

Adaptive linguistic style for an assistive robotic health companion based on explicit human feedback

Hannes Ritschel, Andreas Seiderer, Kathrin Janowski, Stefan Wagner, Elisabeth André

Angaben zur Veröffentlichung / Publication details:

Ritschel, Hannes, Andreas Seiderer, Kathrin Janowski, Stefan Wagner, and Elisabeth André. 2019. "Adaptive linguistic style for an assistive robotic health companion based on explicit human feedback." In *Proceedings of the 12th ACM International Conference on Pervasive Technologies Related to Assistive Environments - PETRA '19, Rhodes, Greece, June 05 - 07, 2019*, edited by Fillia Makedon, 247–55. New York, NY: ACM Press.
<https://doi.org/10.1145/3316782.3316791>.



Adaptive Linguistic Style for an Assistive Robotic Health Companion Based on Explicit Human Feedback

Hannes Ritschel, Andreas Seiderer, Kathrin Janowski, Stefan Wagner and Elisabeth André

Human-Centered Multimedia, Augsburg University

Augsburg, Germany

{ritschel,seiderer,janowski,andre}@hcm-lab.de

ABSTRACT

In the future, an increasing amount of social robots will be found in our domestic environments to support and facilitate everyday life. Especially in the context of assistive, health-related support, more and more robotic products are on the way to the consumer market. While one can observe that many commercial efforts are put into the visual appearance, embodiment, motion and sound of companion robots, this paper focuses on the robot's conversational skills. We investigate how to adapt the robot's linguistic style to the individual user's preferences. This includes two forms of robot persona in the context of information retrieval tasks and games, as well as politeness with regard to recommendations. Therefore, we present an autonomous companion robot, which adapts its spoken language based on explicit human feedback. It provides several functionalities for information retrieval, reminders, communication and entertainment as well as health-related recommendations. Results of the in-situ study with elderly participants indicate that human preferences vary with regard to the robot's employed politeness strategies. Furthermore, the participants preferred assistant persona over companion persona in the information retrieval context.

CCS CONCEPTS

• **Computing methodologies** → **Intelligent agents**; *Reinforcement learning*; • **Human-centered computing** → Haptic devices;

KEYWORDS

social robots; assistive companions; linguistic style; adaptation; persona; politeness; reinforcement learning

ACM Reference Format:

Hannes Ritschel, Andreas Seiderer, Kathrin Janowski, Stefan Wagner and Elisabeth André. 2019. Adaptive Linguistic Style for an Assistive Robotic Health Companion Based on Explicit Human Feedback. In *The 12th Pervasive Technologies Related to Assistive Environments Conference (PETRA '19)*, June 5–7, 2019, Rhodes, Greece. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3316782.3316791>

1 INTRODUCTION

Thanks to technological advances of the last years more and more robotic products appear on the market – not only for scientific research, but especially for assistance and entertainment in users' domestic environments. Similar to the ubiquitous use of mobile devices and smart speakers for information retrieval, communication, notification, entertainment, tracking of everyday activities as well as health monitoring, social robots of the future aim to adopt and complete these functionalities. Consumer products, such as *Jibo*, *Buddy*, *Pillo*, *Mabu*, *ElliQ*, *Lynx* or *Moorebot*, already combine embodiment with essential assistive functionalities, smart home integration, entertainment and communication.

Communication with the human is a crucial part for these devices. Language, gaze, facial expression and gestures make it possible to create more natural interaction by mimicking human behaviors. Many efforts are put into the visual appearance, embodiment, motion and sound of domestic robots (e.g. *Jibo*, *Buddy*) or robotic toys (e.g. *Cozmo*, *Miko*) to portray certain personality stereotypes. However, spoken language in human-human interaction often carries the essential information, supplemented by non-verbal behaviors. Therefore, the robot's spoken language requires special attention and should be chosen carefully. Not only does the formulation impact the perceived personality profile but also the perceived persuasion, for example in the context of recommendations by robotic elderly assistants [7]. This is a major aspect to consider, especially with regard to assistive functionalities related to health like reminding of taking pills.

We explore how an assistive social robot should express itself in terms of linguistic style. The investigation is threefold: we focus on politeness with regard to health-related recommendations, as well as two forms of robot persona in the context of information retrieval tasks and games. Persona, meaning the robot's "fictional personality with varied and stable behavioral and personality patterns" [3], plays a key role to represent a compelling interaction partner. Furthermore, the selective use of politeness is important for the robot to achieve its own goals, e.g. when it has to persuade the user to do exercises or go for a walk.

Robots initially benefit from the novelty effect, but experience has shown that people tend to get bored after a short time. One issue is that the robot usually sticks to the same kind of behavior and is not able to dynamically adapt to the user's preferences. Therefore, social robots should learn how to express themselves in an appealing and convincing manner and adapt their behaviors to the individual human's needs and preferences, in particular when they serve as an assistive companion. In order to learn about the individual user's preferences regarding linguistic style, we present

a robotic companion which employs an autonomous learning approach based on explicit feedback from the human. While providing several applications for entertainment, health-related recommendations, information retrieval and communication, the robot adapts its language based on the user's positive or negative feedback. A hardware control panel is used for interacting with the robot and for providing the feedback signal.

The paper is structured as follows: Section 2 outlines related work covering robotic companions and linguistic style in the context of robots. Section 3 explains in detail the presented adaptation approach, including the communicative strategies, learning and prototype setup. Section 4 describes the results and feedback of an in-situ study with two senior users in their own homes.

2 RELATED WORK

We split up related work in two research areas: (1) robotic companions which serve to entertain, communicate and support the human and (2) linguistic style in the context of social robots. These areas outline important research and recent technology trends, which are combined by our work.

2.1 Robotic Companions

In order to assist humans in their everyday lives social companions provide support and advice to entertain, to help with information-seeking or cognitive tasks, or to improve mental and emotional wellbeing. With technological advances of the last years more and more products are on the way to the consumer market, especially for health care and entertainment. Typical features include information retrieval, communication with family members or medical professionals, calendar management, reminders and games. For example, *Pillo* is a domestic robot which monitors and dispenses medication, reminds the user of his or her care plan and answers questions concerning the nutritional value of food. Usually, these robots have either a physical face (e.g. *Mabu* and *Elliq*), or it is rendered on a (touch) screen (e.g. *Jibo* and *Buddy*). Interaction with the user is achieved with a voice interface, mobile app or touch.

Besides commercial products, companion robots have long been subject of research. The *AlwaysOn* project [28] developed a social companion agent for elderly people living on their own. It offers various activities for entertainment (such as games), for communication with other people, or exercises for the user's wellbeing. A notable feature is also the system's approach to develop a relationship to the user by planning shared activities and avoiding sensitive topics such as health advice in the early stages. It is "always on": the user can interact with it at any time.

Within the *CompanionAble* project, Schroeter et al. [26] present a socially assistive robot companion for older people suffering from mild cognitive impairment. The mobile robot is combined with smart home technologies to provide social and cognitive support. Functionalities include reminders of appointments and activities, video calls, storage of personal items and a cognitive stimulation game addressing the impairments of the target group.

When it comes to health-related robotic applications, a growing body of research in assistive robots is robot-assisted training [32, 33], covering the support of exercises, often combined

with gamification elements to increase enjoyment and intrinsic motivation. For example, Schneider et al. [25] use a NAO robot to give instructions for isometric exercises. Their results show that exercising with a humanoid companion is more effective than exercising alone. Besides post-stroke rehabilitation [30], people suffering from dementia [31] or autism spectrum disorders [13], research also addresses the elderly population: exercises for the user's physical and mental wellbeing include e.g. arm, finger or cognitive training [5]. Robots are used to give instructions, demonstrate the sequence of movements or poses and to provide feedback based on the user's performance. Study results by Fasola and Matarić [5] show a clear preference for robots over virtual agents in terms of social attraction, enjoyableness and helpfulness when performing exercises.

It seems a logical extension to integrate such functionality in domestic robotic companions. Today's consumer-grade domestic robot companions often do not yet have the flexibility to demonstrate and monitor difficult exercises due to their limited physical embodiment, sensors and actuators. However, these devices often have an integrated display which may be used to illustrate instructions. While displays raise the problem of distraction from the robot's embodiment, they are the best trade-off to present additional information. For example, within the *CARE* system [19], exercises for physical and mental wellbeing are presented in the form of text and images on a digital picture frame without a robot. The authors further combine exercises and recommendations with gamification aspects to increase user appreciation [8]. Companion robots offer the potential to integrate such functionality in the future.

2.2 Linguistic Style for Social Robots

With spoken language being a very important communication channel, the robot's spoken language is investigated and manipulated in a range of robot experiments. For example, Aly and Tapus [1] realize user-robot personality matching where both gestures and speech are adapted to the human's personality profile. While the authors employ Natural Language Generation (NLG) to generate the robot's spoken words accordingly Tapus et al. [30] use scripted utterances to express robot personality and explore adaption in post-stroke rehabilitation therapy. NLG is also used in the context of story telling for adapting the robot's linguistic style [23]. Linguistic style, as well as adaptation, furthermore address the robot's spoken language with regard to humor [21, 22, 34] based on Reinforcement Learning and human social signals [20].

In the context of socially assistive robots for the elderly, Bartl et al. [3] investigate the impact on acceptance of robot persona. While persona means the "fictional personality with varied and stable behavioral and personality patterns" [3], robots benefit from several modalities, including language, gaze behavior, facial expression and gestures, to express these characteristics. Bartl et al. use the robot to remind about appointments in a calendar application. The expressed persona results in either prevalent friendly, emotional, enthusiastic and expressive (companion) or more authoritative and formal spoken language (assistant) behavior. The robot expresses its corresponding persona with language (e.g. fillers, words of agreement, pronouns, etc.) as well as facial expression (smiles) and pose (head tilt). Results of their Wizard of Oz (WoZ) experiments indicate that the companion persona is preferred and

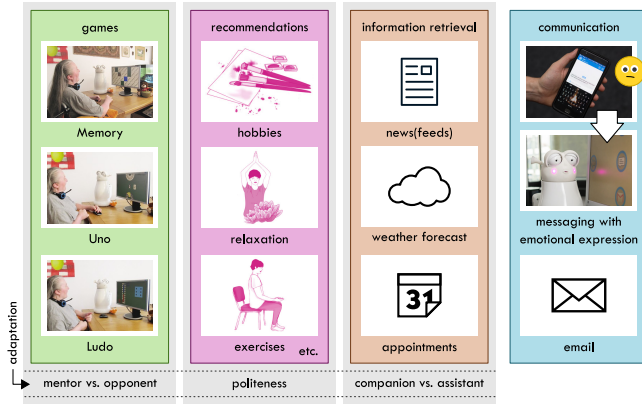


Figure 1: Overview of the robot's applications.

that it has a positive impact on the robot's likability and perceived social intelligence. Moreover, they report on the observed novelty effect as well as some desired functionalities like reminding about events and further recommendations like drinking water.

Huijnen et al. [9] point out that different users have different preferences regarding the role, responsibilities and type of persona expressed by companion robots and that there is a need to adjust the robot's character accordingly. Results of their studies also indicate that apart from the functionality the robot's character and interaction style are significantly more important than its design and physical embodiment.

With interaction becoming more and more natural, e.g. by using natural language to communicate with smart speakers, embodied agents also need to meet increased expectations about their social competence [10]. For example, research results show that politeness impacts the perceived persuasiveness of recommendations by robotic elderly assistants [7]. However, the authors also come to the conclusion that this perception varies greatly between age groups. They also observed that while some participants found polite wordings more persuasive, the opposite was the case for others. This indicates that the robot's use of politeness needs to be tailored to the individual user in order to be most effective.

In order to explore individual human preferences with respect to the linguistic style of assistive companions, we present an adaptive social robot which learns from explicit feedback. While providing assistive functionalities, including entertainment, health-related recommendations, information retrieval and communication, the experiment aims to explore individual human preferences over a longer period of time. Our contribution consists of the adaptation approach for linguistic style as well as insights from an in-situ study. In contrast to the WoZ study by Bartl et al. [3] we extend the robot's assistive functionalities and investigate linguistic style in three respects for a completely autonomous assistive companion. The evaluation is carried out in the field with the target elderly population in their own homes for one week.




Figure 2: User interface and robot expression examples.

3 ADAPTIVE LINGUISTIC STYLE FOR AN ASSISTIVE COMPANION

This paper combines a domestic robotic companion with an adaptation approach for its spoken language in order to explore the user's individual preferences. It offers typical assistive functionalities, such as information retrieval (news, weather forecast, appointments, contacts), health-related recommendations and communication (email, chat), as well as games (Memory, Uno, Ludo) and a set of scripted jokes. Figure 1 presents an overview of the functionalities.

The user can give explicit positive or negative feedback with regard to the robot's language. Based on this data, an autonomous adaptation process learns his or her preferences and tailors the robot's linguistic style to the individual user over time. The goal of this learning process is to optimize the spoken language to be most efficient and pleasing for the individual human in the corresponding application context. Motivated by the research by Bartl et al. [3], Gebhard et al. [6] and Hammer et al. [7], we focus on linguistic style as expressed in two different types of persona and politeness. The former takes two pairs of contrasting robot personas into account: *companion vs. assistant* for information retrieval tasks and *mentor vs. opponent* in the context of games. The latter covers eight politeness strategies from the literature for health-related recommendations.

3.1 Mentor vs. Opponent

The first aspect which is addressed by the adaptation process is the robot's language during gameplay. Based on the ideas of the virtual poker game described by Gebhard et al. [6], the companion robot is able to behave and express itself either as a *mentor* (collaborative) or an *opponent* (competitive). Collaborative robot behavior means that it behaves towards the shared goal of enjoying a social game together with the human. We are especially interested in which robot persona is more compelling in the context of games.

Both personas are configured based on the "Big Five" personality model [14]. In particular, they differ in the dimensions agreeableness and neuroticism. The *mentor* is configured to be more agreeable and less neurotic, which reinforces positive emotions. The *opponent* is less agreeable and more neurotic, and thus negative emotions

are amplified. Game events, such as losing a token or successfully revealing a pair of cards in the Memory game, are appraised based on the OCC model [18] from the perspective of the robot. Depending on the persona the same event elicits different robot emotions. The more agreeable *mentor* is configured to like the user, it shows positive emotions and expresses empathy for the human. The less agreeable *opponent* dislikes the user and therefore experiences resentment or gloating depending on the event.

Ultimately, these reactions are expressed in the robot's language. For example, when it loses the game, it may say "Congratulations! You won!" when it acts as *mentor* while the *opponent* persona can result in "You do not deserve that victory!" (both translated from German). In addition, the robot's comments are emphasized by corresponding non-verbal behaviors (see Figure 2). The virtual game board, cards, tokens, and game statistics are shown on the screen next to the robot.

3.2 Companion vs. Assistant

Bartl et al. [3] investigate the role of persona for social robots in the elderly domain. Based on their definition of *companion* and *assistant* persona, as well as the presented corresponding cues of prototypical behavior, our robot's adaptation mechanism explores the human's individual likings. The goal of the learning process is to identify which of the two personas is preferred in an information retrieval context.

The two types of persona are employed within the scope of several applications. These include news, weather forecast, appointments, contacts and messaging. When the user is interested in news feeds, the robot reads out loud headlines and the abstract of recent articles. When triggering the weather forecast for the current or next seven days, the robot generates a corresponding description and gives appropriate hints (e.g., to take along an umbrella or suitable clothing when it is getting cold). Furthermore, an overview of appointments in the calendar as well as an overview and details of contact information can be retrieved. Automatic reminders are triggered in configurable intervals before appointments.

Based on [3], the *companion* and *assistant* persona is reflected in the robot's language. Within the scope of these applications, the robot's language for instructions, suggestions, reminders, notifications, confirmations and descriptions differs e.g. with respect to fillers ("oh", "ah"), words of agreement ("okay", "alright", "good") or pronouns (informal "we" vs. formal "you"). Texts from external services, such as news article headlines, abstracts or message contents, are not modified.

3.3 Politeness Strategies

An important part of the robot's assistive functionalities is a set of pervasive health-care recommendations. As research by Hammer et al. [7] indicates that politeness impacts the human's perceived persuasion, the robot explores different politeness strategies. The adaptation mechanism selects the strategy based on the user's feedback. Depending on the employed strategy, the formulation changes, but the semantic content remains the same. The employed politeness strategies were identified by Johnson et al. [11] based on the theory by Brown and Levinson [4]. Accordingly, recommendations can be formulated by the robot either as (1) direct commands,

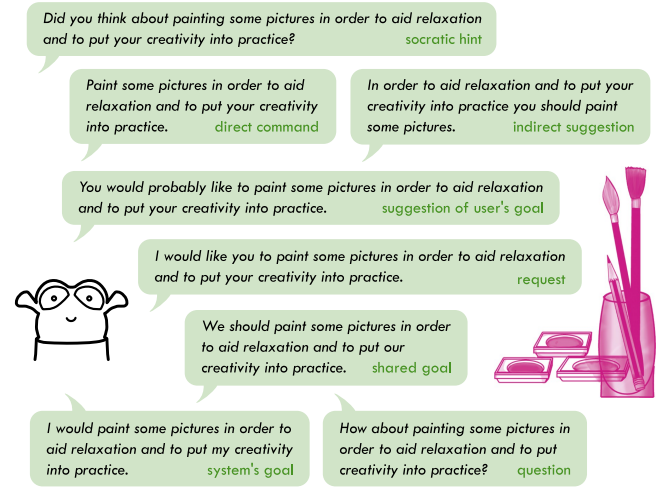


Figure 3: Examples of the same recommendation presented with different politeness strategies and an image (translated from German).

(2) indirect suggestions, (3) requests, (4) actions expressed as the robot's goal, (5) actions expressed as shared goal, (6) questions, (7) suggestions of user goals, or (8) socratic hints. Figure 3 illustrates an example recommendation for each strategy translated from German. Each of the recommendations is prepared as scripted text for each of the politeness strategies. In addition, the content is illustrated with images.

Recommendations target the elderly population to support their independence. They address physical, mental and environmental wellbeing and encourage activities, e.g. with respect to hobbies, relaxation, exercises and much more. They suggest reading books, airing and turning up the heating in the morning, making gymnastics, listening to music, watering the plants, drinking enough, and much more. Contents are adapted from the CARE system [19], where recommendations are presented with a digital picture frame with both text and images on the basis of context information acquired by sensors in the elderly user's domestic environment.

3.4 Adaptation Process

The robot employs Reinforcement Learning (RL) as a machine learning framework for exploring and learning about human preferences with regard to the robot's linguistic style. In RL, the autonomous agent investigates different actions iteratively and learns which one is the best of them based on scalar feedback: the reward signal.

Our learning tasks are defined as k -armed bandit problems [29], a simple form of RL, where the agent's goal is to find the best of k actions (see below). In each time step t , the agent selects an action $A_t \in \mathcal{A}$ from the set of actions \mathcal{A} , executes it and observes the received scalar reward R_{t+1} . The reward signal is used to update the estimated action value $Q_t(A_t)$, which indicates how good or bad A_t performs. Comparing the Q -values of all actions allows to estimate which action is the most effective. Over time and with increasing amount of rewards received for each action, the estimated Q -values become more and more precise. In each step, Upper Confidence




Figure 4: Overview of the adaptation process, involving one k -armed bandit problem for each application category.

Bound (UCB) action selection [29] is used for balancing exploitation and exploration, i.e., the agent's choice of the *greedy* (best) action with the highest Q -value versus exploring another one.

For learning about the robot's linguistic style there are three distinct agents (see Figure 4). Each of them focuses on the language characteristics of interest in the corresponding application context as described in Section 3.1, 3.2 and 3.3. The first one addresses the politeness strategies of recommendations, the second one explores the expression of persona in the context of games and the last one investigates the expression of persona for the information retrieval applications. Accordingly, the set of actions \mathcal{A} is defined as follows.

Within the context of recommendations, the actions correspond to the eight politeness strategies. The learning agent for information retrieval tasks and communication has two actions: it employs either a *companion* or *assistant* persona. During the games, the set of actions represents the *mentor* and *opponent* persona. Actions directly influence the robot's linguistic style: for each action, a set of scripted utterances is prepared for each application. During interaction, the robot chooses its text output from the pool of utterances associated with the last action of the learning agent.

For each of the robot's utterances the reward signal (+1 or -1) can be provided directly via the control panel's feedback buttons. They light up as soon as feedback is required. If the user provides the reward signal within 30 seconds the corresponding agent updates the Q -value of the associated action.

At every point in time only one agent is learning depending on the current application context. When switching to another application (e.g. from recommendations to games), the corresponding learning agent becomes active. Overall, the three agents learn about the individual user's linguistic style preferences based on his or her explicit feedback.

3.5 Hardware and Software Setup

The robotic companion consists of a Reeti robot, a screen, a control panel and a computer for control (see Figure 5). While the robot and screen serve as output medium the control panel is the interface for human input. It was implemented with physical buttons due to issues reported in the literature with respect to touch-based interaction in the context of the elderly domain. For example, Motti et al. [17] pointed out that touch-based interaction may alleviate barriers of getting started for elderly people, but at the same time they detected a number of usability issues, such as the timing of tapping gestures.

3.5.1 Control panel. The interaction with the robotic companion is realized completely throughout a custom built control panel (see Figure 6). It consists of ten physical buttons with integrated LED lights to (1) let the user give explicit positive or negative feedback




Figure 5: Hardware and software setup.

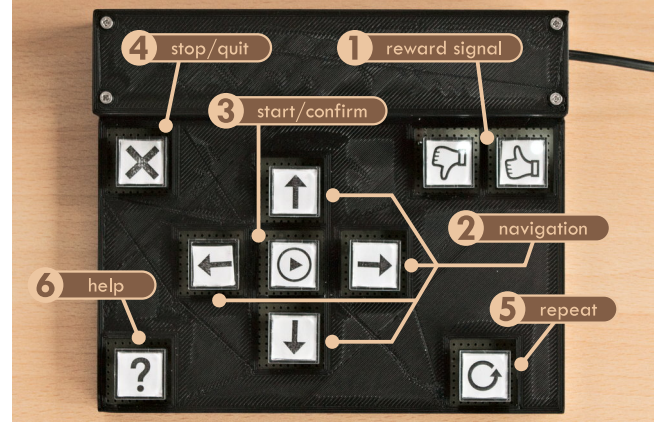


Figure 6: Control panel with physical buttons to interact with the robotic companion.

for the adaptation process, (2) navigate through the menu and choose selections, (3) start an application or confirm selections, (4) stop the robot's text output or quit an application, (5) repeat the robot's last utterance, and (6) provide help.

An Arduino compatible *Micro Pro* microcontroller is connected via USB to interface the buttons and LEDs. It is based on an *ATmega32U4* chip and is running with 3.3 V and a clock of 8 MHz. The controller is able to both send key presses like a regular PC keyboard and act as a serial device at the same time. Key presses are sent from the panel when a button is pressed, the serial connection allows to switch LEDs on or off. The power provided by a USB 2.0 port is sufficient to power all components of the control panel.

Since each of the buttons and LEDs has to be connected independently to the controller the count of I/O pins is not sufficient for the 10 buttons and 10 LEDs. Additionally, if all 10 LEDs are switched on, the current draw could exceed the maximum rating of 200 mA of the *ATmega32U4*. To solve these problems a *SX1509* I/O expander breakout board with 16 I/O ports was added to which all LEDs are

wired to. The board is connected via I²C to the microcontroller. All buttons are connected to the I/O ports of the microcontroller where the internal pull-up resistors are switched on to prevent floating inputs. The buttons are “debounced” by the firmware of the microcontroller. Each button press is sent as keyboard key press to the PC; messages from the PC via serial connection are used to light up the LEDs. The surrounding case is 3D printed and protects the electronics against damage.

3.5.2 Robot behaviors. In order to take advantage of the robot’s physical embodiment, it is equipped with animations, including gaze behavior and blinks. For example, during the games the robot shifts its gaze [15] towards the screen or user and shows basic emotions depending on its game progress (see Figure 2). After one minute of human inactivity the robot stops any idle animations: since it is designed as an “always on” [28] companion, it chooses a sleeping pose with closed eyes to protect the motors. Additionally, the LEDs in the robot’s cheeks light up to indicate that it is in “stand by” mode and can be reactivated as desired. The display turns off after three minutes to save energy. Pressing any button on the control panel reactivates the robot’s animations and display.

4 IN-SITU STUDY

Evaluating an assistive robot companion in a user’s home poses a number of challenges. For instance, people refused to let a robot into their home, which can be explained by the lower affinity with electronic devices that is still prevalent in the elderly population [7, 12]. Furthermore, our application is aimed at senior citizens who would like to have assistance but are still healthy enough to interact with the system, which further limited the number of candidates. Another problem was the fact that our hardware setup (such as the custom-made control panel) could only be placed in one single home at any given time, which constrained the duration of the study.

We therefore conducted a preliminary in-situ study with one female user, aged 64, and one male user, aged 61. Both were German native speakers. The robot stayed at each user’s home for one week. During that time they were free to interact with it as often as they wanted and to pick whatever application they were interested in. The robotic companion was completely autonomous. For example, calendar reminders were created automatically depending on the appointments entered by the user.

Participants were told that they should pay particular attention to the robot’s spoken language. Their feedback should be given depending on whether they like or dislike the way the robot expresses itself, independent of the semantic content or actual information presented in information retrieval tasks. They were not told which aspects of linguistic style were manipulated in detail.

4.1 Individual Preferences

In total, the participants interacted 133 and 135 times with the robot, respectively: Figure 8 plots the number of application invocations for each participant. Throughout these interactions, the female study participant provided feedback 979 times, the male participant 532 times. Based on the users’ feedback, the robot’s adaptation mechanism learned about their individual preferences as follows.




Figure 7: The female study participant plays the Memory game.




Figure 8: Number of application invocations.




Figure 9: Adaptation progress for all game activities (female test person).

4.1.1 Mentor vs. Opponent. Figure 9 illustrates the adaptation progress with respect to mentor/opponent persona in games for the female participant, who provided feedback 464 times. It plots both the values (higher is better) of each action and the reward (received from the control panel) over time. At the end the *mentor* persona emerges as the preference of both female and male participant with regard to the final values. However, if one looks at the complete plot and rewards of the female participant, the adaptation process did not receive enough negative feedback to extract a clear preference. With almost all feedback being positive, independent on the expressed linguistic style, both options have almost the same value. The difference between them is very small, so that the plot uses a logarithmic Y axis to emphasize them. For the male test person, who provided feedback 247 times, the *mentor* persona is also superior.




Figure 10: Adaptation progress for all information retrieval tasks (female test person).

However, there were not enough samples for the *assistant* persona in order to make a clear statement.

4.1.2 Companion vs. Assistant. In the context of the information retrieval tasks, the *assistant* persona turns out to be the favorite for both the female and male test person. Figure 10 illustrates the learning progress exemplary for the female test person, who provided feedback 298 times. It shows an initial exploration phase of roughly 150 learning steps, involving negative and positive feedback for both types of persona. After a longer period of positive feedback for the *assistant* persona, its value emerges as clear favorite for the learning algorithm. Subsequent negative feedback does not necessarily force the robot to switch back to the companion persona at once since there are enough positive samples for the assistive one.

For the male test person, who provided feedback 97 times, the learning agent initially primarily chose the *companion* persona. In the first place, it received both positive and less negative feedback. However, from the point of time when the learning agent selected the *assistant* persona, this one is ranked clearly higher and received only positive rewards from the user. In contrast to Bartl et al. [3], the results indicate that the *assistant* persona was preferred by both the female and male user.

4.1.3 Politeness Strategies. The adaptation results for the robot’s politeness are very different for the two subjects. Figure 11 presents the learning progress for the female test person, who provided feedback 217 times. In the end, recommendations formulated as the system’s goal, direct command or request were preferred most. Indirect suggestions and less direct formulations, such as socratic hints or the suggestion of the user’s goal were the worst performers.

Figure 12 illustrates the learning for the male test person, who provided feedback 188 times. In contrast to the female test person, requests had by far the poorest performance when it comes to received rewards. Similarly, indirect commands, the suggestion of user’s goal and the system’s goal were not rated very highly.

All in all, these results show that both test persons had very different preferences when it comes to politeness strategies. While the female test person preferred in particular short formulations (such as the direct commands or requests) and those which were formulated as the system’s goal, the male test person favored formulations which were not nearly as direct: questions, indirect suggestions and socratic hints.

4.1.4 Perception of Adaptation. To evaluate the adaptation process, we asked the participants after the study whether they had

noticed any changes in the robot’s verbal behavior. Both stated that they had not noticed any changes in the first place, which indicates that the adaptation of the robot’s linguistic style is subconscious. We interpret this as a positive insight due to the fact that the adaptation process itself did not have a directly noticeable negative impact on the perceived user experience. This potential problem was addressed in advance by employing UCB action selection, as this mechanism ensures that the robot does not stick to the greedy action permanently and encourages a certain degree of variety in the robot’s linguistic style in the long run. The only change in the robot’s language was noticed by the male participant, who identified the robot’s switching of informal versus formal pronouns (“Sie” vs. “Du” in German).

4.1.5 Content Bias and Repetitions. From the interaction logs we observed two types of bias which influenced the adaptation process. For the news retrieval application, the female participant’s feedback regarding the robot’s linguistic style was biased. One can observe that headlines with negative polarity lead to a consistently negative feedback in the interaction logs. Despite the instruction to ignore the semantic content of the robot’s language, this confirms that evaluating spoken language is challenging.

Another minor problem was the fact that the robot’s utterances were scripted. Many of them included variations with the same semantic meaning but different formulation. Since the final utterance was selected randomly from the pool of available variations, repetitions could occur occasionally. This also caused negative rewards by the female test person, which can be observed from the interaction logs.

5 FUTURE WORK

In the domestic environment further reward sources for adaptation of social companions should be explored. Especially with regard to health, current sensor technology makes it possible to get information about the user’s environment. With smart home components and body-worn technology becoming cheaper and accessible, these devices open up opportunities for further investigations. For example, in previous research an electronic kitchen scale with custom hardware has been used to acquire information about drink consumption [24]. This allows to explore and provide assistance in the context of nutrition and to improve recommendations by considering the personal and environmental context [27].

Moreover, Natural Language Generation is an important part for generation of individualized, linguistic content on-the-fly. For example, Anselma and Mazzei [2] use NLG in their virtual dietitian app to generate advice according to individual diet plans and food consumption, which is entered via the user interface. In combination with current sensing technology and adaptation, a social domestic companion could also contribute to health-aware food consumption by optimizing appropriate advice based on the user’s selection, reaction or behaviors.

Apart from the robot’s spoken language, the games could be augmented by physical objects in addition to the virtual screen to increase immersion and facilitate interaction. For example, Wrobel et al. [35] use a set of cards with QR codes and a custom hardware response box in their card game; Mehlmann et al. [16] use cards with tracking markers and a Microsoft Surface table in a puzzle




Figure 11: Adaptation progress for the presented recommendations (female test person).




Figure 12: Adaptation progress for the presented recommendations (male test person).

game with a NAO robot. Additional sensors could also be integrated for interactive exercises, e.g., to give the user feedback, active guidance and to provide personalized training. For example, Fasola and Mataric [5] use a Microsoft Kinect to track the user's performance of arm exercises based on the skeleton pose. Herpich et al. [8] propose a body-worn Microsoft Band 2 to record accelerometer data and a pressure-sensitive floor mat for physical exercises.

6 DISCUSSION AND CONCLUSION

In this paper, we presented a socially assistive companion robot, which adapts its linguistic style to the individual user's preferences. A Reinforcement Learning approach is used for exploring and optimizing the robot's language over time, focusing on the expression of two types of persona and politeness in different task contexts.

Building on previous research in Human-Robot Interaction (HRI) and Human-Computer Interaction (HCI), the robot's integrated applications address elderly people in their domestic environment. Inspired by typical features of robotic companions and assistance systems from research and recent robotic consumer products, the autonomous robot offers games, health-oriented recommendations as well as communication and information retrieval functionalities. It consists of a Reeti robot, which talks to the user and shows facial expressions, a display to present information, virtual game boards and the menu, as well as a hardware control panel.

We have also provided valuable insights from a preliminary in-situ study, where the robot was placed in participants' homes for one week. In general, the linguistic adaptation process was unobtrusive. Results indicate that human preferences indeed vary, which we identified with regard to the participants' different ratings for the robot's politeness in the context of health-related recommendations. In contrast, we could not confirm the human preference for robots with companion over assistant persona as reported in a former research study. We observed the opposite as result of the adaptation process within information retrieval tasks.

We also observed two types of bias, such as repetitions and external language contents in information retrieval tasks which contributed to the reward given by the user. While the former can be fixed by providing more scripted text variations or by using Natural Language Generation, the latter suggests being careful when evaluating polarized contents. This points out that the evaluation of spoken language is challenging.

We aim to verify our findings in a future study, building on the current prototype with more participants. All in all, it is possible that users are not consciously aware of which linguistic style they prefer, which would support a learning approach as opposed to manual configuration.

ACKNOWLEDGMENTS

This research was funded by the Bavarian State Ministry for Education, Science and the Arts (StMWFK) as part of the ForGenderCare research association.

REFERENCES

- [1] Amir Aly and Adriana Tapus. 2016. Towards an intelligent system for generating an adapted verbal and nonverbal combined behavior in human-robot interaction. *Auton. Robots* 40, 2 (2016), 193–209.
- [2] Luca Anselma and Alessandro Mazzei. 2018. Designing and testing the messages produced by a virtual dietitian. In *Proceedings of the 11th International Conference on Natural Language Generation, Tilburg University, The Netherlands, November 5-8, 2018*. Association for Computational Linguistics, 244–253.
- [3] Andrea Bartl, Stefanie Bosch, Michael Brandt, Monique Ditttrich, and Birgit Lugrin. 2016. The Influence of a Social Robot's Persona on How it is Perceived and Accepted by Elderly Users. In *Social Robotics - 8th International Conference, ICSR 2016, Kansas City, MO, USA, November 1-3, 2016, Proceedings (Lecture Notes in Computer Science)*, Vol. 9979. 681–691.
- [4] Penelope Brown and Stephen C Levinson. 1987. *Politeness: Some Universals in Language Usage (Studies in Interactional Sociolinguistics)*. Cambridge university press.
- [5] Juan Fasola and Maja J Matarić. 2013. A Socially Assistive Robot Exercise Coach for the Elderly. *J. Hum.-Robot Interact.* 2, 2 (jun 2013), 3–32.
- [6] Patrick Gebhard, Marc Schröder, Marcela Charfuelan, Christoph Endres, Michael Kipp, Sathish Pammi, Martin Rumpfer, and Oytun Türk. 2008. IDEAS4Games: Building Expressive Virtual Characters for Computer Games. In *Intelligent Virtual Agents, 8th International Conference, IVA 2008, Tokyo, Japan, September 1-3, 2008. Proceedings (Lecture Notes in Computer Science)*, Vol. 5208. Springer, 426–440.
- [7] Stephan Hammer, Birgit Lugrin, Sergey Bogomolov, Kathrin Janowski, and Elisabeth André. 2016. Investigating Politeness Strategies and Their Persuasiveness for a Robotic Elderly Assistant. In *Persuasive Technology - 11th International Conference, PERSUASIVE 2016, Salzburg, Austria, April 5-7, 2016, Proceedings (Lecture Notes in Computer Science)*, Vol. 9638. Springer, 315–326.
- [8] Madita Herpich, Thomas Rist, Andreas Seiderer, and Elisabeth André. 2017. Towards a Gamified Recommender System for the Elderly. In *Proceedings of the 2017 International Conference on Digital Health, London, United Kingdom, July 2-5, 2017*. ACM, 211–215.
- [9] Claire A. G. J. Huijnen, Atta Badii, Herjan van den Heuvel, Praminda Caleb-Solly, and Daniel Thiemert. 2011. "Maybe It Becomes a Buddy, But Do Not Call It a Robot" - Seamless Cooperation between Companion Robotics and Smart Homes. In *Ambient Intelligence - Second International Joint Conference on Aml 2011, Amsterdam, The Netherlands, November 16-18, 2011. Proceedings*. 324–329.
- [10] Kathrin Janowski, Hannes Ritschel, Birgit Lugrin, and Elisabeth André. 2018. *Sozial interagierende Roboter in der Pflege*. Springer Fachmedien Wiesbaden, Wiesbaden, 63–87.
- [11] W. Lewis Johnson, Richard E. Mayer, Elisabeth André, and Matthias Rehm. 2005. Cross-Cultural Evaluation of Politeness in Tactics for Pedagogical Agents. In *Artificial Intelligence in Education - Supporting Learning through Intelligent and Socially Informed Technology, Proceedings of the 12th International Conference on Artificial Intelligence in Education, AIED 2005, July 18-22, 2005, Amsterdam, The Netherlands (Frontiers in Artificial Intelligence and Applications)*, Vol. 125. IOS Press, 298–305.
- [12] Katja Karrer, Charlotte Glaser, Caroline Clemens, and Carmen Bruder. 2009. Technikaffinität erfassen – der Fragebogen TA-EG. *ZMMS Spektrum* 29 (09 2009).
- [13] Changchun Liu, Karla Conn, Nilanjan Sarkar, and Wendy Stone. 2008. Online Affect Detection and Robot Behavior Adaptation for Intervention of Children With Autism. *IEEE Trans. Robotics* 24, 4 (2008), 883–896.
- [14] RR McCrae and OP John. 1992. An introduction to the five-factor model and its applications. *Journal of personality* 60, 2 (1992), 175.
- [15] Gregor Mehlmann, Markus Häring, Kathrin Janowski, Tobias Baur, Patrick Gebhard, and Elisabeth André. 2014. Exploring a Model of Gaze for Grounding in Multimodal HRI. In *Proceedings of the 16th International Conference on Multimodal Interaction, ICMI 2014, Istanbul, Turkey, November 12-16, 2014*. ACM, 247–254.
- [16] Gregor Mehlmann, Kathrin Janowski, and Elisabeth André. 2016. Modeling Grounding for Interactive Social Companions. *KI* 30, 1 (2016), 45–52.
- [17] Lilian Genaro Motti, Nadine Vigouroux, and Philippe Gorce. 2013. Interaction techniques for older adults using touchscreen devices: a literature review. In *Proceedings of the 25th IEME conference francophone on l'Interaction Homme-Machine, IHM '13, Talence, France, November 12-15, 2013*. ACM, 125–134.
- [18] Andrew Ortony, Gerald L. Clore, and Allan Collins. 1988. *The Cognitive Structure of Emotions*. Cambridge University Press.
- [19] Thomas Rist, Andreas Seiderer, Stephan Hammer, Marcus Mayr, and Elisabeth André. 2015. CARE - extending a digital picture frame with a recommender mode to enhance well-being of elderly people. In *9th International Conference on Pervasive Computing Technologies for Healthcare, PervasiveHealth 2015, Istanbul, Turkey, 20-23 May 2015*. IEEE, 112–120.
- [20] Hannes Ritschel. 2018. Socially-Aware Reinforcement Learning for Personalized Human-Robot Interaction. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*. International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM, 1775–1777.
- [21] Hannes Ritschel and Elisabeth André. 2018. Shaping a Social Robot's Humor with Natural Language Generation and Socially-Aware Reinforcement Learning. In *Workshop on Natural Language Generation for Human-Robot Interaction at INLG 2018*. Tilburg, The Netherlands.
- [22] Hannes Ritschel, Ilhan Aslan, David Sedlbauer, and Elisabeth André. 2019. Irony Man: Augmenting a Social Robot with the Ability to Use Irony in Multimodal Communication with Humans. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems.
- [23] Hannes Ritschel, Tobias Baur, and Elisabeth André. 2017. Adapting a Robot's Linguistic Style Based on Socially-Aware Reinforcement Learning. In *26th IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN 2017, Lisbon, Portugal, August 28 - Sept. 1, 2017*. IEEE, 378–384.
- [24] Hannes Ritschel, Andreas Seiderer, Kathrin Janowski, Ilhan Aslan, and Elisabeth André. 2018. Drink-O-Mender: An Adaptive Robotic Drink Adviser. In *Proceedings of the 3rd International Workshop on Multisensory Approaches to Human-Food Interaction (MHFI'18)*. ACM, New York, NY, USA, Article 3, 8 pages.
- [25] Sebastian Schneider, Michael Goerlich, and Franz Kummert. 2017. A framework for designing socially assistive robot interactions. *Cognitive Systems Research* 43 (2017), 301–312.
- [26] Ch. Schroeter, Steffen Müller, Michael Volkhardt, Erik Einhorn, Claire A. G. J. Huijnen, Herjan van den Heuvel, Andreas van Berlo, Andreas Bley, and Horst-Michael Gross. 2013. Realization and user evaluation of a companion robot for people with mild cognitive impairments. In *2013 IEEE International Conference on Robotics and Automation, Karlsruhe, Germany, May 6-10, 2013*. IEEE, 1153–1159.
- [27] Andreas Seiderer, Stephan Hammer, Elisabeth André, Marcus Mayr, and Thomas Rist. 2015. Exploring Digital Image Frames for Lifestyle Intervention to Improve Well-being of Older Adults. In *Proceedings of the 5th International Conference on Digital Health 2015, Florence, Italy, May 18-20, 2015*. ACM, 71–78.
- [28] Candace L. Sidner, Timothy W. Bickmore, Bahador Nooraie, Charles Rich, Lazlo Ring, Mahni Shayganfar, and Laura Vardoulakis. 2018. Creating New Technologies for Companionable Agents to Support Isolated Older Adults. *TiS* 8, 3 (2018), 17:1–17:27.
- [29] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning - An Introduction (Second Edition)*. MIT Press.
- [30] Adriana Tapus, Cristian Tapus, and Maja J. Matarić. 2008. User - robot personality matching and assistive robot behavior adaptation for post-stroke rehabilitation therapy. *Intelligent Service Robotics* 1, 2 (2008), 169–183.
- [31] A. Tapus, C. Tapus, and M. J. Matarić. 2009. The use of socially assistive robots in the design of intelligent cognitive therapies for people with dementia. In *2009 IEEE International Conference on Rehabilitation Robotics*. 924–929.
- [32] Konstantinos Tsiakas, Maher Abujelala, and Fillia Makedon. 2018. Task Engagement as Personalization Feedback for Socially-Assistive Robots and Cognitive Training. *Technologies* 6, 2 (2018).
- [33] Konstantinos Tsiakas, Maria Kyrarini, Vangelis Karkaletsis, Fillia Makedon, and Oliver Korn. 2018. A Taxonomy in Robot-Assisted Training: Current Trends, Needs and Challenges. *Technologies* 6, 4 (2018).
- [34] Klaus Weber, Hannes Ritschel, Ilhan Aslan, Florian Lingensfelder, and Elisabeth André. 2018. How to Shape the Humor of a Robot - Social Behavior Adaptation Based on Reinforcement Learning. In *Proceedings of the 2018 on International Conference on Multimodal Interaction, ICMI 2018, Boulder, CO, USA, October 16-20, 2018*. ACM, 154–162.
- [35] Jérémy Wrobel, Ya-Huei Wu, Hélène Kerhervé, Laila Kamali, Anne-Sophie Rigaud, Céline Jost, Brigitte Le Pévédic, and Dominique Duhaut. 2013. Effect of agent embodiment on the elder user enjoyment of a game. In *ACHI 2013-The Sixth International Conference on Advances in Computer-Human Interactions*.