Perception, 2011, volume 40, pages 695 – 704

doi:10.1068/p6900

The perception of humanness from the movements of synthetic agents

James C Thompson, J Gregory Trafton¶, Patrick McKnight
Department of Psychology, George Mason University, 4400 University Drive, MS3F5, Fairfax, VA 22030, USA; e-mail: jthompsz@gmu.edu; ¶ Naval Research Laboratory, Code 5515, 4555 Overlook Avenue, Washington, DC 20375-5337, USA
Received 23 December 2010, in revised form 24 May 2011; published online 28 July 2011

Abstract. As technology develops, social robots and synthetic avatars might begin to play more of a role in our lives. An influential theory of the perception of synthetic agents states that as they begin to look and move in a more human-like way, they elicit profound discomfort in the observer—an effect known as the Uncanny Valley. Previous attempts to examine the existence of the Uncanny Valley have not adequately manipulated movement parameters that contribute to perceptions of the humanness or eeriness. Here we parametrically manipulated three different kinematic features of two walking avatars and found that, contrary to the Uncanny Valley hypothesis, ratings of the humanness, familiarity, and eeriness of these avatars changed monotonically. Our results indicate that, when a full gradient of motion parameter changes is examined, ratings of synthetic agents by human observers do not show an Uncanny Valley.

1 Introduction

Humans have a remarkable capacity to perceive and understand the actions of others. From subtle differences in the way someone moves we can identify qualities such as their gender, emotion, and even their social status (Blake and Shiffrar 2007; Johansson 1973). As virtual and robotic technology has changed, however, we are increasingly faced with the need to interpret the actions of synthetic social agents: computer-generated (CG) avatars and social robots. Synthetic agents have become more common in workplace, social, and entertainment settings, and as such it has become increasingly important to understand the factors that influence our interactions with them. One theory suggests that, as a robot or CG avatar becomes more human-like in appearance, we experience the viewing of the synthetic agent increasingly more pleasant until its appearance reaches a point at which very subtle differences from human-like produces a feeling of profound discomfort in the observer (Mori 1970). This effect has been called ‘bukimi no tani’ or the ‘Uncanny Valley’ (Mori 1970). The unpleasant feeling we experience when we see a synthetic agent that moves very much like a human is suggested to be something akin to encountering a zombie or the undead (Mori 1970). Avoiding the Uncanny Valley in the design of synthetic agents has become an important design principle in robotics and graphic design communities (Fabri et al 2004; Fong et al 2003; Groopman 2009).

While there are a number of psychological accounts that have been proposed to underlie the Uncanny Valley effect (MacDorman and Ishiguro 2006; Steckenfinger and Ghazanfar 2009), actual evidence in favor of the existence of the Uncanny Valley is sparse. One study reported finding an Uncanny Valley in participant ratings of the eeriness and humanness of a sequence of morphed static images that spanned three different parameter spaces: from non-human robot face to human-like robot face, then to human face (MacDorman and Ishiguro 2006). Unfortunately the ‘valley’ was observed at the transition from non-human robot to human-like robot, whereas the Uncanny Valley hypothesis predicted that it should be at the transition from human-like robot to the human face. An evolutionary mechanism contributing to the Uncanny Valley was recently suggested, when it was reported that monkeys spent more time looking at either an 'unrealistic' CG monkey avatar or an actual monkey than they did looking at a 'realistic' CG monkey (Steckenfinger and Ghazanfar 2009). The reduced looking time to the realistic CG monkeys was attributed to macaques perceiving these stimuli as real conspecifics that don't live up to expectations of what a conspecific should look like, although factors such as the perceived (un)attractiveness or threat of the realistic avatar could not be ruled out.

One of the limitations of the few studies of the Uncanny Valley is that there has been no attempt to systematically manipulate a well-controlled parameter and examine the effects of perceived humanness or eeriness. Instead, by morphing between images that vary across a wide range of parameters (MacDorman and Ishiguro 2006), or presenting just a few examples (Steckenfinger and Ghazanfar 2009), it is not possible to isolate parameters that contribute to a synthetic figure or robot appearing eerie or weird. In addition, the previous attempts to determine the existence of the Uncanny Valley have focused on the perception of synthetic faces. Faces appear to be coded in a high-dimensional face space (Lee et al 2000), and isolating the dimension(s) within this space that might produce an Uncanny Valley effect is complex. In one of the few studies that have attempted to present a continuum of morphs from animated image to a photo of a real actor, Hanson et al (2005) found that all levels along the continuum were rated positively.

In the present study we, instead, focused on the role of motion in the perception of humanness and eeriness of animated figures. It has been suggested that adding natural motion captured from humans to robots will make the robots appear more human-like (Matsui et al 2007). Others have used human-like movements and body postures to make robots appear less creepy or threatening (Bethel and Murphy 2006). At the same time, Matsui et al (2007) warn that, owing to the differences in kinematics and joint structures, slight differences in the motion between a human body and a robot could strongly influence the way the robot is perceived, producing an Uncanny Valley effect. Indeed, a key proposal of the Uncanny Valley hypothesis made by Mori (1970) was that motion plays a pivotal function in the generation of the Uncanny Valley: even slight variations from human-like movements are suggested to produce or amplify the 'valley' in ratings of familiarity or eeriness. We created two CG avatars that walked in place, and parametrically manipulated three different kinematic parameters that are important to the perception of gait. We examined the effects of manipulations to these different kinematic parameters to ratings of the humanness, familiarity, and eeriness of the avatars. Using a model comparison approach, we compared the Uncanny Valley account to a model in which preferences for more human-like avatars increased monotonically.

2 Methods

2.1 *Participants*

A total of forty undergraduate students from George Mason University participated in this study in exchange for course credits (twenty-two males, mean age $\pm$ SD $= 19.1 \pm 0.7$ years). All participants had normal or corrected-to-normal vision. All participants provided written informed consent and all procedures were approved by the Human Subjects Review Board at George Mason University. Participants were randomly assigned to either a Mannequin avatar condition or a Whole-body avatar condition.

2.2 *Stimuli*

Stimuli consisted of two avatars, created in Poser 5.0 and animated with the walking motion of an individual human supplied by Vicon Motion Systems (Oxford, UK) and modified in Matlab (Mathworks, Natick, MA). Gait was chosen as the action, as it is familiar and unambiguous, and because there has been considerable research into the

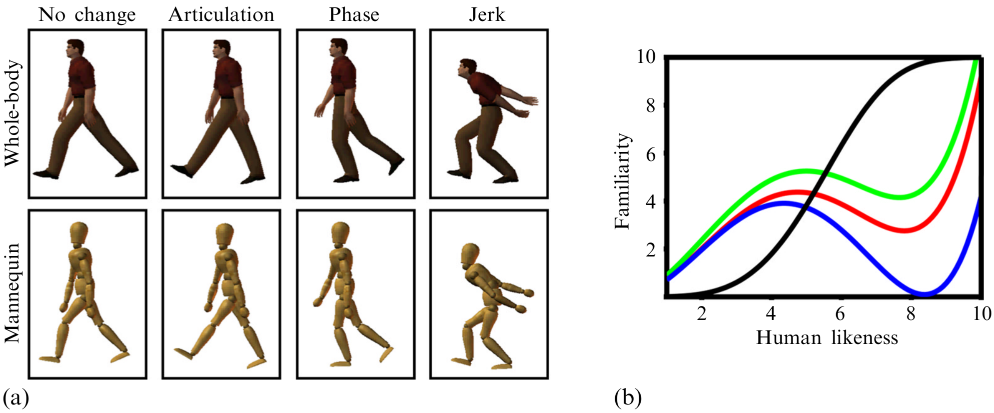

Figure 1. Avatars and the Uncanny Valley as a polynomial function. (a) The two different avatars and sample frames from the unchanged-walker video and from the three kinematic manipulation videos. (b) The effects of the manipulations to humanness on ratings of familiarity (or alternatively, 1/eeriness) modeled as either a 5th-order polynomial function (red, green, or blue) or a Weibull function (black). In the polynomial model, the 'valley' is determined by the 3rd order term, as is demonstrated by the three colored lines. The red and green lines differ only in terms of the 1st order term (red = 0.22; green = 0.44), whereas the red and blue lines differ only in terms of the 3rd order term (red = 0.0105; blue = 0.011). Note that even a small fluctuation in the 3rd order term can amplify the valley, indicating the sensitivity of this model to a possible Uncanny Valley.

parameters that contribute to gait perception. One figure consisted of a Mannequin with 13 joints (figure 1a) that was chosen to provide minimal form cues apart from those consistent with a human body shape. The second figure consisted of a whole-body male walker with 13 joints that was selected in order to examine the possible interaction between the presence of further form cues (hair, face, clothing) and movement in the generation of an Uncanny Valley effect (figure 1a).

We manipulated three different kinematic parameters that are important to the perception of human gait. We reduced the articulation provided by the multiple ball-and-socket joints that comprise the human body by decreasing the joint rotation of the wrists, elbows, knees, and ankles. This made the figure walk with increasingly stiff arms and legs. Articulation has been shown to be an important contributor to the perception of human motion (Aggarwal et al 1994; Beauchamp et al 2002). We introduced a phase offset between opposing limbs during the gait cycle, reducing the fluidity of the gait. Several studies have demonstrated that the phase relationship between opposing limbs, especially the feet, is an important factor in the perception of gait (Bertenthal and Pinto 1994; Casile and Giese 2005). We also introduced a biomechanically implausible jerk action at a random point during the gait cycle (Candidi et al 2008). Still images from the videos are shown in figure 1, and videos of the stimuli are available on the *Perception* website (see http://dx.doi.org/10.1068/p6900). The first two parameters (Articulation, Phase) comprised a manipulation that was constant throughout the phase cycle, while the third parameter (Jerk) was a temporally brief manipulation that occurred at random points across a gait cycle. Morphs between each of these modifications and the original motion files, in 10% increments, were made using a spatiotemporal morphing algorithm modified from Thompson et al (2005). Each motion file was then anchored from least 'natural' (ie least amount of Articulation, most amount of Jerk, largest Phase-offset) to most 'natural' (ie full Articulation, no Jerk, no Phase-offset), where most 'natural' is the original, unaltered motion caption file. The ten levels of the Articulation, Phase, and Jerk manipulations were then imported into Poser and applied to the two different figures with a tracking camera so that the figure walked with no net translation.

2.3 *Procedure*

Stimuli were presented on a PC computer with E-Prime software. Participants sat approximately 60 cm from a 22-inch CRT computer monitor. They were instructed that they would see short videos of an animated figure and would be asked to make some judgments about the figure. Participants were then shown a 4 s (three cycles of gait) video of one of the animated figures walking and were asked "How human did you find the figure?", "How familiar did you find the figure?", and "How eerie did you find the figure?", with each question answered on a seven-point (1 indicating "not at all" to 7 indicating "extremely") Likert scale. These three measures have been used in previous examinations of the Uncanny Valley (MacDorman and Ishiguro 2006; MacDorman et al 2009), as well as the initial proposal by Mori (1970), allowing the direct comparison of the findings of the present study with previous research. Ratings of humanness and familiarity were expected to rise, and eeriness was expected to decline, as stimuli became more natural. Participants responded by pressing the appropriate number (1 to 7) on the keypad of the keyboard, and were encouraged to answer according to their first impression of the stimuli. After the final response, the next trial began. The three different manipulations were presented in separate blocks, with the 10 levels of each manipulation presented in random order within a block. Each participant viewed 8 trials of each of the 10 levels for the three manipulations, presented across 6 blocks.

2.4 *Analysis*

Rather than simply testing if ratings of the walking CG avatars demonstrated a pattern that could be interpreted as either like the hypothesized Uncanny Valley effect or not, we instead used a model-comparison approach. In a model-comparison analysis, if a more parsimonious model can provide as good or better account of observer ratings than a model of the Uncanny Valley effect, it would provide positive evidence against the Uncanny Valley. The Uncanny Valley hypothesis predicts a non-rectilinear relationship between manipulations of the animated figure and the reaction of observers. The relationship between these manipulations and observer ratings can be characterized by a fifth-order polynomial model [figure 1b, see equation (1)]. Initial model fitting with third-, fourth-, or sixth-order polynomial produced poorer model fits than the fifth-order polynomial. The predictor (eg level of Articulation) gets transformed in a standard polynomial sequence ranging from 0 (the intercept) to some integer value k. Note that our model omitted the intercept because the Uncanny Valley would specify an intercept at 0; thus it was implied rather than estimated. Higher values of k allow for the upward and downward shifts expected in the Uncanny Valley.

$$y = b_1 x + b_2 x^2 + b_3 x^3 + b_4 x^4 + b_5 x^5 \ . \qquad (1)$$

In the fifth-order polynomial model, the 'valley' is determined by the third-order term. As is shown in figure 1b, even small fluctuations of the third-order term can amplify the valley, indicating the sensitivity of this model to a possible Uncanny Valley.

An alternative model that eliminates any possibility of a break in continuity is the Weibull (1951) function—a function that necessitates a relationship that cannot decrease from any previous point. The advantage of the Weibull function was that it restricted the expected relationship and did so in the most parsimonious manner, as it had one fewer parameter than the polynomial model. Our Weibull model deviated slightly from the generic equation in that we specified two scaling factors S and C to account for the metric of our response variable [see equation (2)]. These two scaling factors provided the ability to estimate the lower and upper bounds of the curve, while also adjusting the predicted values to fall within the range of the observed reaction ratings.

$$y = \left[1 - \mathrm{e}^{-(x/\lambda)^k}\right] S + C \ . \qquad (2)$$

The parameters of the polynomial and Weibull-function models were estimated using the nls function in R (http://www.r-project.org). The polynomial model had five parameters while the Weibull one had four parameters, making it the more parsimonious of the two models. By using the same estimation algorithm, we were able to directly compare the relative predictions via a standard F test of the residual sums of squares. In total, 18 model pairs were tested by using a 2 model (polynomial versus Weibull) $\times$ 3 predictor (Articulation, Jerk, and Phase) $\times$ 3 observer rating (Humanness, Familiarity, and Eeriness) $\times$ 2 (Whole-body versus Mannequin) design. The model comparison serves as the most important hypothesis since that comparison assesses the possibility of the Uncanny Valley—possible only if the polynomial model fits better than the Weibull model. The advantage of the nested model comparison over traditional hypothesis tests is that we know both models will fit relatively well and certainly better than by chance alone (ie the parameters will all be significant) found from initial model-fit tests; thus, hypothesis testing provided little to no test of the relevant hypothesis. Instead, the nested-model comparison relied solely on the model fit for each function and the better-fitting more parsimonious model would be preferred.

3 Results

The observer ratings of Humanness, Familiarity, and Eeriness for the manipulations of Articulation, Jerk, and Phase-offset, for the Mannequin and Whole-body avatars are shown in figures 2 and 3. These figures also show the empirical fit of the Weibull (blue) and polynomial models (red). The results were plotted so that increases in movement ‘naturalness’ were represented as increases along the x-axis. Ratings of Humanness and Familiarity increased as the gait of the Mannequin and Whole-body avatars increased in Articulation, while ratings of Eeriness decreased as Articulation increased. Similarly, ratings of Humanness and Familiarity increased, and ratings of Eeriness decreased, as the amount of Jerk or the Phase change decreased. As can be seen from the effects of changes to the three separate movement parameters for the Mannequin

Table 1. Nested-model comparison results.

Stimuli	DV	IV	Weib Res SS	Poly Res SS	DF	SS	F value	Pr(> F)
Whole-body	human	Articul	225.23	226.42	1	−1.19	−1.13	1
	familiar	Articul	219.87	222.56	1	−2.68	−2.59	1
	eerie	Articul	276.77	300.59	1	−23.81	−17.03	1
	human	Jerk	319.57	325.4	1	−5.84	−3.86	1
	familiar	Jerk	325.97	332.93	1	−6.96	−4.49	1
	eerie	Jerk	336.16	352.26	1	−16.11	−9.83	1
	human	Phase	336.25	363.85	1	−27.61	−16.31	1
	familiar	Phase	352.42	372.75	1	−20.33	−11.73	1
	eerie	Phase	396.49	397.48	1	−0.99	−0.54	1
Mannequin	human	Articul	305.68	306.05	1	−0.37	−0.21	1
	familiar	Articul	255.99	258.77	1	−2.78	−1.88	1
	eerie	Articul	518.51	529.11	1	−10.59	−3.5	1
	human	Jerk	290.74	290.64	1	0.1	0.06	0.81
	familiar	Jerk	289.24	288.68	1	0.56	0.34	0.56
	eerie	Jerk	507.76	518.44	1	−10.69	−3.61	1
	human	Phase	258.1	273.92	1	−15.82	−10.11	1
	familiar	Phase	285.96	300.06	1	−14.1	−8.22	1
	eerie	Phase	449.97	452.82	1	−2.85	−1.1	1

Note: Weib Res SS = Weibull residual SS, Poly Res SS = polynomial residual SS. DF = degrees of freedom of the model comparison. F = F-value from the model comparison. Pr(> F) = p-value for F given DF from the model comparison.

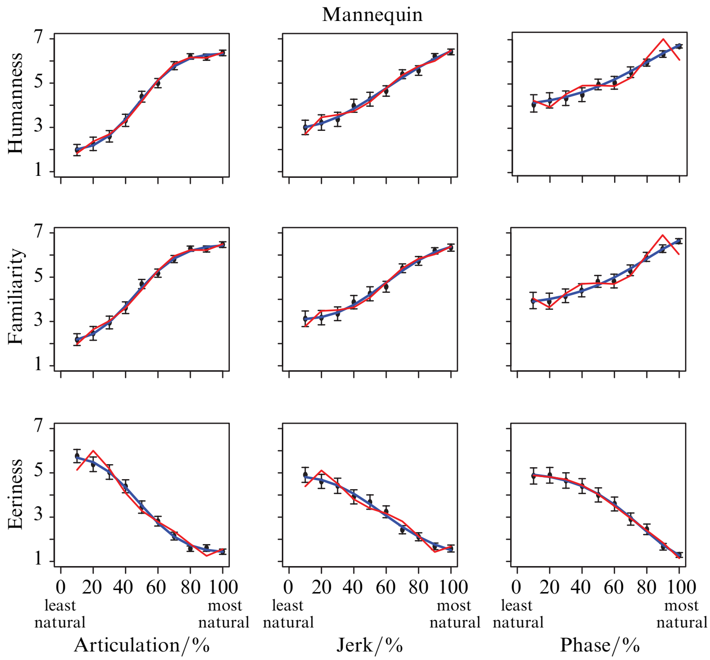

Figure 2. Effects of the kinematic manipulations to the Mannequin avatar to observer ratings of Humanness, Familiarity, and Eeriness. Results are plotted as a function of increasing movement 'naturalness', where 100% natural is the movement from original motion capture file. Circles represent the mean rating, error bars represent ±1 SEM. Red line represents the polynomial function, blue line represents the Weibull function.

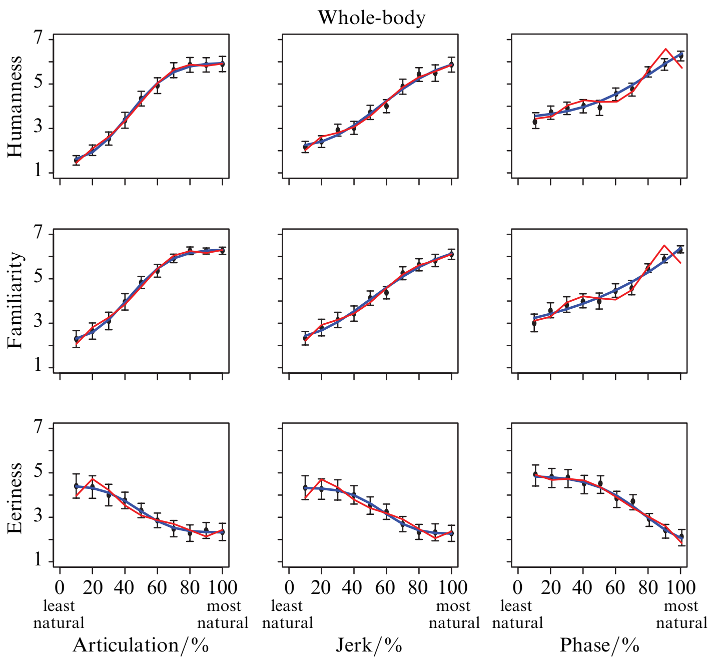

Figure 3. Effects of the kinematic manipulations to the Whole-body avatar to observer ratings of Humanness, Familiarity, and Eeriness. Results are plotted as a function of increasing movement 'naturalness', where 100% natural is the movement from original motion capture file. Circles represent the mean rating, error bars represent ±1 SEM. Red line represents the polynomial function, blue line represents the Weibull function.

(figure 2) and Whole-body (figure 3) conditions, there did not appear to be any sign of a valley in observer ratings. Instead, responses appeared to be characterized by a floor, followed by a rise in ratings, and a ceiling effect.

In all conditions and stimuli, the Weibull and polynomial models fit the data adequately with the expected significant effects for every predictor parameter [all ps < 0.05; results available on the *Perception* website (see http://dx.doi.org/10.1068/p6900, tables S1 & S2)]. A nested-model comparison based upon the ratio of the residual SS indicated the two models did not differ significantly for any manipulation (see table 1). Furthermore, the Weibull models tended to fit better despite having one fewer parameter. Relative fit can be assessed by the residual SS and, since the Weibull consistently produced lower residual SS, we can say with some certainty that the Weibull function may not statistically differ from the polynomial but in all cases it would be considered the preferred model. Visual examination of the empirical fit of the Weibull and polynomial models clearly shows that the two functions are largely overlapping. The instances in which the two functions diverge (eg Eeriness ratings for Articulation and Jerk with both the Whole-body or Mannequin and Humanness ratings for Phase-change with Whole-body or Mannequin) are characterized by the largest differences in residual SS in favor of the Weibull model (table 1).

4 Discussion

Advances in social robots, entertainment, and virtual-reality technology have increased the likelihood that we will see and interact with synthetic agents. While there has been considerable research into how we perceive human movement, little is understood of our perception of the movement of CG agents. Previous theoretical work suggested that as an android or CG avatar becomes closer to human-like, it evokes a sense of profound discomfort in the observer—a phenomenon labeled the Uncanny Valley (Mori 1970). The original description of the Uncanny Valley emphasized that movement plays a vital role in the generation of this effect. While there has been little empirical evidence to support this proposal, it has proven highly influential in robotic design communities (Groopman 2009). In the present study, we examined the contribution of motion to the generation of the Uncanny Valley effect. We systematically altered the articulation, phase relationship, and a biomechanically implausible jerk movement of two different CG avatars and found no evidence in support of the Uncanny Valley. Instead, we found that as the avatars moved in a more human-like manner, observers systematically rated them as more human-like, more familiar, and less eerie. Our findings indicate that parametric manipulation of the kinematics of CG avatars changes how they are perceived by human observers in a monotonic manner, up to a ceiling level.

The three kinematic variables manipulated in the present study have previously been demonstrated to play an important role in the visual perception of human movement. The articulation provided by the multiple ball-and-socket joints that comprise the human body present a complex pattern of motion, and there is evidence that neural mechanisms involved in the processing of human movements are tuned to the articulatory structure of the human body (Beauchamp et al 2002; Jastorff et al 2006). Similarly, the phase relationship between limbs is an important cue for the visual processing of human movement, and disturbing this phase relationship has strong effects of the ability to discriminate different patterns of biological motion (Bertenthal and Pinto 1994; Casile and Giese 2005). It has been demonstrated that the biomechanical plausibility of a movement affects the way that it is perceived (Avenanti et al 2007; Moseley and Brugger 2009). Our data indicate that all three of these factors contribute to the perception of humanness of CG avatars, and that their systematic disruption leads to an avatar being perceived as eerie or unfamiliar.

There have been multiple theoretical accounts of the basis of Uncanny Valley, from pathogen avoidance to evolutionary aesthetics (MacDorman and Ishiguro 2006; Steckenfinger and Ghazanfar 2009). Each of these accounts has been based on the assumption that CG figures or androids look eerie or weird because they too closely resemble real humans. However, the results of the present study would indicate that this assumption is not correct at least as far as the movement of these CG figures is concerned. We found that as CG avatars moved in a more human-like way, they were perceived as less eerie or unfamiliar. Our findings suggest that, at least for motion, there is no Uncanny Valley when synthetic figures show smooth gradient of change from non-human motion to human motion. What then determines why some CG avatars or androids sometimes look strange or eerie? Several studies have indicated that disruption of configural face cues might make CG figures look eerie or unpleasant (Seyama and Nagayama 2007; MacDorman et al 2009). A modified version of the Uncanny Valley hypothesis suggested that the increased realism present in highly rendered CG figures makes the observer less tolerant of minor abnormalities of appearance, such as incorrect skin color or texture (Steckenfinger and Ghazanfar 2009). According to this proposal, more realistic figures are treated as real and thus evoke expectations about how they should look or act. When these expectations are not met, the stimuli are perceived as strange or unusual. Our findings do not support the extension of this proposal to the way we perceive the motion of CG figures, as we found that as avatars moved in a more human-like manner they were perceived as less strange or unusual. A study by Chaminade et al (2007) found that participants were more likely to rate a point-light figure running as ‘biological’ than a rendered, anthropomorphic figure. These authors did not examine manipulations of the motion of the figure, although their findings do suggest that the form and the motion of an avatar might interact to influence how ‘biological’ it is perceived to be. Our findings do not necessarily imply that form and motion cannot interact to influence the perception of avatars—it is not, however, clear that such an interaction would produce an Uncanny Valley.

One could argue that our data did not demonstrate an Uncanny Valley because the motion of the stimuli was not realistic enough, and that ratings of our stimuli fall on the left side of the ‘valley’. Certainly we did not capture the motion of the muscles, movement of the skin or hair, or changes to clothes. However, we find this argument unconvincing for several reasons. First, ratings of humanness and familiarity of our altered stimuli were at or close to ceiling. It is possible that such ratings represent a relative scaling, and that absolute levels of perceived humanness could be achieved by animating body and facial musculature and skin. In turn, such animation could theoretically lead to an Uncanny Valley. Against this claim, robots such as those developed by Hanson et al (2005), which have elastic skin that folds and wrinkles with facial expressions, have avoided the Uncanny Valley. Second, from a practical sense one can conclude from our data that, when using avatars animated by motion-capture data from human actors, one can create stimuli that are perceived as human-like; thus one need not fear straying into the Uncanny Valley with such stimuli. More importantly, however, the weakness of this argument is that for any given stimulus set in which evidence contrary to the Uncanny Valley is reported, one can always argue that it is not realistic enough and one has not reached the ‘valley’ yet. Unless one specifies the parameters that actually produce the Uncanny Valuable, this line of argument renders the hypothesis unfalsifiable.

It is possible that the Uncanny Valley exists for the perception of static stimuli, such as faces, but is not produced by changes to the motion of stimuli. This is not to say that observers do not notice errors in the kinematics of synthetic stimuli—on the contrary, the changes to ratings of the humanness of CG figures following the manipulation of kinematic parameters in the present study would suggest that observers are very sensitive to such errors. However, it is possible that the faces of synthetic agents sometimes might look eerie or weird because of some highly atypical combination of parameters that is easily avoided. From a theoretical standpoint, if the perceived eeriness of a figure results not from its similarity to a real human, but because of some idiosyncratic position in high-dimensional face space, then it is misleading to call these phenomena an Uncanny Valley. This possibility would be consistent with a far simpler account of how we perceive synthetic agents of varying degrees of humanness: that is, as they become more human-like, they appear less eerie or weird. To properly demonstrate the existence of the Uncanny Valley it is important to systematically manipulate basic features that contribute to perceptions of realism or humanness. By using a smooth gradient of changes to a single-parameter space that contributes to perceptions of realism or humanness it should be possible to determine if the weirdness or eeriness of a synthetic figure shows a bump as it becomes very close to human-like in its appearance.

Acknowledgments. We wish to thank Evan Francis and J Malcolm McCurry for assistance with data collection. This work was partially supported by the Office of Naval Research under funding document N0001409WX20173 to JGT. The views and conclusions contained in this document should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Navy.

References

Aggarwal J K, Cai Q, Liao W, Sabata B, 1994 "Articulated and elastic non-rigid motion: A review" *Proceedings of the IEEE Workshop on Motion of Non-Rigid and Articulated Objects* (Austin, TX: IEEE Computer Society)pp 2 – 14

Avenanti A, Bolognini N, Maravita A, Aglioti S M, 2007 "Somatic and motor components of action simulation" *Current Biology* **17** 2129 – 2135

Beauchamp M S, Lee K E, Haxby J V, Martin A, 2002 "Parallel visual motion processing streams for manipulable objects and human movements" *Neuron* **34** 149 – 159

Bertenthal B I, Pinto J, 1994 "Global processing of biological motion" *Psychological Science* **5** 221 – 225

Bethel C L, Murphy R R, 2006 "Affective expressions in appearance-constrained robots" *HRI'06: Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human – Robot Interaction* (Montreal: ACM SIGCHI) pp 327 – 328

Blake R, Shiffrar M, 2007 "Perception of human motion" *Annual Reviews of Psychology* **58** 47 – 73

Candidi M, Urgesi C, Ionta S, Aglioti S M, 2008 "Virtual lesion of ventral premotor cortex impairs visual perception of biomechanically possible but not impossible actions" *Social Neuroscience* **3** 388 – 400

Casile A, Giese M A, 2005 "Critical features for the recognition of biological motion" *Journal of Vision* **5** 348 – 360

Chaminade T, Hodgins J, Kawato M, 2007 "Anthropomorphism influences perception of computer-animated characters' actions" *Social Cognition Affect on Neuroscience* **2** 206 – 216

Fabri M, Moor D, Hobbs D, 2004 "Mediating the expression of emotion in educational collaborative virtual environments: An experimental study" *Virtual Reality* **7** 66 – 81

Fong T, Neurobakhsh I, Dautenhahn K, 2003 "A survey of socially interactive robots" *Robotics and Autonomous Systems* **42** 143 – 166

Groopman J, 2009 "Robots that care: Advances in technological therapy" *New Yorker* November 2, pp 66 – 77

Hanson D, Olney A, Pereira I A, Zielke M, 2005 "Upending the Uncanny Valley" *Proceedings of AAAI National Conference, Pittsburgh, PA* pp 24 – 31

Jastorff J, Kourtzi Z, Giese M A, 2006 "Learning to discriminate complex movements: Biological versus artificial trajectories" *Journal of Vision* **6** 791 – 804

Johansson G, 1973 "Visual perception of biological motion and a model for its analysis" *Perception & Psychophysics* **14** 195 – 204

Lee K J, Byatt G, Rhodes G, 2000 "Caricature effects, distinctiveness, and identification: Testing the face – space framework" *Psychological Science* **11** 379 – 385

MacDorman K F, Green R D, Ho C-C, Koch C, 2009 "Too real for comfort: Uncanny responses to computer generated faces" *Computers in Human Behavior* **25** 695 – 710

MacDorman K F, Ishiguro H, 2006 "The uncanny advantage of using androids in social and cognitive science research" *Interaction Studies* **7** 297 – 337

Matsui D, Minato T, MacDorman K F, Ishiguro H, 2007 "Generating natural motion in an android by mapping human motion", in *Humanoid Robots: Human-like Machines* Ed. M Hackel (Vienna: Itech) pp 351 – 366
Mori M, 1970 "Bukimi no tani" (The Uncanny valley) *Energy* **7** 33 – 35
Moseley G L, Brugger P, 2009 "Interdependence of movement and anatomy persists when amputees learn a physiologically impossible movement of their phantom limb" *Proceedings of the National Academy of Sciences of the USA* **106** 18798 – 18802
Seyama J, Nagayama R S, 2007 "The uncanny valley: Effect of realism on the impression of artificial human faces" *Presence: Teleoperators and Virtual Environments* **16** 337 – 351
Steckenfinger S A, Ghazanfar A A, 2009 "Monkey visual behavior falls into the uncanny valley" *Proceedings of the National Academy of Sciences of the USA* **106** 18362 – 18366
Thompson J C, Clarke M, Stewart T, Puce A, 2005 "Configural processing of biological movements in human superior temporal sulcus" *Journal of Neuroscience* **25** 9059 – 9066
Weibull W, 1951 "A statistical distribution function of wide applicability" *Journal of Applied Mechanics – Transactions of ASME* **18** 293 – 297

ISSN 0301-0066 (print) ISSN 1468-4233 (electronic)

PERCEPTION

VOLUME 40 2011

www.perceptionweb.com