Even Better than the Real Thing: How Imperfection Shapes Trust and Engagement with Digital Humans

Johnny Tang-Ear
University of Sydney
Johnny-tangear@sydney.edu.au

Mike Seymour
University of Sydney
Mike.seymour@sydney.edu.au

Alan Dennis
Indiana University
ardennis@iu.edu

Lingyao Yuan
Iowa State University
lyuan@iastate.edu

Catherine Hardy
University of Sydney
Catherine.hardy@sydney.edu.au

Abstract

This study challenges the assumption that more realism in digital humans always leads to greater trust and engagement. Using eye-tracking and post-exposure surveys, we compared viewer responses to three video presenters: a highly realistic digital human, a real human, and an imperfect altered human, represented by a real presenter altered to have unblinking eye contact. While participants rarely noticed visual imperfections consciously, the human with subtle flaws led to significantly greater trust and willingness to pay. The imperfect video outperformed the fully realistic, unaltered human video, suggesting that perfect realism may not always be best. These findings offer important implications for the design of AI-driven digital humans, highlighting that strategic imperfection can enhance authenticity, trust, and engagement in customer interactions. Moreover, the results contribute new empirical insights into the Uncanny Valley theory, suggesting that user affinity and trust may peak not at perfect realism, but can peak at a point just prior to the full realism.

Keywords: Digital Humans, Human-Computer Interaction, Deepfake, Uncanny Valley.

1. Introduction

Digital humans, characterized by their highly realistic human-like faces and voices, are being deployed as Artificial Intelligent (AI) conversational agents across various applications (Seymour et al., 2023). Their roles range from avatars representing individuals to autonomous artificial intelligence-driven agents. As these technologies rapidly evolve, developers and researchers must critically assess how realistic digital humans need to be to optimize user engagement and trust.

The prevailing view in the field suggests that the design of digital humans should strive for the highest possible photorealism—virtually indistinguishable from an actual human, as subtle imperfections may create a feeling of uncanniness (Wang et al., 2015). However, findings from our experimental study challenge this assumption.

We conducted a lab experiment using eye-tracking technology to measure viewer responses to short video presentations designed to simulate a product sales scenario. Each participant viewed one of three conditions. The first condition involved an unaltered video featuring a real human presenter. The second featured a highly realistic digital human created through state-of-the-art deepfake identity-swapping techniques (Seymour et al., 2023). The third presented a partially altered video of a real human, intentionally designed to be imperfect yet close in appearance to the original unaltered footage.

Eye-tracking provided detailed insights into viewer attention patterns, capturing gaze behaviors throughout the viewing experience. Eye-tracking technology offers a granular look into viewer engagement by monitoring and recording eye movements and gaze patterns, which can detect the viewers' attention (Halszka et al., 2017). In this research, eye-tracking serves as a primary technique to investigate user behavior when watching the three video treatments. By comparing the viewer's gaze patterns from eye trackers between videos, we can gain valuable insights into the differences between the three presentation techniques or appearances.

URI: https://hdl.handle.net/10125/111923
978-0-9981331-9-5

Complementing these quantitative measures, participants completed surveys evaluating their affinity and trust of the presenter and their willingness-to-pay.

This study explores the following research questions:

RQ1: Do viewers exhibit behavioral differences, as measured by eye-tracking data, indicating that they perceive or react differently to variations in visual realism among the three video treatments?

This question investigates whether viewers' gaze patterns, specifically, where viewers look and the duration of their gaze, reflect recognition or response to differences in visual quality and accuracy between the treatments.

RQ2: Do the differences between video treatments affect participants' self-reported willingness to pay, trust, or affinity toward the presenter?

This question forms the core focus of the study. It aims to determine which visual representation has the greatest influence on user perception. To address this, participants' self-reported responses are analyzed in conjunction with the behavioral data recorded via eye-tracking.

This research contributes meaningful insights to the Information Systems (IS) and Human-Computer Interaction (HCI) literature by showing that increased visual realism does not necessarily correlate with enhanced viewer trust or attention. These results underscore the complexity of design choices regarding digital human realism and suggest important trade-offs that must be considered in practical applications.

2. Background

The use of digital humans is rapidly expanding, with numerous companies now commercially offering digital-human solutions (e.g., HeyGen, UneeQ, and D-Id.com) (Seymour et al., 2023). High-quality, realistic digital humans are no longer limited to high-budget media and entertainment productions; organizations and even individuals now have the capability to develop personalized digital avatars. Platforms like Epic Games' MetaHuman Creator and Reallusion iClone allow users to design diverse digital characters, supporting a broad range of racial and physical representations.

These developments suggest that while digital humans must achieve a certain minimum quality to communicate emotional content effectively, perfect photorealism is not necessarily the primary determinant of effectiveness. Industries such as healthcare, education, finance, and corporate communications are increasingly adopting digital humans (Seymour et al., 2023), raising important questions about the role visual quality plays in establishing trust and engagement in customer interactions.

Research suggests that more realistic digital humans are preferred to less realistic ones, such as cartoons (Ma et al., 2024; Seymour et al, 2021). It may be that once digital humans approach high levels of realism, the subtle imperfections might not significantly impact user perception, as users may not explicitly notice these differences, although some subtle visual imperfections may indeed subconsciously influence users' behavior, trust, affinity, and their willingness to pay (Wang et al., 2015).

Digital humans are realistic digital entities controlled either by humans or AI. Their implementation across sectors, including fashion, entertainment, gaming, education, and corporate communications, is driven by advancements in computer graphics, GPU hardware, and neural rendering technologies, particularly deep learning-based "deepfake" techniques (Seymour et al., 2023). Applications range from digital characters in films and commercials to AI-driven virtual agents used as customer service representatives, sales assistants, or social influencers.

Decisions regarding the visual appearance of digital humans vary by application. Some, like Apple's emojis, adopt deliberately stylized designs, whereas many applications strive for highly realistic representations. Previous studies suggest that realism positively correlates with user trust and affinity (Wang et al., 2015); however, questions remain about whether digital humans must achieve perfect realism or merely reach a threshold of realism sufficient to avoid the uncanny valley (Seymour et al., 2021).

The adoption of digital humans began within the entertainment and film industry, where realism is crucial for narrative believability. Technical imperfections in films often break immersion, causing audiences to shift attention from the narrative to the filmmaking process itself. Consequently, film production invests heavily in creating characters realistic enough to maintain audience empathy and narrative credibility.

Recent technological advancements now enable the cost-effective production of near-realistic interactive digital humans. Neural rendering techniques, such as generative adversarial networks (GANs), produce extremely realistic digital faces from training data. Nevertheless, visual imperfections persist and are detectable despite rapid advancements (Seymour et al., 2023).

Industries adopting digital humans for customer-facing roles typically aim for film-industry standards of realism, which can be expensive. A practical question thus arises: is such investment in near-perfect realism justified outside the entertainment context? Our earlier research questioned whether users genuinely prefer highly realistic AI-driven agents, or whether a lower but adequate level of realism would suffice. The current study extends this exploration to assess whether users might be subconsciously influenced by subtle imperfections, even if they do not consciously recognize them.

The "uncanny valley" theory argues that user affinity (natural liking) increases as digital avatars become more realistic, but dramatically drops once avatars are nearly, but not fully, humanlike (Mori et al., 2012; Wang et al., 2015). See Figure 1. After crossing this valley, affinity theoretically rises again, reaching its peak at complete realism. Despite its widespread use, the original uncanny valley theory was largely speculative, lacking empirical validation, particularly concerning the exact trajectory of affinity after surpassing the uncanny point.

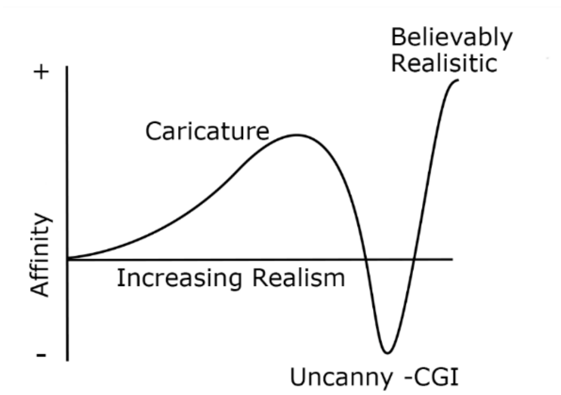

Figure 1. Uncanny Valley

The theory was developed before modern media technologies, such as high-resolution smartphones, significantly shaped user expectations around digital content. As technological advances have transformed user experiences, the notion of realism and acceptable imperfections might have evolved. While prior research primarily addressed overcoming the uncanny valley (Wang et al., 2015), our research explores the subtler space beyond this threshold—where users no longer experience explicit discomfort but might still respond differently to varying levels of imperfection.

From a theoretical perspective, a gap exists between perfect realism and minor imperfections. Traditionally, any departure from perfect realism was assumed inherently detrimental to affinity, trust, and effectiveness. However, our findings suggest a counterintuitive insight: deliberate visual imperfections might enhance viewer trust and willingness to pay, challenging the assumption that maximal realism always yields the best outcomes.

To understand how visual imperfections influence user behavior, we examine three key outcomes in the e-commerce context: user affinity, willingness to pay, and trust. Willingness to pay describes the monetary value users assign to a product based on subjective perceptions of quality and value (Rosen, 1974; Wertenbroch & Skiera, 2002). Previous research has shown that willingness to pay in online contexts can be affected by numerous factors, including product imagery and system design (Yuan & Dennis, 2019) (Tripathi et al., 2009). Visual quality, such as the realism of digital-human presenters in product videos, could thus substantially influence purchasing decisions.

Trust reflects an individual's willingness to be vulnerable based on positive expectations of others' actions (Mayer et al., 1995), and it extends to interactions with virtual agents, avatars, and digital systems (Benbasat & Wang, 2005; Komiak & Benbasat, 2006; Lowry et al., 2008). Affinity, on the other hand, captures how favorably and realistically viewers perceive virtual avatars. Both trust and affinity are critical determinants of user engagement and purchasing behavior in digital interactions (Etemad-Sajadi, 2016).

Therefore, investigating how minor visual imperfections affect these outcomes offers practical guidance for organizations employing digital humans in customer-facing roles, potentially optimizing both effectiveness and efficiency in their deployments.

While the idea that minor visual imperfections might increase the effectiveness of digital humans may initially appear counterintuitive, it can be understood as consistent with broader societal norms regarding interpersonal communication. Three commonly accepted aspects of human interaction provide a possible useful framework for interpreting these findings.

First, prior research has consistently shown that direct eye contact fosters perceptions of connection and trust between interactants. Conversely, individuals who avoid eye contact or whose gaze is averted are often judged as less trustworthy or sincere (Bente et al., 2001). Direct eye contact during conversation fosters a sense of connection and trust between individuals. Eye contact is generally perceived as a key component of meaningful, honest interaction, reinforcing the social bond between speaker and listener.

Second, consider the widespread use of eye makeup. While often said to enhance natural features, it is equally plausible to argue that practices such as applying eyeliner or eyeshadow primarily serve to draw attention to the eyes. These enhancements subtly direct focus toward the eyes without being consciously noticed by the viewer. Importantly, they highlight key facial features without making the individual's eyes appear distracting from the overall impression of authenticity.

Figure 2: An example of the brain's separate and specialist facial processing.

Figure 3. The Same Image, just rotated 180 Degrees

Third, it is a common belief that individuals who avoid eye contact, or whose gaze appears "shifty", are perceived as less trustworthy. Colloquial expressions such as "look me in the eye and tell me that" underscore the strong association between steady eye contact and perceived honesty. Thus, if a digital human's eyes are slightly exaggerated in their directness, appearing to "stare" more than a perfectly realistic human might, it may unintentionally enhance perceptions of trustworthiness. This subtle imperfection, while deviating from perfect realism, could reinforce positive judgments about the digital human's sincerity.

Additionally, research in cognitive psychology provides further support for this interpretation. Established phenomena such as the Thatcher Effect demonstrate that the human brain processes key facial features, particularly the eyes and mouth, separately from the overall facial structure (Seymour et al, 2018). This suggests that visual discrepancies in the eyes' appearance relative to the face may not be inherently unsettling. Figures 2 and 3 present the same image rotated 180 degrees. Despite extreme distortions in eye and mouth positioning, viewers typically do not perceive the face as disturbing when presented upside down, highlighting the brain's isolated processing of these critical features when orientation cues are altered. By analogy, subtle gaze-related imperfections (e.g., unblinking eyes) in our altered condition may go unnoticed yet still influence trust.

Together, these insights suggest that minor imperfections, particularly those related to eye presentation, may not disrupt, and may even enhance, the effectiveness of digital humans in fostering trust, affinity, and persuasive engagement.

H1: Users will display greater visual attention to imperfections in a) the eyes, mouth, and face of the digital human, and b) the eyes of the altered human.

H2: Compared to a real human presenter, participants in the digital human presenter condition will report lower a) affinity, b) trust, and c) willingness to pay.

H3: Compared to a real human presenter, participants in the altered human presenter condition will report lower a) affinity, b) trust, and c) willingness to pay.

3. Methodology

3.1. Participants

We recruited 75 adult participants from an Australian university, all fluent in English. Participants were recruited via promotional flyers that did not disclose the monetary value of incentives to prevent coercion and ensure ethical compliance. After completion of the study, each participant received a $45 movie voucher as a token of appreciation, proportional to the approximately 15-minute time commitment required for participation. A power analysis using G*Power (Faul et al., 2007) determined that the design had a power of .57 to detect a medium effect.

3.1. Experimental Design

The study employed an experimental design conducted in a controlled laboratory environment at the university. Participants viewed one of three video stimuli (each about 1 minute) promoting a new Apple iPad product and completed a survey (see Figure 4):

- Control Condition: An unaltered video of a real human presenter
- Digital Human Condition: A video of a highly realistic digital human created using advanced neural rendering and identity-swapping techniques
- Eye Altered Condition: A partially altered video of the same real human presenter by incorporating subtle visual imperfections to the eyes (i.e., eye fixation without blinking)

The size of the presentation on screen was sized to the approximate size of a mobile device, of an iPad, to provide a plausible presentation on a large high resolution computer monitor.

We conducted randomization checks across the three treatments and found no significant differences in gender, race, education, English ability, income, or attitude to Apple (lowest p-value was $p=0.330$), so we conclude there is no evidence of a randomization failure.

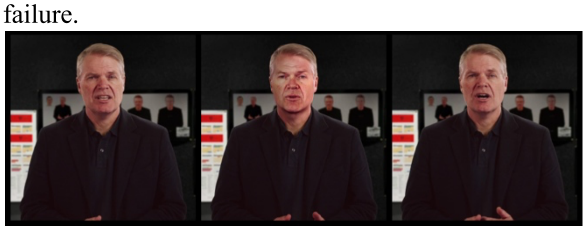

Figure 4: The three treatments, Eye altered (left), digital human (middle), and the control (unaltered).

3.2. Procedure

Upon arrival, each participant was briefed about the study objectives and procedure as well as provided informed consent via a Participant Consent Form. The experiment comprised four sequential phases: 1) a calibration of eye-tracking equipment, 2) viewing of video stimuli under one randomly assigned experimental condition, 3) completion of a post-study survey, and 4) completion of the post-study demographic survey. Each session concluded with a participant debriefing.

3.2.1. Calibration. The calibration of the eye-tracking experiment involved participants being asked to adjust the height of the chair to around the optimal height of 50 centimeters to 55 centimeters, as well as ensuring that participants were approximately 60 centimeters away from the screen where adjustments were made to the computer screen to account for height differences. During the calibration, participants were asked to follow a white moving dot on the computer screen to ensure that the eye-tracking hardware was able to accurately detect the participants' eye movement and gaze in order to validate the data (Djamasbi, 2014). Furthermore, as participants are focused upon a white moving dot with a series of static dots as checkpoints, this enables the eye tracking hardware to determine and assess the participant's pupil data on the computer screen (Djamasbi, 2014).

3.2.2. Post-study Survey. Immediately after viewing the assigned video, participants completed a structured post-study survey (see Appendix A) designed to measure subjective experiences, trust, affinity, and willingness to pay using 7-point Likert scales. The survey instrument is designed to capture overall perceived trustworthiness rather than to differentiate explicitly between competence, benevolence, and integrity. Although validated trust scales exist (Mayer et al., 1995; Benbasat & Wang, 2005), they are lengthy and suited to ongoing interactions, not a one-minute video pitch. We therefore adapted items for ecological validity in our sales-agent context while still covering core trustworthiness dimensions and confirmed robustness through internal consistency and validity. All were reliable, and an HTMT analysis showed they had discriminant validity. We also included open-ended demographic questions (gender, ethnicity, education year, native English speaker, income), as well as the participants' attitudes towards Apple products.

3.2.3. Validity. A series of factors were taken into consideration prior to and during the experiment to mitigate the different types of biases encountered. These include a measurement bias (Page & Henderson, 2008) via standardization of data collection through having a consistent 7-point Likert scale to ensure that all participants' responses were recorded throughout the post-study survey. Other biases include sampling bias, which was mitigated through random allocation of the different video stimuli provided to participants and random selection of participants (Popovic & Huecker, 2023) where they were recruited through QR code flyers, as well as confirmation bias, which was alleviated by balancing the survey with a range of open and closed-ended questions and statements to reduce the amount of influence on how questions are asked. The surveys also include a range of attention checks to ensure that participants were paying attention to the questions being asked.

Figure 5. Room setup, with eye tracking below the monitor

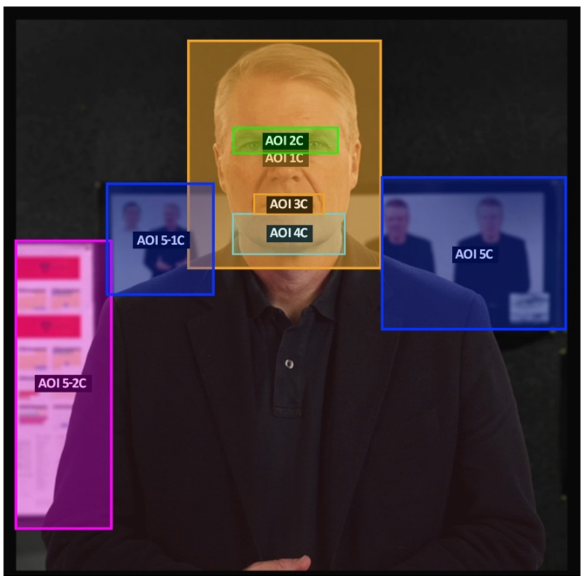

Figure 6. Areas of interest (AOI), shown as overlays.

3.3. Materials

The primary materials included Participant Information Statements, Participant Consent Forms, pre- and post-study surveys, video stimuli, and eye-tracking equipment.

Eye-tracking data were collected using the iMotions eye tracker, positioned at the base of the monitor (see Figure 5), which captured detailed eye movement and temporal gaze distribution metrics. The eye tracker provided precise insights into participants' visual attention and behavioral responses to each video treatment. Eye-tracking data were collected to examine participants' eye movements and gaze patterns while viewing the experimental videos. The primary eye-tracking metrics used in this analysis included Areas of Interest (AOIs) and Heat Maps. An Area of Interest (AOI) refers to predefined regions within a visual stimulus from which specific eye-tracking metrics can be extracted (Farnsworth, 2023). AOIs themselves are not direct metrics but rather regions designated to quantify key visual behaviors.

4. Results

We began with a heat map analysis to understand overall gaze patterns. Figure 7 shows the distribution and intensity of gaze points across the stimulus for the three treatments. Areas marked in red represent high concentrations of gaze points, signifying greater visual attention and frequency of gaze fixation. Yellow and green areas indicate moderate to low levels of attention, respectively, while uncolored regions were largely unattended (Farnsworth, 2024). The gaze patterns were similar across the treatments.

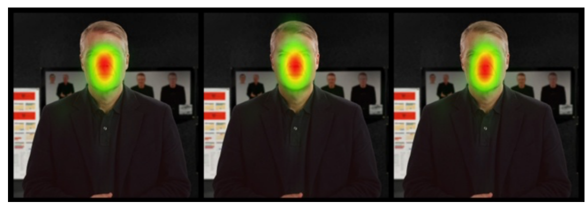
Figure 7: Heat maps show very similar patterns

We assessed the treatment effects with regressions using the controls of gender (female or not), Ethnicity (Caucasian or not), native English speaker (or not), education level, and annual income. There were no significant differences between the digital human and the real human for any of the four outcomes in Table 1. Surprisingly, the altered human treatment had *greater* affinity (p=.046) and *greater* trust (p=.017) than the real human treatment, but there were no differences for the bid amount or video quality. H2 and H3 are not supported.

Table 1. Treatment Statistics

		Real Human (Control)	Digital Human	Altered Human
Bid	Mean	1716.72	1607.88	1516.54
	St Dev.	389.13	423.84	573.71
Affinity	Mean	3.87	3.89	4.34
	St Dev.	1.05	1.12	0.95
Trust	Mean	4.00	4.30	4.45
	St Dev.	1.00	1.04	1.00
Video Quality	Mean	5.25	4.63	5.28
	St Dev.	1.41	1.39	1.22

We tested whether eye gaze influenced the outcomes using the Hayes PROCESS macro (Hayes, 2013), a regression-based approach that estimates direct and indirect effects in mediation and moderation models. This allowed us to examine whether participants' gaze patterns mediated the relationship between the experimental treatments and the outcome variables. The results are presented in Table 2, with significant relationships highlighted in grey. We included controls in the model but omitted them from Table 2 to improve readability (the only significant control was *Apple* attitude).

Interestingly, the digital human treatment and the altered human treatment did not affect the percentage of time participants spent looking at different AOIs compared to the control treatment of the real, unaltered human. H1 is not supported. However, the percent of time spent looking at the eyes (AOI2) significantly increased the amount bid for all treatments. Likewise, the percent of time looking at the mouth (AOI3) increased affinity, while looking at one background area (AOI5) reduced affinity. No other AOIs were significant, and no treatment by AOI interaction was significant, which means the effects were consistent across treatment (i.e., looking at the distorted eyes in the altered human treatment had the same effect as looking at the normal eyes in the real human treatment).

We also observed direct positive effects from the altered human treatment to affinity and trust (i.e., unrelated to eye gaze), as matching the results of the tests of H2 and H3 above. This suggests that the unexplained positive results are not related to eye gaze.

Table 2. Statistical Results

Relationship	Beta	t-test	P value
DH -> AOI1 Face	.301	.968	.333
DH -> AOI2 Eyes	.010	.029	.977
DH -> AOI3 Mouth	-.067	.299	.765
DH -> AOI4 Chin	-.117	.461	.645
DH -> AOI5	.021	.117	.907
DH -> AOI51	-.268	.994	.320
DH -> AOI52	-.196	.627	.531
AH -> AOI1 Face	.296	1.103	.270
AH -> AOI2 Eyes	-.150	.580	.562
AH -> AOI3 Mouth	.346	1.078	.281
AH -> AOI4 Chin	-.054	.222	.824
AH -> AOI5	.484	1.596	.110
AH -> AOI51	.247	.710	.478
AH -> AOI52	.220	.752	.452
AOI1 Face -> Affinity	-.354	1.505	.132
AOI1 Face -> Bid	-.040	.293	.770
AOI1 Face -> Trust	-.129	.611	.541
AOI1 Face -> VideoQuality	-.102	.461	.645
AOI2 Eyes -> Affinity	.351	1.140	.254
AOI2 Eyes -> Bid	.616	2.307	**.021**
AOI2 Eyes -> Trust	.142	.539	.590
AOI2 Eyes -> VideoQuality	-.038	.237	.813
AOI3 Mouth -> Affinity	.359	2.239	**.025**
AOI3 Mouth -> Bid	.298	1.763	.078
AOI3 Mouth -> Trust	.167	1.038	.299
AOI3 Mouth -> VideoQuality	-.071	.503	.615
AOI4 Chin -> Affinity	.053	.412	.680
AOI4 Chin -> Bid	-.110	.935	.350
AOI4 Chin -> Trust	.063	.549	.583
AOI4 Chin -> VideoQuality	-.114	.733	.463
AOI5 -> Affinity	.192	.809	.419
AOI5 -> Bid	.054	.305	.760
AOI5 -> Trust	-.008	.112	.911
AOI5 -> VideoQuality	.056	.383	.702
AOI51 -> Affinity	-.356	2.123	**.034**
AOI51 -> Bid	-.010	.321	.748
AOI51 -> Trust	-.225	1.277	.202
AOI51 -> VideoQuality	-.480	1.786	.074
AOI52 -> Affinity	-.247	1.020	.308
AOI52 -> Bid	.064	.319	.750
AOI52 -> Trust	-.109	.285	.776
AOI52 -> VideoQuality	.082	.042	.967
DH x AOI1 Face -> Affinity	.130	.411	.681
DH x AOI1 Face -> Bid	.455	1.386	.166
DH x AOI1 Face -> Trust	.089	.226	.821
DH x AOI1 Face -> VideoQuality	-.214	.423	.672
DH x AOI2 Eyes -> Affinity	-.108	.124	.902
DH x AOI2 Eyes -> Bid	-.367	1.106	.269
DH x AOI2 Eyes -> Trust	.168	.512	.608
DH x AOI2 Eyes -> VideoQuality	.514	1.793	.073
AH x AOI1 Face -> Affinity	-.274	.679	.497
AH x AOI1 Face -> Bid	.368	.890	.373
AH x AOI1 Face -> Trust	-.469	1.177	.239
AH x AOI1 Face -> VideoQuality	-.219	.335	.738
AH x AOI2 Eyes -> Affinity	.211	.443	.658
AH x AOI2 Eyes -> Bid	-.751	1.221	.222
AH x AOI2 Eyes -> Trust	.055	.078	.938
AH x AOI2 Eyes -> VideoQuality	.140	.213	.831
Affinity -> Bid	.231	1.362	.173
Trust -> Bid	-.125	.606	.545
VideoQuality -> Bid	-.101	.839	.401
DH -> Affinity	.046	.211	.833
DH -> Bid	-.192	.598	.550
DH -> Trust	.331	1.153	.249
DH -> VideoQuality	-.487	1.670	.095
AH -> Affinity	.681	2.119	**.034**
AH -> Bid	-.524	1.330	.184
AH -> Trust	.775	2.432	**.015**
AH -> VideoQuality	.188	.309	.757

Note: DH stands for Digital Human Treatment; AH stands for Altered Human Treatment; AOI stands for Area of Interest (AOI)—percentage of time spent on the area of interest. Please refer to Figure 6 for different AOI.

5. Discussion and Limitations

This study set out to explore how subtle imperfections in real human presenters and digital human representations influence user trust, affinity, and purchasing behavior in an e-commerce context. Contrary to prior assumptions in both research and practice that equate visual perfection with optimal outcomes, our findings reveal three key insights that challenge this logic.

First, our results indicate that the digital human presenter performed comparably to the unaltered real human in all measured outcomes—trust, affinity, perceived video quality, and willingness to bid. Participants did not react negatively to the digital human, suggesting that in controlled, short-form content such as product pitches, realistic digital humans can be as effective as real human presenters. This finding aligns with the growing adoption of AI-driven agents in commercial contexts, demonstrating that realism can be achieved without penalties.

Second, and more unexpectedly, we found that the subtly imperfect real human presenter, with minor visual anomalies such as unblinking eye contact, generated significantly higher trust and affinity than the unaltered real human. While video quality ratings and bid amounts did not differ, participants attributed greater emotional connection and credibility to the imperfect presenter. This suggests that minor imperfections could humanize the presenter, potentially activating subconscious cues associated with sincerity and attentiveness. These results offer empirical support for rethinking the assumption that visual perfection always equates to better outcomes in human-agent interaction.

In the original Uncanny Valley theory, the x-axis is most often translated as 'increasing realism', with the highest affinity presumed at 100% real human likeness. Within this framing, both the digital human and the altered human fall short of perfect realism and thus should, by definition, elicit lower affinity. Our results, however, suggest otherwise: participants attributed greater trust and affinity to the "imperfect altered human" than to the unaltered real human. This indicates that the altered condition represents not merely another point on the Uncanny graph, but a qualitatively distinct position where subtle imperfection can enhance perceived trustworthiness, this contradicts the accepted expectation that maximum affinity aligns with maximum realism.

Third, eye-tracking data revealed that gaze behavior significantly influenced participants' responses. Specifically, increased visual attention to the presenter's eyes predicted higher bid amounts, while increased gaze toward the mouth was associated with stronger feelings of affinity. Introducing imperfections, subtly draw attention to the parts of the presenter's face in ways that drive trust in users and people in general. Importantly, these gaze patterns were consistent across all three treatment conditions, suggesting that viewer engagement is more strongly shaped by spontaneous visual attention than by the nature of the presenter (real or artificial). This unscores that eye gaze itself is a powerful mediator of perceived trust and engagement.

5.1. Implications for Future Research

While the study offers robust contributions, it is not without limitations. Most notably, the visual imperfection we tested, lack of blinking—was subtle and context-specific. Future research could extend and repeat the approach to validate replicability with different subjects. In doing it could also explore a broader range of 'imperfections' (e.g., slight facial asymmetries around the eyes or micro-expression anomalies) to understand which elements enhance or diminish perceived authenticity.

Additionally, the presenter in all conditions was Caucasian and spoke without a strong accent. Cultural familiarity and demographic similarity are known to influence trust and affinity. Future studies should investigate how race, gender, language, and other identity cues interact with imperfection and realism to shape user perceptions.

Future research could investigate the boundaries of imperfection in human and digital representations specifically, the extent to which visual alterations can be introduced before crossing a perceptual threshold where trust and authenticity are undermined. We suspect this threshold is highly content-dependent and will vary by individual differences, making it unlikely that a single standard of "acceptable imperfection" exists across contexts. Identifying and mapping these boundaries would not only clarify when subtle flaws enhance perceived sincerity but also when distortions risk triggering distrust. This is a point to explore further with attention to how these dynamics may shift across modalities (e.g., visual versus vocal imperfections) and situational contexts.

Moreover, although sound quality was controlled and held constant in this study, it was not independently manipulated. Given the role of audio in shaping perceptions of professionalism and competence, future experiments could examine if imperfections in voice synthesis, tone, or sync also influence viewer trust?

Finally, the experiment was conducted in a laboratory setting using a brief product pitch. While valid for many digital scenarios, future work could seek to replicate these findings in more natural or commercial environment where actual sales are involved, such as digital commerce or actual online shopping scenarios.

5.2. Implications for Practice

These findings hold practical significance for organizations deploying digital human presenters. The assumption that greater realism always leads to better outcomes should be reconsidered. Instead, strategic imperfection, particularly in subtle facial cues like eye gaze, may enhance perceived authenticity and trustworthiness, especially in contexts where emotional engagement is critical.

Our findings suggest a counterintuitive effect: the "less than real" presenter was perceived as more trustworthy than the fully real human. This challenges a core assumption of the Uncanny Valley framework, where 100% real humans are expected to yield the highest levels of affinity. In contrast to prior research, our experiment demonstrates that perfect realism is not always the point of maximum trust or affinity.

For digital human developers and marketing professionals, this suggests a shift in design priorities: rather than striving for flawless realism, attention might be better focused on how intentional, human-like imperfections can be used to simulate natural behaviour. This could lead to more effective, emotionally resonant interactions without the high costs associated with ultra realistic digital human production.

Importantly, as digital humans become increasingly pervasive in commercial and organizational contexts, ethical and regulatory frameworks must evolve to address the complexities of human-agent interaction. While our findings suggest that subtle visual imperfections can enhance trust and perceived authenticity, they also blur the line between genuine and synthetic representation. This ambiguity introduces ethical risks, particularly when users are unaware that they are interacting with AI-driven agents. Developers and deploying organizations have a responsibility to be transparent about the artificial nature of digital humans, especially in scenarios involving persuasion, influence, or customer decision-making. Trust should not be manipulated covertly, and users must be able to distinguish between real and synthetic communicators in ways that respect their autonomy and consent.

Moreover, the effectiveness of "strategic imperfection" raises important questions about the ethics of designing for subconscious influence. Our results indicate that small visual anomalies, such as unblinking gaze, can increase trust without participants' conscious awareness. This creates the potential for ethically ambiguous design practices, especially in domains such as finance, healthcare, or customer service, where users may make decisions under the impression of interacting with sincere, trustworthy human agents. While such design strategies may improve engagement, they must be implemented with caution. Clear ethical guidelines and design accountability are essential to ensure that emotional resonance is not used to exploit user vulnerability, and that the deployment of digital humans prioritizes transparency, fairness, and responsible interaction.

6. References

Benbasat, I., & Wang, W. (2005). Trust in and adoption of online recommendation agents. *Journal of the Association for Information Systems*, *6*(3), 72 - 101. https://doi.org/10.17705/1jais.00065

Bente, G., Krämer, N. C., Petersen, A., & De Ruiter, J. P. (2001). Computer animated movement and person perception: Methodological advances in nonverbal behavior research. Journal of Nonverbal Behavior, 25(3), 151–166. https://doi.org/10.1023/A:1010690525717

Djamasbi, S. (2014). Eye Tracking and Web Experience. *Transactions on Human-Computer Interaction*, *6*(2), 37 - 54. https://doi.org/10.17705/1thci.00060

Etemad-Sajadi, R. (2016). The impact of online real-time interactivity on patronage intention: The use of avatars. *Computers in Human Behavior*, *61*, 227 - 232. https://doi.org/10.1016/j.chb.2016.03.045

Farnsworth, B. (2024). Eye tracking: The complete pocket guide. Retrieved 14/05/2025, from https://imotions.com/blog/learning/best-practice/eye-tracking/

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*(2), 175–191. https://doi.org/10.3758/BF03193146

Halszka, J., Holmqvist, K., & Gruber, H. (2017). Eye tracking in educational science: Theoretical frameworks and research agendas. *Journal of eye movement research*, *10*(1), 1 - 18. https://doi.org/10.16910/jemr.10.1.3

Hayes, A. F. (2013). Introduction to mediation, moderation, and conditional process analysis: A regression-based approach. Guilford Press.

Komiak, S. Y. X., & Benbasat, I. (2006). The effects of personalization and familiarity on trust and adoption of recommendation agents. *MIS Quarterly*, *30*(4), 941 - 960. https://doi.org/10.2307/25148760

Lowry, P. B., Vance, A., Moody, G., Beckman, B., & Read, A. (2008). Explaining and predicting the impact of branding alliances and web site quality on initial consumer trust of e-commerce web sites. *Journal of Management Information Systems*, *24*(4), 199 - 224. https://doi.org/10.2753/MIS0742-1222240408

Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review*, *20*(3), 709 - 734. https://doi.org/10.2307/258792

Mori, M., MacDorman, K. F., & Kageki, N. (2012). The Uncanny Valley. *IEEE Robotics & Automation Magazine*, *19*(2), 98 - 100. https://doi.org/10.1109/MRA.2012.2192811

Page, L. A., & Henderson, M. (2008). Appraising the evidence: what is measurement bias? *Evidence Based Mental Health*, *11*(2), 36 - 37. https://doi.org/10.1136/ebmh.11.2.36

Popovic, A., & Huecker, M. R. (2023). Study Bias. In *StatPearls [Internet]*. Treasure Island (FL): StatPearls Publishing. https://www.ncbi.nlm.nih.gov/sites/books/NBK574513/

Rosen, S. (1974). Hedonic prices and implicit markets: product differentiation in pure competition. *Journal of Political Economy*, *82*(1), 34 - 55. https://doi.org/10.1086/260169

Seymour, M., Riemer, K., Yuan, L., & Dennis, A. R. (2023). Beyond deep fakes. Communications of the ACM, 66(10), 56–67. https://doi.org/10.1145/3584973

Seymour, M., Yuan, L. I., Dennis, A., & Riemer, K. (2021). Have We Crossed the Uncanny Valley? Understanding Affinity, Trustworthiness, and Preference for Realistic Digital Humans in Immersive Environments. *Journal of the Association for Information Systems*, *22*(3). https://doi.org/10.17705/1jais.00674

Wang, S., Lilienfeld, S. O., & Rochat, P. (2015). The Uncanny Valley: Existence and Explanations. *Review of general psychology*, *19*(4), 393 - 407. https://doi.org/10.1037/gpr0000056

Wertenbroch, K., & Skiera, B. (2002). Measuring consumers' willingness to pay at the point of purchase. *Journal of Marketing Research*, *39*(2), 228 - 241. https://doi.org/10.1509/jmkr.39.2.228.19086

Yuan, L. I., & Dennis, A. R. (2019). Acting Like Humans? Anthropomorphism and Consumer's Willingness to Pay in Electronic Commerce. *Journal of Management Information Systems*, *36*(3), 450 - 477. https://doi.org/10.1080/07421222.2019.1598691

Appendix A – Survey Questions

Bid Amount	Suppose this tablet was for sale on eBay or Facebook Marketplace. The suggested retail price is A$2199. How much would you bid for the tablet?
Affinity Alpha=.870	I have affinity with the sales agent. I feel a closeness with the sales agent. I feel a likeness with the sales agent. I feel rapport with the sales agent. The sales agent was friendly.
Trust Alpha=.900	Overall, the sales agent was trustworthy. I trust the sales agent. I can rely on the sales agent. The sales agent can be trusted with selling the product. I have confidence with the sales agent. I feel confident about the sales agent's skills. The sales agent had integrity.
Video Quality Alpha=.906	The video has good quality. I can view the video clearly. There is no issue with the quality of the video. The video is high quality.
Apple Attitude Alpha=.796	I will recommend Apple products to someone who seeks my advice. Next time I will buy Apple products. I will switch to other brands if I experience a problem with Apple products. I will stay with Apple products even if I experience problems.
Demographics Open-ended	What is your gender? What is your race/ethnicity? What is your highest level of schooling that you have completed? Is English your native language? What is your annual gross income?