Towards a Computational Approach for Proactive Robot Behaviour in Assistive Tasks

Ilenia Cucciniello
University of Naples Federico II
Naples, Italy
il.cucciniello@studenti.unina.it

Antonio Andriella
Pal Robotics
Barcelona, Spain
antonio.andriella@pal-robotics.com

Silvia Rossi
University of Naples Federico II
Naples, Italy
silrossi@unina.it

ABSTRACT

While most of the current work has been focused on developing adaptive techniques to respond to human-initiated inputs (what behaviour to perform), very few of them have explored how to proactively initiate an interaction (when to perform a given behaviour). The selection of the proper action, its timing and confidence are essential features for the success of proactive behaviour, especially in collaborative and assistive contexts. In this work, we present the initial phase towards the deployment of a robotic system that will be capable of learning what, when, and with what confidence to provide assistance to users playing a sequential memory game.

CCS CONCEPTS

• **Human-centered computing** → **User centered design**; *User studies*; *User models.*

KEYWORDS

Proactive Behaviour, Socially Assistive Robots, Human-Robot Interaction, Adaptive Behaviour

ACM Reference Format:
Ilenia Cucciniello, Antonio Andriella, and Silvia Rossi. 2023. Towards a Computational Approach for Proactive Robot Behaviour in Assistive Tasks. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction (HRI '23 Companion), March 13–16, 2023, Stockholm, Sweden.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3568294.3580139

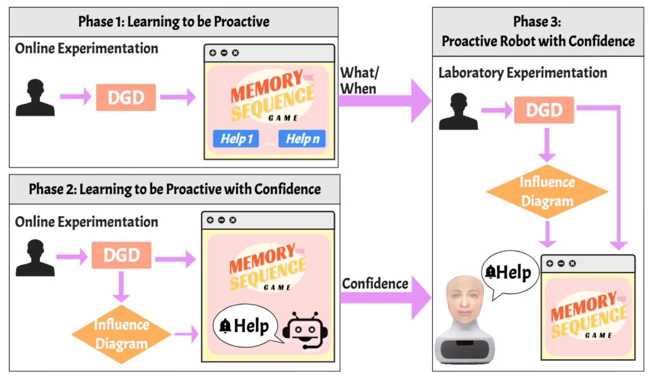

Figure 1: Illustration of the learning pipeline in three phases. Note that DGD is the questionnaire we use to assess the player's profile.

1 INTRODUCTION

Socially assistive robots aim to improve the quality of life of their users through social interactions [11, 16] which implies that the interaction is focused on the users to help them to achieve specific goals. Perceiving the user's needs and intentions and acting proactively are fundamental to social interaction.

With respect to learning what behaviour the robot shall perform according to the user's preferences and needs, most of the research in the field has aimed at solving this problem by employing different AI methodologies: symbolic task planning [2], Bayesian Network (BN) [17], Reinforcement Learning (RL) [5], and Inverse Reinforcement Learning (IRL) [1], among others. Nonetheless, such approaches result in robots that learn only to react to humans' actions or to given events. In the real world, robots are also required to proactively take the initiative [18]. They need to identify the requirements of a situation and gather information to decide *what*, but also *when* to perform a given action in an anticipatory way. Indeed, time plays a crucial role in the quality of collaborative tasks [8]. If a robot intervenes when it is not needed, it can negatively impact the interaction, especially trust [12]. And so, if the system does not intervene when the human expects it to take an action, their engagement might wipe off over time [9].

To address this issue, in recent years, researchers have started exploring new techniques for making robots proactive. Nonetheless, most of the approaches have focused on learning the levels of proactivity of the robot, which is controlled in a Wizard-of-Oz (WoZ) fashion ([10, 15, 19]). On the other hand, very few have focused on developing AI reasoning capabilities to learn the correct timing for robot behaviour [4, 6, 7]. Moreover, none of the previous approaches proposed a computational approach to learn the proactivity directly from humans.

We focus on defining a computational approach that can be used for learning (i) when humans require assistance and (ii) how confidently the robot should take control and intervene. Such a framework is tested on a sequential memory task. Here, we devised a learning pipeline consisting of three main phases, aiming at learning an Influence Diagram (ID) [14] that fits the participant's assistive needs (see Fig. 1). In the first phase, namely, *Learning to be Proactive*, we request participants to play alone with the possibility of asking for assistance when necessary. In the second phase, namely, *Learning to be Proactive with Confidence*, participants play

HRI '23 Companion, March 13–16, 2023, Stockholm, Sweden

ACM ISBN 978-1-4503-9970-8/23/03…$15.00
https://doi.org/10.1145/3568294.3580139

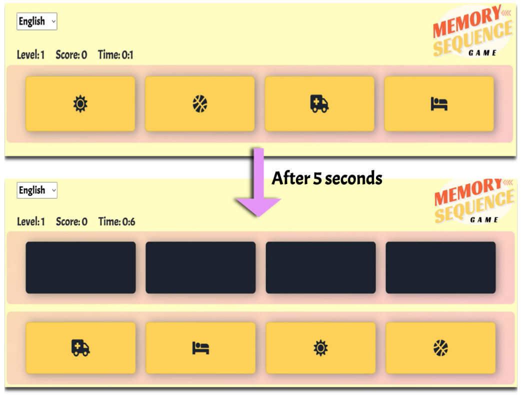

Figure 2: Example of game sequence of 4 icons: the sequence shown to the user (top), the sequence requested to repeat to the user after 5 secs (bottom).

with the assistance of a proactive virtual-screen robot trained with the data collected in the previous phase. At this point, we learn how confident the robot is in offering assistance. Finally, in the third phase, namely, *Proactive Robot with Confidence*, we endow the Furhat robot with the model fine-tuned in the second phase, aiming to assess whether and to what extent the proactivity has an impact on participants' acceptability and performance.

In this work, we present the details of the first phase (see Fig. 1), describing the preliminary results.

2 PHASE 1: LEARNING TO BE PROACTIVE

2.1 Methodology

The objective of this first phase was to learn both the reactive behaviour of the robot, that is, what assistive action to offer, and the timing at which such assistance needs to be provided. Towards such a goal, we designed an experimental study in which participants were requested to play a sequential memory game with increasing levels of difficulty (Fig. 2). Participants were told to play the game with the assistance of a virtual robot that was introduced as a peer from whom they could seek help if necessary. Hence, when requested, the robot could then unlock one of the degrees of assistance available in that state. According to [10], this strategy is referred to as a "low level of proactivity" where the user can only explicitly request help. Specifically, the user could click on the robot icon to unlock suggestions and decide *when* and *what* to apply.

To avoid a one-size-fits-all policy for all the participants, we profiled them by using the Demographic Game Design model (DGD) [3]. In this way, we could learn the proactive robot behaviour that best fits a given player's profile.

2.2 Assistive Task

The task was a sequential memory game in which, for each level, the participant was requested to remember a sequence of random icons (e.g., anchor, bicycle, cloud, sun) and repeated them in the correct order. The sequence of icons was shown for 5 seconds, after that cards were flipped and the participant was requested to compose the entire sequence. The next level was accessible only if the current one was completed.

At any time, the participant could decide to end the game in order to avoid experiencing negative feelings, such as discouragement or frustration, in the case of non-completion.

The goal of the game was to pass as many levels as possible while trying to minimise the number of mistakes and requests for assistance. Points were assigned depending on whether the participant placed the correct card (+10 correct card / -5 wrong card) and on the assistance requested (See Sec. 2.3). Each level had a different number of cards to remember and a different deck of cards to choose from. Regarding the number of cards in the sequence to remember, it ranged from 4 up to 7, according to Miller's Law [13] which stated that the number of objects an average human mind can remember while running is 7 ± 2. Concerning the deck of cards, they ranged from 4 up to 15 cards.

2.3 Assistive Behaviour

We defined four degrees of assistance that the participant could request (see Fig. 3):

- *Hide Card* (-1 point), allows for the temporary removal of a card from those to be selected;
- *Suggest Position* (-2 points), allows for the temporary removal of cards to the right or left of those to be chosen;
- *Indicate Position* (-3 points), allows for the temporary removal of all the cards except the correct one;
- *Review Sequence* (-4 points), allows one to review the entire sequence of cards for 5 seconds.

2.4 Procedure

For the online experimentation, we developed a website[1] with Flask[2] and hosted it on PythonAnywhere[3] for free. Before starting to play the game, a tutorial with game instructions about how to score points and request assistance was shown. Next, the participant was offered the possibility to try the game, with up to three levels. After the user agreed to participate, they were asked to first fill out a demographic questionnaire about age, gender, and level of education, and then the DGD questionnaire. Finally, the user could start playing the game. After passing a level, the player was provided with their current score and ranking position with respect to the others in order to keep them engaged. The game ended in two cases: when the user pressed the end button or when they achieved the maximum level.

2.5 Evaluation Measures

Concerning subjective measures, the DGD Questionnaire was used to profile the participants. The questionnaire provides four non-exclusive types of players (i.e., Conqueror, Manager, Wanderer, and Participant).

Concerning the objective measures, the highest level achieved in the game and the total score were considered the participant's performance. Furthermore, the number of mistakes and hints requested were included as dimensions to classify the user's playing

[1] http://qcdeveloper.pythonanywhere.com/
[2] http://flask.pocoo.org/
[3] https://www.pythonanywhere.com/

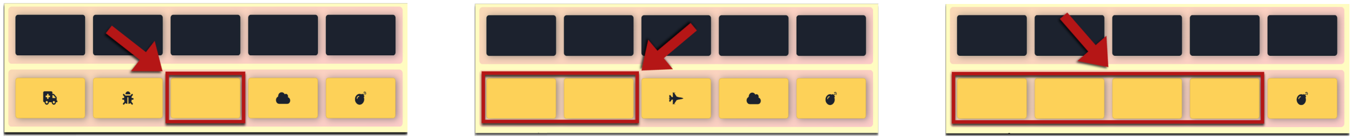

(a) *Hide card* (-1 point) **(b) *Suggest position* (-2 points)** **(c) *Indicate position* (-3 points)**

Figure 3: Examples of the assistance provided by the virtual robot

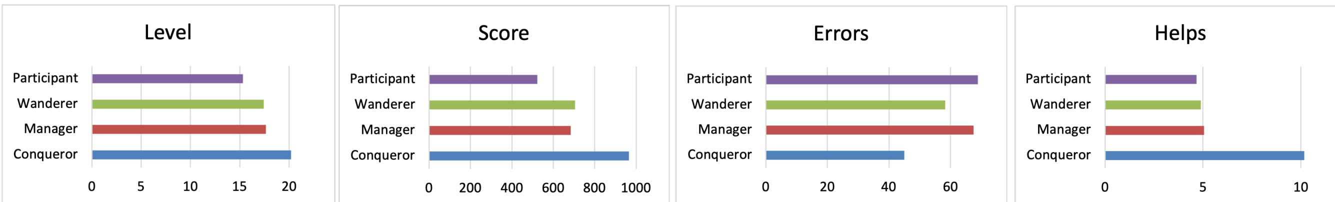

Figure 4: Average values for each profile of the following discriminative variables: a) *Level* achieved b) *Score* c) *Mistakes* and d) requested *Helps*.

style. In order to evaluate when the user requested a hint, we calculated the "help time" which is the time difference between the request for help and the previous action. Finally, to evaluate what hints were requested, we calculated the number of occurrences of each of them.

3 PRELIMINARY RESULTS

We divided the participants into four groups based on the DGD questionnaire. However, some participants might belong to more than one class; hence, the results presented in this section involve statistics with repeating players. From those, we excluded those participants who did not achieve level 10, which was the first with 7 cards to remember and 7 cards in the deck. The data of 83 participants were analysed, 45 of whom belong to the Conqueror class, 20 to the Manager, 9 to the Wanderer, and finally 9 to the Participant. For each class, we identified the following discriminative variables: number of levels achieved, score, number of mistakes, and assistance requested (See Fig. 4).

From the results, it emerged that the Conqueror played longer on average and got the highest score, asking for more help than others with the result of making fewer mistakes. On the other hand, the Manager asked for less assistance and consequently made more mistakes, lowering the total score. Given the smaller number of participants belonging to the Wanderer and Participant classes, we could not find any relevant differences in the data. However, when we consider the *what*, that is, the degree of assistance that participants requested during the game (see Fig. 5), we can notice how participants who belonged to the Participants class switched almost equally between "Hide card" and "Review sequence" while the others requested the maximum level of assistance most of the time. It is noteworthy to observe that a kind of pattern in the selection of the degrees of assistance could be inferred.

We speculated that Conquerors were those who challenged themselves to finish the game or at least achieve the maximum level according to their capacity. Indeed, they asked for the amount of help they needed to pass to the next level and keep their score higher.

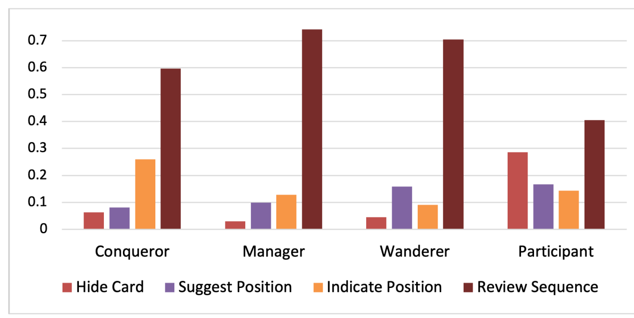

Figure 5: Probability of *what* level of assistance for each profile was request.

On the other hand, the Manager seemed to prefer the highest level of assistance only when they started to make several mistakes in a row. Their strategy did not pay in terms of score, and they ended up with a score quite far from those belonging to the Conqueror. Concerning the Wanderer, it can be observed that the strategies for requiring assistance were different. Lower levels were requested, not incrementally as for the Conqueror and Manager, and the highest level was requested to avoid losing the game and moving on to the next level. Finally, for those belonging to the Participant class, we can hypothesise that they did not use the assistance for getting to the next level, on the contrary, they seemed to play for the sake of enjoying the game. Indeed, most of the participants who belonged to the Participant class were those who requested less assistance, committed more mistakes, and ended up finishing the game before the minimum level.

With respect to the time they requested help, we can notice some interesting patterns (see Fig. 6). Indeed, the participants who belonged to the Conquer profile, on average, waited for the longest before requesting help. A similar pattern can be observed for those belonging to the Manager profile. However, in this case, their timings were slightly lower than those of the Conqueror. Contrarily, those who belonged to the Participant and Wanderer profiles showed a completely different pattern. Indeed, they seemed not to hesitate when they requested the "Hide Card" and "Indicate Position" helps, while for the "Suggest Position" and "Review Sequence" they delayed a bit more their request.

Figure 6: Average time (s) for each profile to request a given level of assistance.

These results confirm our main idea of profiling the participants according to the DGD questionnaire. Indeed, we could observe different patterns for the four classes in terms of "what" assistance (see Fig. 5) and "when" it was requested (see Fig. 6).

4 NEXT STEPS

With the data collected in the first phase, we have built an ID whereby the system can learn at what moment it is more likely to provide a given degree of assistance. The ID will be used in phases 2 and 3. Note that the initial classes returned by the DGD questionnaire might be increased to consider the intra-classes relationship (e.g., a person might be both a Conquer and a Manager). In the following, we describe how we envision the next two phases.

4.1 Phase 2: Learning to be Proactive with Confidence

In the second phase, the goal is to learn the level of confidence by means of which the robot will take control of the task. According to [10], we defined four increasing levels: *None*, in which subjects could only explicitly request help *Notification*, in which the user is informed of a solution *Tip*, in which a solution is directly suggested to the user *Action*, in which the system selects an option for the user. Note that the *None* level has already been trained in the first phase, in which participants could request assistance when they needed it. Here, we focus on the remaining three levels. As in the first phase, participants will be assisted by a virtual robot, which, according to the ID model of the class they belong to, will propose the action of assistance that best fits their needs. Aiming to train the level of confidence, the robot would suggest the action with a random confidence level, and the responses of participants would be recorded. In this way, we extend our ID, which will return not only what action to take and when to take it, but also with what confidence to recommend it to the user.

4.2 Phase 3: Proactive Robot with Confidence

In the last phase, we aim at assessing whether a proactive robot would have an impact on participants' performance, trust, and overall experience. To do so, we will endow a Furhat robot with proactive capabilities based on the ID trained in the first two phases. Specifically, each participant will be profiled with the DGD and requested to play with the robot endowed with an ID that reflects their profile.

5 ACKNOWLEDGMENT

This project was partially funded by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 801342 (Tecniospring INDUSTRY), by the Italian MUR and by the European Union under the CHIST-ERA project COHERENT (PCI2020-120718-2).

REFERENCES

[1] Antonio Andriella, Carme Torras, Carla Abdelnour, and Guillem Alenyà. 2022. Introducing CARESSER: A framework for in situ learning robot social assistance from expert knowledge and demonstrations. *User Modeling and User-Adapted Interaction* (2022).
[2] Antonio Andriella, Carme Torras, and Guillem Alenyà. 2020. Cognitive System Framework for Brain-Training Exercise based on Human-Robot Interaction. *Cognitive Computation* (2020).
[3] Chris Bateman and Richard Boon. 2005. *21st Century Game Design (game development series).* Charles River Media, Inc.
[4] Sera Buyukgoz, Jasmin Grosinger, Mohamed Chetouani, and Alessandro Saffiotti. 2022. Two ways to make your robot proactive: Reasoning about human intentions or reasoning about possible futures. *Frontiers in Robotics and AI* 9 (2022).
[5] Caitlyn Clabaugh, Kartik Mahajan, Shomik Jain, Roxanna Pakkar, David Becerra, Zhonghao Shi, Eric Deng, Rhianna Lee, Gisele Ragusa, and Maja Matarić. 2019. Long-Term Personalization of an In-Home Socially Assistive Robot for Children With Autism Spectrum Disorders. *Frontiers in Robotics and AI* 6 (2019).
[6] Anaís Garrell, Michael Villamizar, Francesc Moreno-Noguer, and Alberto Sanfeliu. 2017. Teaching Robot's Proactive Behavior Using Human Assistance. *International Journal of Social Robotics* 9 (04 2017).
[7] Jasmin Grosinger, Federico Pecora, and Alessandro Saffiotti. 2016. Making Robots Proactive through Equilibrium Maintenance. In *Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI).*
[8] Guy Hoffman, Maya Cakmak, and Crystal Chao. 2014. Timing in human-robot interaction. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction.* 509–510.
[9] Bahar Irfan, Aditi Ramachandran, Samuel Spaulding, Dylan F. Glas, Iolanda Leite, and Kheng Lee Koay. 2019. Personalization in Long-Term Human-Robot Interaction. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI).* 685–686.
[10] Matthias Kraus, Nicolas Wagner, and Wolfgang Minker. 2020. Effects of Proactive Dialogue Strategies on Human-Computer Trust. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20).* 107–116.
[11] Gianpaolo Maggi, Elena Dell'Aquila, Ilenia Cucciniello, and Silvia Rossi. 2021. "Don't Get Distracted!": The Role of Social Robots'Interaction Style on Users'Cognitive Performance, Acceptance, and Non-Compliant Behavior. *International Journal of Social Robotics* 13, 8 (2021), 2057–2069.
[12] Daniel C McFarlane and Kara A Latorella. 2002. The scope and importance of human interruption in human-computer interaction design. *Human-Computer Interaction* 17, 1 (2002), 1–61.
[13] George Armitage Miller. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review* 63, 2 (1956), 81.
[14] Zdzisław Pawlak. 2004. Decision Networks. In *Rough Sets and Current Trends in Computing.* Springer Berlin Heidelberg, 1–7.
[15] Zhenhui Peng, Yunhwan Kwon, Jiaan Lu, Ziming Wu, and Xiaojuan Ma. 2019. Design and Evaluation of Service Robot's Proactivity in Decision-Making Support Process. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19).* 1–13.
[16] Silvia Rossi, Marwa Larafa, and Martina Ruocco. 2020. Emotional and Behavioural Distraction by a Social Robot for Children Anxiety Reduction During Vaccination. *International Journal of Social Robotics* 12 (2020), 765–777.
[17] Thorsten Schodde, Kirsten Bergmann, and Stefan Kopp. 2017. Adaptive Robot Language Tutoring Based on Bayesian Knowledge Tracing and Predictive

Decision-Making. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction (HRI '17)*. 128–136.
[18] Chapa Sirithunge, AG Buddhika P Jayasekara, and DP Chandima. 2019. Proactive robots with the perception of nonverbal human behavior: A review. *IEEE Access* 7 (2019), 77308–77327.
[19] Hao Tan, Ying Zhao, Shiyan Li, Wei Wang, Ming Zhu, Jie Hong, and Xiang Yuan. 2020. Relationship between social robot proactive behavior and the human perception of anthropomorphic attributes. *Advanced Robotics* 34, 20 (2020), 1324–1336.