From Robots to Racism –

Attempting To Replicate The Uncanny Valley Effect With Biological Faces And Investigating Individual Differences

Jana Westermann, s2085887

University of Twente, July 2021

Bachelor Thesis

First Supervisor: Dr. M. Schmettow

Second Supervisor: Dr. Simone Borsci

Introduction

Even if you do not know its name, you are sure to have come across an example of the Uncanny Valley (UV) effect at some point in your life: a character in an animated movie or a video game, or maybe even just a store mannequin that looked a little too human-like, eventually you will encounter something that looks to be just on the wrong side of artificially human and the unsettling feeling that comes with it. This is the Uncanny Valley effect, which was first described by Masahiro Mori in 1970 in a japanese essay entitled “Bukimi No Tani”. This translates more accurately to “Valley of Eeriness”, but the phrase was first translated as “Uncanny Valley” in 1978, and the translation stuck around (Hsu, n.d.)

In his essay, Mori hypothesized that there was a point on the scale of artificial to human where an artificial face was almost, but not perfectly, human-like, and that this would provoke a reaction of uneasiness and revulsion in any human who saw it. This sudden drop in affinity was what Mori described as the Bukimi No Tani, or Uncanny Valley (Mori, 1970/2012).

The UV effect has since been proven to exist in practice in a number of studies – Palomäki et al. (2018) found that it was replicable with four out of the six stimuli sets they tested, suggesting that the effect relied on photorealistic, and pre-evaluated images, and was less reliable when using morphed images. Löffler et al. (2020), meanwhile, showed that the effect was not limited only to humanoid robots, but could also be found using zoomorphic ones: people preferred robots that were either high or low in animal likeness over ones mixing realistic and unrealistic features. In 2016, Mathur and Reichling also showed that the effect was replicable on a large set of various stimuli including both real life robot faces and a controlled set of edited faces, and that the effect occurred regardless of presentation time. Replications of the Mathur and Reichling study with slight modifications further proved that the UV effect would always occur (Slijkhuis, 2017; Koopman, 2019).

Explaining the Uncanny Valley Effect

However, no matter how well the UV effect can be proven and replicated across studies, there is still the question of *why* it occurs in the first place. As MacDorman et al. (2009) put it, the various theories put forth to explain the UV can be generally divided into two groups: theories involving automatic, fast processing, and those that involve slower, more general cognitive processes that occur later.

Fast Processing Theories

One of the theories in the area of fast processing would be that of threat avoidance: even in his original essay (Mori, 1970/2012), Mori hypothesised that the feeling of eeriness we experience when viewing faces in the uncanny valley is an instinct of self-preservation. Human-like, but not entirely human faces would suggest corpses, or sickness, and therefore we would want to avoid them, causing a reaction of disgust. In their review of this Pathogen Avoidance hypothesis, Wang et al. (2015) note that, while the hypothesis has never been tested directly, there have been studies showing that high levels in disgust sensitivity predict high levels of uncanny valley sensitivity.

Other automatic processing theories, as listed by Wang et al. (2015), would be the Mortality Salience or Terror Management theory, which suggests that we have a negative reaction to uncanny stimuli because they remind us of our own eventual death and trigger a fear or being replaced by android doppelgangers, and the Evolutionary Aesthetics theory, which suggests that we react negatively to faces that do not fit our idea of aesthetically pleasing faces, as we are primed to want to procreate with people who have good genes and will ensure healthy, aesthetic offspring. There was an attempt to test the latter theory by Hanson (2005), who compared faces to ones modified to be more aesthetically pleasing. Hanson found that the modified faces provoked a lower UV effect. However, as Wang et al. point out, there were no clear conclusions to be drawn as the modified faces could have a reduced UV effect because of the modification itself and not the increased aesthetic value.

Slow Processing Theories

When it comes to theories related to slower cognitive processing, Wang et al. (2005) again list a number of possible explanations: the Violation of Expectation hypothesis, the Category Uncertainty hypothesis, and the Mind Perception hypothesis.

The Violation of Expectation hypothesis is also one that was proposed in Mori's original essay (1970/2012): he suggested the eerie feeling we get from certain artificial stimuli occurs because we think they are organic at first glance, but then realise that they are in fact not upon further exposure. This mismatch between our expectations and the perceived reality causes us to have a negative reaction towards the stimulus as it gives us conflicting perceptual cues, causing cognitive stress.

The Category Uncertainty theory is similarly based on the way uncanny stimuli cause cognitive stress: here, the proposed reason for the eerie feeling we get from certain stimuli is due to the fact that we are unsure what category to place them in – dead or alive, human or non-human, and so on. This was partially tested, such as by Palomäki et al. (2018), whose study showed that using morphed faces was not as successful in producing the UV effect as pre-selected stimuli, showing that this is an unlikely explanation, and by Yamada et al. (2012), who found that indeed stimuli that were more difficult to neatly categorise were seen as less likeable and trustworthy by participants.

Finally, there is the theory of Mind Perception, which suggests that the reason we experience eeriness at uncanny faces is that they look human, but lack the capacity to experience pain or emotions, something which we expect of humans. This disconnect between looking human-like, but violating our criteria by which we categorise something as human is what causes the eerie feeling. This was tested and shown to be a plausible explanation by Gray and Wegner (2012), who found that anything that we perceive to lack the capacity of experience while seemingly human or human-like provokes a feeling of uncanniness.

What the theories in either category have in common is the fact that they have either not been tested at all or been tested only in studies that were ultimately inconclusive. Another, more promising, commonality of these theories is that the proposed reasons for the UV effect are not tied to the fact that

the stimuli in question are artificial. Concepts such as pathogen avoidance, selecting for evolutionary aesthetics, or expecting something to be alive and being disturbed when it is not, are all things that, if true, would be more evolutionary safeguards designed to appropriately react to biological stimuli than anything related to androids. This suggests that it should be possible to find an Uncanny Valley with exclusively biological stimuli if any of these theories have any weight, which they seem to have, even though none of them has been conclusively proven.

Beyond the Artificial Uncanny Valley

The theory that the UV is an evolutionary phenomenon originally unrelated to androids or other technological advances has been put forth before, such as by Moosa and Ud-Dean (2009), but to our knowledge never been actually tested, which is what this study aims to do: To test the hypothesis that it is possible to recreate the Uncanny Valley phenomenon with a collection of exclusively biological stimuli. If it is possible to prove this, it would have a wide range of fascinating implications. For instance, it would mean that modern day social structures which lead to the marginalisation of everyone who does not fit the widely accepted norms, are a natural consequence instead of something we are taught either directly or by observing the people around us.

Therefore**,** if we suppose that there is an observable UV effect on biological faces due to evolutionary instincts to protect us from a number of possible dangers, it opens up a whole new field of investigation regarding the societal and political consequences. If it is possible to show a connection between the Uncanny Valley and contemporary discrimination of minorities and anyone who is “other”, it would open up an entirely new way of looking at these issues, and how to best solve them.

Morton (2019) puts forth a similar idea in his essay on the inherent racism and ableism in the idea of an Uncanny Valley: If we assume that on the one end of the scale are healthy human beings and on the other hand there are “cute”, unthreatening non humans, this “opens up a forbidden zone of uncanny beings that reside scandalously in the Excluded Middle region” – including human beings who have

disabilities, or who have been subjected to the kind of racist ideologies that likens them to animals more than their fellow humans.

A similar idea is brought up by Woods and Hare (2019) who bring up the concept of simianisation – the enduring racist belief that black people are more like great apes than humans, dating back to the beginnings of the transatlantic slave trade. Woods and Hare further explore this by looking at how, if actual Great Apes fall into the Uncanny Valley, being close to human, but not quite, this ties back directly to the way we as a society treat black people, and how this makes Great Apes and the UV a tool that can be used to further the dehumanization and devaluation of Black people in western society.

Williams (2019) touches on the ableist side of the Uncanny Valley in his essay, in which he explores the way in which science-fiction has historically used cyborgs as a form of talking about disabilities and disabled people. He goes on to cite Robertson (2019), who in her book talked about the fact that the Uncanny Valley is not an objective fact, but rather "a highly subjective thing, based on factors such as 'physical and cognitive abilities, age, sex, gender, sexuality, ethnicity, education, religion, and cultural background;'". He goes on to mention how, in everyday life, not everyone who is disabled or sick immediately becomes something eerie and uncanny to the people around them, but rather that it is possible for humans to adjust and even embrace the unfamiliar.

In her essay, Rhee (2013), similarly connects the Uncanny Valley to science fiction, in this case Phillip K. Dick's *Do Androids Dream of Electric Sheep?*. In her critique of Mori's theory, she argues that the Uncanny Valley should not be seen as a naturally ingrained fact of humanity, but rather as a normativising tool that mirrors the way society splits its members into healthy and unhealthy, perfectly human and other. She goes further to say that if we view the Uncanny Valley as something natural and ingrained, "we risk uncritically and erroneously privileging certain definitions of humanness as 'fixed and inherent' holding up conceptions of the human that hierarchize and exclude on the basis of this presumption".

These essays may make compelling arguments about the Uncanny Valley theory as something that is impacted by and impacts society and societal power structures, however they lack empirical proof

of their points – which is hard to find elsewhere as well. The closest to a study on the societal aspect of the UV effect was done by MacDorman and Entezari in 2015: they investigated how individual differences in people affect their sensitivity to the Uncanny Valley effect. Their findings showed that there were in fact a number of personality traits that directly or indirectly increased sensitivity to the effect, among them Neuroticism and Religious Fundamentalism. However, their study was still done with artificial stimuli, and as such fails to connect the Uncanny Valley effect directly to human societal structures – proving that religious fundamentalists are more easily spooked by humanoid robots does not prove that they are more likely to be hostile towards minorities. In order to close this gap, it is necessary to investigate individual differences as they apply to the Uncanny Valley effect in relation to biological faces.

Research Focus

As the main focus on this study remains on proving an Uncanny Valley effect exists in regards to biological stimuli in the first place, an in-depth investigation of individual differences is beyond the scope of this project. Therefore, investigating how different personality traits relate to the UV effect will be limited to a number of short personality scales. This can be seen as a partial replication of the MacDorman and Entezari study, as it will be using the same measure for Uncanny Valley sensitivity, but using different personality traits. The scales used are for Authoritarianism, which is in line with what MacDorman and Entezari found to be correlated, a Need For Closure scale to explore a trait more related to fast processing, and a Big Five inventory, as this provides a decent overview of a number of universal personality traits. With this exploratory approach this study will attempt to discover if the MacDorman and Enterzari study's findings can be replicated with biological stimuli and a number of other personality traits.

Methods

Procedure

Participants of the online questionnaire were first faced with a welcome page, explaining that the aim of the survey was to assess emotional response to a number of stimuli. Afterwards, they were asked to read and agree to an informed consent form.

The online survey consisted of two main parts: stimuli rating and personality scales. Before the stimuli rating section began, they were presented with a selection of the stimuli in order to help anchor participants in their ratings. For the stimuli rating section, participants were shown each of the 100 stimuli for 2 seconds, before automatically being redirected to the next page, on which they were asked to rate the stimulus on two randomly selected scales out of the total six. After each set of 25 stimuli ratings, participants were directed to take a short break before proceeding to the next set.

Following all four blocks of 25 stimuli, participants were directed to a page telling them they had finished with the rating part of the survey and were now asked to fill out a number of short questionnaires, followed by the Need For Closure Short Form Scale, Very Short Authoritarianism Scale and Short Big Five Inventory.

Finally, participants were debriefed by informing them about the true purpose of the study, investigating the Uncanny Valley, and given the option to leave comments on the survey itself.

Design

The main hypothesis of the study was investigated using a within-subjects design which had each participant rate all 100 stimuli based on their eeriness. The dependent variable was eeriness, and the independent variable human likeness.

For the exploratory second part of the study, a correlation questionnaire survey design was used in order to collect data on the strength of a variety of personality traits. This data was then used as

independent variables, with strength of the Uncanny Valley effect being the dependent variable, in a linear regression model.

Measures and Materials

Stimuli

Adapting the criteria used by Marthur & Reichling (2016), stimuli of various biological faces were gathered. The adapted criteria required the stimuli to be pictures showing the individual's entire face, from the front, in a reasonably high quality. The pictures should be of real species and not computer generated, as well as not depicting famous persons. Through this, a total of 111 pictures were collected, including pictures of modern day humans of various races found on license-free stock photo sites, a variety of monkey and great ape faces, as well as busts depicting human ancestors suchs as Neanderthals. All pictures were cropped to be 450x450 pixels, and cut out as to only display the face itself in front of a white background, not showing other body parts.

Next, four researchers individually rated each stimulus on its human likeness (from 1-100), emotional valence (from -100 to 100) and perceived emotional expression (categories were neutral, happiness, sadness, anger, surprise, fear and disgust). Based on these ratings, the Intraclass Correlation Coefficient could be calculated in order to see how much the raters agreed on the various factors. The stimuli with markedly low levels of interrater agreement were excluded from the study. Furthermore, a number of Marthur and Reichling's robot faces were selected to be included in the stimuli in order to provide a comparison group. For this, the 10 stimuli with the lowest level of general error across the range of emotional reaction as found in Koopman (2019)'s study were chosen. The final set of stimuli consisted of exactly 100 images.

Survey

The 100 stimuli were uploaded into a Qualtrics online survey. The survey was designed to show participants all 100 images in randomised groups of 25 each, providing short breaks between each set of

25. Each image was set to be shown for a total of two seconds before automatically proceeding to the next page, where two scales were displayed. These were randomly selected out of a total of six, which measured eeriness and likability. In order to ensure the most accurate results, these scales were available in Dutch, German, and English, using the translations created by Koopman (2019) for the eeriness scale, or created for this study for the likability scale. The rest of the survey was presented in English. Also included were a number of short form personality scales, namely the Very Short Authoritarianism Scale (Bizumic & Duckitt, 2018), the Big Five Short Form Scale (Schupp & Gerlitz, 2014), and the Short Form Need For Closure Scale (Roets & Van Hiel, 2011). Short form scales were purposely chosen to account for the fact that participants had already completed the lengthy first half of the survey.

Measures

The Eeriness scale used was created by Ho and MacDorman in 2011, who found that it had “excellent psychometric properties", with a high internal reliability, having an overall Crohnbach’s Alpha of .74 for Eeriness, and applicability to stimuli outside of humanlike robots. It consists of 8 items, which are each rated on a scale of 0 - 100.

To measure likability, a continuous visual analog scale (VAS) ranging from -100 to 100 was used, with the two extremes being “this face seems less friendly, more unpleasant, creepy” and “this face seems more friendly and pleasant, less creepy”, as used by Marthur and Reichling (2016). The scale was translated into German and Dutch by native speakers for this study.

The Very Short Authoritarianism Scale (VSA) was developed and tested by Bizumic and Duckitt (2018). It consists of 6 items which are answered on a 9-item Likert scale, ranging from “very strongly agree” to ”very strongly disagree”. The scale showed satisfactory internal validity and reliability.

The Short Form Need For Closure Scale (NFC-SF), created by Roets and Van Hiel (2011), consists of 15 items taken from the original long form version (Kruglanski et al., 2013). It showed comparable, and thus sufficient, validity and reliability to the long form scale.

Finally, the short form of the Big Five Inventory (BFI-S) was created by Gerlitz and Schupp (2005) and consists of a total of 15 items taken from the original long form version. It had sufficient internal reliability considering its low number of items, and showed validity for all five internal dimensions measured by the inventory. The English versions of the items were taken from Lang et al.'s (2011) paper on the applicability of the scale in a wide range of settings.

Participants

The online survey was completed by 84 participants. Out of these, 22 surveys had to be excluded due to missing or invalid answers or due to the fact that they failed the attention check question. Therefore the final number of participants is 62. Participants were gathered via the university's test subject pool SONA, as well as the researchers' friends and family. In the name of privacy, the survey did not collect any demographic information of the participants, as this did not seem relevant to measuring the Uncanny Valley effect.

Statistical analysis

For the statistical analysis of the data, a polynomial regression between human likeness and eeriness of the stimuli was performed in order to produce a curve with one shoulder and a trough for each participant. The second part of the study used linear regression in order to test for correlation between personality traits and Uncanny Valley sensitivity.

Results

The Uncanny Valley Effect in Biological Faces

In order to investigate whether an Uncanny Valley effect can occur using a set of biological stimuli, four polynomial models were estimated, following the methods described in Schmettow's book on statistics (2021).

The final model being used is a multi-level polynomial model using a Bayesian estimate with MCMC sampling and a beta-distribution. Here, the multi-level polynomial model is used to account for the fact that each participant has their own ratings, as averaging the results would not show whether all participants experience the UV effect. This is similar to Koopman's (2019) analysis on robotic faces. Using a Beta regression for the model allows us to minimise the effect of the participants' different response styles by adjusting the variance in order to account for anchoring (Schmettow, 2021).

Figure 1 shows the cubic polynomial curves adjusted for a beta distribution with random effects standard deviations. As we can see, the beta distribution has decreased the variance in responses to the point where the individual curves are mostly parallel, with the individual troughs overlapping. The only shift between curves is vertical, due to the curves having different centers.

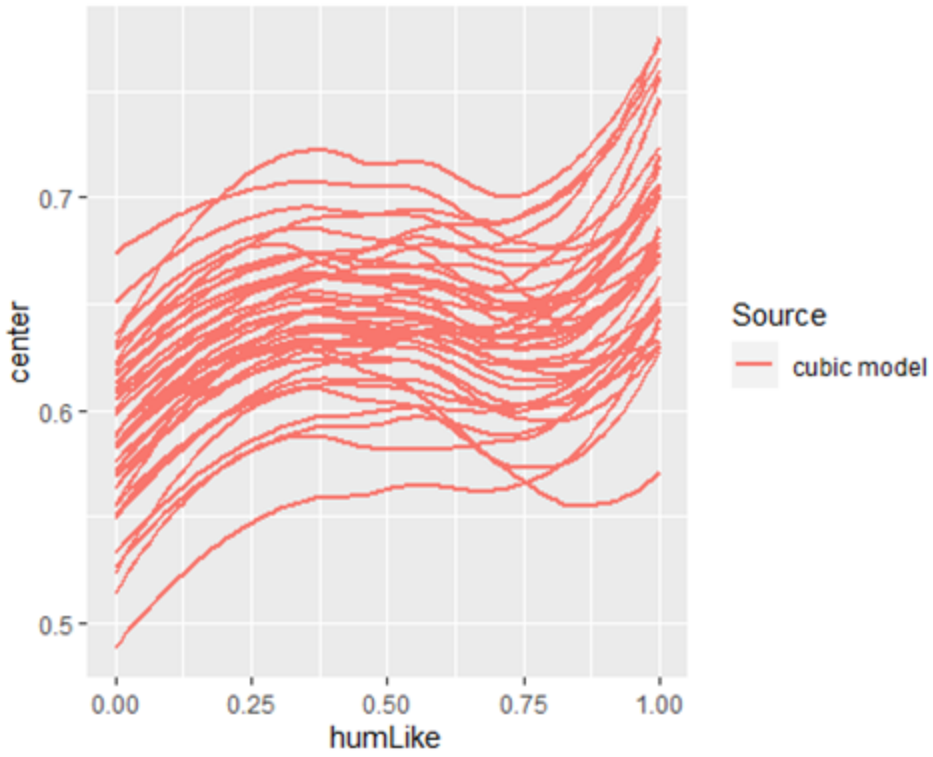

Figure 1. Spaghetti plot of a cubic polynomial model using a beta distribution to adjust for anchoring showing a number of participants' curves for human likeness ratings.

Finally, the universality of the Uncanny Valley effect across the sample had to be investigated. This was done according to the method described in Section 6.4 of Schmettow (2021). Figure 4 shows the individual positions of shoulders and troughs of participants' ratings. Looking at it, it quickly becomes clear that they are rather similar. This is further supported by figure 5, which shows the probability of an uncanny valley effect across participants, which is close to or reaches 100% for almost all participants, with the few outliers still achieving at least 79.5% probability.

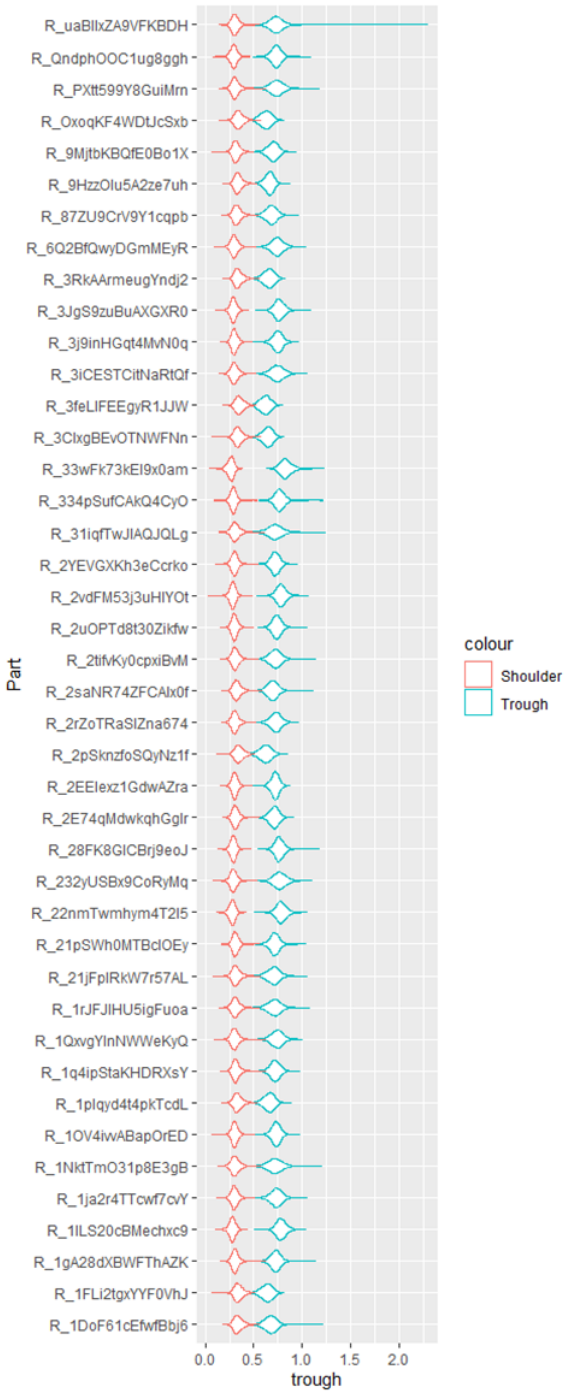

Figure 4. Shoulders (in red) and throughs (in blue) of the individual participants' rating curves showing the similarity across participants

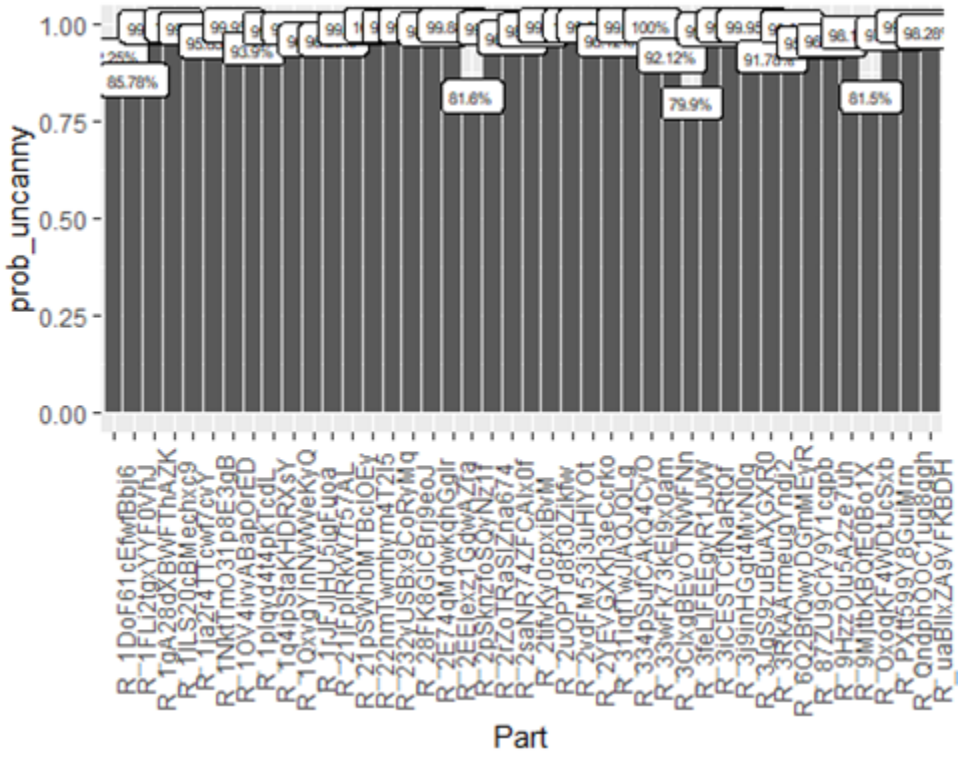

Figure 5. Probability of an Uncanny Valley effect for a selection of participants with participant ID on the X-axis and probability of UV effect on the Y-axis.

As the data shows, there is in fact an observable uncanny valley effect using biological faces. Considering the fact that the participants' curves are largely similar, it seems unlikely that there would be any individual differences due to personality traits. However, it is still worth investigating how and why MacDorman and Entezari found significant correlations in their 2015 paper. Therefore, we will proceed with the second half of this research paper, namely whether certain personality traits can be used to predict sensitivity to the Uncanny Valley effect.

Investigating personality traits in relation to the Uncanny Valley

For this, the results of the various personality scales had to be calculated, as well as sensitivity to the Uncanny Valley effect. In the name of partially reproducing their work, this used each participant's average eeriness rating across all biological stimuli, which is how sensitivity was computed in MacDorman and Entezari's 2015 study. Table 1 shows the descriptive statistics of the resulting variable – a higher score suggesting the participant found the stimuli generally more eerie and therefore indicating higher sensitivity to the Uncanny Valley effect. It also shows the descriptive statistics of the results of each of the personality scales. For all three scales, mean scores across all items were calculated. For the VSA, scores could range from -4 to +4, with higher scores indicating higher levels of authoritarianism. The BFI-SF works similarly, with possible scores for each of the five traits ranging from 1 to 5, with higher scores again indicating high level of the respective trait. Lastly, the possible scores for the NFC-SF ranged from 1 to 6, with 6 indicating the highest level of need for closure.

Table 1

Descriptive Statistics of the Variable Sensitivity to the Uncanny Valley effect, as well as the scores for the Very Short Authoritarianism Scale (VSA), The Big Five Inventory Short Form (BFI-SF), and the Need For Closure Scale Short Form (NFC-SF). (n=62)

	Min	Max	SD	Mean
Sensitivity	38.75	65.77	5.55	52.93
VSA	-4	.83	1.01	-2.2
BFI-SF				
Neuroticism	1	5	.99	3.06
Extraversion	2	5	.86	3.7
Openness	1.67	5	.75	4.1
Agreeableness	2	5	.76	3.9
Conscientiousness	1.67	5	.72	3.62
NFC-SF	2.2	5.4	.72	3.8

After this, the Pearson Correlation for each score in relation to the Sensitivity score was computed, as well as a linear regression analysis performed. Table 2 shows the correlations, as well as their 95% Confidence Intervals. The data shows slight negative correlations between UV sensitivity and Authoritarianism, Need for Closure, as well as Agreeableness and Neuroticism out of the Big Five. There is little to no correlation between Sensitivity and Openness, and slight positive correlations to Conscientiousness and Extraversion. All of the correlations fall within their respective 95% Confidence Intervals, suggesting that they are plausible results. However, all of these 95% CIs also include 0, which means it is also entirely possible there is no correlation at all.

Table 2

Pearson Correlations between UV Sensitivity and the various personality traits, as well as their 95% Confidence Intervals.

Trait	Pearson Correlation to UV Sensitivity	95% Confidence Interval	
		Lower	Upper
Authoritarianism	-.195	-.422	.056
Need For Closure	-.145	-.379	.107
Openness	.001	-.247	.248
Neuroticism	-.014	-.261	.235
Conscientiousness	.177	-.074	.407
Extraversion	.122	-.130	.359
Agreeableness	-.189	-.417	.061

The sensitivity scores, as well as the scores for the various personality scales were successfully visually checked for normality. One notable occurrence was that the scores for Authoritarianism skewed slightly to the left, indicating especially low levels of the trait among the group of participants.

It would at this point be possible to attempt to rectify the mistakes of the MacDorman and Entezari study by investigating correlations between these personality traits and a different measure for sensitivity, namely subtracting the position of each participant's first shoulder from the position of their trough. With this measure (Table 3), the sensitivity scores would no longer be impacted by the individual rating styles. However, if we look at the resulting variable, the range of the values is extremely close, with a standard deviation of .07, further supporting the finding that the Uncanny Valley effect presents in different people with little to no individual differences. Therefore, further investigating the relationship between this newly calculated variable and the various personality traits would very likely not result in any notable findings.

Table 3

Range, Mean and Standard Deviations for UV Sensitivity calculated by subtracting Shoulder from Trough of each participant, in order to ignore variance due to rating styles.

Min	Max	Mean	SD
.27	.57	.41	.07

Discussion

The first of the research questions in this study, which aims to answer whether it is possible to show an Uncanny Valley effect with biological stimuli as opposed to artificial ones, can be answered with a clear yes. This, in combination with the fact that the through and shoulder of the participants' curves are largely identical, only differing in individual rating sensitivity, suggests that there is a clear evolutionary aspect to the UV effect – it seems to be a universal emotional reaction that is not exclusive to robotic or otherwise artificial stimuli. This consequently opens up a large variety of future research questions: it might be valuable to investigate the UV with more specific subsets of biological faces such as exclusively human stimuli, measuring for reactions to different ethnicities, genetic disabilities, scarring or burns, plastic surgery and countless others. Furthermore, this lends additional credibility to the evolutionary explanations for the UV effect described in the introduction of this paper, which opens up further areas to investigate – which of the evolutionary explanations is the most likely one? Finally, the fact that there was a dataset of 100 biological stimuli created for this study, and that it provably works to show the Uncanny Valley effect, means that there are also possible future research options that could make use of the stimuli set, such as replicating other studies on the Uncanny Valley that were done with artificial stimuli in order to see if there are differences in the ways people react to the UV in artificial vs. biological faces.

When it comes to answering the second research question, whether certain personality traits can predict Uncanny Valley sensitivity, we can already see this is untrue based on the first half of the data analysis – the overall shape of the curve is very similar for all participants and numerical differences in the ratings are most likely due to response style bias. This is in direct disagreement with the 2015 study by MacDorman and Entezari, mentioned earlier, who found a correlation using robot faces – the implication being that their results were erroneous, and in fact due to response style bias, which describes how different participants choose different ratings to express the same emotional responses. One participant might rate a certain level of eeriness as 60 and another as 75, while both experience the exact same reaction. This would mean that the results of the previous study were not in fact proof that any personality traits predict sensitivity to the Uncanny Valley – their value for sensitivity was just an artifact

in the data due to differences in rating style of each participant. This phenomenon is also known as anchoring (Schmettow, 2021), and explains how each participant chooses a different range for their ratings and sticks to it going forward. This variance in response styles is wholly irrelevant to explaining the Uncanny Valley effect or any sensitivity to it. Assuming that there are different levels of sensitivity to the UV effect based on the different anchoring points of the participants is a wholly spurious correlation, and any relationships between this "sensitivity" value and personality traits is merely a coincidence. Therefore it can be concluded that the UV is actually more or less the same for every person, with no variance between different personality types other than the exact number a person chooses to rate a certain stimuli. This is further supported by the recalculated sensitivity value excluding variance due to anchoring, which shows that the actual value barely differs between participants. This is a noteworthy finding, as it shows that there is no individual variance in experiencing the Uncanny Valley effect, but it is in fact a universal, deeply ingrained instinctual response owing to evolutionary development. Therefore, no further research is needed in the field of personality traits and other possible predictors for strength of the Uncanny Valley effect, and past research showing correlations to that effect should be re-inspected for their methodical accuracy as it is highly likely that any conclusive results in fact only showed the participants' differing response style and nothing more.

The implications of the texts cited in the introduction remain important, however – if there is an ingrained Uncanny Valley effect we all experience which leads us to shy away from faces that are familiar, but slightly different, what does that mean for the visibly marginalised among us? One can only hope that, as humanity evolves, we continue to gain enough awareness of our instinctive reactions and urges to be able to control them – as well as being open to expanding our concept of what is familiar and what is "other" as society, and the people we see represented in media, continues to diversify and expand on a global scale.

For all of the results it is worth noting that several participants commented saying they found the survey to be very long, indicating that as the study went on, results may have gotten more inaccurate,

even if the included attention check was completed successfully. Future studies should aim for a lower number of stimuli, or provide the study in separate parts to be completed on different days.

References

Bizumic, B. & Duckitt, J. (2018). Investigating Right Wing Authoritarianism With a Very Short Authoritarianism Scale. *Journal of Social and Political Psychology*, 6(*1*), 129–150. doi:0.5964/jspp.v6i1.835

Gerlitz, J. & Schupp, J. (2005). Zur Erhebung der Big-Five-basierten Persönlichkeitsmerkmale im SOEP. In *DIW Research Notes,* 4. Berlin, Germany: DIW Berlin.

Gray, K., & Wegner, D. M. (2012). Feeling robots and human zombies: Mind perception and the uncanny valley. *Cognition*, 125(*1*), 125–130. doi:10.1016/j.cognition.2012.06.007

Hanson, D. (2005). Expanding the aesthetic possibilities for humanoid robots. *IEEE-RAS international conference on humanoid robots*. Retrieved from: https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.472.2518&rep=rep1&type=pdf

Ho, C.-C., & MacDorman, K. F. (2010). Revisiting the uncanny valley theory: Developing and validating an alternative to the Godspeed indices. *Computers in Human Behavior*, 26(*6*), 1508–1518. doi:10.1016/j.chb.2010.05.015

Hsu, J. (n.d.), *Robotics' Uncanny Valley Gets New Translation.* Live Science. https://www.livescience.com/20909-robotics-uncanny-valley-translation.html

Koopman, Robbin (2019). The Uncanny Valley as a universal experience : a replication study using multilevel modelling. Retrieved from: https://essay.utwente.nl/77172/

Kruglanski, A. W., Atash, M. N., De Grada, E., Mannetti, L. & Pierro, A. (2013). Need for Closure Scale (NFC) . *Measurement Instrument Database for the Social Science*. Retrieved from

Löffler, D., Dörrenbächer, J. & Hassenzahl, M. (2020). The Uncanny Valley Effect in Zoomorphic Robots: The U-Shaped Relation Between Animal Likeness and Likeability. *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20),* 216-270. doi:10.1145/3319502.3374788

MacDorman, K. F., & Entezari, S. O. (2015). Individual differences predict sensitivity to the uncanny valley. *Interaction Studies Social Behaviour and Communication in Biological and Artificial*

Systems, 16(*2*), 141–172. doi:10.1075/is.16.2.01mac

MacDorman, K. F., Green, R. D., Ho, C.-C., & Koch, C. T. (2009). Too real for comfort? Uncanny responses to computer generated faces. *Computers in Human Behavior*, 25(*3*), 695–710. doi:10.1016/j.chb.2008.12.026

Mathur, M. B., & Reichling. (2016). Navigating a social world with robot partners: A quantitative cartography of the Uncanny Valley. *Cognition*, 146(*1*), 22-32. doi:0.1016/j.cognition.2015.09.008

Moosa, M. & Ud-Dean, M. (2009). Danger Avoidance: An Evolutionary Explanation of Uncanny Valley. *Biological Theory,* 5(*1*), 12–14. doi:10.1162/BIOT_a_00016

Mori, M., (2012). The Uncanny Valley [From the Field]. (MacDorman, K., & Kageki, Trans.) *N.IEEE Robotics & Automation Magazine*, 19(*2*), 98–100. doi:10.1109/mra.2012.2192811 (Original work published 1970)

Morton, T. (2019). Racism, Ecology and The Uncanny Valley. Encounters with the Posthuman: Materiality, Vitality, Narrativity, 38-52. Retrieved from: http://aihumanities.org/ko/wp-content/uploads/mangboard/2019/07/16/F272_Racism,%20Ecology%20and%20The%20Uncanny%20Valley_Timothy%20Morton.pdf

Palomäki, J., Kunnari, A., Drosinou, M., Koverola, M., Lehtonen, N., Halonen, J., Repo, M. & Laakasuo, M. (2018). Evaluating the replicability of the uncanny valley effect. Helyonen, 4(11). doi: 10.1016/j.heliyon.2018.e00939

Rhee, J. (2013). Beyond the Uncanny Valley: Masahiro Mori and Philip K. Dick's Do Androids Dream of Electric Sheep? *Configurations*, 21(*3*), 301–329. doi:10.1353/con.2013.0022

Roets, A., & Van Hiel, A. (2011). Item selection and validation of a brief, 15-item version of the Need for Closure Scale. *Personality and Individual Differences*, 50(*1*), 90–94. doi:10.1016/j.paid.2010.09.004

Schmettow, M. (2021) New statistics for design researchers (Draft Version). Retrieved from: https://schmettow.github.io/New_Stats/index.html

Slijkhuis, P.J.H. (2017). *The Uncanny Valley Phenomenon: A replication with short presentation times.* Retrieved from: https://essay.utwente.nl/72507/

Wang, S., Lilienfeld, S., & Rochat, P. (2015). The uncanny valley: Existence and explanations. *Review of General Psychology*, 19(*4*), 393–407. doi:10.1037/gpr0000056

Williams, D. (2019). Heavenly Bodies: Why It Matters That Cyborgs Have Always Been About Disability, Mental Health, and Marginalization. doi: 10.2139/ssrn.3401342

Woods, V., & Hare, B. (2019). Chapter 7: Uncanny valley of the apes. *Why We Love and Exploit Animals: Bridging Insights from Academia and Advocacy*. Abingdon: Routledge.

Yamada, Y., Kawabe, T., & Ihaya, K. (2012). Categorization difficulty is associated with negative evaluation in the "uncanny valley" phenomenon. *Japanese Psychological Research*, 55(*1*), 20–32. doi:10.1111/j.1468-5884.2012.00538.x

Appendix A

SPSS Syntax

```
DATASET ACTIVATE DataSet1.
COMPUTE sensitivity=MEAN(S1_2B_1 TO S99_2E_1).
EXECUTE.
RECODE BFI_1 BFI_2 BFI_4 BFI_5 BFI_7 BFI_8 BFI_9 BFI_11 BFI_12 BFI_13 BFI_15 ('Strongly '+
    'disagree'='1') ('Somewhat disagree'='2') ('Neither agree nor disagree'='3') ('Somewhat '+
    'agree'='4') ('Strongly agree'='5').
EXECUTE.
RECODE BFI_3 BFI_6 BFI_10 BFI_14 ('Neither agree nor disagree'='3') ('Strongly disagree'='5')
    ('Strongly agree'='1') ('Somewhat agree'='2') ('Somewhat disagree'='4').
EXECUTE.
RECODE NFCSF_1 NFCSF_2 NFCSF_3 NFCSF_4 NFCSF_5 NFCSF_6 NFCSF_7 NFCSF_8
NFCSF_9 NFCSF_10 NFCSF_11 NFCSF_12 NFCSF_13 NFCSF_14 NFCSF_15 ('Moderately
Agree'='5') ('Moderately Disagree'='2')  ('Slightly agree'='4') ('Slightly disagree'='3') ('Strongly agree'='6')
('Strongly disagree'='1').
EXECUTE.
COMPUTE bfi_n=MEAN(BFI_1,BFI_2,BFI_3).
EXECUTE.
COMPUTE bfi_e=MEAN(BFI_4,BFI_5,BFI_6).
EXECUTE.
COMPUTE bfi_o=MEAN(BFI_7,BFI_8,BFI_9).
EXECUTE.
COMPUTE bfi_a=MEAN(BFI_10,BFI_11,BFI_12).
EXECUTE.
```

```
COMPUTE bfi_C=MEAN(BFI_13,BFI_14,BFI_15).
EXECUTE.
RECODE VSA_2 VSA_3 VSA_6 ('very strongly disagree'='-4') ('strongly disagree'='-3') ('slightly '+
    'disagree'='-1') ('slightly agree'='1') ('somewhat agree'='2') ('strongly agree'='3') ('very '+
    'strongly agree'='4') ('somewhat. disagree'='-2') ('unsure/ neutral'='0').
EXECUTE.
RECODE VSA_1 VSA_4 VSA_5 ('unsure/ neutral'='0') ('very strongly disagree'='4') ('strongly '+
    'disagree'='3') ('slightly disagree'='1') ('slightly agree'='-1') ('somewhat agree'='-2')
    ('strongly agree'='-3') ('very strongly agree'='-4') ('somewhat. disagree'='2').
EXECUTE.
COMPUTE VSA_score=MEAN(VSA_6,VSA_1,VSA_3,VSA_2,VSA_4,VSA_5).
EXECUTE.
DESCRIPTIVES VARIABLES=sensitivity nfcscore bfi_n bfi_e bfi_o bfi_a bfi_C VSA_score
  /STATISTICS=MEAN STDDEV RANGE MIN MAX.
CORRELATIONS
  /VARIABLES=sensitivity bfi_n bfi_e bfi_o bfi_a bfi_C nfcscore VSA_score
  /PRINT=TWOTAIL NOSIG FULL
  /CI CILEVEL(95) BIAS(TRUE)
  /MISSING=PAIRWISE.
COMPUTE sens=trough-shoulder.
EXECUTE.
DESCRIPTIVES VARIABLES=sens
  /STATISTICS=MEAN STDDEV MIN MAX.
```