$).

7. Discussion

H1: Based on the collected data, it was verified that the level of trust participants had in Roboldo was influenced by the condition in which they were exposed. A decrease in Roboldo's competence and benevolence was observed, and statistical differences were found in condition C4. In this condition, the agent provided false clues to the players, leading to no matches in the cards. This result makes sense because participants had a certain level of trust before interacting with the agent based solely on its image. However, when the interaction began, and the agent provided false hints, the *a priori* level of trust decreased because the agent was hindering rather than helping. It was noticeable that no difference was found in condition C5, where the agent also provided incorrect information. The difference between the two conditions was the emotional behavior of the agent, with the expression of emotions in C4 but not in C5. Hence, it can be inferred that incorrect suggestions can affect a person's level of trust in an agent, particularly when the agent expresses emotions.

H2: The responses regarding the perception of the person-agent relationship revealed that in the conditions where the agent provided correct information, participants perceived a significant difference (before and after interaction) in the dimension measuring the agent's intelligence. This change in perception is valid since the agent provided card hints that helped the players complete the game quickly and easily. Therefore, players may have perceived the agent as more intelligent after the interaction compared to their perception before the interaction. In the condition where the agent provided false hints and simulated emotions (C4), a statistical difference was found before and after the interaction in the dimension of anthropomorphism. It can be observed that the agent's simulation of emotions in this condition led the participants to perceive the agent as more "alive" after the interaction than before. There were no differences in some dimensions, likely due to the small sample size in each condition. It is believed that a larger sample is necessary to observe the behavior in human-agent perception across the dimensions of the Godspeed questionnaire.

H3: A statistical difference was found when analyzing the data from all conditions regarding the game experience (positive and negative affects). This provides evidence that the developed scenario positively and negatively influenced the players. When splitting the data by condition, statistical differences were observed between the participants who did not interact with the agent before and after regarding positive and negative affects. A statistical difference was found in positive affects in the second and third study conditions. Conversely, in conditions C4 and C5, the differences were observed in negative affects. The presence of the agent providing true or false clues influences the players' game experience. However, further studies with a larger sample size are needed for more significant evidence.

8. Conclusion

In conclusion, this study validates the hypothesis that virtual agents become less valuable when they fail to provide helpful information to the user. Our findings demonstrate that deliberate errors in the agent's responses have a negative impact on users' confidence, perception, and experience in the task. In contrast, helpful tips and information have a positive effect.

However, we acknowledge that the small sample size and limitations in simulating the agent's emotions may have influenced the study's results. Therefore, we recommend conducting future studies with a larger sample size and developing agents with improved expression capabilities to enhance user accuracy and connection.

Despite these limitations, this research provides valuable insights into how virtual assistants can influence human decision-making and trust. Our findings can inform the development of more effective virtual assistants that promote greater user engagement and trust. Overall, we hope that this study contributes to a better understanding of how virtual agents can be designed to provide helpful assistance and improve the user experience.

Supplementary Material

`https://osf.io/qr8vk/?view_only=b4099caae9944c24b4ce35c7613a0717`.

References

Bartneck, C., Kulić, D., Croft, E., and Zoghbi, S. (2009). Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International journal of social robotics*, 1(1):71–81.

Benbasat, I. and Wang, W. (2005). Trust in and adoption of online recommendation agents. *Journal of the association for information systems*, 6(3):4.

Chowanda, A., Flintham, M., Blanchfield, P., and Valstar, M. (2016). Playing with social and emotional game companions. In *International Conference on Intelligent Virtual Agents*, pages 85–95, Springer. Cham.

Darwin, C. and Prodger, P. (1998). *The expression of the emotions in man and animals*. Oxford University Press, USA.

Gass, R. H. and Seiter, J. S. (2018). *Persuasion: Social influence and compliance gaining*. Routledge.

Hancock, P., Billings, D., and Schaefer, K. (2011). Can you trust your robot? *Ergonomics in Design: The Quarterly of Human Factors Applications*, 19:24–29.

Hashemian, M., Paradeda, R., Guerra, C., and Paiva, A. (2019). Do you trust me? investigating the formation of trust in social robots. In *EPIA Conference on Artificial Intelligence*, pages 357–369, Springer. Cham.

Johnston, O. and Thomas, F. (1981). *The illusion of life: Disney animation*. Disney Editions, New York.

Moyer-Gusé, E. (2008). Toward a Theory of Entertainment Persuasion: Explaining the Persuasive Effects of Entertainment-Education Messages. *Communication Theory*, 18(3):407–425.

Paradeda, R. B. (2020). *Personality-based Persuasion by an Interactive Social Robot Storyteller*. PhD thesis, Instituto Superior Técnico, Universidade de Lisboa.

Poels, K., de Kort, Y., and Ijsselsteijn, W. (2007). D3. 3: game experience questionnaire.

Ragni, M., Rudenko, A., Kuhnert, B., and Arras, K. O. (2016). Errare humanum est: Erroneous robots in human-robot interaction. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 501–506.

Rosenthal-von der Pütten, A. M., Krämer, N. C., and Herrmann, J. (2018). The effects of humanlike and robot-specific affective nonverbal behavior on perception, emotion, and behavior. *International Journal of Social Robotics*, pages 1–14.

Salem, M., Eyssel, F., Rohlfing, K. J., Kopp, S., and Joublin, F. (2013). To err is human(-like): Effects of robot gesture on perceived anthropomorphism and likability. *International Journal of Social Robotics*, 5:313–323.

Shapiro, S. S. and Wilk, M. B. (1965). An analysis of variance test for normality (complete samples). *Biometrika*, 52(3/4):591–611.

Shiban, Y., Schelhorn, I., Jobst, V., Hörnlein, A., Puppe, F., Pauli, P., and Mühlberger, A. (2015). The appearance effect: Influences of virtual agent features on performance and motivation. *Computers in Human Behavior*, 49:5–11.

Torre, I., Carrigan, E., McDonnell, R., Domijan, K., McCabe, K., and Harte, N. (2019). The effect of multimodal emotional expression and agent appearance on trust in human-agent interaction. In *Proceedings of the 12th ACM SIGGRAPH Conference on Motion, Interaction and Games*, pages 1–6.

Türkgeldi, B., Özden, C. S., and Aydoğan, R. (2022). The effect of appearance of virtual agents in human-agent negotiation. *AI*, 3(3):683–701.

Yang, Y., Ma, X., and Fung, P. (2017). Perceived emotional intelligence in virtual agents. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pages 2255–2262.