Robot Errors and Human Teachers: The Effects of Personality and Patience During Learning Tasks

Cory J. Hayes[1] and Laurel D. Riek[1]

Abstract— **Robots are expected to become a ubiquitous technology in the near future, where different people from a wide variety of backgrounds may have daily interactions with robots. These co-present humans will expect to be able to customize robot behavior to suit their preferences and needs through intuitive methods, such as learning from demonstration (LfD). However, these interactions will vary depending on the personal qualities of the user, and to the best of our knowledge, no work has explored how these qualities affect the teaching of a robot during an LfD task. This paper introduces a novel experimental design for testing patience during LfD instruction, and explores the relationship between personality and patience while teaching an autonomous DARwIn-OP robot. Another contribution of this paper is the introduction of systematic error simulation during LfD.**

I. INTRODUCTION

People have different backgrounds that can affect their attitudes towards technology. Some individuals may come from cultures where robots are treated as living entities with respect, while others may come from cultures that view robots more as tools [20]. People also have a wide range of cognitive and physical abilities that can affect how they perceive, interact with, and accept robots. One way to help address the concern of universal acceptance is by allowing people to customize robot behavior to more ideally suit their preferences and needs.

Learning from demonstration (LfD) is one viable way for non-technical members of society to accomplish this task of developing and modifying custom robot behavior. In LfD, a learner creates a mapping between states and actions by watching a teacher perform the task and attempting to replicate the action [2], [4], [6], [23].

The main benefit of LfD is that it is an intuitive way for people to teach others and does not require the teacher to have highly specialized knowledge. Researchers have noted its significant advantages in robot control and it has already been explored to a large degree in field of automatic robot programming [3]. LfD could also play a crucial role in the prevalence of robots in general society in the near future by allowing non-robotics experts to create custom behavior for robots [2].

In addition to allowing people who are not technically inclined to customize their robot's behavior, LfD also addresses a major concern from robotics and human-robot interaction (HRI) researchers. Developers cannot explicitly program for all situations a social robot may be in given the wide range of possibilities varying complexity of each situation. Learning from demonstration allows for the robot to adapt to new situations by learning new behaviors on the fly. By allowing for the creation of robot behavior that adapts to the needs of its user, a robot will become individualized making it different from any other robot as well as personalized, since will reflect the needs of its surrounding environment as well as the actors in that environment [8].

In regards to unique human attributes, personality in particular plays a critical role in HRI. In both the HRI and human-computer interaction (HCI) literature, personality traits have been shown to affect how people interact with different technologies [1], [5], [8], [19], [22], [28]. In HRI, personality is largely discussed in the domain of robotic anthropomorphism where people project and/or attribute human like qualities to robots [13], [14], [21].

Personality has also been shown to affect the willingness of people to accept a new technology and their willingness to adapt to it. The Technology Acceptance Model (TAM) is the most widely recognized model of technology acceptance, and is used to predict and understand how people may come to adapt a technology [9]. Devaraj et. al [10] explored the TAM along with the previously discussed five personality traits and showed how they relate to perceptions of a technology and willingness to adapt, with the exception of "openness" not being supported by findings.

However, to our knowledge, there has been little to no work that explores the possible relationship of human personality and a person's patience when dealing with robots. This is an important area to research since given the wide variety of tasks that social robots will be asked to accomplish [25], robots will definitely make mistakes. While some robotic mistakes could have drastic consequences, such as in the areas of healthcare and socially assistive robots, it is also important to explore the effects that minor annoyances, such as a robot not performing in an expected way, could have on user patience and overall acceptance.

This paper explores how personality types factor into the degree of patience a person has when teaching and correcting an error-prone robot during a learning from demonstration task. In our work, we use the widely-accepted five factor model of personality, which focuses on the factors of openness, conscientiousness, extraversion, agreeableness, and neuroticism [11].

Openness describes the appreciation of intellectual creativity, exploration, and the predisposition of novelty. Conscientiousness describes the degree to which people exhibit self-discipline and responsibility; people who are conscientious

[1]C.J. Hayes and L.D. Riek are with the Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN, 46556 USA. e-mail: chayes3@nd.edu, lriek@nd.edu.

are goal oriented and rely on planned behavior.

Extraversion focuses on the degree to which a person is outgoing and exhibits positive emotions in the company of other people. Agreeableness describes the compassion and cooperativeness of a person; people who score highly on agreeableness are well tempered. And finally, neuroticism describes emotional instability such as the propensity of anger and sadness and is scored on a negative scale (i.e., a high degree of neuroticism is a negative aspect, whereas for all other dimensions a high score is a positive aspect).

This motivates our primary research question: *do certain personality traits affect an individual's patience when interacting with a robot during a learning from demonstration task?*

This leads us to five hypotheses:

H1: People who score higher on conscientiousness will be less patient with a robot than those who who score lower on this scale.

H2: People who score lower on emotional stability will be less patient with a robot than those who score higher on this scale.

H3: People who score higher on agreeableness will be more patient with a robot than those who score higher on this scale.

H4: People who score higher on extraversion will be more patient with a robot than those who score higher on this scale.

H5: People who score higher on openness will be more patient than those who score lower on this scale.

These hypotheses are motivated by the defining qualities of each trait. People who score higher on conscientiousness are more goal-oriented and organized, so it is reasonable to assume that such a person will not positively respond to unexpected robot behavior that inhibits progress towards reaching a goal.

People who score lower on emotional stability are more prone to experience negative emotions such as anger and anxiety, and unwanted robot behavior should more easily frustrate a person with lower emotional stability. The reason we use the term "emotional stability" rather than "neuroticism" is due to the specific personality test we employed (see Section II-B).

A high agreeableness score indicates that the individual is more well-tempered and cooperative than those who score lower on this scale, so such a person would be more tolerable and patient with a robot.

The hypothesis for those who score higher on extraversion is a less directly related to patience than the other traits, but we hypothesize that since those who score high on extraversion exhibit more positive emotions than those who score lower, they would in turn be more patient with a robot.

A high score on openness signifies that a person is open to new experiences and is intellectually curious, so their response to unexpected robot behavior would be more positive than one who scores lower on this scale.

In Section II, we discuss our experimental methodology, and report our results in Section III. In Section IV, we discuss the implication of these findings for the robotics community, particularly for roboticists interested in building robots that will work alongside people.

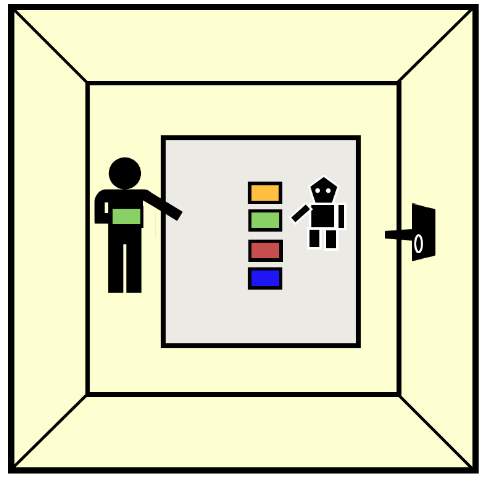

Fig. 1. This diagram shows the experiment room setup for our mock LfD task. A participant (left), presents a color card to the robot (right), and this action is captured by a Kinect behind the robot.

II. Methodology

In order to measure the degree of patience a person has when dealing with a robot that behaves in an undesirable and potentially frustrating way, we designed an experiment that uses a mock LfD approach. In this experiment, participants were given the task of teaching four colors to a small autonomous humanoid robot positioned on top of a table. To teach the robot a new color, participants 1) hold up a color card in front of the robot, 2) state its color while simultaneously 3) pointing to the corresponding color card on the table in front of the robot, and 4) wait for the robot to point to the same table card while stating the color.

This study was advertised as one where the purpose was to determine how effectively a person, regardless of technical background, could teach a robot through a LfD task. Though participants were told that the robot would be learning new behavior from them, the robot did not actually learn anything from the participant and was programmed to autonomously interact with participants in a predetermined manner. We decided not to implement actual robot learning at this point in order to precisely control for robot errors, and a setup that does involve true LfD would be one immediate followup to this study.

A. Programming and Setup

To program this experiment, we combined the capabilities of the humanoid robot and a computer running ROS Electric on Ubuntu 11.10. The robot used in the experiment was a DARwIn-OP, developed by ROBOTIS [16]. DARwIn-OP has a height of 17.89 inches and features a built in personal computer, a microcontroller, multiple high performance servo motors, and a variety of sensors such as a gyroscope and accelerometer that enables it to balance while walking.

Although the DARwIn-OP contains microphones and a camera, both of which are applicable to this experimental

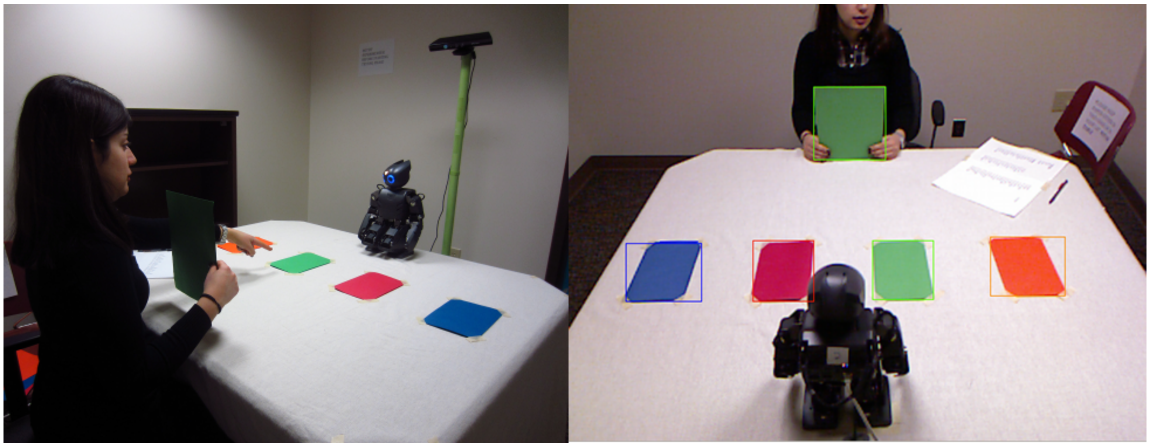

Fig. 2. Viewpoint of the room from a perspective similar to that of the participant's (left) and the Kinect's viewpoint (right) with blob detection enabled.

setup, we did not utilize these sensors. Instead, we used ROS to create a custom program that employed the blob detection capabilities of the Microsoft Kinect to identify the color cards a participant held up in front of the robot. When the program recognizes the color of the held up card, it then determines the actions of the robot and forwards a command through an Ethernet connection. This command signifies what the DARwIn-OP says, what card it points to, and also controls its head movement.

Figure 1 demonstrates the experimental set up. The participant is situated at a table opposite from the robot, see the left image of Figure 2 for a viewpoint of the room similar to that of the participant's. The DARwIn-OP is positioned on top of the table facing towards the participant. In front of the robot are four color cards that it can point to; the participant holds up slightly larger color cards to present to the robot. Behind the robot and positioned at human height is the the Microsoft Kinect, see the right image of Figure 2 for the Kinect's viewpoint of the room.

B. *Preliminary Tasks*

Prior to performing the LfD portion of the experiment, each participant completed an online survey primarily featuring a personality test. This personality test was taken from the 100-item unipolar personality test developed by [15], known as the Trait Descriptive Adjectives (TDA). This test is noted for being a considerably robust and reliable test for the five personality traits [27]. The traits targeted are surgency (equivalent to extraversion), agreeableness, conscientiousness, emotional stability (the inverse of neuroticism), and intellect (equivalent to openness). Each of these five traits has a total of twenty items devoted to negative and positive attributes. Surgency, agreeableness, conscientiousness, and intellect each have 10 items devoted to both the positive and negative attributes. However, emotional stability has 6 traits devoted to its positive attributes and 14 traits devoted to its negative attributes.

At the end of this online survey, participants were instructed to designate times during which they were available to visit our lab to complete the LfD portion of the experiment.

Once a participant arrived to the experiment room, he/she filled out various forms and was given instructions for the experiment. In addition to an instruction form, participants also watched a tutorial video in which the same robot is taught a set of four different colors: pink, purple, brown, and yellow. The manner in which the person in the video taught and tested the robots ability to identify colors is identical to what the participants did for this experiment, with the only difference being the colors used. Participants were told that due to the scenario shown in the tutorial video, the robot has already learned the color "yellow", and that color is used to transition between the three stages of the experiment.

C. *Learning from Demonstration Task*

The mock LfD portion consists of three stages: pre-training, training, and testing.

1) Pre-training: In pre-training, the participant simply holds up a card in front of the robot to practice presenting a card to the robot and reacting to its response. To follow the LfD narrative, the robot only states "I see a card" when the participant holds up any card other than yellow, since it has not yet been taught any color by the participant. To move on to the training stage of the experiment, a participant holds up the yellow card twice consecutively. In return, the robot informs the participant that it is ready to learn new colors.

2) Training: In the training phase, the participant teaches the robot a color by holding up a card, stating its color, and pointing to the respective color card in front of the robot. To simulate learning, the robot pauses briefly before looking at the card on the table that the participant points to, and then points and repeats the color of the card. After pointing to the card, the robot looks back up at the participant and states that it is ready to learn the next color. Once training is done for all four colors, the participant signifies they are ready move on to the timed testing phase by once again holding up the yellow card twice. The robot then informs the participant it is ready to test its new color knowledge and that it is ready for the first card. It is important to note that for both the training and testing phases, the participants follow a provided sequence of colors to present to the robot.

3) Testing: In the testing phase, participants are given the task of testing and correcting the robot to the point where it can correctly identify five colors in a row within seven minutes. A correct identification happens when the participant holds up the card, without stating or pointing to the color, and the robot points to and states the correct color on its first attempt. When the participant presents a color card to the robot, it mimics making at a decision by first scanning the cards in front of it twice and then pointing to a card. The testing phase is the portion of the experiment where we measure participant patience by manipulating the expected behavior of the robot.

While the program is autonomous, it is designed to intentionally make numerous mistakes during this testing phase. In this case, a mistake is where the robot states the color held up by the participant while pointing to the incorrect color. Participants are warned that mistakes may happen during the testing phase since the robot is using a simple machine learning algorithm that considers 1) the color the participant holds up, 2) the color the participant states, and 3) the participants pointing action during the training phase of the experiment.

When the robot misidentifies a color, the participant can choose to correct the robot by repeating the learning process for that one color, or they can choose to move on to the next color in the sequence by raising the yellow card once. Due to this simple machine learning algorithm, there is a stated trade-off proposed to the participants. This trade-off is between the participants making sure the robot correctly identifies a color through multiple re-teachings if necessary, or deciding to move on to the next color in a sequence and hoping that the robot later correctly identifies the color after correctly identifying other colors.

However, the robot makes mistakes on a very frequent basis during the testing phase that requires participants to correct the robot a relatively large amount of times before it eventually makes a correct identification for a color. The number of times the robot makes a mistake for any given color in the sequence is predetermined by a randomly generated list. The program uses this random list for all participants so that each participant experiences the same robot behavior of correct/incorrect identifications. While the number of mis-identifications for any given color is random, the robot does not mistake a certain color less than five times or more than ten times before eventually correctly identifying it. The program is designed to make the goal of five correct identifications in a row impossible to reach, regardless of how long the person interacts with the robot.

After seven minutes, the LfD task ends and the participant leaves the experiment room. Participants are then asked to provide any feedback they may have regarding the experiment on an online form. The experimenter then debriefs the participant to inform them of the true objective of the experiment Finally, the participant is compensated for their time with a $5 gift card before leaving.

D. Measurements

We employed the TDA personality assessment instrument as noted by [27]. Participants were asked to rate themselves for each of the 100 adjectives on a scale from 1 to 9, a score of 1 signifying that an adjective was "extremely inaccurate" and 9 signifying that it was "extremely accurate. Each adjective listed in the personality test corresponds to either a positive or negative attribute for a specific personality dimension. For example, the item "pleasant" is a positive attribute for the dimension of Agreeableness, while "rude" is a negative attribute for the same dimension. Positive attributes are summed normally depending on the response on the 9-point Likert scale. However, negative attributes are reflected (i.e. a "1" would become a "9" and vice-versa) before being summed with the positive attributes.

We created our own measurement to calculate a participant's degree of patience with the robot. This measurement was calculated by dividing the number of times a person actually corrected the robot by the total number of possible times the person could correct the robot. To illustrate this, consider a scenario where the robot would intentionally misidentify a certain color in the sequence five times before finally identifying the correct color. If the participant gave up re-teaching that color after two attempts, the ratio of attempted corrections out of possible corrections for this particular instance would be 2/5.

We used a ratio instead of counting the number of times a participant corrected the robot to account for amount of times participant attempted to correct the robot before moving on. For example, a participant who was patient enough to correct robot mis-identifications until the robot got it correct would have a noticeably higher patience ratio than a person who had the same total number of attempted corrections but moved on to the next color in a sequence after two mis-identifications.

To determine how personality corresponds to the degree of patience the participant had with the robot, each dimension in TDA instrument was compared separately to this patience ratio through correlation analysis.

III. Results

We recruited 39 participants on the University of Notre Dame campus, via through mailing lists and physical flyers. All participants were native English speakers. We targeted native English speakers due to concerns raised during beta testing regarding some of the unclear meanings of adjectives used in the personality test. There were a total 18 female participants and 21 male participants. The average age of the participants was 21.76 with a standard deviation of 3.11.

Before running correlation analysis for the relationship of patience and any of the five personality dimensions, we first verified the normalcy of participant patience scores. This check was done through the use of a parametric test. This test resulted in z-scores for kurtosis and skewness of 1.208 and 1.506 respectively. Since these scores fell below the threshold of 1.96 for small data sizes, this showed that the patience scores followed a normal distribution. After confirming this normality, we then performed correlation analysis using Pearson's correlation coefficient *r*. None of these relationships were significantly correlated ($p < .05$), though we report them below.

Conscientiousness and Patience resulted in a Pearson correlation coefficient of r = -.144 and a significance value of .190. The coefficient of determination R^2 was .021, meaning that Conscientiousness could only account for 2.1% of the variation in Patience and 97.9% of the variability is accounted for by other variables.

Emotional Stability and Patience had a Pearson correlation coefficient of r = -.073 and a significance value of .329. R^2 was .005, meaning that Emotional Stability could only account for .5% of the variation in Patience and 99.5% of the variability is accounted for by other variables.

Agreeableness and Patience had a Pearson correlation coefficient of r = -.174 and a significance value of .144. R^2 was .030, meaning that Agreeableness could only account for 3% of the variation in Patience and 97% of the variability is accounted for by other variables.

Surgency and Patience had a Pearson correlation coefficient of r = -.022 and a significance value of .447. R^2 was 4.93×10^{-4}, meaning that Surgency could only account for .049% of the variation in Patience and 99.051% of the variable is accounted for by other variables.

Intellect and Patience had with a Pearson correlation coefficient of r = -.093 and a significance value of .288. R^2 was .009, meaning that Intellect could only account for 0.9% of the variation in Patience and 99.1% of the variability is accounted for by other variables.

IV. Discussion

We did not find support for any of our five hypotheses regarding the direct relationship between personality traits and patience while teaching a robot during an LfD task. This lack of support may be due to several reasons.

Firstly, our patience measure may not be well-suited to this particular experimental paradigm. We sought to reward participants who spent more time correcting the robot during the testing phase, which may have led to an uneven weighting. For example, we saw a few extreme cases in our data, with patience scores of 100 and 4.1. The former score represents a participant who attempted to correct the robot every time it made a mistake, while the latter represents one who very rarely did so. A larger sample would help determine the robustness of the measure.

A second interpretation of the findings might be that participants deduced the true objective of the LfD task, i.e., they realized that the robot was not actually attempting to learn from them, and so they acted differently during the experiment. However, given participant verbal and written comments after the LfD task, we do not believe this is the case. No participant indicated that they had any notion of the true purpose of the experiment beyond what was described to them at the onset of the study, and most were surprised at the reveal when they were debriefed.

Third, it may be possible that there is no clear relationship between personality and patience when teaching a robot. However, given the strong evidence in the literature correlating multiple personality attributes with patience across a range of tasks [7], [12], [17], it seems reasonable to expect such a finding would hold for robot teaching as well. It is possible our experiment did not leave sufficient time to truly tax participants' patience [18], [24], or that novelty effects of the robot played a role.

Indeed, this novelty effect could be impactful enough to nullify the effect that personality plays in a LfD task. The qualitative comments provided by participants hint at this possibility. Many participants stated that while the robot was ”remarkably unintelligent”, they still enjoyed interacting with the robot because it was either their first time or one of the few times they had done so. One possible additional factor that could be used in a patience measurement for an LfD task is a definitive measurement of user familiarity with robots.

Though our results do not find support for a relationship between personality and patience during LfD tasks, we still believe that this work presents a useful contribution to the robotics community. We implemented a fully autonomous system for HRI, which is notable since a large percentage of HRI work is implemented through Wizard of Oz approaches [26]. While it may have exaggerated what a person could experience while trying to create new robot behavior using LfD, this design was a convincing replication of a task that may be encountered in future real-world scenarios.

Furthermore, our participants reported believing that they were actually teaching a robot and capable of modifying its behavior during our experiment. This suggests that other researchers can use and modify this experimental design to explore other potential effects that may arise when a person uses LfD techniques with a robot.

Personality underlies the way people behave, how they perceive the world, and both their short-term and long-term mental states. It has been linked to multiple human qualities such as age, culture, gender, and educational level. Additionally, personality remains relatively stable throughout a person's life. In essence, personality defines all of the possible different ways a user may interact with a robot. This encompasses apparent differences in behavior (e.g. the participants mentioned earlier with patience scores of 4.1 and 100) as well as more subtle differences such as two users with closer scores, the latter reflecting the motivation of this study to determine how personality affects patience.

Teaching interactions with a robot will have a trade-off between robot behavior accuracy and interaction time requirements. Longer interactions with a robot will lead to more accurate robot behaviors, but there are differences in the amount of time people will devote to teach a robot new behaviors. Therefore, we argue that it is vital to focus on quantifying observable human behavior that stems from underlying user personality during interactions with robots. We focus on patience, but this extends to other kinds of human behavior.

From a computing standpoint, being aware of a user's personality will allow a robot to effectively balance the accuracy-time trade-off during a LfD interaction. For example, a robot that knows a user has a certain personality type and will not devote much time for teaching might employ specific strategies of active learning. These strategies would lead to the shortest possible interaction before achieving what the user considers to be "acceptable" robot behavior. On the other hand, a robot that knows a user has a personality type that corresponds to a more patient person may employ active learning techniques that lead to a longer interaction. This would result in new behavior that is significantly more accurate and personalized than simply being "acceptable."

The creation of these metrics will allow robots to be able to dynamically analyze human behavior and adjust their own behavior to better suit their user based on personality-centered models. These models will not be perfect from the start since there will be differences in human behavior across different types of learning tasks, and will require fine-tuning through repeated interactions. However, these metrics would eventually allow robots to evolve into highly personalized machines, paving the way for a future where most, if not all, people interact with robots on a daily basis.

V. Acknowledgements

This material is based upon work supported by the National Science Foundation Grant No. IIS-1253935. The authors also thank Tariq Iqbal, Maryam Moosaei, Michael Gonzales, Lucinda Krahl, Elise Eiden, and Samantha Rack for their help and support.

References

[1] E. Abrahamian, J. Weinberg, M. Grady, and M. Stanton. The Effect of Personality-Aware Computer-Human Interfaces on Learning. *Journal of Universal Computer Science 10*, 1, 2004.

[2] B. Argall, S. Chernova, and M. Veloso. A Survey of Robots Learning from Demonstration. *Robotics and Autonomous Systems*, 57, 2009.

[3] C. Atkeson and S. Schaal. Robot Learning from Demonstration. *Machine Learning: Proceedings of the 14th International Conference*, 3, 1997.

[4] S. Calinon and A. Billard. Incremental learning of gestures by imitation in a humanoid robot. *In proc. of the 2nd International Conference on Human Robot Interaction*, 2007.

[5] O. Celiktutan, E. Sariyanidi, and H. Gunes. Let me tell you about your personality: Real-time personality prediction from nonverbal behavioural cues. In *11th IEEE International Conference on Automatic Face and Gesture Recognition*, 2015.

[6] S. Chernova and A. L. Thomaz. Robot Learning from Human Teachers. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 2014.

[7] O. Curry, M. Price, and J. Price. Patience is a Virtue: Cooperative People Have Lower Discount Rates. *Personality and Individual Differences*, 44, 2008.

[8] K. Dautenhahn. Robots We Like to Live With? A Developmental Perspective on a Personalized, Life-Long Robot Companion. *13th IEEE International Workshop on Robot and Human Interactive Communication*, 2004.

[9] F. Davis. User Acceptance of Information Technology: System Characteristics, User Perceptions, and Behavioral Impacts. *International Journal of Man-Machine Studies*, 38, 1993.

[10] S. Devaraj, R. Easley, and J. Crant. Research Note How Does Personality Matter? Relating the Five-Factor Model to Technology Acceptance and Use. *Information Systems Research*, 19, 2008.

[11] J. Digman. Personaloty Structure: Emergence of the Five-Factor Model. *Annual Review of Psychology*, 41, 1990.

[12] S. B. G. Eysenck and H. J. Eysenck. The Place of Impulsiveness in a Dimensional System of Personality Description. *British Journal of Social and Clinical Psychology*, 16, 1977.

[13] T. Fong, I. Nourbakhsh, and K. Dautenhahn. A Survey of Socially Interactive Robots. *Robotics and Autonomous Systems*, 42, 2003.

[14] S. Fussel, S. Kiesler, L. Setlock, and V. Yew. How People Anthropomorphize Robots. *In proc. of the 3rd International Conference on Human Robot Interaction*, 2008.

[15] L. Goldberg. The Development of Markers for the Big-Five Factor Structure. *Psychological Assessment*, 4, 1992.

[16] I. Ha, Y. Tamara, H. Asama, and D. Hong. Development of Open Humanoid Platform DARwIn-OP. *In proc. of 2011 SICE Annual Conference*, 2011.

[17] D. Hanna, I. Sneddon, and V. Beattie. The Relationship between the Stockpersons Personality and Attitudes and the Productivity of Dairy Cows. *Animal*, 3, 2009.

[18] C. J. Hayes, C. Crowell, and L. D. Riek. Automatic Processing of Irrelevant Co-Speech Gestures with Human but not Robot Actors. *In proc. of the 8th International Conference on Human Robot Interaction*, 2013.

[19] K. Isbister and C. Nass. Consistency of Personality in Interactive Characters: Verbal Cues, Non-verbal Cues, and User-characteristics. *International Journal of Human-Computer Studies*, 53, 2000.

[20] N. Kitano. Roboethics - A Comparative Analysis of Social Acceptance of Robots between the West and Japan. *The Waseda Journal of Social Sciences*, 1997.

[21] K. Lee, W. Peng, S. Jin, and Y. Chang. Can Robots Manifest Personality?: An Empirical Test of Personality Recognition, Social Responses, and Social Presence and Human-Robot Interaction. *Journal of Communication*, 56, 2006.

[22] S. Pesty and D. Dahaut. Acceptability in Interaction From Robots to Embodied Conversational Agent. *In proc. of the International Conference on Computer Graphics Theory and Applications*, 2011.

[23] J. Peters and S. Schaal. Reinforcement learning of motor skills with policy gradients. *Neural Networks*, 21, 2008.

[24] R. W. Picard and S. B. Daily. Evaluating affective interactions: Alternatives to asking what users feel. In *CHI Workshop on Evaluating Affective Interfaces: Innovative Approaches*, pages 2119–2122, 2005.

[25] C. Ray and F. Mondada. What Do People Expect from Robots? *In proc. of 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2008.

[26] L. D. Riek. Wizard of Oz Studies in HRI: A Systematic Review and New Reporting Guidelines. *Journal of Human Robot Interaction*, 2012.

[27] G. Saucier. Mini-Markers: A Brief Version of Goldberg's Unipolar Big-Five Markers. *Journal of Personality Assessment*, 63, 1994.

[28] A. Tapus, C. Tapus, and J. Mataric. Hands-Off Therapist Robot Behavior Adaptation to User Personality for Post-Stroke Rehabilitation. *In proc. of the IEEE International Conference on Robotics and Automation*, 2007.