Infinite Personality Space for Non-fungible Robots

Daniel H. Grollman(✉)

Sphero, Inc., 4772 Walnut St, Suite 206, Boulder, CO 80301, USA
dan@sphero.com

Abstract. We outline a novel method for defining robot personality for the purposes of individual differentiation. Rather than a designer-developed set of behaviors where a users' preferences are learned and inserted into pre-written scripts, our approach allows for each robot to have and express a unique personality. This uniqueness reduces the fungibility of the robots, which may lead to increased user engagement.

Keywords: Human-robot interaction · Robot personality

1 Introduction

Robots that co-exist with non-specialized, untrained users are more and more frequently being developed to leverage 'social' capabilities to smooth their interactions. Beyond the basic technologies such as natural-language processing, gaze tracking, and theory of mind, there are two concepts of interest that have proven popular in the industry. The first, *personality*, seeks to imbue a social robot with a coherent entity-hood, described with vague but human-understandable terms such as 'helpful,' 'whimsical,' 'sassy,' or 'sparkling.' This goal is often achieved by hand-crafting robot behaviors based on a designer's understanding of how the desired personality would be expressed, and this personality and the robot's expression thereof is fixed for the life of the robot.

To be more enticing to a wider audience, multiple different personalities may be developed, and consumers enabled to select from within this fixed set. To further differentiate between individual robots (which may share the same physical form and base personality), the second concept, *personalization*, aims to let individual users make the robot's form and behavior more unique. Users are often able to (and do) modify the surface characteristics of their robot via paint, stickers, markers, etc., or even via use and wear. The software on the robot is generally not as malleable, but developers often allow for some user customization by selecting similarly 'surface' characteristics such as voice, gender pronoun, graphical avatar, etc. The available options define a restricted, often discrete space of robot 'characters,' which can still be somewhat easily replicated between users. (We ignore here the relatively small, but robust, 'hacker' or 'maker' communities that delve much deeper and change the hardware and software of robots in ways unintended or unimagined by the producing company.)

© Springer International Publishing AG 2016
A. Agah et al. (Eds.): ICSR 2016, LNAI 9979, pp. 94–103, 2016.
DOI: 10.1007/978-3-319-47437-3_10

These available characters are often further personalized over the the life of interaction between a user and their robot, by having the robot learn preferences of the former and adjust its behavior. For example, a user's affinity for pizza could be learned (by mining past food orders) and pizza could be suggested as a solution when the user expresses hunger. We characterize these sorts of adaptations as 'slotting in' the user's preference into a pre-scripted response (offering preferred food when hungry).

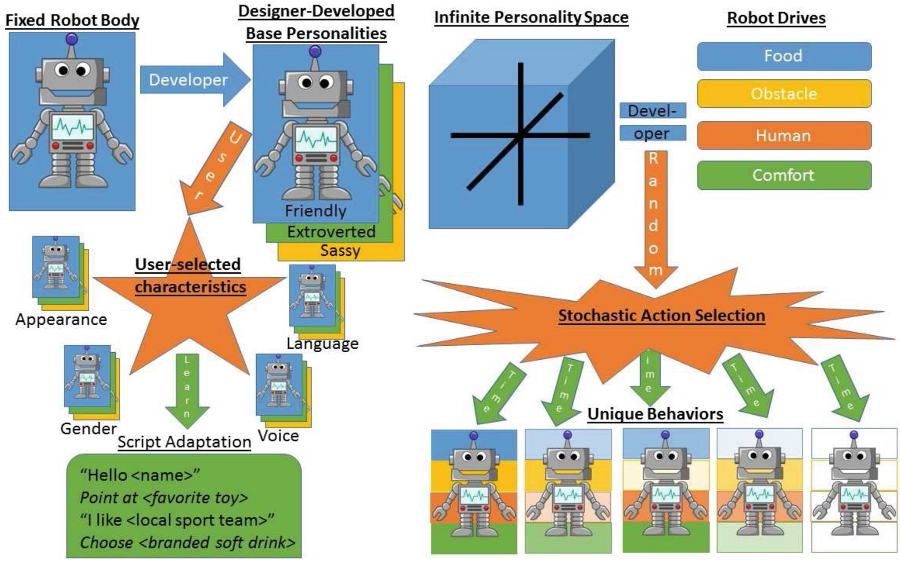

(a) Standard Approach (b) Our Approach

Fig. 1. Differing approaches to robot personality and personalization. In the standard approach (left), a set of robot personalities are developed by designers with explicit customization options that can be selected by the user. Additionally, verbal or behavioral 'scripts' can have slots for utilizing learned preferences. In our approach (right), a developed set of robot 'drives' defines an infinite personality space. Drive parameters are used in the stochastic selection of robot actions, resulting in differentiated, unique behavior over time.

We claim that these approaches (illustrated in Fig. 1a) are non-scalable due to their reliance on developer and designer time to create both the personalities and the myriad scripts required for personalization. While adaptive learning can lead to differentiated robots, these robots aren't truly unique, as they are still following the same script, with only shallow changes. Instead, we propose a method (outlined in Fig. 1b) for defining an infinite personality space that a robot can occupy, leading to truly unique robots. This approach is complementary to current personalization techniques, and could be even made adaptive, where the robot's personality adjusts over time in a way not currently possible.

2 Related Work

As robots move beyond the traditional niches of industrialized automation, where they are often segregated from humans, and military uses, where the users can be heavily trained, they are increasingly expected to operate side by side with untrained, naive users, who may have minimal exposure to technology, let alone autonomous robots. Several strands of research have indicated that social interaction cues can be used to facilitate interaction between these humans and robots, by enabling robots to conform to expected roles and processes familiar from human-human, or human-animal interaction [10]. For example, exaggerated motions, which may be sub-optimal from a robot energy expenditure point of view, can aid humans in better recalling, enjoying, and understanding the intent of a robot's reaching gesture [3]. This approach is related to the concepts of transparency, legebility and predictability [2], which states that it is important that a robot make its internal state (goals, 'thinking process') intelligible to its human counterparts, in order to speed collaboration.

Likewise, attending to human social norms is an area of active work. Several groups are investigating ways to have robots learn the 'rules of the road' that govern human navigation, both vehicular [11] and pedestrian [7]. Note that these rules are often implicit, hard to articulate, and culturally dependent, and so cannot be pre-programmed, and instead are often learned from demonstration. Some work has instead looked at recognizing human reactions to violations of these rules, in an effort to have a robot self-adjust [13].

Beyond collaborative benefits such as improved task performance [4], social behaviors on robots have been shown to impact a human's trust and compliance with a robot [6]. Robots with social behaviors such as politeness are also deemed more intelligent, capable and approachable than their non-social counterparts [9]. Thus, these behaviors can be used to counteract a general negative perception of robots due to their portrayal in popular media [12], as well as compensate for errors in their actual behaviors [1].

Socially active robots have further been shown to have positive impacts on the humans themselves, beyond any functional task the robot is made to perform. Sometimes, this impact is the point, such as in diet-assistance [5] or fitness coach robots, whose whole purpose is to use social skills to help people modify their behavior. Similarly, social companion robots have been shown to have positive health benefits such as stress reduction and increased tolerance for pain [14].

While the full range of impact of robot social behaviors on human interactants is still unknown, the commercial robot industry has embraced this approach whole-heartedly. The number of robots purporting social interaction capabilities on the market has soared, with the most recognizable perhaps being Jibo (Fig. 2a) and Pepper (Fig. 2b), both of which have been marketed more as friends and companions than appliances. Other products, such as the ZARO hospital assistance robot are built upon common platforms, such as the Nao (Fig. 2c). The possibility of multiple companies developing different applications and personalities for the same robot raises the specter of physically identical systems behaving in wildly different ways, and no research that we are aware of

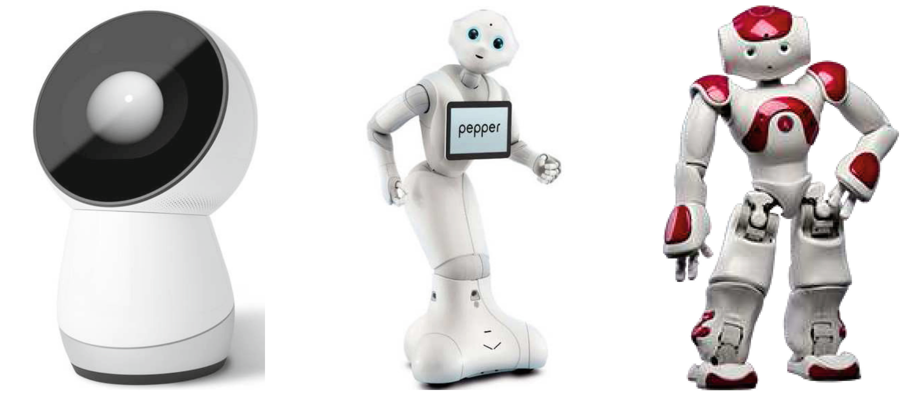

(a) Jibo (b) Pepper (c) Nao

Fig. 2. Forerunners of the coming social robot deluge. Jibo (left) is one of a number of home robot companions marketed almost more for their character and personality than their functional utility. Pepper (center) is marketed as the first robot that recognizes and reacts to its users emotional needs. The Nao robot (right) is used as the base for a number of different social robots.

addresses how humans may react to this. To a lesser extent, company-developed social robots may face similar issues, as their personalities are tweaked to meet cultural and geographic norms for different markets.[1]

In contrast to research robots, which generally only interact with humans for relatively short periods of time, commercial robots are designed to have a long-term presence in the user's life. Industry development has adopted the concept of personality to encompass all of the social capabilities of a robot, beyond the mere functional ones. While academic work on robot personality is somewhat thin, the concept of personality (and associated concepts of attitudes, emotions, and moods) are well studied in the psychological literature, although a generally agreed-upon unified model is still lacking [8].

We take 'personality' as commonly used to mean a sense of a unifying gestalt behind an entities' behavior. Industrial robot personalities will likely be heavily influenced by those already in use in the gaming industry, where non-player characters are often designed to be engaging and social. These interactions are highly scripted and require many hours of designer, developer, and potentially actor time and effort. More dynamic behaviors are achieved via hand-crafted behavior trees, a variant of Finite State Machines, where several underlying behaviors are switched between based on context and user input. Similar approaches will likely be used to develop new, embodied robot characters, but note that all of the resulting characters are fixed. Research has yet to be carried out to examine human reaction to long-term interaction with the resulting characters, but anecdotal evidence suggests that without massive amounts of programmed variation

[1] "Pepper, the emotional robot, learns how to feel like an American" Wired, 6/7/16.

and adaptability, users will find repeated interaction at best dull, and at worst annoying.

3 Approach

Most social robot systems are developed by taking a desired, perhaps already implemented, functionality and layering social behaviors on top of it. We take an opposing view and argue that in order for the robot's personality to be really unified and 'shine through,' it needs to be developed first, and functional utility added later. Accordingly, we focus here on developing the core personality system of a robot, and leave functional utility for future work.

Our basic model considers a robot that has some set of continuous sensors (S) and actuators (A), that together determine what the robot *can* do in the world. We concern ourselves with a model for the robot's personality (P), which determines what the robot *opts* to do. In order for the robot to do anything at all, we must consider some drives (D) that define what the robot *needs* to do.

To achieve infinite diversity in personality, we consider a continuous, bounded personality space. A robot's personality is represented by a point in this space, and as it is infinite, all robots can have different (albeit perhaps similar[2]) personalities. The robot's personality is then used to drive the robot's decision making and behaving. Adaptation could be achieved by moving the robot's personality in this space, which will in turn change how the robot reacts to changes in its environment. For convenience, we take $P \in \{0, 1\}^K$, where K is the dimensionality of the personality space.

We define a drive D as a behavior that takes in a state of the world and a potential action and produces an *acceptability* of performing that action in that state, dependent upon the robot's personality. That is $D(S, A, P) \rightarrow [0, 1]$, where 0 indicates that the action is *not* acceptable to this drive in this state with this personality, and 1 indicates that it is, with differing acceptabilities in-between.

Given a set of drives ($\{D_k\}_{k=1}^K$) and a current state s_t, the total acceptability (α) of a proposed action (a) is

$$\alpha_a = \prod_{k=1}^{K} D_k(s_t, a, p_k) \tag{1}$$

which defines a pseudo-distribution (values in $[0, 1]$, un-normalized) over the entire action space of the robot. Even without the normalization constant, we can sample from the underlying distribution using rejection sampling (with a uniform proposal distribution) to find an action that is more-or-less acceptable to all of the robot's drives. The use of sampling (rather than a MAP estimate) is deliberate, as it brings randomness into the robot's behavior, which makes it seem more 'alive.'

[2] An open question is how different two personalitites must be in order to be *perceived* as different by humans, we leave this for future work.

Note that the number of drives and the dimensionality of the personality space are the same, K. That is, the drives implicitly define the personality space of the robot. In essence, each drive defines a continuum of behaviors, dependent on the personality parameter p_k that smoothly changes the drive's behavior between two extremes as we will show in the next section.

3.1 Implementation

We implement our personality system on a simple robot, shown in Fig. 3a. The robot has three time-of-flight sensors and two sonar range finders facing forward to detect obstacles, and measures the ambient light level at three locations (again forward facing). A color camera on a tilt motor is used to locate human faces in front of the robot (range, bearing and height), and a custom IR board provides the range and bearing to the charging dock, as well as its current battery charge. The robot has a treaded drive system controlled by linear and angular velocities, and can tilt the camera. The input space is variable-dimensional ($11+3N$, $N =$ number of visible humans) and the action space is 3D.

We implement four drives on this system, with an associated 4-dimensional personality space. Each drive defines acceptability as a Gaussian distribution over the action space ($D_k(s_t, p_k, a) = \mathcal{N}(a|\mu, \Sigma)$) with μ_l, μ_a, μ_t being the centers of the distribution in linear, angular, and camera tilt space, and $\sigma_l, \sigma_a, \sigma_t$ being the corresponding entries in the diagonal covariance matrix (Σ). For simplicity, we do not consider cross-covariance terms in this work, and leave out scaling constants in the following.

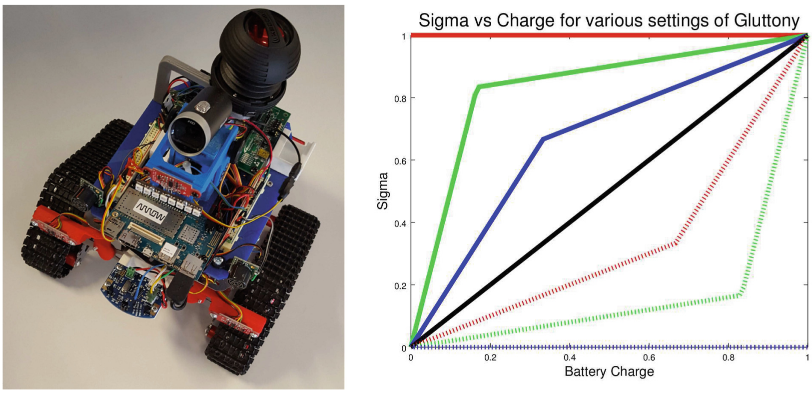

(a) Marz (b) PiecewiseLinear

Fig. 3. Left: Our robot platform senses obstacles in front of it with sonar and time-of-flight and can drive via a treaded system. The camera tilts, and is used to locate human faces. Right: The piecewise linear function maps battery charge and gluttonous-ness to variance in the food drive.

Food Drive. The food drive serves to keep the robot charged by placing the center of acceptability on linear and angular velocities that will drive the robot towards the charger. It considers the range and bearing to the charger (r_c, b_c) and the current charge level (c) and computes $\mu_l = r_c \cos(b_c), \mu_a = b_c$. The robot's personality space for this drive runs from food-seeking or *gluttonous* $(p_{\texttt{food}} = 1)$ to food-ignoring $(p_{\texttt{food}} = 0)$ and is reflected in the computed variances $\sigma_l = \sigma_a = pl(c, p_{\texttt{food}})$, where pl is the piecewise-linear function in Fig. 3b.

Comfort Drive. Depending on personality, the comfort drive makes the robot seek out and stay in comfortable, well-lit areas. It takes in the three ambient light levels (l_l, l_c, l_r - left, center, right) and sets $\mu_l = 1 - \max(l_r, l_c, l_l), \mu_a = l_l - l_r$ to slow the robot as brightness reaches a maximum, and turn towards the brighter side. Again, we use the personality to set the variance, where $\sigma_l = \sigma_a = 1 - p_{\texttt{comfort}}$. As the robot's *laziness* increases, it tends to more often seek out and bask in the light.

Obstacle Drive. While the robot will not deliberately collide with obstacles, the distance to which it is willing to approach them depends on the personality dimension of *cautiousness*. Considering the three time of flight sensors (t_l, t_c, t_r) and the two sonar sensors (s_l, s_r), the obstacle drive sets $\mu_l = (1 - p_{\texttt{obstacle}})\texttt{min}(t_l, t_c, t_r, s_l, s_r)$ to slow the robot as it approaches an obstacle, and $\mu_a = \texttt{sign}(t_l - t_r)(1 - \texttt{min}(t_l, t_c, t_r, s_l, s_r))$ to turn the robot towards the freer side, faster when it is closer to an obstacle. Note that this drive uses the personality to change the mean of the distribution (slowing down faster as cautiousness increases), and the variance is set $\sigma_l = \sigma_a = \min(t_l, t_c, t_r, s_l, s_r)$ to decrease as an obstacle is neared, to ensure the robot does not collide.

Human Drive. The only drive to consider camera tilt, the human drive guides the robot to approach humans and look them in the face. Given the range, bearing, and height of the N visible people $(\{r_h^{(n)}, b_h^{(n)}, h_h^{(n)}\}_{n=1}^N)$, the drive considers each human individually and returns the average acceptability $\alpha_{\texttt{human}} = \frac{1}{N}\sum_{n=1}^N \mathcal{N}(A|\mu^{(n)}, \Sigma^{(n)})$ where $\mu_t^{(n)} = h_h^{(n)}, \mu_l^{(n)} = r_h^{(n)}, \mu_a^{(n)} = b_h^{(n)}$, and the variances depend on the robots *friendliness*, as $\sigma_t^{(n)} = \sigma_l^{(n)} = \sigma_a^{(n)} = 1 - p_{\texttt{human}}$.

4 Experiments and Results

Our experiments aimed at determining whether or not our infinite personality space and drive-centric system gave rise to recognizable and measurable differences in robot behavior. To do so we not only interviewed humans who interacted with our physical robot platform, but also replicated the robot's functionality in a web-based simulator to examine longer-term behavioral differences. The personalities we examined were hand-picked to highlight the differences achievable with this system.

4.1 Quantitative

We examine here the impact of one personality dimension on robot behavior. Specifically, with other personality traits held constant, we expect the personality trait of gluttony to impact the amount of time the robot spends charging, with more gluttounous robots spending more time accumulating charge. In our multi-robot simulator, we simulate several identical robots with the same initial conditions (location, orientation, and charge) that only differ in their value of $p_{\texttt{food}}$ and track the number of times they dock, and the total amount of time they spend docked over several hours.

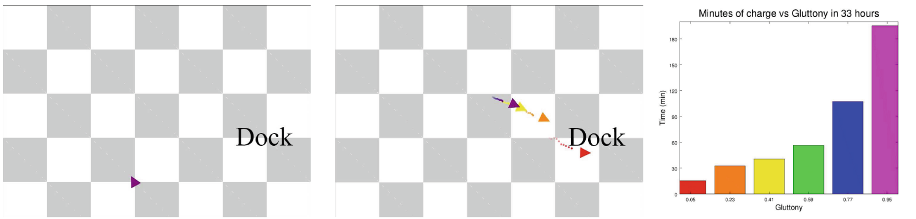

(a) Initial (b) Divergent behavior (c) Time spent charging

Fig. 4. Effect of gluttonous personality trait on charging time and frequency. (a) All robots start with the same initial conditions but quickly diverge by choosing to approach the dock or not (b). After 33 h, differences in behaviors are apparent, as gluttony directly impacts total time spent charging (c).

Initially (Fig. 4a) all of the robots are at the same location, but as they begin to get range and bearing readings on the dock, they quickly diverge (Fig. 4b). After 33 h of simulated time, the differences in behavior have become apparent, as shown in Fig. 4c. The most gluttonous robots spend around 12.5 times more time charging than the least gluttonous. Additionally, while all robots began with 5 % state of charge (to stimulate charging), during the simulation the least gluttonous robot ($p_{\texttt{food}} = 0.05$) was observed to keep its battery at 1 %, while the most gluttonous ($p_{\texttt{food}} = 0.95$) increased its to 96 %.

4.2 Qualitative

While our qualitative results indicate that changes in personality do, in fact, lead to changes in behavior, we also wish to examine the perception of the robot's personality by interacting humans. To do so we performed a series of informal demonstrations for naive users ($N < 20$, not part of the team that developed the robot) comparing various personalities. The robot was exhibited in both our office space and a dedicated 'living room' environment, with a couch, chair, lamps, etc. While no statistical conclusions can be drawn from such a casual study, different personalities were anecdotally visible, as described below:

- Robots with 'friendliness' turned down were seen to be indifferent to the presence of humans, while those with 'friendliness' turned up were seen as more engaging, and elicited more interaction.
- Robots with 'cautiousness' turned down were seen as less skilled, due to their increased likelihood of getting stuck in corners
- Robots with 'laziness' turned up were seen as "falling in love" with the lamp, as they would approach the light and stop, while those with 'laziness' turned down would ignore it.

Note that users often did not interact with the robot long enough for it to charge, so differences in behavior related to 'gluttony' are not discussed, but were covered in Sect. 4.1.

5 Future Work

There are some limitations to our current approach that can be investigated in future work. While the system does scale to additional drives (the personality space grows linearly), our use of rejection sampling to find acceptable actions may not be a tenable solution in higher dimensions. Even with only 4 dimensions our robot was, at times, unable to find an acceptable action in the time allotted. This issue becomes particularly acute when one drive has low acceptability over much of the action space (i.e., when near a wall, the obstacle drive only accepts a small portion of available actions). Likewise, as the dimensionality of the output space grows, the computational limits of our system may be taxed.

We specifically worked with a deliberately simple robot system, in order to focus on our ability to represent personality and demonstrate differences via behavior. For example, we did not utilize any memory or time-extended actions, and built an entirely reactive system. However, there is nothing in our framework that precludes these capabilities from being included, and doing so will undoubtedly be necessary to achieve functional utility.

On that note, our robot is personable and entertaining, but as yet serves no functional goal. While there are markets and use cases for purely entertainment robots, greater acceptance may be achieved by having the robot have some functional utility. In our framework, these uses may take the form of drives (to deliver mail, for example), which would then interact with the other drives and personality to give rise to a unique, functional *and* personable robot.

Lastly, the work presented here focused on defining an infinite personality space that can give rise to an infinite number of unique robots. Still, however, we take the personality as fixed for the lifetime of the robot. An interesting possibility is, however, to allow the personality of the robot to change over time, perhaps through interaction with a human. For example, reinforcement learning techniques could be used to reward observed behavior, which could then be used to change the robot's personality to make the good behavior more likely to occur.

6 Conclusions

By taking a personality-first view of robot behavior and operating in an infinite personality space, we have defined a novel way of developing a social robot. Our main goal of developing non-fungible robots that truly differ is achieved, as each robot's personality can be unique, and will result in idiosyncratic behavior. These differences in behavior are both measurable and observable to humans.

Acknowledements. Many thanks to Michael Gielniak for writing the simlutator and running studies, and to Dave Hygh, Quentin Michelet, and Patrick Martin for making the robots and keeping them running.

References

1. Cha, E., Dragan, A.D., Srinivasa, S.S.: Perceived robot capability. In: Robot and Human Interactive Communication, pp. 541–548 (2015)
2. Dragan, A.D., Lee, K.C.T., Srinivasa, S.S.: Legibility and predictability of robot motion. In: International Conference on Human-Robot Interaction, pp. 301–308 (2013)
3. Gielniak, M.J., Thomaz, A.L.: Enhancing interaction through exaggerated motion synthesis. In: International Conference on Human-Robot Interaction, pp. 375–382 (2012)
4. Jung, M.F., Lee, J.J., DePalma, N., Adalgeirsson, S.O., Hinds, P.J., Breazeal, C.: Engaging robots: easing complex human-robot teamwork using backchanneling. In: Computer Supported Cooperative Work, pp. 1555–1566 (2013)
5. Kidd, C.: Designing for long-term human-robot interaction and application to weight loss. Ph.D. thesis, MIT (2008)
6. Kiesler, S., Goetz, J.: Mental models of robotic assistants. In: CHI 2002 Extended Abstracts on Human Factors in Computing Systems, pp. 576–577 (2002)
7. Kirby, R.: Social robot navigation. Ph.D. thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, May 2010
8. Moshkina, L.V.: An integrative framework of time-varying affective robotic behavior. Ph.D. thesis, Atlanta, GA, USA, AAI3464090 (2011)
9. Mumm, J., Mutlu, B.: Human-robot proxemics: physical and psychological distancing in human-robot interaction. In: International Conference on Human-Robot Interaction, pp. 331–338 (2011)
10. Phillips, E.K., Schaefer, K., Billings, D.R., Jentsch, F., Hancock, P.A.: Journal of Human-Robot Interaction **5**(1) (2016)
11. Sheh, R.K.M., Hengst, B., Sammut, C.: Behavioural cloning for driving robots over rough terrain. In: International Conference on Robotic Systems, pp. 732–737 (2011)
12. Sundar, S.S., Waddell, T.F., Jung, E.H.: The hollywood robot syndrome: media effects on older adults' attitudes toward robots and adoption intentions. In: International Conference on Human Robot Interaction, pp. 343–350 (2016)
13. Sutcliffe, A., Grollman, D., Pineau, J.: Estimating people's subjective experiences of robot behavior. In: AAAI Fall Symposium on AI for HRI (2014)
14. Wada, K., Shibata, T.: Robot therapy in a care house - its sociopsychological and physiological effects on the residents. In: International Conference on Robotics and Automation (2006)