

Measuring Users’ Responses to Humans, Robots, and Human-like Robots with Functional Near Infrared Spectroscopy

Megan Strait¹ and Matthias Scheutz¹

Abstract—The *Uncanny Valley Hypothesis* (UVH) describes the sudden change in a person’s affect from affinity to aversion that is evoked by robots that border a human-like appearance. The portion of the human-likeness spectrum in which such aversion is posited to occur is referred to as the “uncanny valley”. However, evidence in support of the UVH is primarily based on subjectively assessed evaluations. Thus it remains an open question as to whether there are behavioral or neurophysiological manifestations of uncanny valley effects. To address this gap in literature, we investigated the activation of the anterior prefrontal cortex (PFC) – a region of the brain associated with emotion regulation – in response to a series of robots with varying human-likeness. We hypothesized that highly human-like robots – which have been found to receive negative subjective attributions – will also elicit increased activity in the PFC versus humans or robots with lesser degrees of human-likeness in accordance with the UVH. Our results show a “valley” in brain activity in the PFC corresponding to the valley observed via subjective measures alone, thus suggesting one neural manifestation (the PFC) of uncanny valley effects and further supporting the affective response (aversion) posited to occur by the UVH. However, the results also reveal a *second* “uncanny valley” in prefrontal hemodynamics, which suggests that the effects (and the contributing factors) are more complex than previously understood.

I. INTRODUCTION

People tend to evaluate agents more positively the more they look and act like a human (e.g., [1], [2]). Thus, as robots and computer agents are increasingly intended for social interactions with people (e.g., [3]), the development of human-like agents has become a focus for researchers in the Human-Robot Interaction (HRI) community (e.g., [4]). However, according to the Uncanny Valley Hypothesis, a person’s affinity towards the agent can change to aversion as the agent nears – but does not attain – a human-like appearance [5], [6]. The portion of the “human-likeness spectrum” in which this change occurs is referred to as the *uncanny valley*.

While there is substantial evidence showing that highly human-like robotic agents receive more negative attributions than humans or agents with less human-like appearances (e.g., [7], [8]), the findings are primarily based on subjective evaluations and the high variability of such evidence has left the Uncanny Valley Hypothesis open to numerous critiques (e.g., [9], [10], [11], [12]). For instance, although many researchers have found that highly human-like agents elicit more negative ratings than their lesser human-like counterparts, others have found the opposite: that people prefer more

human-like agents and rate them more positively (e.g., [13], [?]). Moreover, some suggest that the UVH is not worth much consideration, as there have been findings that participants affective responses to uncanny agents become more positively-valenced following repeated and/or prolonged exposure to the stimuli (e.g., [11], [14]). Given the contrasting results, other approaches to investigating the UVH, such as those targeting the potential neurophysiological or behavioral manifestations of the uncanny valley effects, may be of utility to further corroborate or refute the UVH.

Recent neurophysiological investigations have uncovered numerous differences in the perception of humans versus human-like robots (e.g., [15], [16]). For example, Chaminade and colleagues found increased activity in the occipital and posterior temporal cortices in response to robot versus human stimuli, suggesting additional visual processing is elicited when perceiving a non-human anthropomorphic agent [15]. Others have found that robotic movement engages the action observation network in the parietal cortex (e.g., [17], [18]). This has led Saygin and colleagues to suggest that prediction error – that is, an agent that is very human-like in appearance might result in expectation of more human-like movement than what it exhibits – might underly the resulting aversion posited by the UVH [18]. However, movement error prediction cannot account for the observation of uncanny valley effects in studies which employ image-based (non-moving) stimuli. Moreover, the parietal cortex is not associated with the limbic system of the brain, and thus does not account for the affective response (aversion) hypothesized to accompany the uncanny valley.

We thus attempt to address these shortcomings in an exploratory investigation in which we employ both survey-based subjective methods *and* brain-based objective measures (using functional near infrared spectroscopy). Specifically, we measured the activation of the anterior prefrontal cortex (PFC) – a functional region of the brain associated with emotion regulation – to investigate whether there are neural correlates of an affective response to highly human-like robots versus humans and less human-like robots. Our findings confirm a “valley” in brain activity in PFC corresponding to that observed via subjective measures, which further supports the affective response posited by the UV. However, the results also reveal a *second* “uncanny valley” in the prefrontal hemodynamics, which suggests that the effects are more complex than previously thought.

¹Tufts University, 200 Boston Avenue, Medford MA USA
{megan.strait, matthias.scheutz}@tufts.edu

II. RELATED WORK

Two primary hypotheses are proposed to underly the effects stemming from the uncanny valley – the *atypical feature hypothesis* and the *category-conflict hypothesis*. The atypical feature hypothesis states that the effects may be a function of a mismatch in features, specifically the presence of one or more atypical features (e.g., [7]). This hypothesis may account for the effects that arise in response to agents that are fairly human-like except for one or two features (e.g., a human face paired with a robotic voice [19]). The second hypothesis posits that on a continuum of two categories (e.g., morphs of a human and a non-human agent linearly interpolated), stimuli that are on the border of two categories will be perceived as ambiguous, and as a result, elicit an uncanny valley response (e.g., [9]). Evaluation of uncanny valley effects are thus primarily assessed in terms of the agent’s appearance and the standard indices characterize agents along a pairing of human-likeness (x-axis) and eeriness (y-axis) [20]. However, others have approached the uncanny valley as a form of mind perception, hypothesizing that greater capacities (e.g., capacity to feel pain) are ascribed to agents with greater human-likeness (e.g., [1], [21]).

Regarding differences in perceptions of human versus human-like agents, there have been a number of recent brain-imaging studies which also may help illuminate the mechanisms involved in uncanny valley effects. Such investigations have shown both differentiable activity levels as well as distinct neural substrates (e.g., [15], [16], [18], [22]). For example, increased responses to robot (compared to human) stimuli have been observed in the occipital and posterior temporal cortices ([15]), which suggests additional visual resources are dedicated when perceiving a human-like agent. In particular, this work showed greater activation bilaterally in the fusiform face area - an area implicated to have human face-specificity – which indicates the perception of human-like facial features of the robotic agent, and further suggests that unfamiliar or uncanny faces may require additional processing of the visual input. Conversely, in perceiving emotional gestures, activity in Brocas area, the left anterior insula, and the orbitofrontal cortex, is reduced for humanoid (versus human) agents, which in contrast to the greater activity observed in the occipital cortices, suggests lesser emotional rapport or resonance with the robotic agent performing the actions [15].

While the aforementioned investigations point to various differences in the perception of human versus human-like agents, there are a number of limitations to their interpretation. Specifically, (1) the subjective assessments of agent characteristics could be interpreted as non-emotional observations of poor aesthetics, (2) eye-gaze behaviors such as low fixation duration could be re-interpreted as disinterest rather than aversion, and (3) the related fMRI work does not investigate the recruitment of the limbic system in particular, nor does it employ particularly uncanny agents.

Recent work on the perception of humans and robots in negatively-valenced contexts has, however, suggested a link

between fronto-cortical hemodynamics and the emotional value of the situation ([21], [22]). Specifically, responses to humans versus robots placed in moral dilemmas shows significant increases prefrontal hemodynamics [21]. In addition, an investigation of people’s affinity/aversions to non-human-like and human-like robots also indicated a relation between increased prefrontal activity and emotion, with increases in hemodynamics corresponding to severe reductions in participants’ interest in interacting with a robot rated as highly eerie ([23], [24]). In general, the prefrontal cortex shows a correspondence with negatively-valenced/highly-arousing stimuli (e.g., [25]) and is directly coupled with the amygdala as part of the limbic system and as an essential functional area in emotion regulation (e.g., [26]). Thus here to explore the relationship between neurophysiological indicators of affect and the standard subjective indices of the uncanny valley, we specifically targeted the prefrontal cortex for its role in processing highly human-like robots versus humans and robots with lesser degrees of human-likeness.

III. MATERIALS AND METHODS

We imaged the anterior prefrontal cortex using a two-channel NIRS instrument (ISS Imagent) while participants viewed a series of 80 images depicting either human or robotic agents. We hypothesized that, in response to highly human-like robots, participants’ prefrontal hemodynamics would show greater change in activity than responses to humans or robots with lesser human-likeness (H1). We further hypothesized that this activity would be associated with self-report measures of high arousal and negative attributions (indexed via ratings of eeriness) to the robot (H2).

A. Stimuli

As human-like characteristics are widely variable between robotic agents and thus impossible to systematically manipulate, we instead collated a set of 80 images (20 human and 60 robot) intended to serve as a spectrum of human-likeness from very mechanical to very human. As a manipulation check, human-likeness was assessed and used as a measure for post-hoc categorization on a nine-point scale. Subjective ratings indicated that approximately half of the images fell into the bottom third of the scale and that the images of humans tended to be perceived as most human-like (see




Fig. 1. Distribution of human-likeness ratings (in blue: human agents).




Fig. 2. Left: average ratings of eeriness as a function of human-likeness for each of the 80 images. Right: average ratings of arousal as a function of eeriness. Blue indicates an image depicting a human agent. Red indicates the best fit model.

Figure 1). The trial structure here consisted of each image being presented to participants for six seconds (preceeded by 3s during which a fixation cross was presented). The resolution of the images was 72dpi with dimensions of 400x400px. The images were presented at the center of a 13-inch monitor on a white background and the ordering of the images was randomized. After the participant viewed all images, the images were then re-presented to obtain self-paced ratings of each image on several dimensions.

B. Measures

We sampled subjective perceptions of the image-based stimuli on a five-point Likert scale regarding two dimensions of agent characteristics – human-likeness and eeriness – each operationalized by three survey items [20]. Based on prior speculation that the eeriness dimension corresponds to a visceral response of participants [27], we also assessed subjective arousal using the Self-Assessment Manikin commonly employed for measuring emotion [28].

Subjective. We sampled subjective perceptions of the image-based stimuli on a five-point Likert scale regarding two common dimensions of agent characteristics – human-likeness and eeriness – each operationalized by three survey items [20]. Based on prior speculation that the eeriness dimension corresponds to a visceral response of participants [20], we also assessed subjective arousal using the Self-Assessment Manikin commonly employed for measuring emotion [28].

Objective. Functional near infrared spectroscopy (NIRS; also known as fNIRS or fNIR) was used to measure participants’ prefrontal neural activity (indexed by hemodynamic changes). A two-channel NIRS oximeter (ISS Imagent; $TR = 11Hz$) was used to collect data from participants’ anterior medial prefrontal cortices (PFC) bilaterally. Based on prior findings that increased oxygenated hemoglobin corresponds to increased aversion ([23]), we expected to observe a positive correspondance between increased hemodynamic activity and higher ratings of eeriness.

C. Population and Procedure

Twenty-six students and staff (16 female) were recruited via an affiliate University website and paid \$10/hour for their participation. All subjects reported being healthy, right-handed, and having no history of brain trauma. Average participant age was 20.8 years old ($SD=2.9$). Upon the receipt of informed, written consent, participants were fitted with the NIRS equipment using a black cap to secure the two sensors on the left and right PFC. A five-minute baseline measurement was then sampled for post-hoc conversion of the raw NIRS data into units of hemoglobin. The participant then proceeded to view the 80 images. After the viewing was complete, participants were then re-presented with each of the images and instructed to rate the images on the two agent-based dimensions (human-likeness, eeriness) and the degree of emotion-induction (arousal).

D. Signal processing

Noise reduction. Prior to analysis, we preprocessed the NIRS data using several standard techniques. Specifically, NIRS data were first converted from raw light attenuation to units of hemoglobin and then filtered to reduce systemic artifacts. Low frequency artifacts were removed by subtracting the timeseries resulting from a 1st degree Savitzky-Golay low pass filter with a cut-off frequency of .08Hz from the original data. A second filter (1st degree Savitzky-Golay, cut-freq. of .5Hz) was then applied to reduce cardiac pulsations. Conversion and filtering yields two signals – oxygenated hemoglobin (HbO) for the left and right PFC – per trial.

Inter-trial variability reduction. The data is then further processed to reduce within-subject variability: each trial is zeroed using the three-second fixation period preceding the stimulus onset and then all trials with the same subjective rating are averaged. For example, if a given participant rated 20 of the 80 images a 1 out of 5 in terms of human-likeness, the corresponding NIRS data (e.g., left oxy-hemoglobin) for those 20 trials would be averaged, resulting in one signal




Fig. 3. Mean baseline-corrected AUC (\pm SD; in $Mol/L \times 10^{-3}$) as a function of arousal (left) and human-likeness (right).

(mean hemodynamic response, HDR). The average of that signal is then computed to yield one statistic (mean AUC).

Baseline correction. Lastly, in order to make inferences regarding the significance of the observed hemodynamic activity, a summary statistic of baseline activity is computed and subtracted out of mean AUC values. Here we used the hemodynamic activity in response to the human stimuli ($N=20$) as the baseline activation. We perform the above reduction of inter-trial variability and then subtract the mean baseline AUC from the experimental AUC. Due to the inherent and wide variability within-subjects of hemodynamic activity (e.g., [29]), the NIRS data were analyzed on a pairwise basis. Statistical inferences were thus made using matched pair t-tests on the resulting (baseline-corrected and normalized) AUC measures, with the null hypothesis that the resulting measure (i.e., difference between the control and experimental condition) is zero.

IV. RESULTS

Using human-likeness as our x-axis, we first analyzed conscious perceptions to characterize our stimulus set based on the aforementioned agent characteristics. We then investigated the relation of the objective NIRS measurements to the emotion dimension and the correspondance of hemodynamic changes as a function of human-likeness.

A. Manipulation Check

To characterize the image set, for each image we averaged self-report ratings across all participants ($N = 26$). This yielded one average rating of arousal, human-likeness, and eeriness per image. The general trend between human-likeness and eeriness was analyzed using Spearman’s correlation coefficient. Exploratory polynomial regression was then used to identify best fit models of the non-linear relationships depicted in Figure 2.

Eeriness. Analysis of the linear trends between eeriness and human-likeness showed a slight, but significant correlation ($r=-.2314$, $p=.0389$). Polynomial regression showed a fourth-order polynomial to best fit the relationship ($r=.6121$). Specifically, along the first half of the scale eeriness ratings increase as ratings of human-likeness increase. Eeriness then peaks around the midpoint of the scale, and then the trend reverses such that eeriness falls as human-likeness continues to increase (see Figure 2, center).

Arousal. We examined the general correspondance between eeriness and arousal using Spearman’s correlation coefficient. The analysis showed a positive, strongly significant linear relationship ($r=.8842$, $p<.0001$; see Figure 2, right). Specifically, perceptions of induced arousal increase with increasing eeriness of an agent.

B. Hypothesis Testing

H1: Do highly human-like robots elicit greater hemodynamic change versus humans or robots with lesser degrees of human-likeness? To understand the relationship between an agent’s perceived human-likeness and the corresponding prefrontal hemodynamic activity, we re-processed the NIRS data averaging oxygenated hemoglobin across all trials with the same rating of *human-likeness* which yielded mean AUC as a function degree of human-likeness, per participant. Figure 3 (right) depicts the baseline-corrected AUC averaged across all participants, by degree of human-likeness.

The results of matched pair t-tests (two-tailed, $\alpha = .05$, $df = N_{obs} - 1$) – testing whether the difference is zero between AUC at each level of arousal versus baseline activity – are shown in Table I. Again, as not all participants utilized the full human-likeness scale (the distribution of ratings of the robot-based stimuli was left-skewed), the number of observations (N) per level of arousal ranged between 12 and 26, hence the critical values for t again vary slightly. All contrasts are statistically significant at the 95% confidence interval (with a Bonferroni-correction of $\alpha = .0056$ for nine contrasts). Here, the results indicate are two substantial peaks (increases) in oxygenated hemoglobin, associated with

TABLE I
MEAN AUC ($Mol/L \times 10^{-3}$) BY DEGREE OF HUMAN-LIKENESS.

Rating	N_{obs}	$T_{\alpha=.0056}$	M	SD	T_{obs}
1	26	± 3.03	-.4077	.0530	-39.22*
1-2	25	3.04	.7114	.0505	70.43*
2	21	3.10	.4469	.1068	19.17*
2-3	25	3.04	.1189	.1297	4.58*
3	19	3.14	-.9835	.1616	-26.53*
3-4	18	3.16	-.5482	.1859	-12.51*
4	15	3.26	.2349	.1884	4.82*
4-5	13	3.36	.5596	.0817	24.71*
5	12	3.43	.2771	.0621	15.44*

a human-likeness rating, h , in the range $[1 < h \leq 2]$ and again, at $[4 < h \leq 5]$. Furthermore, there is a significant valley in oxy-hemoglobin in the range of $[2 < h \leq 3]$

H2: Are the changes elicited in response to highly human-like robots associated with aversion? To understand the relationship between arousal induction and the measures of prefrontal hemodynamic activity, we preprocessed the NIRS data using the aforementioned procedures and analyzed matched pair differences. Here, we averaged oxygenated hemoglobin across all trials with the same rating of *arousal* induction which yielded mean AUC as a function of arousal level, per participant. Each participant’s AUC statistic was then corrected for baseline activity to yield mean paired (within subjects) difference in AUC. Figure 3 (left) depicts the mean baseline-corrected AUC averaged across all participants, by degree of arousal induced.

The results of matched pair t-tests (two-tailed, $\alpha = .05$, $df = N_{obs} - 1$) – testing whether the difference is zero between AUC at each level of arousal versus baseline activity – are shown in Table II. As not all participants utilized the full arousal scale, the number of observations (N) per level of arousal ranged between 18 and 25, hence the critical values for t vary slightly. All contrasts are statistically significant at the 95% confidence interval (with a Bonferroni-corrected $\alpha = .01$), indicating that the increased prefrontal activity associated with the upper half of the arousal scale is significantly greater than the activity in response to human-based stimuli. Furthermore, the decreased prefrontal activity associated with the center of the arousal scale is significantly less than that associated with the human stimuli.

V. DISCUSSION

The results of this exploratory investigation into the uncanny valley effects (UVEs) on prefrontal hemodynamics contribute two primary implications to current understanding of both UVEs as well as neurological differences in how people perceive humans and robots. Furthermore, this investigation yields a characterization of the largest set of image-based stimuli to-date for exploring UVEs.

We first confirmed the utility of our stimulus set via subjective perceptions of agent-based characteristics (*human-likeness, eeriness, and arousal*). While the continuum of agents with increasing human-likeness is strongly skewed left, pairing with both arousal and the eeriness dimension show a “valley” consistent with prior work (e.g., [20]).

We then looked at the relationship between subjective perceptions of an agent’s eeriness, introspective ratings of the corresponding arousal induced, and prefrontal hemodynamic

activity – which indicated increases in oxy-hemoglobin reflect some degree of negative emotion induction. Specifically, we found a strongly significant correlation between ratings of eeriness and arousal ($r = .88$), consistent with prior speculation that the eeriness dimension captures visceral reactions to uncanny agents (e.g., [20]), which suggests that there is indeed an emotional response (e.g., aversion) evoked. To identify the directionality of the relation, we examined the NIRS data as a function of arousal, which showed that an increase in oxy-hemoglobin corresponds to high arousal (and hence by transitivity, high eeriness). This is consistent with [23] who found an increase in HbO corresponded with significant reductions in affinity for a robot perceived as highly eerie. The data also show a substantial decrease mid-scale in hemoglobin, which we speculate may reflect interest.

Lastly, we examined the relationship between subjective perceptions of human-likeness and prefrontal hemodynamics, which suggested the existence of *two* uncanny valleys. Specifically, we found two peaks in hemodynamic activity: one near the ceiling of the scale (closest to human), consistent with prior indications, but also a peak in the range of $[1 < h \leq 2]$. The correspondance between hemodynamic activity and arousal suggests that these two peaks in hemodynamic activity reflect the induction of some degree of negative emotion, supporting both of the prominent hypotheses. Specifically, the peak in the range of human-likeness closest to human stimuli supports the category conflict elicited by agents bordering the category of human (versus non-human). Whereas inspection of images that received average ratings of human-likeness in that lower range revealed robots that have one or two primary human-like features (e.g., eyes and/or fingers) but lack other key features (e.g., mouth, arms, etc.), including, for example: the Wakamaru, iCub (without casing), Robothespian, Eccerobot, and Albert Hubo robot. Hence, the peak lower in the human-likeness spectrum shows support for the elicitation of UVEs due to an atypical feature or conflict in features (e.g., eyes) combined with an otherwise non-human appearance (e.g., lack of mouth, arms, etc.).

Hence, our findings suggest that there is indeed an emotional response to robotic agents with eerie appearances and that this response, moreover, is evoked not only for highly human-like robots, but *also* for robots with low ratings of overall human-likeness (i.e., robots with ratings of $[1 < h \leq 2]$) but that have some key human-like features. Furthermore, analysis of the prefrontal hemodynamics suggests that the subjective measures alone do not fully capture the relationship between human-likeness and the uncanny valley. Hence, our experimental results also have methodological consequences for the field of HRI in that they point to the need for additional measures such as imaging of prefrontal neural activity for understanding the uncanny valley and untangling the complex factors contributing to its effects.

VI. LIMITATIONS

The aim of this investigation was to explore the neurophysiological manifestations of uncanny valley effects and understand how they relate to a person’s affinity/aversion to

TABLE II
MEAN BASELINE-CORRECTED AUC ($Mol/L \times 10^{-3}$) BY AROUSAL.

Rating	N_{obs}	$T_{\alpha=.01}$	M	SD	T_{obs}
1	24	+/-2.80	.0616	.0411	7.34*
2	25	2.79	-.0933	.0951	-4.90*
3	25	2.79	-.6923	.1209	-28.63*
4	21	2.84	.7357	.1257	26.82*
5	18	2.89	.2771	.0621	18.93*

a robotic agent. Despite the encouraging results, several remaining limitations require additional experimentation. First and foremost, this study is limited by the stimuli. Specifically, the stimuli used here were comprised of still images of human and robotic agents. While this addresses the effects that aesthetics and design of robotic agents may have on human observers, it does not measure any aspect of human-robot interaction (e.g., effects of voice/appearance mismatch, [19]). Secondly, although the sampling of prefrontal hemodynamics constitutes an objective measure, the implications for regional activation and correspondance between emotion induction and directionality of oxygenated hemoglobin are inferential. Specifically, a two-channel NIRS instrument does not allow for specific localization of the activation source (e.g., differentiating contributing activity from the medial versus lateral prefrontal cortex). Moreover, the correspondance between emotion and oxy-hemoglobin above was indexed via an introspective measure of emotion, which although correlated with a number of physiological indicators, is still inferential. Further investigation employing a number of confirmatory measures such as fMRI to localize activation or neurophysiological measures (e.g., galvanic skin response, pupil dilation, facial corrugator activation) to corroborate affect induction from uncanny stimuli is necessary.

VII. CONCLUSIONS

The uncanny valley describes the sudden change in a person's affect from affinity to aversion evoked by robots that border a human-like appearance. Most prior evidence for this effect was based on subjective evaluations. And while several neurophysiological investigations point to underlying differences in the perception of humans versus human-like robots, the relation of such perceptual differences to the emotion induction posited by the UV remains unaddressed, begetting the question of whether the effects are mainly observational. We have addressed this shortcoming in an exploratory evaluation in which we employed both subjective *and* objective measures. Our findings confirm a "valley" in prefrontal hemodynamics corresponding to the uncanny valley observed via subjective measures, thus supporting the affective response posited by the UV. However, the results also reveal a *second* "uncanny valley" in prefrontal hemodynamics, which suggests that the effects (and the contributing factors) are more complex than previously understood and thus warrant additional investigations.

REFERENCES

- [1] K. Gray and D. M. Wegner, "Feeling robots and human zombies: mind perception and the uncanny valley," *Cognition*, 2012.
- [2] M. Walters, D. Syrdal, K. Dautenhahn, R. Te Boekhorst, and K. Koay, "Avoiding the uncanny valley: robot appearance, personality and consistency of behavior in an attention-seeking home scenario for a robot companion," *Autonomous Robots*, 2008.
- [3] C. Lisetti, "Believable agents, engagement, and health interventions," in *Human-Computer Interaction International*, 2011.
- [4] H. Ishiguro, "Android science," in *Cogsci*, 2005.
- [5] K. F. MacDorman and N. Kageki, "The uncanny valley," *IEEE Robotics and Automation*, 2012.
- [6] M. Mori, "The uncanny valley," *Energy*, 1970.

- [7] K. F. MacDorman, P. Srinivas, and H. Patel, "The uncanny valley does not interfere with level 1 visual perspective taking," *Computers in Human Behavior*, 2013.
- [8] K. F. MacDorman, "Subjective ratings of robot video clips for human likeness, familiarity, and eeriness: An exploration of the uncanny valley," in *CogSci*, 2006.
- [9] T. J. Burleigh, J. R. Schoenherr, and G. L. Lacroix, "Does the uncanny valley exist? an empirical test of the relationship between eeriness and the human likeness of digitally created faces," *Computers in Human Behavior*, 2013.
- [10] C. Bartneck, T. Kanda, H. Ishiguro, and N. Hagita, "Is the uncanny valley an uncanny cliff?" in *RO-MAN*, 2013.
- [11] H. Brenton, M. Gillies, D. Ballin, and D. Chatting, "The uncanny valley: does it exist?" in *Animated Characters Interaction*, 2005.
- [12] A. Tinwell and M. Grimshaw, "Bridging the uncanny: an impossible traverse?" in *Mindtrek*, 2009.
- [13] E. Broadbent, V. Kumar, X. Li, J. Sollers, R. Stafford, B. MacDonald, and D. Wegner, "Robots with display screens: a robot with a more humanlike face display is perceived to have more mind and a better personality," *PLoS one*, vol. 8, no. 8, p. e72589, 2013.
- [14] H. Chen, R. Russell, and K. Nakayama, "Crossing the 'uncanny valley': adaptation to cartoon faces can influence perception of human faces," *Perception*, 2010.
- [15] T. Chaminade, M. Zecca, S. Blakemore, A. Takanishi, C. D. Frith, S. Micera, P. Dario, G. Rizzolatti, V. Gallese, and M. A. Umiltà, "Brain response to a humanoid robot in areas implicated in the perception of human emotional gestures," *PLoS one*, 2010.
- [16] M. Cheetham, P. Suter, and L. Jncke, "The human likeness dimension of the uncanny valley hypothesis: Behavioral and functional mri findings," *Frontiers in Human Neuroscience*, 2011.
- [17] E. S. Cross, R. Liepelt, A. F. de C. Hamilton, J. Parkinson, R. Ramsey, W. Stadler, and W. Prinz, "Robotic movement preferentially engages the action observation network," *Human Brain Mapping*, 2012.
- [18] A. Saygin, T. Chaminade, H. Ishiguro, J. Driver, and C. Frith, "The thing that should not be: predictive coding and the uncanny valley in perceiving human and humanoid robot actions," *Social Cognitive and Affective Neurosci*, 2012.
- [19] W. J. Mitchell, K. A. Szerszen, A. S. Lu, P. W. Schermerhorn, M. Scheutz, and K. F. MacDorman, "A mismatch in the human realism of face and voice produces an uncanny valley," *i-Perception*, 2011.
- [20] C.-C. Ho and K. F. MacDorman, "Revisiting the uncanny valley theory: Developing and validating an alternative to the godspeed indices," *Computers in Human Behavior*, 2010.
- [21] M. Strait, G. Briggs, and M. Scheutz, "Some correlates of agency ascription and its role in emotional and non-emotional decision-making," in *Affective Computing and Intelligent Interaction*, 2013.
- [22] A. M. R. von der Putten, F. P. Schulte, S. C. Eimler, L. Hoffmann, S. Sobieraj, S. Maderwald, N. C. Kramer, and M. Brand, "Neural correlates of empathy towards robots," in *HRI*, 2013.
- [23] M. Strait, C. Canning, and M. Scheutz, "Let me tell you! Investigating the effects of robot communication strategies in advice-giving situations based on robot appearance, interaction modality, and distance," in *Human-Robot Interaction (HRI)*, 2014.
- [24] M. Strait and M. Scheutz, "Using near infrared spectroscopy to index temporal changes in affect in realistic human-robot interactions," in *Physiological Computing Systems (PhyCS)*, 2014.
- [25] D. Heger, R. Mutter, C. Herff, F. Putze, and T. Schultz, "Continuous recognition of affective states by functional near infrared spectroscopy signals," in *Affective Computing and Intelligent Interaction*, 2013.
- [26] H. L. Urry, C. M. van Reekum, T. Johnstone, N. H. Kalin, M. E. Thurow, H. S. Schaefer, C. A. Jackson, C. J. Frye, L. L. Geischar, A. L. Alexander, and R. J. Davidson, "Amygdala and ventromedial prefrontal cortex are inversely coupled during regulation of negative affect and predict the diurnal pattern of cortisol secretion among older adults," *Journal of Neuroscience*, 2006.
- [27] C.-C. Ho, K. F. MacDorman, and Z. Pramono, "Human emotion and the uncanny valley: a GLM, MDS, and isomap analysis of robot video ratings," in *HRI*, 2008.
- [28] M. Bradley and P. Lang, "measuring emotion: The self-assessment manikin and the semantic differential," *Journal of Behavioral Therapy and Experimental Psychiatry*, 1994.
- [29] C. Canning and M. Scheutz, "Functional near-infrared spectroscopy in human-robot interaction," *Journal of Human-Robot Interaction*, 2013.