Drink-o-mender: an adaptive robotic drink adviser

Hannes Ritschel, Andreas Seiderer, Kathrin Janowski, Ilhan Aslan, Elisabeth André

Angaben zur Veröffentlichung / Publication details:

Ritschel, Hannes, 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: Boulder, CO, USA — October 16 - 20, 2018*, edited by Anton Nijholt, Carlos Velasco, Marianna Obrist, Katsunori Okajima, and Charles Spence, 3. New York, NY: ACM Press. https://doi.org/10.1145/3279954.3279957.

Drink-O-Mender: An Adaptive Robotic Drink Adviser

Hannes Ritschel, Andreas Seiderer, Kathrin Janowski, Ilhan Aslan and Elisabeth André
Human-Centered Multimedia, Augsburg University
Augsburg, Germany
{ritschel,seiderer,janowski,aslan,andre}@hcm-lab.de

ABSTRACT

Social robots become increasingly important in the domain of healthcare and maintenance. Nutrition is not an exception: research robots are used in several experiments to teach people about healthy nutrition. Moreover, robotic products, whose task is to keep track of the user's nutrition, to provide tips, reminders, and to recognize anomalies in health-related behaviors are on the way to market. To convince users of a robot's recommendations, speech is an important interaction modality. However, automatically adapting a robot's spoken advise depending on users' behaviors is still a challenge. We address this issue by building Drink-O-Mender, an interactive installation, which includes a Reeti robot augmented with additional sensing and adaptation abilities. The robot is designed to offer drinks in a social setting. It aims to convince users of consuming healthy drinks while adapting its spoken advices depending on the users' selected beverage choices. The installation is equipped with custom hardware including a smartscale to sense the type and quantity of consumed drinks. We describe the interactive installation in detail and demonstrate feasibility of generating adaptive spoken advices by reporting on insights gained from exhibiting the installation during a public event and observing interactions of 78 users with the robot.

ACM Reference Format:
Hannes Ritschel, Andreas Seiderer, Kathrin Janowski, Ilhan Aslan and Elisabeth André. 2018. Drink-O-Mender: An Adaptive Robotic Drink Adviser. In *3rd International Workshop on Multisensory Approaches to Human-Food Interaction (MHFI'18), October 16, 2018, Boulder, CO, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3279954.3279957

1 INTRODUCTION

The popularity of health and well-being related applications in app stores suggest that people are increasingly interested in integrating technological solutions to monitor their individual everyday activities and to reflect on their overall health and well-being. A variety of specialized apps and devices can already be utilized to track personal and behavioral data, including apps which apply machine learning techniques to improve the quality of nutrition data logging on smart watches [5]. These developments open up new possibilities for healthier and more sensible diets, particularly with regard to the problem of obesity among western populations and studies showing that, for example sugar contents in some drinks marketed to children are unacceptably high [3].

While diets may be considered a personal choice, food and drink consumption is often a social activity, celebrated in public and private spaces. Thus, tips and recommendations with regard to healthy nutrition in homes and public gatherings need to be provided through adequate and socially acceptable interfaces for multiuser settings. We believe that social companions, such as embodied social robots or contemporary voice assistants may be appropriate for the task. In order to explore the idea of an intelligent social robot as a nutrition adviser we have implemented a prototype whose aims is two-fold, (i) to shed light on the actual nutritional value of different drinks, and (ii) to learn from interactions with a heterogeneous multitude of people how to convince the user to select a more healthy one.

Drink-O-Mender consists of a social robot, which is placed on a table together with different fruit juices and iced tea, and custom hardware, which allows to sense the user's selected beverage and quantity. This data is used to optimize the robot's information presentation strategy with the aim of supporting the user's healthy nutrition. We use a *Reeti* robot with an expressive face and text-to-speech output since it has been shown that a three-dimensional embodied avatar has some advantages in a three-dimensional environment when it comes to human decision making [29]. In addition, initial insights on the adaptation approach, the user acceptance and usability of the prototype are provided, which we gained from exhibiting the system for four hours at a public scientific event in the field and observing interactions with more than 70 users.

The paper is structured as follows: Section 2 presents related work involving social robots in the nutrition context as well as adaptation based on Reinforcement Learning (RL). Section 3 gives an overview of the interaction, learning, hardware and implementation. Section 4 describes first insights drawn from a scientific exhibition where the prototype was shown to the general public.

2 RELATED WORK

We split up related work in two research areas: (1) research about social robots in the nutrition context, where they are employed to support the user's health (awareness) (2) Reinforcement Learning for personalization of social robots, which adapt their behaviors to improve or facilitate Human-Robot Interaction (HRI). Both areas summarize important research and technology trends, which are combined by our work.

2.1 Social Robots and Nutrition

Providing nutrition information is one of the applications provided by current health assistance devices. The recently-developed domestic robot *Pillo* [32] acts as a health manager for its family. Its main purpose is to dispense pills to family members, who are identified by face recognition, and to ensure that none of them miss their scheduled doses. It also manages health-related appointments such as sports training and supports video communication with medical professionals or emergency calls to family members. Furthermore, it is supposed to answer questions about the amount of calories in certain foods. *Mabu* is another commercial robot with similar goals which is currently under development.

Kidd and Breazeal [15] presented a robotic coach, which assists people with weight loss and maintenance for long-term, supportive care. Their anthropomorphic robot with head and eye movement is able to carry out a short conversation regularly to talk about daily goals for diet and exercise. Within the conversation, the user has to report about the nutritional intake and amount of exercises conducted, while the robot gives feedback, advice or suggestions and makes small talk to evaluate whether the user enjoys the interaction or not. In [16], the authors compare the talking robotic interface with a touch-screen computer (same software but neither head nor voice) and a paper log for tracking a user's calorie intake and amount of exercise. Their results show that the long-term motivation of subjects who used the social robot was significantly higher than that of the other two groups and they also formed the strongest "working alliance" with their assistive device.

Rich et al. [25] developed an *Always-On* companion agent to support elderly people living on their own. The activities which could be initiated by the agent also include counseling the user for diet and exercise, which is a very personal topic. Therefore, the agent is not allowed to bring it up before the relationship to the user has reached a certain stage. The appropriate stage for each activity is defined by the authors of the content plug-ins, as is the increase in closeness gained by completing them. The system's relationship manager then schedules health-related conversations for a time after the bond has been strengthened by less sensitive activities.

Cruz-Maya and Tapus [4] evaluated the role of embodiment and voice in a multimedia learning scenario about nutrition and healthy eating tips. In their experiment, the authors investigate whether the robot's physical embodiment improves learning in comparison with a tablet only. Information about healthy diet is presented with text and images as slides on a tablet user interface embedded in a Kompai robot, accompanied with synthesized or recorded voice. Their results did not show an improvement in the user's learning performance when using a robot, which according to the authors could be attributable to the fact that participants concentrated only on the tablet with the presented content.

Similarly, results by Hammer et al. [9] showed equal persuasion for health-related recommendations (including nutrition) presented either by a robotic elderly assistant or a tablet PC. However, the robot was perceived as more usable, less complex and easier to learn. Participants felt slightly more confident and had a slightly stronger willingness to use it at home.

Learning about healthy nutrition is also subject of the research by Short et al. [30]. Results of their study with 26 first-grade children show an increase in engagement and high levels of enjoyment when interacting with a *DragonBot* robot. The authors conduct a Wizard of Oz (WoZ) experiment, where a teleoperator provides dialog selection while the overall interaction flow is controlled autonomously. Three nutrition topics are part of different learning sessions: "packing lunch boxes", which also includes selection of non-sugary drinks, "choosing after-school snacks" and "building balanced meals".

Kruijff-Korbayová et al. [20] investigated the role of off-activity talk as part of a system designed for long-term support of diabetic children. Their results show that off-activity talk impacts the motivation of participants to play again with a NAO robot.

All in all, social robots provide the opportunity to increase human awareness about health and nutrition in a world with growing number of digital assistants and companions. They can provide important related information in an engaging, multi-modal manner in a diversity of environments.

2.2 Adaptive Social Robots

Social robots, which adapt their behaviors to human users, are used in a variety of settings. This section outlines existing works based on Reinforcement Learning [31], an on-line, autonomous machine learning framework. Similar approaches include for example the TAMER framework [19] and POMDPs, which are also used for assistance in healthcare [13].

There are multiple options to provide the RL feedback signal (reward) in HRI scenarios. Apart from task-related information like user performance (e.g. in exercises/games), human social signals can be used, including smile and gaze [8, 12, 21], laughter [35], tactile [2] or prosodic [17] feedback, interaction distance, gaze meeting, motion speed timing [23], gesture and posture [24, 27], or gaze direction [6]. Another option is to use physiological data from ECG [22] or EEG [34].

Application scenarios include for example learning of social behavior [2], student tutoring [8], the assistive domain including post stroke rehabilitation [33] and intervention for children with autism spectrum disorder [22]. Similar to the scenario at hand, which focuses on the adaptation of presented information through speech, entertainment is one research area of social robots, where speech and content presentation are of central importance, too. Scenarios include Japanese Manzai [11], standup comedy [14, 18], joke [35] and story telling [27], where presented contents or their delivery (e.g. animation, sound or voice parameters) are optimized in real-time for an individual user or a larger audience. RL is also used to learn the best strategy for maintaining long-term user engagement, for example when playing games [21].

Similar to the aforementioned examples, our adaptive approach tries to assist the human with healthy nutrition. In contrast to Kidd and Breazeal [15], we do not focus on tracking consumed food but on optimizing the content of presented information via speech. We use a social robot as interaction partner since additional social behaviors expressed by the robot provide the ability to increase the human's interest [20] and embodied agents might also have advantages in a three-dimensional environment particularly with regard to human decision making [29].

To facilitate the process of sensing consumed drinks, we developed hardware which weighs the drinking vessel automatically with a smartscale, so that it becomes no longer necessary to explicitly report this data via speech or text input. Moreover, the user's selected beverage is sensed by custom hardware, too. Both serve as feedback signal in terms of reward for the adaptation process.

3 AN ADAPTIVE ROBOTIC DRINK ADVISER

Social robots provide a range of possibilities for communication, including speech, gaze or gestures. In the context of supporting the user's nutrition, the most obvious of them is speech. To convince the human of healthy nutrition, our social robot does not only present information, but is also equipped with an autonomous learning process. It adapts the spoken content in terms of scripted statements over time depending on the users' selection of drinks. Our prototype, which was shown at a public scientific exhibition, optimizes the learned behavior not for individual users in terms of personalization but for the average population. For this event, the robot is equipped with a set of recommendations/statements in different categories: they either put emphasis on the high amount of calories in drinks with high nutritional value or draw attention to those with the lowest amount of calories.

Figure 1 illustrates the general idea: a social robot presents information about several drinks to the user. Its overall goal is to encourage healthy behavior by making the human aware of the nutritional value. This information is stored in a database. It includes the amount of calories, which is addressed in the robot's utterances presented to the human user. Moreover, it is used to calculate rewards for the learning process, which selects the robot's statements depending on the user's beverage selection and quantity.

Custom hardware was developed to obtain this information and to make Drink-O-Mender completely autonomous. It includes vessel holders to identify the user's selection and a smartscale, which is able to act both as sensor and as actuator. While the ability to light up and animate the built-in LED ring is mainly included for facilitating the interaction to attract attention and to clarify the internal state of the application, the sensor for weighing the cup is essential for the adaptation process. Together with the information from the vessel holders, this data serves as a basis to adapt the robot's strategy (see Section 3.1). The approach was implemented in a prototype, which combines the adaptive behavior and hardware with a Reeti robot, including non-verbal behaviors.

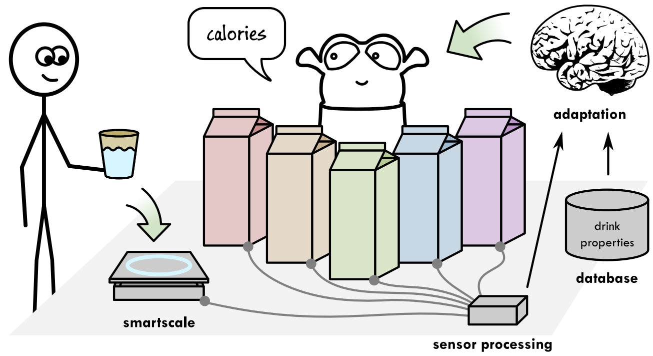

Figure 1: Conceptual prototype setup.

3.1 Adaptation Process

Task of the robot is to convince the user of healthy nutrition by exploring different formulations and their influence on the user's decision. For this purpose, a set of *actions* in the form of scripted texts was prepared. Each of them has several alternatives with the same semantic content, but different formulation to make the interaction more varied. Goal of the adaptation process is to find the most convincing action so that the user chooses a healthy drink.

Finding the most profitable action from a set of actions $\mathcal{A}$ over a sequence of learning steps can be formalized as n-armed bandit problem [31], a reduced form of Reinforcement Learning, as there is no notion of state but only a set of actions to evaluate. For each time step t, the bandit executes an action $a_t \in \mathcal{A}$ and receives a reward r_t, which is used to calculate the action's current value $Q_t(a)$. Actions are selected by either choosing a random one (exploration) or the *greedy* action a^* (exploitation), which is the one with the highest estimated value $Q_t(a^*) = \max_{a \in \mathcal{A}} Q_t(a)$ at time t.

At the exhibition, content was focused on the amount of calories. In our learning task, actions correspond to different categories of scripted statements. Those in the first category try to discourage the visitor from choosing drinks with high nutritional value by explicitly highlighting this fact right away (e.g. "Grape juice has a relatively high amount of calories. Apple and orange juice are in the mid-field range."). Conversely, the robot tries to draw attention to those drinks with the lowest amount of calories in the second category. More options can be included easily by adding more statements or information about drinks for further experiments.

Each interaction starts with ϵ-greedy action selection, i.e. either execution of a random action with probability ϵ or the greedy action with probability $1-\epsilon$. After talking about the drinks, the user makes his or her choice, selects a vessel, fills up a cup and drinks up. The reward $r_t \in [0; 1]$ for learning is calculated based on the quantity and drink properties from the database. It takes the quantity q_t, maximum filling level q_{max} of the cup (both in milliliters) and nutritional value $kcal_t$ of the drink into account, with $kcal_{max}$ being the maximum amount of kilo-calories of all drinks:

$$r_t = 1 - \frac{kcal_t}{kcal_{max}} \cdot \frac{q_t}{q_{max}}$$

While the prototype optimizes with regard to the amount of calories only (for example, drinking large quantities with high nutrient concentration results in a low reward), reward calculation can be easily adjusted for further experiments. At the end of the interaction, the value $Q_t(a)$ of the selected action a is updated with k_a being the number of times that a has been executed so far:

$$Q_t(a) = \frac{k_a - 1}{k_a} \cdot Q_{t-1}(a) + \frac{1}{k_a} \cdot r_t$$

In general, adaptation allows to optimize behavior either for individual users which interact over a longer period of time, or for multiple users. The prototype was presented in a non-laboratory setting where people could not be identified and did not return to the stand. Thus, we were interested in whether the robot adapts to the rewards based on the data collected throughout all interactions with a variety of persons (see Section 4.1).

3.2 Hardware Setup

In order to provide a reliable data source for reward calculation of the adaptation process, custom sensor and actuator devices were developed. A hardware overview is visible in Figure 2, details in 3.

3.2.1 Vessel Detection. First of all we needed a robust way to detect which vessel was chosen by the user. For this task we 3D-printed five

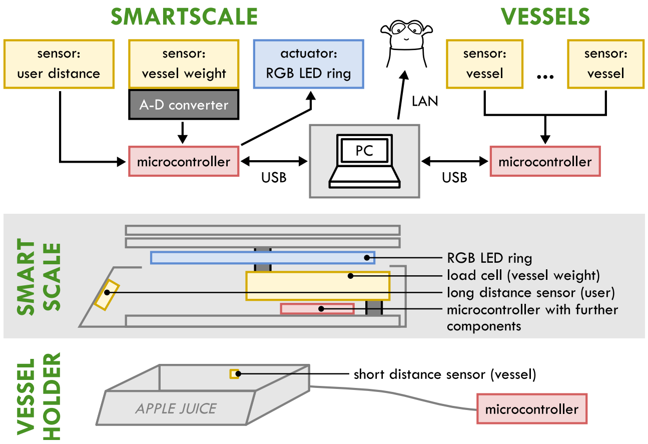

Figure 2: Hardware setup.

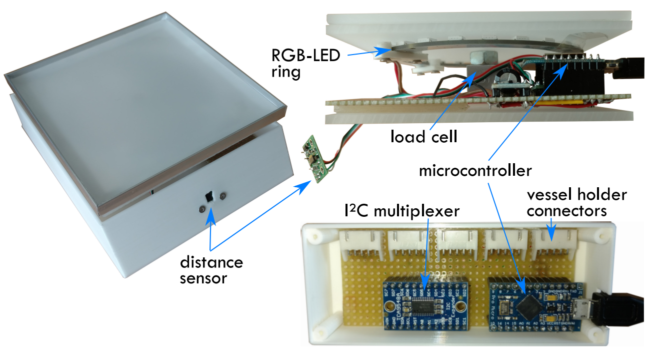

Figure 3: Electronics of the custom hardware prototypes. Left / above two views of the scale: with and without the 3D-printed surrounding case. Below a top-view of the micro-controller unit of the vessel holders with removed case top.

vessel holders, each containing an optical near-distance presence sensor (APDS-9960[1]). This type of sensor is also used in smartphones, for example to detect whether there is an object or ear in front of the phone so that the display can be switched off.

The sensors are connected via five pin JST connectors to an Arduino Pro Micro 3.3 V (ATmega32U4) micro-controller. Since all sensors use the I²C bus and have the same address an I²C multiplexer is required. We used the eight channel TCA9548A I²C multiplexer breakout board from Adafruit[2]. The Arduino is connected via USB to the PC. We implemented a firmware for the Arduino that is able to switch between the sensors via the multiplexer and detects a vessel with a maximum distance of about three centimeters. The internal sample rate is reduced to about 10 Hz. Whenever the state of a vessel holder changes an event is sent to the PC using the USB serial connection.

[1]https://www.broadcom.com/products/optical-sensors/integrated-ambient-light-and-proximity-sensors/apds-9960
[2]https://learn.adafruit.com/adafruit-tca9548a-1-to-8-i2c-multiplexer-breakout/overview

3.2.2 Smartscale. We developed a smartscale that is able to provide information about the amount of fluid (weight) and the user presence in front of the system, which triggers or stops the interaction. Additionally, it can be used as an actuator as it is possible to control a RGB-LED (WS2812B) ring with 24 LEDs. It can be regarded as an extended stationary version of the mobile smartscale presented in [28].

To be able to measure the weight of a vessel placed on the scale a loadcell with a maximum load of 5 kg is used. The loadcell is connected to a HX711 analog-digital-converter with I²C. To detect a user in front of the system we integrated an I²C Time-of-Flight laser distance sensor breakout board from Pololu (VL53L0X[3]) which is able to measure distances of up to 200 cm. The sensors and LED ring are connected to a WeMos D1 mini ESP8266 micro-controller. For this setup we did not require the Wifi capability of the board but as this controller provides much higher performance and memory than an AVR-based Arduino it was easily possible to integrate several LED color animations with a smooth refresh rate (30 Hz) while still being able to provide sensor data of the loadcell and distance sensor of about 30 Hz.

The brightness of the LEDs is reduced by software so that the power (max. 500 mA) of the USB 2 connection is enough to provide a bright light of the LEDs which are easily visible even in a bright environment. Short but high current peaks of the LEDs and micro-controller are handled by an 1000 µF capacitor. WS2812B LEDs require 5 V TTL logic so that a 3.3 V / 5 V level-shifter was required to allow stable operation with the ESP8266.

The micro-controller is connected to the PC via USB and uses a serial connection for communication. The LED color of the whole ring and different animations can be set via the serial connection.

To provide good stability the loadcell was screwed between two acrylic milk glass plates with a thickness of 3 mm. A second acrylic milk glass plate was glued with double-sided tape on top of the upper plate so that the screws are hidden and the light of the RGB-LEDs gets more diffused. The surrounding case of the lower part was 3D-printed and the front plate, to which the distance sensor is screwed to, has a slope of 30° so that a nearby human in front of the smartscale can be detected. To protect the electronics from accidentally spilled fluid a heightened rim out of an aluminum profile was glued to the rim of the top plate.

3.2.3 Firmware. Sensor data processing is partly done by the micro-controller but also by the sensors themselves. The calibrated weight data is filtered with a moving average of four samples at a sample rate of 30 Hz by the ESP8266. A threshold is used to detect the removal or reset of a vessel in the holders. Only state changes are communicated by the ATmega32U4 firmware.

3.3 Data Processing and Interpretation

Further data processing and communication is done on the PC which communicates with the sensors via serial connection. For this task we developed a lightweight generic event stream processor called *Eventerpretor*. It allows the creation of multiple parallel stream pipelines consisting out of threaded nodes which are connected using *LinkedBlockingQueues* with timeouts to prevent locks.

[3]http://www.st.com/en/imaging-and-photonics-solutions/vl53l0x.html

The following abstract base node types exist: input, transform and output. For this setup we used for example serial input nodes, moving average transform nodes, threshold transform nodes with delays and postgresql and serial output nodes.

To communicate with a serial connection in both directions, which is required for the smartscale, it is necessary to open just a single connection. Multiple connections are usually not allowed by the OS. To handle this problem a global manager was implemented which is able to share the connections.

We added additional options to a threshold node, since, for example, the distance sensor used for user detection should not immediately switch to *present* and *absent* whenever the distance reached or fell below a specific value. This node allows us to control the delay for outputting the presence or absence. This was especially required as it was expected that users just pass by the system or shortly move out of the detection range of the distance sensor during system interactions. Additionally, we used this node type to prevent wrong weight values by shortly delaying the weight output because the weight plate of the scale can slightly swing after placing or removing a weight.

To communicate with other programs we decided to utilize the postgresql database connection which is used to send notifications and receive them without polling since the database was already used to store drink data. This mechanism also offers an easy way to add logging of the exchanged data in the future.

3.4 Interaction

The interaction itself was created with *Visual SceneMaker* [7], which allows modeling agent behavior as parallel, hierarchical finite state machines. It communicates with the hardware via postgresql notifications to trigger animations on the smartscale and to receive user presence, current weight as well as the state of the vessel holders. A specialized software module translates the SceneMaker's command messages for speech and animation actions to the Reeti API and reports their execution status back to the application.

Reeti's behavior is split up in two main states: *idle* and *interaction*. The former triggers random comments, animations and gaze behavior on the robot when no visitor is present. When a person stands in front of the setup, the latter handles both low-level dialog behavior (pronouncing and interrupting planned utterances) and the procedure of small talk, presentation of drinks/offering them, instructions, handling user inputs, hardware actuation and learning.

Figure 4 shows the part of the state machine which models the high-level interaction flow. The interaction with the visitor is grouped into several sections, each of which is represented by a super-node containing the fine-grained actions. Most of these super-nodes are executed in a linear manner, but we also included branches for error handling. For example, we did not allow the users to mix different juices because this would interfere with the reward calculation for the learning process. Therefore, if the visitor happened to pick up more than one juice vessel at the same time, the robot would tell them to use only one.

After the greeting phase, Drink-O-Mender chooses a recommendation strategy based on what has been learned from previous interactions. According to the chosen strategy, Reeti either draws attention to the two juices with the lowest amount of calories or to

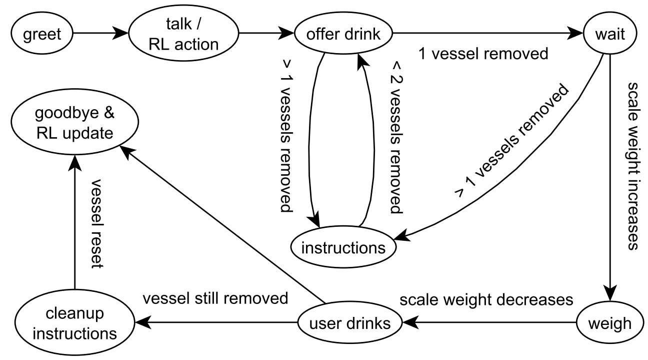

Figure 4: Simplified state machine of the high-level interaction flow.

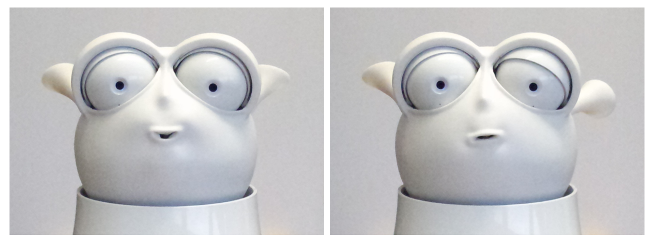

Figure 5: Examples for Reeti's facial expressions.
***Left:* happy. *Right:* confused.**

the one with the highest amount, before informing the visitor that the remaining two juices have a medium amount of calories). This phase (labeled *talk / RL action* in Figure 4) has to take place before the user picks up a juice vessel because otherwise the recommendation would not be able to influence their choice.

3.5 Nonverbal Signals

To make the interaction with the robot more natural, we made use of Reeti's movable head and expressive face. The Reeti robot has three degrees of freedom in the neck which enables it to look at different points of interest, such as the scale or the visitor. During the idle phase, the gaze target is chosen at random. In the interaction phase, it is set explicitly at certain points in the dialog in order to show Reeti's attention focus. For example, Reeti looks at the visitor when greeting him or her to show that it is aware of their presence. However, when it tells the visitor to pick one of the juice vessels, it turns its head down in order to draw attention to the items in question. These movements are executed in parallel to the dialog by means of a separate *gaze* state machine.

Furthermore, we use emotional facial expressions to illustrate Reeti's mood. For this purpose, we play back predefined animation sequences for the servo motors in Reeti's face. These motors then move Reeti's ears and eyelids or deform the robot's silicon skin around the mouth. During the interaction phase Reeti is mostly smiling in order to make the visitor feel welcome and respected. In case of error handling it would switch to a confused expression. Its emotional state is more varied in the idle phase in order to attract curious passers-by. During that phase, Reeti makes various

Figure 6: Prototype setup at the exhibition with Reeti robot, vessels, vessel holders and smartscale (front).

comments about the situation to express excitement, nervousness or impatience about its upcoming performance, surprise about the number of visitors, or neutral musings about colleagues who might visit as well. The commands for the matching facial expressions are embedded in the script for Reeti's speech and executed while the sentence is being spoken. Two expressions are shown in Figure 5.

4 LESSONS LEARNED

The prototype was presented in May 2018 at the "Lange Nacht der Wissenschaft" ("Long Night of Science"), a popular scientific exhibition in Augsburg, Germany, which took place for the very first time with more than 2000 visitors (see Figure 7). Due to the non-laboratory setting with noise and poor lighting (which precluded the use of speech recognition or face tracking), the time limit, huge rush and many factors beyond our control, this demonstration was not intended as a formal evaluation. It did not include a questionnaire about the visitors' background or their detailed perception of our system. The interaction itself was designed to be relatively short to accomodate the busy schedule of the visitors. We used this opportunity to gain first subjective insights about the acceptance and usability of the system.

4.1 Adaptation Results

Five drinks with different nutritional value were selected for the exhibition (see Figure 6): tomato juice and iced tea (least calories), apple and orange juice (medium calories) as well as grape juice (most calories). In four hours, 78 visitor interactions were logged. While the actual amount of people interacting with the robot was higher, some of these interactions were not recorded because people went away without finishing the interaction. Moreover, some of the recorded interactions were not consistent or successful. For example, some people did not understand the robot completely because of noise or they did not wait until it completed its instructions and mixed them up. Sometimes several people tried to interact at the same time, which was not identifiable for the robot.

Figure 7: A visitor interacts with the robot.

Figure 8 plots the drink quantities and rewards for 78 recorded interactions. It illustrates the connection between drink type, quantity and reward, which is the measurement of "success" from the algorithmic perspective: the higher the nutritional value, the lower the reward. When looking at the robot's "success" in terms of persuasion, i.e. whether the user selected a drink with low nutritional value, the results are as following: tomato juice and iced tea were chosen 21 times, apple and orange juice 25 times and grape juice 32 times. This indicates that spoken information about the calorie intake alone might not be sufficient.

Ten of the interactions were not carried out in the intended form, for example, some people put the much heavier vessel on the smartscale instead of the cup, which is attributable to intentional or unintentional misuse of the system (see Section 4.3). This can be seen in the plot at any learning step where the measured quantity exceeds the maximum capacity of the cup (100 milliliters), i.e. in interaction 0, 2, 14, 15, 32, 43, 53, 58, 67 and 71. Depending on the drink type, this caused the minimum reward value zero for grape juice in interaction 14, 15 and 58 (see Section 3.1), since this was the drink with the largest amount of calories.

During the exhibition, the exploration rate of the n-armed bandit was set to $\epsilon = 0.2$. As a result of the adaptation process, the robot's favored action was the one which explicitly emphasized the high amount of calories of grape juice. The experiment showed that the data from the hardware can serve as a reliable source for reward calculation of the adaptation process, however, the argumentation of the robot must be improved to assess whether the adaptation approach really works.

4.2 Misleading Affordances

Most visitors assumed that the robot was capable of speech recognition, which was not implemented because we had anticipated that there would be too much background noise in this setting. One reason for this expectation could be that the robot addressed them using speech, so they tried to respond in kind. Another may be that the robot appeared as a social being with a face, and therefore the visitors tried to interact in a natural manner. For similar reasons, many visitors expected the robot to use the eye cameras for face tracking or tried to use gestures for indicating their selected beverage. These observations show that the robot needs to make better use of human-like communication behavior in the future.

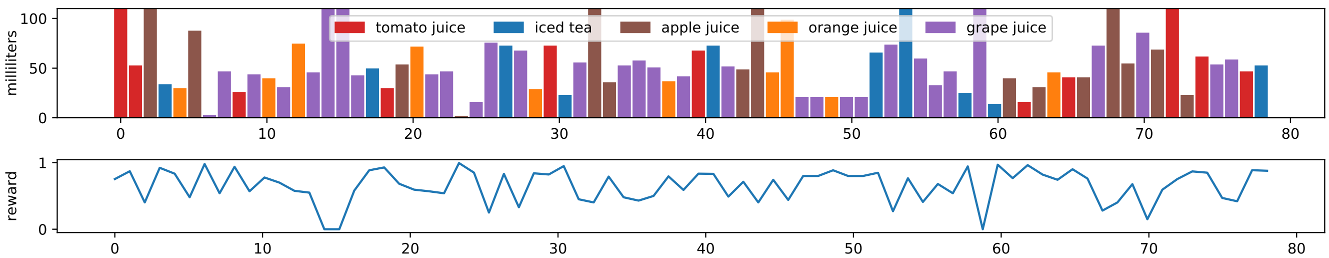

Figure 8: Consumed drinks and rewards for all interactions.

There was also confusion about the scale: some users tried to place the selected juice pack on it. One possible explanation could be that we did not allow the cup to be filled on top of the scale in order to avoid damage by spilled liquids. Therefore, the scale was not occupied when a visitor picked up the juice pack, while its prominent position and LED animation implied that something needed to be placed on it. Since this affordance was visible long before the cup was to be placed on the scale, it appears that the visitors chose to interact with the scale using the only movable item at that time.

4.3 Acceptance of the Robot

Overall, the visitors appeared to like the robot. Many of them called Reeti cute (mostly women, but also at least two men), making excited comments like "this one's adorable. Oh god!" or "oh, there's a little robbie! Oh, that's cute." One person stated "I'll take him with me now, he's funny". A third visitor pointed out that Reeti appeared non-threatening and rather submissive. These observations are in line with the findings reported by Wu et al. [36]. According to them, elderly people who were interviewed about assistive robots preferred robot models which were small and discrete, especially those they described as "lovely", "cute" or "funny", with a creature-like appearance that was not too close to that of a human.

At least one visitor saw potential in asking Reeti for nutrition information as opposed to manually searching via smartphone app, and another one stated their desire to have something like that in their own home. However, others were less enthusiastic about the idea. One visitor said it would be "weird to set up such a thing", and another one believed that nobody would listen to the robot's advice because it was merely "a toy" without authority. Yet another chose to interact with Reeti despite hating robots, but was more interested in messing with the system to test its reaction. The latter two observations may be caused by Reeti's submissive appearance and interaction style. To avoid patronizing the user, we had decided that Reeti would only make polite suggestions, but would not criticize the user for making a suboptimal choice. Furthermore, although many people liked Reeti for its cute and non-threatening appearance (see above), this factor may have contributed to its lack of authority, leading them to take the robot less seriously.

In order to motivate the users, the robot needs to be both appealing and persuasive. Besides the physical appearance, finding the proper balance between politeness and persuasiveness is still an open challenge, since it depends on a multitude of situational and individual factors [10]. It is likely that this aspect of our prototype can be improved by adapting not only the content of the recommendation, but also the manner in which it is presented.

Some users also wished for more information regarding the drinks. One criticized the choice of juices, complaining about the conventional farming methods used in their production, and stated "he's not informed enough for my liking, unfortunately." Others expected more detailed information about the actual content of the filled cup. This implies that the user's acceptance of the application can be greatly improved by taking more than just the calories into account. For example, the system could reason about vitamins or a user's preference for organic farming, and adding a speech interface would enable Reeti to answer questions about those properties of the available drinks.

5 CONCLUSION

We have outlined that a social robot capable of speech and facial expression would be a promising interface to provide spoken nutrition advice. To explore our assumption, we built Drink-O-Mender, an adaptive system, consisting of a Reeti robot and additional sensing hardware which enables the robot to adapt its spoken advice to the type and quantity of drinks consumed by human users.

We have described the integrated adaptation mechanism to reward or punish the robot's behavior, which is based on a reduced form of Reinforcement Learning. Our motivation for the specific mechanism was to have the robot learn advising users in a manner which would result in users choosing more healthy drinks, for example, through steering the user away from high-calorie drinks and drawing their attention to those with fewer calories. The presented approach has several benefits, including (1) real-time adaptation in the background, which allows to adapt the robot behavior immediately after each interaction, (2) easy to use hardware and robotic interface, as the user only picks up the vessel and fills up the cup, which is weighed automatically, and (3) a small footprint of the smartscale and vessel holders. Furthermore, the system implements a sensing option for consumed drinks based on custom hardware instead of video or audio processing, which may be apt for privacy sensitive studies including "in the wild" studies.

We have also provided valuable insights gained from observing users interacting with the robotic drink adviser during a public event, such as how the robot's (perceived) cuteness aroused interest, but its lack of (attributed) authority was problematic. In our future work, we aim to address the prototype's learning process, which should be expanded to explore and learn about different kinds of robot behaviors and their influence on the users' drink selection. One option is to incorporate politeness strategies for recommendations [10] or an interactive dialog system with different states of information presentation [26] to optimize the robot's strategy. We also plan to explore alternative sensing techniques for the tracking of different drink vessels, such as computer vision. Furthermore, we intend to conduct a formal evaluation to measure the performance of the learning algorithm and to confirm the subjective observations made during this first public presentation.

ACKNOWLEDGMENT

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] 2014. *The 23rd IEEE International Symposium on Robot and Human Interactive Communication, IEEE RO-MAN 2014, Edinburgh, UK, August 25-29, 2014*. IEEE.

[2] Rémi Barraquand and James L. Crowley. 2008. Learning polite behavior with situation models. In *Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction, HRI 2008, Amsterdam, The Netherlands, March 12-15, 2008*. ACM, 209–216.

[3] Jane Boulton, Kawther M Hashem, Katharine H Jenner, Ffion Lloyd-Williams, Helen Bromley, and Simon Capewell. 2016. How much sugar is hidden in drinks marketed to children? A survey of fruit juices, juice drinks and smoothies. *BMJ Open* 6, 3 (2016).

[4] Arturo Cruz-Maya and Adriana Tapus. 2016. Teaching nutrition and healthy eating by using multimedia with a Kompai robot: Effects of stress and user's personality. In *16th IEEE-RAS International Conference on Humanoid Robots, Humanoids 2016, Cancun, Mexico, November 15-17, 2016*. IEEE, 644–649.

[5] Simon Flutura, Andreas Seiderer, Ilhan Aslan, Chi Tai Dang, Raphael Schwarz, Dominik Schiller, and Elisabeth André. 2018. DrinkWatch: A Mobile Wellbeing Application Based on Interactive and Cooperative Machine Learning. In *Proceedings of the 2018 International Conference on Digital Health, DH 2018, Lyon, France, April 23-26, 2018*. ACM, 65–74.

[6] Pierre Fournier, Olivier Sigaud, and Mohamed Chetouani. 2017. Combining artificial curiosity and tutor guidance for environment exploration. In *Workshop on Behavior Adaptation, Interaction and Learning for Assistive Robotics at IEEE RO-MAN 2017*.

[7] Patrick Gebhard, Gregor Mehlmann, and Michael Kipp. 2012. Visual SceneMaker - a tool for authoring interactive virtual characters. *J. Multimodal User Interfaces* 6, 1-2 (2012), 3–11.

[8] Goren Gordon, Samuel Spaulding, Jacqueline Kory Westlund, Jin Joo Lee, Luke Plummer, Marayna Martinez, Madhurima Das, and Cynthia Breazeal. 2016. Affective Personalization of a Social Robot Tutor for Children's Second Language Skills. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*. AAAI Press, 3951–3957.

[9] Stephan Hammer, Kim Kirchner, Elisabeth André, and Birgit Lugrin. 2017. Touch or Talk?: Comparing Social Robots and Tablet PCs for an Elderly Assistant Recommender System. In *Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction, HRI 2017, Vienna, Austria, March 6-9, 2017*. ACM, 129–130.

[10] 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*. Springer International Publishing, Cham, 315–326.

[11] Kotaro Hayashi, Takayuki Kanda, Takahiro Miyashita, Hiroshi Ishiguro, and Norihiro Hagita. 2008. Robot *Manzai*: Robot Conversation as a Passive-Social Medium. *I. J. Humanoid Robotics* 5, 1 (2008), 67–86.

[12] Jacqueline Hemminghaus and Stefan Kopp. 2017. Towards Adaptive Social Behavior Generation for Assistive Robots Using Reinforcement Learning. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction, HRI 2017, Vienna, Austria, March 6-9, 2017*. ACM, 332–340.

[13] Jesse Hoey, Craig Boutilier, Pascal Poupart, Patrick Olivier, Andrew Monk, and Alex Mihailidis. 2012. People, sensors, decisions: Customizable and adaptive technologies for assistance in healthcare. *TiiS* 2, 4 (2012), 20:1–20:36.

[14] Kleomenis Katevas, Patrick GT Healey, and Matthew Tobias Harris. 2015. Robot Comedy Lab: experimenting with the social dynamics of live performance. *Frontiers in psychology* 6 (2015).

[15] Cory D. Kidd and Cynthia Breazeal. 2007. A Robotic Weight Loss Coach. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, July 22-26, 2007, Vancouver, British Columbia, Canada*. AAAI Press, 1985–1986.

[16] Cory D. Kidd and Cynthia Breazeal. 2008. Robots at home: Understanding long-term human-robot interaction. In *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*. IEEE, 3230–3235.

[17] Elizabeth S Kim and Brian Scassellati. 2007. Learning to refine behavior using prosodic feedback. In *Development and Learning, 2007. ICDL 2007. IEEE 6th International Conference on*. IEEE, 205–210.

[18] Heather Knight. 2011. Eight Lessons Learned about Non-verbal Interactions through Robot Theater. In *Social Robotics - Third International Conference, ICSR 2011, Amsterdam, The Netherlands, November 24-25, 2011. Proceedings (Lecture Notes in Computer Science)*, Vol. 7072. Springer, 42–51.

[19] W. Bradley Knox and Peter Stone. 2009. Interactively shaping agents via human reinforcement: the TAMER framework. In *Proceedings of the 5th International Conference on Knowledge Capture (K-CAP 2009), September 1-4, 2009, Redondo Beach, California, USA*. ACM, 9–16.

[20] Ivana Kruijff-Korbayová, Elettra Oleari, Ilaria Baroni, Bernd Kiefer, Mattia Coti Zelati, Clara Pozzi, and Alberto Sanna. 2014. Effects of off-activity talk in human-robot interaction with diabetic children, See [1], 649–654.

[21] Iolanda Leite, André Pereira, Ginevra Castellano, Samuel Mascarenhas, Carlos Martinho, and Ana Paiva. 2011. Modelling Empathy in Social Robotic Companions. In *Advances in User Modeling - UMAP 2011 Workshops, Girona, Spain, July 11-15, 2011, Revised Selected Papers (Lecture Notes in Computer Science)*, Vol. 7138. Springer, 135–147.

[22] 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.

[23] Noriaki Mitsunaga, Christian Smith, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. 2008. Adapting Robot Behavior for Human–Robot Interaction. *IEEE Trans. Robotics* 24, 4 (2008), 911–916.

[24] Anis Najar, Olivier Sigaud, and Mohamed Chetouani. 2016. Training a robot with evaluative feedback and unlabeled guidance signals. In *25th IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN 2016, New York, NY, USA, August 26-31, 2016*. IEEE, 261–266.

[25] Charles Rich, C Sidner, Bahador Nooraei, and William Coon. 2012. Operating in a hierarchy of time scales for an always-on relational agent. In *Workshop on Real-Time Conversations with Virtual Agents, Santa Cruz, CA (September 2012)*.

[26] Verena Rieser and Oliver Lemon. 2011. *Reinforcement Learning for Adaptive Dialogue Systems - A Data-driven Methodology for Dialogue Management and Natural Language Generation*. Springer.

[27] 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.

[28] Andreas Seiderer, Simon Flutura, and Elisabeth André. 2017. Development of a Mobile Multi-device Nutrition Logger. In *Proceedings of the 2Nd ACM SIGCHI International Workshop on Multisensory Approaches to Human-Food Interaction (MHFI 2017)*. ACM, New York, NY, USA, 5–12.

[29] Kazuhiko Shinozawa, Futoshi Naya, Junji Yamato, and Kiyoshi Kogure. 2005. Differences in effect of robot and screen agent recommendations on human decision-making. *Int. J. Hum.-Comput. Stud.* 62, 2 (2005), 267–279.

[30] Elaine Short, Katelyn Swift-Spong, Jillian Greczek, Aditi Ramachandran, Alexandru Litoiu, Elena Corina Grigore, David Feil-Seifer, Samuel Shuster, Jin Joo Lee, Shaobo Huang, Svetlana Levonisova, Sarah Litz, Jamy Li, Gisele Ragusa, Donna Spruijt-Metz, Maja J. Mataric, and Brian Scassellati. 2014. How to train your DragonBot: Socially assistive robots for teaching children about nutrition through play, See [1], 924–929.

[31] Richard S. Sutton and Andrew G. Barto. 1998. *Reinforcement learning - an introduction*. MIT Press.

[32] Vincent Tao, Kristy Moy, and Vibhuti Arya Amirfar. 2016. A little robot with big promise may be future of personalized health care. *Pharmacy Today* 22, 9 (Sep 2016), 38.

[33] Adriana Tapus, Cristian Tapus, and Maja J. Mataric. 2008. User - robot personality matching and assistive robot behavior adaptation for post-stroke rehabilitation therapy. *Intelligent Service Robotics* 1, 2 (2008), 169–183.

[34] Konstantinos Tsiakas, Maher Abujelala, and Fillia Makedon. 2018. Task Engagement as Personalization Feedback for Socially-Assistive Robots and Cognitive Training. *Technologies* 2 (May 2018), 49.

[35] Klaus Weber, Hannes Ritschel, Florian Lingenfelser, and Elisabeth André. 2018. Real-Time Adaptation of a Robotic Joke Teller Based on Human Social Signals. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*. International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM, 2259–2261.

[36] Ya-Huei Wu, Christine Fassert, and Anne-Sophie Rigaud. 2012. Designing robots for the elderly: Appearance issue and beyond. *Archives of Gerontology and Geriatrics* 54, 1 (2012), 121 – 126.