

Interactive Realistic Digital Avatars - Revisiting the Uncanny Valley

Michael Seymour
University of Sydney
mike.seymour@sydney.edu.au

Kai Riemer
University of Sydney
kai.riemer@sydney.edu.au

Judy Kay
University of Sydney
judy.kay@sydney.edu.au

Abstract

The 40-year-old Uncanny Valley theory is influential in the discussion surrounding acceptance of realistic graphical agents. This theory was formulated by observing robots. While it has been shown to be valid when observing digital characters, little has been studied about acceptance when people interact with avatars, rather than simply observe a recording.

The emerging technology that will soon be able to create realistic avatars challenges the conventional view built on this theory, that affinity is a function of ‘appearance’, necessitating a reevaluation of the dimensions of the problem. We introduce a broader theoretical foundation with an additional dimension, namely interactivity.

Studies that have validated the Uncanny Valley used still images or clips from existing media, but did not explore interactivity. In this study we suggest that interactivity operates on an independent, orthogonal dimension to ‘appearance’, and that interaction can ‘overcome the valley’ in affinity due to matching and common human non-verbal cues. We hypothesize that these cause the user to process the avatar differently. We contribute to the literature a new way to theorize the relationship between avatar realism and affinity, including both avatar appearance and interaction, and outline a research design to study this relationship.

1. Introduction

We are on the cusp of realistic interactive human awareness and responsiveness in computers, the creation of embodied conversation agents (ECA) with realistic human faces [4]. While the realization of such agents involves various aspects such as responsiveness, appropriateness to task, accuracy in user intent, and quality of computer responses, one key aspect is the *visual presentation* of the agent itself. In this paper we study the effect of interactive digital avatars with photo-realistic human faces, based on technology and techniques that are only just becoming available and that originate from the field of feature film effects [17].

The dominant theory to capture the acceptance of photo-realistic face technology is the so-called *Uncanny Valley* [14]. This theory, now over 40 years old, predicts that acceptance increases steadily as realism in face presentation increases, before dropping off sharply and becoming negative. It then rises strongly again as realism approaches perfect reproduction (see Figure 1 below). The resulting ‘valley’ captures metaphorically the effect that faces can be almost real but feel ‘wrong’ in subtle, yet important ways. These close, but not accurate, faces make the outcome worse before it becomes better in a very non-linear fashion. Furthermore, the theory states that, movement “changes the shape of the uncanny valley graph by amplifying the peaks and valleys” [27].

While the theory has been successfully applied in the context of computer animation, gaming and digital representations of human faces [43], it is important to note firstly that the theory has its origins in robotics, not in digital technology, and that secondly most applications of the theory since have focused on pre-rendered stills and pre-animated film clips [22,25]. In general, the focus of the theory has been on the *appearance* of the digital or robotic artifact as judged by an observer.

Given recent advances in real-time rendering technology, in this study we move beyond the study of mere appearances to considering the effects of *interaction*, when the user is involved with the photo-realistic human avatar, not merely judging its appearance. Given that movement was stated in the original theory to amplify the ‘uncanny’ effect the prevailing assumption has been that interactivity makes the effect even more pronounced. However, we have initial, anecdotal evidence that suggests that the opposite happens, and that the uncanny valley might be mitigated in interactive situations.

Hence, we set out to review the current theory, by taking account of the influence of appearance and of interaction as two separate effects on the acceptance of digital avatars. We present the design and setup of an ongoing study to test the effect of the two variables by employing different digital avatars with varying degrees of photo-realism.

Regarding the uncanny valley we reason that as we introduce interactivity, we are no longer merely exploring the effect of creating better animations, but an entirely different phenomenon: the simulation of human interaction with an agent communicating via their face. We suggest that interaction becomes the main phenomenon that drives the effect, with appearance receding into the background. As a consequence we suggest that the original theory needs to be modified as it fails to capture this important dimension of interactivity.

The findings will be significant as technology is close to enabling graphical faces with the potential to bridge the Uncanny Valley. We can thus expect to see many more applications of this technology, for example the creation of graphical faces for ECAs such as Apple’s Siri, Google’s EVA, Amazon’s Alexa, or similar technology. There is still much to be done to provide an interface, but the implications for specialist areas and applications in health, education and special needs are potentially far-reaching.

We begin by discussing the significance of faces, the theory of the Uncanny Valley and how the brain processes faces specifically. We then outline recent advances in technology and discuss why interactivity affects the Uncanny valley. We present a research model, outline a corresponding research design and discuss potential implications of our research.

2. Background

2.1. The Uncanny Valley

Until recently, it was not possible to produce a realistic rendering of a human face. Many would argue that it is still almost impossible to reliably produce an animated face that escapes looking ‘eerie’. Typically rendered faces suffered from looking ‘unnervingly real’ due to rendering limitations, even when rendered using computationally expensive and time consuming computer graphics approaches. Masahiro Mori’s 1970 paper [27] termed this limitation “bukimi no tani”, commonly translated as “uncanny valley”.

As seen in Figure 1, Mori’s theory maps a function of increasing realism against an axis of affinity. As the realism of a human representation increases, the affinity increases, but only to a point. Once realism exceeds a particular point, there is a sharp drop in acceptance or affinity. In the theory, the state of most affinity is explicitly formulated to be that of a completely believable or real person, where the difference is impossible to detect. But it also postulates a quite high level of affinity with moderate realism, just before the valley.

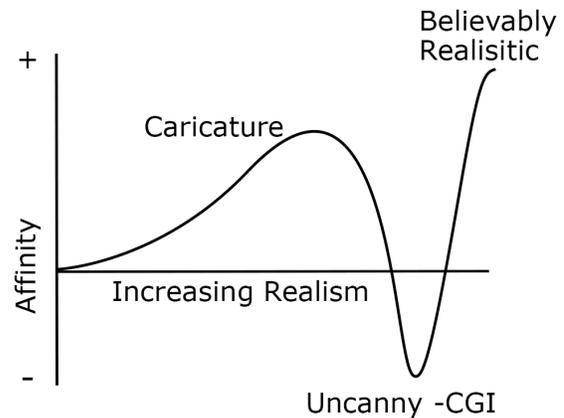


Figure 1: The Uncanny Valley, updated to reflect its application to CGI

Determining what can trigger a negative response in viewers is not a matter of a simple equation. There appears to be no single cause that makes a digital face seem uncanny [40]. The original theory speculated that this ‘eerie sensation’ when viewing a close approximation of human reality “is probably a form of instinct that protects us from proximal, rather than distal, sources of danger. Proximal sources of danger include corpses, members of different species, and other entities we can closely approach” [27].

Formulated in Japan in reference to robotics and prosthetics, the non-linear response has been shown to be useful in explaining our view of computer graphics images (CGI) of faces [31]. The hypersensitivity to facial representation (see below) has been the basis of research aiming to extend the original Uncanny Valley theory from robotics to computer graphics. Most research explores the purely aesthetic variations in combination with physiological theories, but without any emphasis on interactive engagement [23,31,41].

In this paper we seek to explore the nature of ‘affinity’ as the defining value in the original Uncanny Valley theory. However, there is no clear measure or metric for the notional ‘affinity’: it is not a dependent variable against which one can test with some independent variable [27]. Importantly, the word itself is a translation from the original Japanese word *Shinwakan* (親和感), and thus is open to interpretation.

While affinity is the currently accepted translation, other English translations have also been used, such as; *familiarity*, *rapport*, and *comfort level* to describe the theory’s reactionary vertical axis [12].

Several previous studies have judged ‘affinity’ by perceived familiarity. These studies asked participants to rate avatars on a 9-point scale in terms of how strange or familiar they judged it to be from 1 (very strange) to 9 (very familiar) [21,42].

2.2 Recent Advances in Technology

Technically, it is still challenging to create a realistic digital avatar as a user interface element. Achieving a realistic human face in a non-interactive environment is only just being addressed at the high end of feature film effects [17]. Excellent work has been done in producing high quality facial rendering and animation [37], but most cannot be rendered in real-time. Figure 2 shows an example of how closely Computer Graphic Images (CGI) avatars can resemble real faces, when rendered away from the constraints of real time interactivity.



Figure 2: CGI avatar vs. real image of source person. Image rendering credit: USC-ICT.

Recently, gaming technology has allowed advancements in real time facial rendering on stand-alone PC systems, but they still require high end graphics cards and powerful processing. Steady advances in computational power and algorithmic innovation means that there is a reasonable expectation that such work will be widely available on accessible computing platforms, including mobile devices, in the near future. This makes it timely to build a framework that supports effective design and use of avatars that are interactive and at a suitable level of realism.

2.3 On the Significance of Faces

Humans are hard wired to interpret human faces. From birth a child responds and learns from their parents' faces. These interpretations are fundamental for the successful growth and functioning of humans [36]. As such the brain has developed the ability to read faces far more specifically and with greater fidelity than any other object we see. Evolution has left us with both the ability to see a face in a few pen strokes of a cartoon or a puffy cloud (Pareidolia) but also to identify and reject those artificial faces which are only approximately close to realistic [26]. Not only can we detect these inferior renditions but we react to them far less favorably than a much simpler caricature.

A simple experiment known as the Thatcher illusion using a still image can illustrate how the brain (of both humans and some primates) processes faces in very specialized ways [1]. Figure 3 shows an inverted face, which seems slightly odd but not greatly alarming to most viewers. When the same image (see Figure 4 below) is rotated 180 degrees, it seems 'grossly off-putting'. This is true even when, as is shown here, the image is only roughly composited and little effort has been made to blend the images seamlessly and match shading or correct position. (*You may need to rotate the page to see that the below is the same image as in Figure 4*).



Figure 3: An example of the brain's separate and specialist facial processing. Compare this image to Figure 4 below.

We propose that this sharp change of response is due to different parts of the brain processing different aspects of the image of the face. It is accepted wisdom

that reading human facial expressions is so central to human communication that we have developed dedicated neurological skills for reading faces. This is also why very small changes in expression can lead to interpreting very different signals from those around us. Thus this illusion suggests that, unlike other routinely seen items, we have specialist eye, mouth and head silhouette neurological processing pathways. When one sees the head inverted in the first image, one part of the brain acknowledges it is upside down. Another part of the brain focuses on the eyes and mouth sections of the face, checking them for expression. As the eyes and mouth look normal, the face itself seems only slightly wrong. Invert the same image and the head is seen as orientated the right way but the eyes and mouth are horribly misshaped. This results in repulsion.

There are numerous other examples of functional Magnetic Resonance Imaging (fMRI) studies showing facial recognition is associated with very specific parts of the brain [26]. We suggest that such specialist face processing directly affects the Uncanny Valley response as the levels of realism changes.



Figure 4: The same image as Figure 3, rotated 180 degrees – note how ‘crude’ the composite seems when seen this way up, but there was no additional processing.

2.4 Neurological Processing of Faces

We have suggested that there are parts of the brain that specialize in processing parts of the face. There are

many other parts of the brain involved in processing facial interactions and some of these are quite relevant and virtually unexplored in terms of the Uncanny Valley theory.

Humans have dedicated neurological pathways for facial recognition but also processing and communicating back to other people. The ‘decision’ to react to a face with a smile comes to the facial nucleus from a network of structures that include the amygdala and multiple interconnected cortical and subcortical motor areas [11]. These constantly both process faces and drive a person’s own facial expressions without deliberate thought.

MacDorman et al [24] state that strong “anecdotal evidence indicates the Uncanny Valley suppresses empathy”, citing many such studies. The authors theorise that this is related to mirror neurons. They outline a series of experiments to investigate if higher realism enhances one’s ability to take the perspective of a character. They suggest this might be so, because of increased activation of mirror neurons [24].

Mirror neurons were first discovered in the ventral premotor region F5 of the macaque monkey [29]. The part of the brain that contains mirror neurons shows fMRI activity both when an action is done and observed. This is because “the stimulus feature encoded in mirror neurons is repeated irrespective of whether the action is observed or executed” [36]. Furthermore, this brain activity occurs when an individual is performing the action or when they are observing someone else performing the same action.

The MacDorman et al study made a base assumption from Lombardo et al [20] that mirror neurons enable inferring the intentions of others based on one’s own intentions in a similar situation. In their study they rejected their hypothesis that the less human or more eerie a character looks, the more it interferes with taking a character’s different perspective. In other words, the effects of the Uncanny Valley do not appear to make the viewer less empathetic. This study however did not explore interaction. The experiment used the same male individuals at different levels of realism but it was based on non-interactive stills.

The paper did highlight other research that had linked neuroimaging to higher photorealism [9]. For example, studies noted a linear relation between a robot’s human photorealism and cortical activation in the medial frontal cortex and the right temporal parietal junction [18]. There was also a significant neural response in the participants’ left premotor cortex when watching an action performed by a human being but not when watching the same action performed by a robot [39]. This suggests potential for strong affinity when observing a person but not a robot.

The inference MacDorman et al [24] make is that these findings suggest that a character's human photo-realism does not affect brain processes and interfere with the ability to see another's perspective. This effectively tested *causes* of the Uncanny Valley. We will suggest a different relevance of mirror neurons; - their role in mitigating the Uncanny Valley during interaction.

3. Theorizing Interactivity

3.1. Interactivity and the Uncanny Valley

The previous neurological studies into mirror neurons are noted to be consistent with Press [32] who found that our "action observation network," which translates observed actions into motor codes, is biologically tuned and not engaged by non-biological agents". We propose that, rather than such neurological activity causing the Uncanny Valley response when observing different faces, it is mirror neurons that lead to empathy, which explains how interaction is able to mitigate the Uncanny Valley effect. Mirror Neurons [15,16,35] have been credited with enabling empathetic mirroring responses.

We label these empathetic mirroring responses as emotional matching and it is a common human response in face-to-face communication. Tight facial emotional mirroring loops with a high fidelity/quality avatar appears to be one promising way to operationalise current theory to influence the Uncanny Valley response.

The term Affective Computing encompasses both the reading of the user's emotional state and reflecting that in the user interface. By matching emotional signals on an avatar to those of a user, we may affect affinity independently of varying realism. By simulating the emotional matching thought to be caused by mirror neurons in the brain we provide not only more believable animation, but interactively we stimulate empathy in the user. By enriching the combined experience of the interaction in this way we circumvent certain aesthetic issues that might otherwise produce a negative affinity in the user.

Matching is not just a facial property. It is also the chameleon effect that refers to unconscious mimicry of postures, mannerisms, and also facial expressions such that one's behavior passively and unintentionally changes to match that of another [6].

This technique of emotionally matching someone in behavior and tone communicates an awareness of their emotional state, acknowledging it by emotionally matching it. Bi-directional relationships between humans and computers were always at the core of the

original work defining Affective Computing [30]. Much of the research since then has been focused on creating systems that can interpret user emotions. Less work has explored emotional content mimicked back via an avatar to the user. This was due to prior limitations of rendering emotionally communicative and subtle facial representations of an ECA. As already noted, advances in computer graphics are making it possible to use subtle human facial characteristics, with the potential to significantly impact social acceptance of an ECA in high awareness interactions.

It is the duality of interaction that offers unexplored positive benefits and a different way of viewing a computer interface. We assert that there are situations where something may be considered unacceptable as an animatronic but could be quite acceptable when combined with interactivity, even though it might still be recognised as artificial (i.e. something less than imperceptibly real).

What we are striving for in creating an avatar is not to make its appearance just 'better' but to give it certain features that we observe in human interaction. These are properties that have been shown in studies of pre-recorded clips to improve an avatar's level of engagement [13].

In the coming together of the user and the avatar the situation allows an emergence of a presence, which is not just a causal relationship, but also a combined deeper and richer connection.

We propose that interactivity and responses, such as mirroring, form an important dimension that is independent of visual realism.

3.2. Research Questions and Variables

The research objective is to explore the issues and technology that would be required to allow for adoption and use of the next generation of human-computer interfaces in the context of education, with related applications in business, health and other areas.

We begin by exploring what interaction is and ask the question: if a realistic avatar is not only moving but interacting - is it still unsettling? Or does interactivity change the problem? In doing so, perception of interactivity will be a combination of factors, such as:

- Responsiveness
- Visual emotional matching
- Understanding of verbal comments
- Intelligent natural responses
- Subtle non-verbal active listening signals
- An avatar voice that is pleasing and expected

Our research is designed to address some of these points and restrict their variability in the testing.

We question if it is possible for a substandard appearance (according to the original metric) to be, at

least partially, mitigated by plausible human interaction. We achieve this by separating and testing independently the effects of appearance and of interaction that involves such aspects as mirroring.

4. Study Overview

4.1. Study Design

We present here the design and foundation work for a study to test the influence of appearance and of interaction as two separate effects on the acceptance of digital avatars. We do this by employing different digital avatars with varying photo-realistic appearance.

The method we propose to use is a ‘Wizard of Oz’ experiment. Participants will be seated with a computer, which has both an avatar displayed on a screen and a webcam. A person ‘behind the curtain’ will drive the avatar’s interaction with the participants. This approach allows for natural responses and removes the variable of having to create artificial intelligent response systems or the influence of having to use a synthetic voice for the avatar. This design is a means to study interfaces that are still currently very technically demanding or yet to be created.

The participant will then interact with three different avatars; the process will be a within-subjects design using a Latin square.

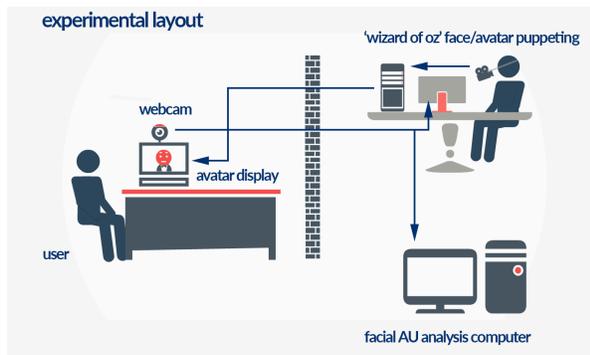


Figure 5. The experimental layout illustrating the ‘Wizard of Oz’ researcher behind a wall controlling the avatar.

The method for exploring these research questions needs to account for many factors. As noted, studies in the past have sampled different isolated examples of human representation to explore if indeed the Uncanny Valley could be mapped. The efficacy of such studies was often limited by not showing examples based on the same face. We will thus use variation utilizing the same actual face at different levels of realism (Figure 6). Since interaction needs to be examined in terms of realism, we will also deploy a non-realistic cartoon

version with similar interaction but without a realistic appearance. We have prototypes of this working for different levels of realism, to validate the approach, and the preliminary testing indicates this will be achievable. The technical details of the face tracking and rendering are outside this scope of this paper.



Figure 6. The same face at different levels of realism (all digital).

The ‘Wizard of Oz’ style interaction will allow the researcher driving the facial interface to be hidden from view, but able to see and hear the participant via their webcam. Figure 7 shows the head mounted display that registers the researcher’s reactions and drives the on-screen avatar. The participants will be interacting unaware, until after completion, that the avatars were driven in real time as digital puppets.



Figure 7. An example of a face input digital puppet system worn by the researcher.

4.2. Research Method and Sampling

The analysis of the study will be qualitative, with open questions, as this is early work in exploring the research questions. As such, we seek to qualify the participant’s affinity with each avatar using:

- an open standard questioning of subjects and
- emotional connection to the avatar.

We will record and analyze the participant’s webcam feed. The analysis will use a system built on

the Facial Action Coding System (FACS) [7] to explore the emotions of the participant during testing.

The research will start with a small sample trial to validate the data collection process and analysis before proceeding. The expected sample size of the primary qualitative research group is expected to be either 18-21, (a multiple of 3, given our 3 types of faces).

Our aim is to reduce the number of variables and have the participants drawn from a population with a similar cultural background as the avatar in order to avoid issues of cross-racial variance. While this aids with reducing factors, it is noted that this places a limitation on the richness of the experience. Introducing a wider range of avatars presents an opportunity for further research.

4.3. Experimental Procedure

The procedural setup of the experiment is to have participants explore a subject while we modify the experience through interactivity and separately realism, measuring for affinity.

There are nine test cases (3x3) as a combination of *Still – Recorded video – Interactive*
Crossed by type:
Cartoon – Avatar – ‘Real’ Person. (as in Figure 6)

The task for the participants is to plan a holiday in London, with the aid of the computer assistant. The task is chosen to avoid strong recall of previous emotions (Which may be hard to separate from emotions and responses due to the avatar), and yet still offer a task that someone can imagine themselves doing which is neither unpleasant nor routine.

During the interaction the video and audio feed of the participant will be used to inform the researcher driving the avatar and be recorded for analysis using a system such as Affectiva software. This software does a FACS based analysis of users. It uses emotional analytics using computer vision and deep learning to provide emotional data points. The FACS system will also be used for the avatar digital puppet system.

For each experiment there will be an open-ended question using a Attrakdiff UX eSurvey style model¹ for evaluation data on how much affinity the participants feel towards the avatar in terms of usability and appearance.

To gauge if there is basic emotional attachment to the avatar, part of the experiment will involve exploring a baseline empathy produced without a facial avatar representation, building on the work of Clifford Nass and others [19,28,33].

For the experiment it will be explained to the participants that they will be evaluating a new avatar system with varying degrees of fidelity. We will observe the impact of their choice and their expressions, to gauge emotional attachment to the different avatars.

4.4. Analysis

The original Uncanny Valley maps affinity vs. realism. We aim to produce a 3D mapping of interaction vs. realism vs. affinity (See Figure 8). For each of the 9 test cases listed above, they will be mapped to a positive or negative degree of affinity.

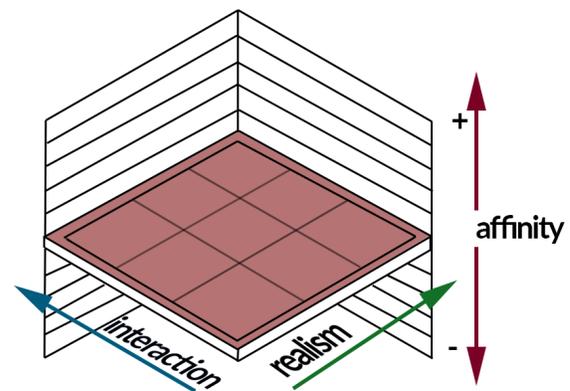


Figure 8. Aim to map (3x3) each of the interactive vs. realism tests with positive or negative affinity.

The baseline of a still image is provided to address basic likeability and appearance. The digital puppet version involves advanced digital real time face tracking and rendering. Interactive digital puppet work will involve emotional matching, something a pre-recorded video cannot do.

5. Discussion

We propose that the simple two-dimensional relationship between affinity and realism in the Uncanny Valley theory does not encapsulate the complexity of a situation where interactivity is introduced. This is both a reflection of the age of the original theory, coupled with advances in technology that offer previously unseen applications.

We propose the metaphor of “flooding” the Uncanny Valley - as it implies that we are no longer merely crossing the Uncanny Valley, but exploring a phenomenon of a different kind, one that in fact is not encompassed by the original theory. The original theory implies that realism is the primary determinate

¹ <http://www.uid.com/en/services/user-research/ux-evaluation.html>

of affinity. We argue that this fails to capture any additional effects of interacting with the avatar and how this can change the relationship between realism and affinity. One could argue that interaction is just an extension of realism but this would imply that they are part of the same axis. We propose that there is an additional dimension that is not merely realism. This new dimension invokes different cognitive processes than just aesthetics; it affects sense making and may well involve different neurological processing.

5.1. Implications

Sproull et al noted in 1996, that “because the human face is such a powerful signal of human identity, adding human faces to interfaces holds promise for interface designers to make interfaces more human-like” [38].

If we are able to more reliably produce avatars that achieve affinity in users then it may be possible to use them effectively in a range of applications such as aged care, elderly companions or assistants. Educational applications include coaching (trial subjects), teaching applications both as mentor/teacher and as study companion [5]. Other applications include numerous medial applications [3] such as better computer interfaces for the disabled, change agents and motivational multiple selves [10].

A promising area of application for avatar research has been education. A higher degree of affinity in this area could translate to improved understanding and measurably higher levels of engagement. This educational application would be enhanced by non-verbal facial communication via an interface. It would allow users to more readily adopt the computer as an extended part of their problem solving or education options, if they could interact beyond simple text and desktop metaphors; if users could encompass the computer as something to be communicated with more naturally, using natural expressions and facial signaling.

5.2. Limitations and future research

Findings of our study will initially be restricted, given the actual subject of the visuals is just one individual, of one particular sex, age and appearance. While there are many general lessons to be learnt from restricting the avatar to one person with a limited set of variables, in reducing the number of variables, it also invites further research with a wider range of faces.

A wider selection of avatars, especially comparing male and female avatars and of different ages would be an obvious area of further research. It is also important to consider comparative ages, for example someone of

parental age with a child avatar may produce parental influences, similarly there is reason to believe that children may respond very differently to peer avatars than parentally aged avatars.

In studies such as tailored educational delivery it has been noted that the persistently lower test scores of students of color as compared to their Euro-American peers have been attributed in part to dialectal differences between students [8]. Hence, research into pairing of appearance and language in avatars to their users may reveal further insights.

The use of such technology is not limited to pedagogical situations where the agent is teacher, there is a strong argument to be made that such agents would work as learning companion systems, as outlined initially by Chan [5].

There is reason to believe that results may not vary over time with extended exposure, yet only a longitudinal study design will reveal if any such effect exists.

As part of the research implementation we seek to also build on the body of research on trust, and in future iterations we aim to explore *Avatar-mediated trust* and *social presence* [2][34].

Finally, further work will be needed to establish in the context of the Uncanny Valley the nature of the engagement around mirror neurons and their proposed role in empathy and affinity.

6. Conclusion

This study aims to contribute to the body of knowledge on human computer interfaces by showing how interaction with digital photo-realistic avatars affects the efficacy of the Uncanny Valley as a theory.

This study intends to show that rather than being a simple two dimensional problem, the way the brain reacts to tight accurate facial interaction adds a new dimension to the users sense of affinity and thus engagement. Applied in a variety of contexts, this could allow a much richer application of avatars because the technology enables a more believable and accurate representation of human faces.

The Uncanny Valley has provided valuable insights since it was first proposed in the field of robotics over forty years ago. The results of this work may reveal a better understanding of affinity with regard to the Uncanny Valley and high fidelity interactive avatars. It raises the possibility of advances in supportive emotional engagement and better human computer interfaces with wide-ranging applications in fields such as education and learning, health care, or advisory, such as in tourism or public services.

7. References

- [1] Adachi, I., Chou, D.P., and Hampton, R.R. Thatcher Effect in Monkeys Demonstrates Conservation of Face Perception across Primates. *Current Biology* 19, 15 (2009), 1270–1273.
- [2] Bente, G., Rüggenberg, S., Krämer, N.C., and Eschenburg, F. Avatar-mediated networking: Increasing social presence and interpersonal trust in net-based collaborations. *Human Communication Research* 34, 2 (2008), 287–318.
- [3] Bickmore, T.W. and Paasche-Orlow, M.K. The Role of Information Technology in Health Literacy Research. *Journal of Health Communication* 17, sup3 (2012), 23–29.
- [4] Cassell, J. Embodied Conversational Interface agent. *Communications of the ACM* 43, 4 (2000), 70–78.
- [5] Chan, T.-W. Learning Companion Systems. 1989.
- [6] Chartrand, T. and Bargh, J. The chameleon effect. *Journal of Personality and Social Psychology* 76, 1999, 893–910.
- [7] Ekman, Paul, and Erika L. Rosenberg, E. *What the face reveals: Basic and applied studies of spontaneous expression using the Facial Action Coding System*. Oxford University Press, 1997.
- [8] Finkelstein, S., Yarzebinski, E., Vaughn, C., Ogan, A., and Cassell, J. The effects of culturally congruent educational technologies on student achievement. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 7926 LNAI, (2013), 493–502.
- [9] Gazzola, V., Rizzolatti, G., Wicker, B., and Keysers, C. The anthropomorphic brain: The mirror neuron system responds to human and robotic actions. *NeuroImage* 35, 4 (2007), 1674–1684.
- [10] Ginsberg, M. The society of mind. Marvin Minsky. *Artificial Intelligence* 48, 1991, 335–339.
- [11] Gothard, K.M. The amygdalo-motor pathways and the control of facial expressions. *Frontiers in Neuroscience* 8, 8 MAR (2014), 1–7.
- [12] Ho, C.-C. and MacDorman, K.F. Revisiting the uncanny valley theory: Developing and validating an alternative to the Godspeed indices. *Computers in Human Behavior* 26, 6 (2010), 1508–1518.
- [13] Kang, S., Feng, A.W., Leuski, A., Casas, D., and Shapiro, A. Smart Mobile Virtual Humans : “ Chat with Me !” im, 1–4.
- [14] Kawaguchi, J. Robocon founder Dr . Masahiro Mori. *Japan Times*, 2011, 1–5.
- [15] Keysers, C. and Gazzola, V. *Mirror Neuron Systems*. 2009.
- [16] Kilner, J.M. and Lemon, R.N. What we know currently about mirror neurons. *Current Biology* 23, 23 (2013), R1057–R1062.
- [17] Klehm, O., Rousselle, F., Papas, M., et al. Recent Advances in Facial Appearance Capture. *Computer Graphics Forum* 34, 2 (2015), 709–733.
- [18] Krach, S., Hegel, F., Wrede, B., Sagerer, G., Binkofski, F., and Kircher, T. Can machines think? Interaction and perspective taking with robots investigated via fMRI. *PLoS ONE* 3, 7 (2008).
- [19] Lee, K.M. and Nass, C. Designing social presence of social actors in human computer interaction. *Proceedings of the conference on Human factors in computing systems - CHI '03*, 5 (2003), 289.
- [20] Lombardo, M. V., Chakrabarti, B., Bullmore, E.T., et al. Shared neural circuits for mentalizing about the self and others. *Journal of cognitive neuroscience* 22, 7 (2010), 1623–35.
- [21] MacDorman, K. Subjective ratings of robot video clips for human likeness, familiarity, and eeriness: An exploration of the uncanny valley. *ICCS/CogSci-2006 long symposium: Toward ...*, January 2006 (2006), 26–29.
- [22] Macdorman, K.F. and Entezari, S.O. Individual differences predict sensitivity to the uncanny valley. *Interaction Studies*, (2015), 1–47.
- [23] MacDorman, K.F., Green, R.D., Ho, C.-C., and Koch, C.T. Too real for comfort? Uncanny responses to computer generated faces. *Computers in Human Behavior* 25, 3 (2009), 695–710.
- [24] MacDorman, K.F., Srinivas, P., and Patel, H. The uncanny valley does not interfere with level 1 visual perspective taking. *Computers in Human Behavior* 29, 4 (2013), 1671–1685.
- [25] McDonnell, R. and Breidt, M. Face Reality : Investigating the Uncanny Valley for virtual faces. *ACM SIGGRAPH ASIA 2010 Sketches*, (2010), 41.
- [26] Meng, M., Cherian, T., Singal, G., and Sinha, P. Lateralization of face processing in the human brain. *Proceedings of the Royal Society B: Biological Sciences* 279, 1735 (2012), 2052–2061.
- [27] Mori, M., Kageki, N., and MacDorman, K.F. The Uncanny Valley. *IEEE Robotics & Automation Magazine*, June (2012), 98–100.
- [28] Nass, C. and Yen, C. *The Man Who Lied to His Laptop: What We Can Learn About Ourselves from Our Machines*. Penguin Group US, 2010.
- [29] di Pellegrino, G., Fadiga, L., Fogassi, L., Gallese, V., and Rizzolatti, G. Understanding motor events: a neurophysiological study. *Experimental Brain Research* 91, 1 (1992), 176–180.
- [30] Picard, R.W. Affective computing: challenges. *International Journal of Human-Computer Studies* 59, 1–2 (2003), 55–64.
- [31] Pollick, F.E. In Search of the Uncanny Valley. *Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering* 40, (2010), 69–78.
- [32] Press, C. Action observation and robotic agents: Learning and anthropomorphism. *Neuroscience and Biobehavioral Reviews* 35, 6 (2011), 1410–1418.
- [33] Reeves, B. and Nass, C. *The media equation: how people treat computers, television, and new media*. Stanford, Calif.: Center for the Study of Language and Information; Cambridge, 1996.
- [34] Riedl, R., Mohr, P.N.C., Kenning, P.H., Davis, F.D., and Heekeren, H.R. Trusting Humans and Avatars: A Brain Imaging Study Based on Evolution Theory. *Journal of Management Information Systems* 30, 4 (2014), 83–114.

- [35] Rizzolatti, G. and Craighero, L. The Mirror-Neuron System. *Annu.Rev. Neurosci.* 27, (2004), 169–192.
- [36] Sagar, M., Bullivant, D., Efimov, O., et al. Embodying models of expressive behaviour and learning with a biomimetic virtual infant. *International Conference on Development and Learning and on Epigenetic Robotics*, (2014), 62–67.
- [37] Seymour, M. The art of digital faces at ICT – from Digital Emily to Digital Ira. *fxguide.com*, (2013), 1–56.
- [38] Sproull, L., Subramai, M., Kiesler, S., Walker, J.H., and Waters, K. When the interface is a face. *Human-computer interaction* 11, 2 (1996), 97–124.
- [39] Tai, Y.F., Scherfler, C., Brooks, D.J., Sawamoto, N., and Castiello, U. The Human Premotor Cortex Is ‘Mirror’ only for Biological Actions. *Current Biology* 14, 2 (2004), 117–120.
- [40] Tinwell, A. Uncanny as usability obstacle. *Online Communities and Social Computing*, (2009).
- [41] Tinwell, A., Abdel, D., and Charlton, J.P. Computers in Human Behavior Perception of psychopathy and the Uncanny Valley in virtual characters. *Computers in Human Behavior* 29, 4 (2013), 1617–1625.
- [42] Tinwell, A., Grimshaw, M., Nabi, D.A., and Williams, A. Facial expression of emotion and perception of the Uncanny Valley in virtual characters. *Computers in Human Behavior* 27, (2011), 741–749.
- [43] Tinwell, A., Grimshaw, M., Nabi, D.A., and Williams, A. Facial expression of emotion and perception of the Uncanny Valley in virtual characters. *Computers in Human Behavior* 27, 2 (2011), 741–749.