

Ascending from the valley: Can state-of-the-art photorealism avoid the uncanny?

Darragh Higgins
higgind3@tcd.ie
Trinity College Dublin
Dublin, Ireland

Donal Egan
doegan@tcd.ie
Trinity College Dublin
Dublin, Ireland

Rebecca Fribourg
rebecca.fribourg@tcd.ie
Trinity College Dublin
Dublin, Ireland

Benjamin R. Cowan
benjamin.cowan@ucd.ie
University College Dublin
Dublin, Ireland

Rachel McDonnell
ramcdonn@tcd.ie
Trinity College Dublin
Dublin, Ireland



Figure 1: (top) High fidelity characters 1–6 used in experiment, (bottom) Low fidelity characters. Males shown on the left and female on the right.

ABSTRACT

Advancements in real-time rendering technology have continued to develop rapidly over the course of the last decade. Consequently, human likenesses have been represented virtually with increasingly impressive detail. There is evidence that this increased resemblance to real humans has an observable and wide-ranging set of effects on human perception, cognition and action in situations that involve digital characters. Studies that seek to advance the science of synthetic animated people have consistently aimed to measure and quantify changes in perceived emotional content mediated through artificial human likenesses. The present study has been built off of this work. The experiment outlined here was constructed to define and measure responses from human participants towards

state-of-the-art photorealistic virtual humans under affected conditions. In particular, we sought evidence for changes in perceptions of human likeness, eeriness and attractiveness that could be observably dependent on conditions of photorealism and character representation.

CCS CONCEPTS

• **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability.

KEYWORDS

Virtual Humans, Uncanny Valley, Animation, Perception



This work is licensed under a Creative Commons Attribution International 4.0 License.

SAP'21, September 16–17, 2021, Rennes, France
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8663-0/21/09.
<https://doi.org/10.1145/3474451.3476242>

ACM Reference Format:

Darragh Higgins, Donal Egan, Rebecca Fribourg, Benjamin R. Cowan, and Rachel McDonnell. 2021. Ascending from the valley: Can state-of-the-art photorealism avoid the uncanny?. In *SAP '21: ACM Symposium on Applied Perception, September 16–17, 2021, Rennes, France*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3474451.3476242>

1 INTRODUCTION

The theoretical notion of the uncanny valley [Mori 1970; Mori et al. 2012] has pervaded scientific discussions of humanoid robots and virtual humans since the end of the 20th century. The term describes a state of unease induced by a realistic artificial representation of a human that occurs at a level of realistic portrayal where those representations differ only subtly from real humans.

This proposed valley begins at the point where people typically begin to report negative feelings towards a virtual or robotic human. The valley is considered to end at the point where realism and positive impressions resume their correlative increase. This proposed correlation implies that there is an upper boundary of realism where both likeability and realism are at an optimized level, although Mori originally supported the idea of designing to stop short of the valley rather than traversing it through maximal realism [Draude 2011].

While sometimes considered to be a theory lacking in substantiation [Blow et al. 2006], there has been growing evidence for the Uncanny Valley effect fostered by studies which have implicated the realism of characters as a governing influence on impressions of virtual characters [Seyama and Nagayama 2007].

Moreover, investigations have highlighted the roles of character appearance and behaviour [Bartneck et al. 2007], individual subjective traits [MacDorman and Entezari 2015] and facial expressions [Tinwell et al. 2011] in shaping the experience of virtual humans, in efforts to quantify a location of the uncanny valley on the scale of possible human representation.

The virtual humans used in the present study represent some of the highest fidelity ‘photorealistic’ characters currently in circulation. While these characters can be displayed and animated in real-time, rendering them at their highest-fidelity requires high levels of computation and expensive GPUs. We aim to determine if these additional computational costs are justified and if such characters are perceived in a positive manner, compared to lower fidelity counterparts. Our findings suggest that current state-of-the-art virtual humans are no longer stuck in the uncanny valley and are rated more attractive, humanlike, and less eerie than lower quality representations.

2 RELATED WORK

With the proliferation of high quality animation technology, however, the correlative desynchronization of likeability and appearance realism has proven to be extremely relevant for researchers seeking to contribute to animated characters in films, cartoons and games [McDonnell 2012], where audience reactions bare importance for the design of virtual humans.

For instance, stylized cartoon renderings of virtual characters have been rated more favourably than more realistic but sickly looking counterparts [Zibrek and McDonnell 2014]. Realistic virtual humans that exhibit low behavioural realism appear to provoke similar effects, compared to high-realism conditions [Guadagno et al. 2007]. There is also evidence that more realistic animations of humans may evoke less trust than cartoon versions [McDonnell and Breidt 2010], and are more easily categorised by humans than semi-realistic characters [Kätsyri et al. 2017].

Many such studies involve full-body representations of virtual humans, although weaker uncanny valley effects have also been measured with screen-mediated images of computer generated

faces [Kätsyri et al. 2019], which provided our research with the motivation to test for this effect using expressive animations of human-like faces.

It has been posited that any uncanny valley study must include emotional evaluation metrics as well as scales for familiarity and human-likeness [Lay et al. 2016]. Both appearance and motion variability among virtual characters have been used to demonstrate observable effects on emotional valence in immersive VR contexts [Mousas et al. 2018].

Higher levels of negative valence and arousal have been observed in people experiencing the uncanny valley effect, where avatar faces can be shown to increase negative valence compared to human faces [Cheetham et al. 2015]. However, there is ambiguity in these empirical claims, with other studies [Mäkäräinen et al. 2015] contending that uncanny characters can be linked to experiences of amusement ahead of negative valence emotions.

This ambiguity was some motivation for our focus on valence intensity ratings in this study, which sought clarity on this important issue of classifying the emotional experience of the uncanny valley. The absence of emotional expressiveness has been measured to heighten perceived eeriness of digital characters in face-only stimuli presentations [Tinwell et al. 2011] such as ours, and so the inclusion of positive and negative valence was also intended to help define a scale of eeriness at emotional dimensions beyond neutrality.

Numerous studies have sought to model human behaviour to methodically describe the effects denoted by the uncanny valley and their perceptual basis [Schindler et al. 2017]. Some proponents of the theory point to evolutionary preferences for aesthetics [Hanson et al. 2005] as a culprit in the generation of negative character impressions. An evolutionary perspective can be supported with evidence that the uncanny valley has been observed in the behaviour of monkeys presented with stimuli representing real and near-real synthetic characters [Steckenfinger and Ghazanfar 2009]. Considering uncanny valley effects from an evolutionary perspective can be useful for developing guiding principles in the domain, but does not provide a framework for measuring perceptual effects of animation realism, as we attempt to do with the present study.

Another approach proposes inconsistent levels of realism between features of a virtual human as the factor which may explain why negative impressions begin to arise at high accuracy representation [MacDorman and Chattopadhyay 2016; Mitchell et al. 2011], defined by perceptual mismatches. There is strong support for the occurrence of uncanny valley under perceptual mismatch conditions [Kätsyri et al. 2015]. However, this approach was not considered appropriate for framing effects of valence on studying the uncanny valley.

A more applicable framework can be derived from studies which consider concept categorisation as the prime culprit for perceived uncanniness [Tondou and Bardou 2011]. The effects denoted by the uncanny valley have been proposed to arise from category errors [Feldman et al. 2009] where difficulties discerning the difference between a human and a replica version may lead to a state of unease. This notion of category error or category uncertainty has been corroborated by studies which have compared category sorting tasks to negative impressions of animated faces [Wang et al. 2015].

Given the salient difference in quality between respective levels of detail (LODs) with our stimuli, we considered category uncertainty to be the approach which best enabled us to frame our predictions. We hypothesised that the higher-fidelity character would be rated more favourably on scales of 'eeriness' than the lower-fidelity, which could potentially be explained by a difference in object classification on behalf of participants. It was expected that this fidelity bias would be reflected in all of our questionnaire data.

3 EXPERIMENT DESIGN

The method outlined here was developed to examine perceptions of highly realistic virtual humans. The experimental format was constructed with 2 (Appearance fidelity) $\times$ 2 (Emotional Valence) conditions, with 3 character variations and 2 character gender conditions for a within subject design. Our main hypothesis was based on appearance fidelity, as we considered it likely that the maximal realism of the highest level of detail to be sufficient for traversing the uncanny valley in regards to virtual humans. We expected levels of uncanny ratings to differ between valence conditions. The experimental paradigm developed here was geared towards exploring potential evidence for and against these hypotheses.

3.1 Stimuli

The stimuli were created in Unreal Engine 4.26 (UE4.26)¹. The avatars used in the stimuli are MetaHumans created using the Unreal Engine online application MetaHuman Creator². MetaHumans are currently the state-of-the-art for realistic digital humans. Each MetaHuman is available in 8 levels of detail (LODs), with LOD0 having the highest level of fidelity and LOD7 having the lowest. The differences in fidelity between LODs arise from differences in geometry complexity (for example, 24000 facial vertices at LOD0 versus 1300 at LOD4), texture complexity, and the use of strand-based hair (LOD0) versus card-based hair (LOD4). Each MetaHuman has a full facial and body rig for animation, with the complexity of these rigs also varying across the different LODs. For example, the LOD0 facial rig consists of 713 joints and 669 blendshapes, whereas the LOD4 facial rig has only 84 joints and no blendshapes. An identical rigging setup across MetaHumans means that a given animation can be used on multiple MetaHumans.

For each gender, we used three MetaHumans chosen on three levels of character material variation, representing three different ethnicities. This decision was motivated by previous studies investigating the uncanny valley [Macdorman et al. 2008] which have ascribed low ratings of eeriness to photorealistic virtual faces but have described the limitation of using only a single style of possible human variation. The character material and gender variation were also included to account for any potential biases in the participant sample. We used these six MetaHumans with LOD0 for our high fidelity condition (Figure 1 (top)) and LOD4 for our low fidelity condition. For our purposes, LOD 4 was considered an appropriate choice for low fidelity, representing the standard of previously gold standard renderings of human likenesses, in terms of both polygon count and appearance fidelity. LOD4 (Figure 1 (bottom)) is also considerably less humanlike than LOD0.

3.1.1 Valence. For each emotional valence, two trained actors' (one female, one male) facial performances were captured via a desk-mounted video camera placed in front of their faces. The valence expressions were derived from literature outlining particular phrases that may be used to represent positive or negative valence emotions [Ben-David et al. 2011]. Each recorded performance was then tracked using the Faceware Studio facial performance tracking software³ and streamed into UE4.26 using the Live Client Plugin for UE4.26⁴. The facial performance was then mapped to the MetaHuman facial animation rig and the resulting animation recorded using UE4.26's Sequence Recorder feature. The commonality of the facial rig across MetaHumans meant that for each emotional valence a single recorded animation could be used across all female characters and similarly for the male characters.

The different components of the stimuli (MetaHuman, animation, lighting, virtual camera, and audio) were pieced together using UE4.26's Sequencer feature. The same virtual camera setup and standard three-point lighting setup was used for all stimuli. For each gender and emotional valence, the same audio track and animation was used, meaning the only difference between stimuli was the actual MetaHuman itself. Several post-processing effects such as ray-tracing and anti-aliasing were used on the final stimuli which were rendered using the UE4.26 Movie Render Pipeline. Each stimuli was rendered at 30 FPS and at a resolution of 2048×1536 and downsampled to 1024×768 . A Nvidia GeForce RTX™ 3090 GPU was used for the rendering.

3.2 Method

The measurements we employed for quantification of the uncanny valley were derived from a revised alternative [Ho and MacDorman 2010] to the Godspeed indices [Bartneck et al. 2007] which can be described as central to the development of experimental techniques for measuring the uncanny valley. The three indices defined in this literature that we used for our measurements were *Human-Likeness*, *Eeriness* and *Attractiveness*. The valence scales were derived from work similar to ours [Mäkäräinen et al. 2014] which aimed to quantify perceived emotional intensity in virtual characters.

Prior to stimuli exposure, participants filled in sheets on informed consent and demographic information, which gathered information on age, gender and video game experience. For the stimuli exposure stage, participants were presented with clips of animated virtual humans expressing positive or negative valence conditions. For each clip, participants were asked to rate the intensity of emotion expressed by the virtual human, on slider scales measuring responses from 1-7 with 1 being 'Not at all happy/sad' and 7 being 'Very happy/sad'. These sliders were used to gauge perceived emotion but also employed as a means to check whether unsupervised participants were actually watching the clips as the experiment progressed. The same scales were used to measure differences with our uncanny indices, with the scales reading *Natural-Artificial*, *Reassuring-Eerie* and *Unattractive-Attractive*. These scales have been subject to scientific validation [Ho and MacDorman 2010] but were also chosen for their unambiguous wording compared to other means of providing uncanny scales to experimental subjects.

¹<https://www.unrealengine.com>

²<https://metahuman.unrealengine.com>

³<http://support.facewaretech.com/studio>

⁴<https://glassboxtech.com/products/live-client>

Since there were three types of character material and two character genders included with the differences in LOD and valence, participants experienced 24 different clips, presented in randomized order. After experiencing all 24 clips, participants were asked for feedback on their decision making processes via a text input box.

3.3 Participants

20 participants (10 male, 10 female) were recruited from a combination of university mailing lists and Prolific participant recruitment⁵. These recruitment methods have been validated as appropriate and useful for perceptual research [Jonell et al. 2020]. In consideration of the complexity of ethical standards around paid participant research [Ripley 2006] we paid the national minimum wage to Prolific recruits, adjusted to the duration of the experiment. The participants were provided with a link to complete the experiment in their web browsers, using Qualtrics survey software⁶. The duration of time taken to complete the experiment was considered as a means of rejection; the maximum accepted timespan was 48 minutes. This was to ensure insofar as possible that participants engaged in the experiment consistently throughout.

4 RESULTS

We initially conducted a 3-way Analysis of Variance to assess the influence of within group factors *Fidelity* (High/Low), *Character* (1–6) and *Clip* (Positive/Negative) on reported levels of uncanniness and emotion intensity. However, there was no evidence in the data to suggest Character had any statistically significant effect on uncanniness and emotional intensity and was thus removed from the analysis to improve interpretability. A Shapiro-Wilk test for normality was performed on the data. The normality assumption was violated ($p > .05$), so we applied an Aligned Rank Transformation to the data as a means to enable our analysis of variance. For the scales of emotion intensity, we averaged results over character and ran 2 way analysis of variance on ratings of emotion intensity for factors of *Fidelity* and *Clip* as there we were not predicting individual differences from characters on emotion intensity. Tukey's post-hoc tests ($\alpha = .05$) were conducted to check significance for pairwise comparisons. In addition, post-hoc tests were corrected using Bonferroni correction.

- *Human-Likeness*: We observed a main effect of Fidelity on human-likeness ($F_{1,437} = 117.96, p < .001$), indicating that participants considered high fidelity characters to look more human ($M = 4.12, SD = 1.82$) than the lower fidelity characters ($M = 2.87, SD = 1.58$) ($p < .0001$).
- *Eeriness*: A main effect of Fidelity was observed on questionnaire responses for Eeriness ($F_{1,437} = 30.97, p < .001$), with low fidelity characters rated as more eerie ($M = 4.05, SD = 1.66$) across the board than high fidelity characters ($M = 3.39, SD = 1.65$) ($p < .0001$).
- *Attractiveness*: A main effect of Fidelity on Attractiveness was established ($F_{1,457} = 120.6, p < .001$). Participants rated Attractiveness to be higher in high fidelity characters ($M = 4.74, SD = 1.41$) than low ($M = 3.6, SD = 1.41, p < .0001$), ($p < .0001$). A main effect of Character on Attractiveness

was found ($F_{1,437} = 7.98, p < .001$). As expected, some characters were rated more attractive than others, with a tendency of character 1 and 2 (Figure 1) to be less attractive than others and character 4 to be more attractive ($p < .0001$).

- *Happiness Intensity*: A main effect of Fidelity was found ($F_{1,19} = 14.45, p = .001$) on happiness intensity ratings, where ratings were higher for high fidelity characters than low ($p = 0.0012$). A main effect of Clip was found ($F_{1,19} = 137.96, p < .001$), where happiness intensity was rated higher in the positive than the negative clip, as expected ($p < .0001$). An interaction between Clip and Fidelity was found ($F_{1,19} = 5.81, p = .03$), where post-hoc tests highlighted that the difference in ratings between low and high fidelity of characters was present in the positive clip ($p < .0001$), but not in the negative clip.
- *Sadness Intensity*: A main effect of Fidelity was found ($F_{1,19} = 6.93, p = .016$) on sadness intensity ratings, where ratings were lower for low than for high fidelity characters ($p = 0.02$). A main effect was also found for Clip ($F_{1,19} = 214.38, p < .001$), where sadness ratings were higher in the negative compared to the positive clip.

5 DISCUSSION

Our results indicate that state-of-the-art photorealism should be considered a positive choice for virtual human applications as it was found to be more attractive, less eerie, more human-like, and happier than lower fidelity depictions. Our results are in line with Zibrek et al. [2019] where photorealism was found to be a positive choice in virtual reality, however other studies had conflicting results on the topic [Cheetham et al. 2015; Mäkäräinen et al. 2015]. Additionally, previous work has shown that characters with skin textures that are less detailed and show less blemishes and other skin irregularities are rated more attractive [Zell et al. 2015], which is consistent with research in the cosmetics industry which shows that blurring skin texture can increase attractiveness [Fink and Matts 2008]. Our low fidelity characters had much smoother skin textures with less detail than our high fidelity ones, however they were not rated more attractive, which we believe was due to the additional reduction in polygon detail counteracted any increase in attractiveness that the skin texture might have afforded.

In our study, our *low level of fidelity* condition was created by lowering the facial rig complexity, and the geometric and texture complexity to mimic characters currently seen across many applications, while retaining the same identity, motion, voice, lighting and rendering conditions as the high fidelity version.

While eeriness ratings were low for our high fidelity characters, it is possible that ratings might reach even lower levels with higher quality facial animation. We used real-time facial tracking and solving with only minimal clean-up in this study, and previous work has shown that lower quality facial animation can appear more eerie on realistic characters [McDonnell 2012]. Future work should investigate the effect of facial tracking and solving.

Interestingly, we found that higher fidelity characters were rated to be happier when expressing a positive emotion. Previous work has shown that brighter lighting conditions can increase happiness ratings on virtual characters [Wisessing et al. 2020], however, our

⁵<https://www.prolific.co/>

⁶<https://www.qualtrics.com/uk/>

study is the first evidence that higher geometric and texture detail can also cause this effect, since our lighting was kept constant.

Our results also indicate that there was no effect of character appearance on uncanny ratings, which implies that future studies could perhaps use less variety in character appearance when studying these effects. We also found no effect of emotion on eeriness ratings, implying that future research might consider testing more neutral expressions to determine it would heighten uncanny ratings [Tinwell et al. 2011].

ACKNOWLEDGMENTS

This work was conducted with the financial support of the Science Foundation Ireland Centre for Research Training in Digitally-Enhanced Reality (d-real) under Grant No. 18/CRT/6224 and under the ADAPT Centre for Digital Content Technology (Grant No. 13/RC/2106_P2) and RADICAL (Grant No. 19/FFP/6409).

REFERENCES

- Christoph Bartneck, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. 2007. Is the uncanny valley an uncanny cliff? *Proceedings - IEEE International Workshop on Robot and Human Interactive Communication*, 368–373. <https://doi.org/10.1109/ROMAN.2007.4415111>
- Boaz M. Ben-David, Pascal H.H.M. Van Lieshout, and Talia Leszcz. 2011. A resource of validated affective and neutral sentences to assess identification of emotion in spoken language after a brain injury. *Brain Injury* 25 (2011), 206–220. Issue 2. <https://doi.org/10.3109/02699052.2010.536197>
- Mike Blow, Kerstin Dautenhahn, Andrew Appleby, Chrystopher L. Nehaniv, and David C. Lee. 2006. Perception of Robot Smiles and Dimensions for Human-Robot Interaction Design.
- Marcus Cheetham, Lingdan Wu, Paul Pauli, and Lutz Jancke. 2015. Arousal, valence, and the uncanny valley: psychophysiological and self-report findings. *Frontiers in Psychology* 6 (7 2015). <https://doi.org/10.3389/fpsyg.2015.00981>
- Claude Draude. 2011. Intermediaries: Reflections on virtual humans, gender, and the Uncanny Valley. *AI and Society* 26 (11 2011), 319–327. Issue 4. <https://doi.org/10.1007/s00146-010-0312-4>
- Naomi H. Feldman, Thomas L. Griffiths, and James L. Morgan. 2009. The Influence of Categories on Perception: Explaining the Perceptual Magnet Effect as Optimal Statistical Inference. *Psychological Review* 116 (10 2009), 752–782. Issue 4. <https://doi.org/10.1037/a0017196>
- Bernhard Fink and Paul J. Matts. 2008. The effects of skin colour distribution and topography cues on the perception of female facial age and health. *Journal of the European Academy of Dermatology and Venereology* 22, 4 (2008), 493–498.
- Rosanna E. Guadagno, Jim Blascovich, Jeremy N. Bailenson, and Cade McCall. 2007. Virtual humans and persuasion: The effects of agency and behavioral realism. *Media Psychology* 10 (2007), 1–22. Issue 1. <https://doi.org/10.1080/15213260701300865>
- David Hanson, Andrew Olney, Ismar A. Pereira, and Marge Zielke. 2005. Upending the Uncanny Valley.
- Chin Chang Ho and Karl F. MacDorman. 2010. Revisiting the uncanny valley theory: Developing and validating an alternative to the Godspeed indices. *Computers in Human Behavior* 26 (11 2010), 1508–1518. Issue 6. <https://doi.org/10.1016/j.chb.2010.05.015>
- Patrik Jonell, Taras Kucherenko, Ilaria Torre, and Jonas Beskow. 2020. Can we trust online crowdworkers?: Comparing online and offline participants in a preference test of virtual agents. *Proceedings of the 20th ACM International Conference on Intelligent Virtual Agents, IVA 2020*. <https://doi.org/10.1145/3383652.3423860>
- Jari Kätysyri, Beatrice de Gelder, and Tapio Takala. 2019. Virtual Faces Evoke Only a Weak Uncanny Valley Effect: An Empirical Investigation With Controlled Virtual Face Images. *Perception* 48 (10 2019), 968–991. Issue 10. <https://doi.org/10.1177/0301006619869134>
- Jari Kätysyri, Klaus Förger, Meeri Mäkräinen, and Tapio Takala. 2015. A review of empirical evidence on different uncanny valley hypotheses: Support for perceptual mismatch as one road to the valley of eeriness. *Frontiers in Psychology* 6 (2015). Issue MAR. <https://doi.org/10.3389/fpsyg.2015.00390>
- Jari Kätysyri, Meeri Mäkräinen, and Tapio Takala. 2017. Testing the ‘uncanny valley’ hypothesis in semirealistic computer-animated film characters: An empirical evaluation of natural film stimuli. *International Journal of Human Computer Studies* 97 (1 2017), 149–161. <https://doi.org/10.1016/j.ijhcs.2016.09.010>
- Stephanie Lay, Nicola Brace, Graham Pike, and Frank Pollick. 2016. Circling around the uncanny valley: Design principles for research into the relation between human likeness and eeriness. *i-Perception* 7 (11 2016). Issue 6. <https://doi.org/10.1177/2041669516681309>
- Karl F. MacDorman and Debaleena Chattopadhyay. 2016. Reducing consistency in human realism increases the uncanny valley effect; increasing category uncertainty does not. *Cognition* 146 (1 2016), 190–205. <https://doi.org/10.1016/j.cognition.2015.09.019>
- Karl F. MacDorman and Steven O. Entezari. 2015. Individual differences predict sensitivity to the uncanny valley. *Interaction Studies. Social Behaviour and Communication in Biological and Artificial Systems/Interaction Studies / Social Behaviour and Communication in Biological and Artificial Systems/Interaction Studies* 16 (11 2015), 141–172. Issue 2. <https://doi.org/10.1075/is.16.2.01mac>
- Karl F. MacDorman, Robert D. Green, Chin-Chang Ho, and Clinton T. Koch. 2008. Too real for comfort? Uncanny responses to computer generated faces. <http://www.macdorman.com>
- Rachel McDonnell. 2012. LNCS 7660 - Appealing Virtual Humans. <http://gv2.cs.tcd.ie/mcdonner/>
- Rachel McDonnell and Martin Breidt. 2010. Face Reality: Investigating the Uncanny Valley for Virtual Faces. In *ACM SIGGRAPH ASIA 2010 Sketches* (Seoul, Republic of Korea) (SA '10). Association for Computing Machinery, New York, NY, USA, Article 41, 2 pages. <https://doi.org/10.1145/1899950.1899991>
- Wade J. Mitchell, Kevin A. Szerszen, Amy Shiron Lu, Paul W. Schermerhorn, Matthias Scheutz, and Karl F. MacDorman. 2011. A mismatch in the human realism of face and voice produces an uncanny valley. , 10–12 pages. Issue 1. <https://doi.org/10.1068/i0415>
- M. Mori. 1970. The Uncanny Valley. *Energy* 7, 4 (1970), 33–35.
- Masahiro Mori, Karl F. MacDorman, and Norri Kageki. 2012. The uncanny valley. *IEEE Robotics and Automation Magazine* 19 (2012), 98–100. Issue 2. <https://doi.org/10.1109/MRA.2012.2192811>
- Christos Mousas, Dimitris Anastasiou, and Ourania Spantidi. 2018. The effects of appearance and motion of virtual characters on emotional reactivity. *Computers in Human Behavior* 86 (9 2018), 99–108. <https://doi.org/10.1016/j.chb.2018.04.036>
- Meeri Mäkräinen, Jari Kätysyri, Klaus Förger, and Tapio Takala. 2015. The funnycanny valley: A study of positive emotional reactions to strangeness. *ACADEMICMINDTREK 2015 - Proceedings of the 19th International Academic Mindtrek Conference*, 175–181. <https://doi.org/10.1145/2818187.2818292>
- Meeri Mäkräinen, Jari Kätysyri, and Tapio Takala. 2014. Exaggerating Facial Expressions: A Way to Intensify Emotion or a Way to the Uncanny Valley? *Cognitive Computation* 6 (12 2014), 708–721. Issue 4. <https://doi.org/10.1007/s12559-014-9273-0>
- Elizabeth B. D. Ripley. 2006. A Review of Paying Research Participants: It's Time to Move beyond the Ethical Debate. *Journal of Empirical Research on Human Research Ethics* 1, 4 (2006), 9–19. <https://doi.org/10.1525/jer.2006.1.4.9> arXiv:<https://doi.org/10.1525/jer.2006.1.4.9> PMID: 19385834.
- Sebastian Schindler, Eduard Zell, Mario Botsch, and Johanna Kissler. 2017. Differential effects of face-realism and emotion on event-related brain potentials and their implications for the uncanny valley theory. *Scientific Reports* 7 (3 2017). <https://doi.org/10.1038/srep45003>
- Ichiro Seyama and Ruth S. Nagayama. 2007. The Uncanny Valley: Effect of Realism on the Impression of Artificial Human Faces. , 337–351 pages. Issue 4.
- Shawn A. Steckenfinger and Asif A. Ghazanfar. 2009. Monkey visual behavior falls into the uncanny valley. www.pnas.org/cgi/doi/10.1073/pnas.0910063106
- Angela Tinwell, Mark Grimshaw, Debbie Abdel Nabi, and Andrew Williams. 2011. Facial expression of emotion and perception of the Uncanny Valley in virtual characters. *Computers in Human Behavior* 27 (3 2011), 741–749. Issue 2. <https://doi.org/10.1016/j.chb.2010.10.018>
- Bertrand Tondou and Nicole Bardou. 2011. A new interpretation of Mori's uncanny valley for future humanoid robots. *International Journal of Robotics and Automation* 26 (2011), 337–348. Issue 3. <https://doi.org/10.2316/Journal.206.2011.3.206-3348>
- Shensheng Wang, Scott O. Lilienfeld, and Philippe Rochat. 2015. The uncanny valley: Existence and explanations. *Review of General Psychology* 19 (12 2015), 393–407. Issue 4. <https://doi.org/10.1037/gpr0000056>
- Pisut Wisessing, Katja Zibrek, Douglas W. Cunningham, John Dingliana, and Rachel McDonnell. 2020. Enlighten Me: Importance of Brightness and Shadow for Character Emotion and Appeal. *ACM Trans. Graph.* 39, 3, Article 19 (April 2020), 12 pages. <https://doi.org/10.1145/3383195>
- Eduard Zell, Carlos Aliaga, Adrian Jarabo, Katja Zibrek, Diego Gutierrez, Rachel McDonnell, and Mario Botsch. 2015. To Stylize or Not to Stylize?: The Effect of Shape and Material Stylization on the Perception of Computer-generated Faces. *ACM Transactions on Graphics* 34, 6, Article 184 (2015), 184:1–184:12 pages.
- Katja Zibrek, Sean Martin, and Rachel McDonnell. 2019. Is Photorealism Important for Perception of Expressive Virtual Humans in Virtual Reality? *ACM Trans. Appl. Percept.* 16, 3, Article 14 (Sept. 2019), 19 pages. <https://doi.org/10.1145/3349609>
- Katja Zibrek and Rachel McDonnell. 2014. Does render style affect perception of personality in virtual humans? *Proceedings of the ACM Symposium on Applied Perception, SAP 2014*, 111–115. <https://doi.org/10.1145/2628257.2628270>