

The Illusion of Performance: The Effect of Phantom Display Refresh Rates on User Expectations and Reaction Times

Esther Bosch
German Aerospace Center, Institute
for Transportation Systems
Braunschweig, Germany
esther.bosch@dlr.de

Robin Welsch
Aalto University
Espoo, Finland
robin.welsch@aalto.fi

Tamim Ayach
HU Berlin
Berlin, Germany
mohammed.tamim.ayach@student.hu-berlin.de

Christopher Katins
HU Berlin
Berlin, Germany
christopher.katins@hu-berlin.de

Thomas Kosch
HU Berlin
Berlin, Germany
thomas.kosch@hu-berlin.de



Figure 1: Users conducted a target selection reaction time task with different display refresh rates. Users were instructed beforehand about the refresh rate. However users received instructions about using a superior or inferior refresh rate while using the opposite or narrated refresh rate. Users who were allegedly using high refresh rates expected a higher performance. In contrast, users who obtained a low refresh rate expected a degraded performance even when using a high refresh rate.

ABSTRACT

User expectations impact the evaluation of new interactive systems. Increased expectations may enhance the perceived effectiveness of interfaces in user studies, similar to a placebo effect observed in medical studies. To showcase the placebo effect, we conducted a user study with 18 participants who performed a target selection reaction time test with two different display refresh rates. Participants saw a stated screen refresh rate before every condition, which corresponded to the true refresh rate only in half of the conditions and was lower or higher in the other half. Results revealed successful priming, as participants believed in superior or inferior performance based on the narrative despite using the opposite refresh rate. Post-experiment questionnaires confirmed participants

still held onto the initial narrative. Interestingly, the objective performance remained unchanged between both refresh rates. We discuss how study narratives influence subjective measures and suggest strategies to mitigate placebo effects in user-centered study designs.

CCS CONCEPTS

• **Human-centered computing** → **User studies**; *HCI theory, concepts and models*; *Empirical studies in HCI*.

KEYWORDS

Refresh Rates, Placebo Effect, User Expectations, Placebo, User Studies, Human-AI Interfaces

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
CHI EA '24, May 11–16, 2024, Honolulu, HI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0331-7/ 24/05
<https://doi.org/10.1145/3613905.3650875>

ACM Reference Format:

Esther Bosch, Robin Welsch, Tamim Ayach, Christopher Katins, and Thomas Kosch. 2024. The Illusion of Performance: The Effect of Phantom Display Refresh Rates on User Expectations and Reaction Times. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24), May 11–16, 2024, Honolulu, HI, USA*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3613905.3650875>

1 INTRODUCTION & BACKGROUND

A display’s refresh rate determines how often a computer can show a new image per second. Nowadays, manufacturers produce consumer displays with refresh rates ranging between 60 Hertz and 500 Hertz (Hz). Higher refresh rates let content appear earlier on a computer display, reducing reaction times between users and the displayed content. Hence, computer displays with high refresh rates reduce the latency between user input and the viewed content, a critical gaming aspect. In this context, higher refresh rates reduce the overall reaction time [10, 17]. Yet, the refresh rate of a computer display is latent. For comparison, Halbhuber et al. [9] showed that phantom network latencies affect players’ subjective expectations and objective performance. Inspired by this work, this paper investigates if and how narrated but not placebic refresh rates influence user expectations and performance.

Interfaces in Human-Computer Interaction (HCI) tend to conceal the complexity of intelligent interface behavior. However, this can change users’ expectations towards systems, potentially leading to biased study results [13] or facilitating increased risk-taking when users believe that technologies improve their capabilities [24]. Parallel to medicine, research indicates that an individual’s confidence in the efficacy of a treatment is enough to bring about enhancements in their physical or mental well-being, even when the treatment lacks genuine effectiveness. This positive change results from the patient anticipating the treatment will positively impact their health. This occurrence, where a person experiences benefits from an inactive intervention, is recognized as the placebo effect [15]. A prerequisite is the expectation and user’s belief in the efficacy of the treatment. A placebo, as evidenced by studies [1, 11, 14], can alleviate pain or contribute to the treatment of illnesses, providing an effective medical intervention despite lacking a specific mechanism for a particular ailment. The effectiveness of the placebo hinges on the patient’s anticipation of its efficacy [2], resulting in positive post-treatment evaluations [16, 21], and sometimes even objective physiological changes [7, 8, 20]. Consequently, the placebo effect complicates the assessment of new medical treatments, regardless of their actual utility. To address this challenge, medical studies incorporate placebos as controls during the evaluation of novel treatments. The effectiveness of a treatment is only acknowledged if the benefits surpass the improvements observed in participants treated with a placebo control condition. This placebo control approach is commonly applied in various scientific fields assessing human responses, such as psychological treatment [2], sports science [18], and visualization research [3]. However, this methodology is not standard when evaluating the efficacy of new interfaces in HCI.

HCI research has demonstrated that placebos can enhance usability and user experience without a functional system. Thus, controlling for user expectations in user studies gains importance in the HCI community [25]. In gaming, fictitious power-up elements that do not impact gameplay and deceptive descriptions of AI adaptation have been found to boost self-reported game immersion [4, 6]. Social media, offering control settings for prioritizing items in a news feed, can lead to higher subjective ratings of user satisfaction, even when these settings have no actual influence [22]. Studies by Kosch et al. [13] revealed that a non-functional supportive interface can induce a placebo effect related to perceived performance gains and

	Stated 60 FPS	Stated 240 FPS
True 60 FPS	Low (H1, H3) Low (H5)	High (H1, H3) Low (H5)
True 240 FPS	Low (H1, H3) High (H5)	High (H1, H3) High (H5)

Table 1: Expected subjective assessments before and after the target selection reaction time task according to our hypotheses. We conducted a full factorial experimental design. Red: Subjective assessment before and after interaction. Blue: Objective reaction time.

reduced workload measures. In the context of participant beliefs about AI, Pataranutaporn et al. [19] conducted a study where individuals interacting with a mental health AI chatbot were informed about different AI characteristics. Despite all participants engaging with the same AI model, those told that the AI was benevolent reported significantly higher levels of trustworthiness, empathy, and effectiveness in providing mental health advice than those primed to believe it was neutral or manipulative. Examining user expectations in human-AI interaction, Kloft et al. [12] found that increased expectations, irrespective of the actual presence of a supportive interface, improve performance due to placebo effects. Interestingly, negative AI descriptions do not bias performance expectations. This dynamic can adversely affect human behavior when utilizing technologies believed to be enhancing. In addition, Vicente and Matute [23] showed that placebo effects manifest even after the interaction. Villa et al. [24] explored the placebo effect within the context of human augmentation technologies, demonstrating a sustained belief of improvement after using a sham augmentation system and an increased willingness to take risks associated with heightened expectancy.

Previous work showed that placebo effects exist when manipulating network latencies, showing positive effects when participants believe in playing a computer game with a low network latency [9]. Inspired by this research, we investigate the users’ performance under placebic refresh rates compared to their actual refresh rates (see Figure 1). Participants were assigned to undergo a target selection reaction time task using a computer display that displayed refresh rates at 60 Hz or 240 Hz. Throughout four rounds, participants completed a reaction time task while being primed with a prior current refresh rate in use. However, the participants used either the true described refresh rate or a placebic stated refresh rate throughout the reaction time task. The results show that participants expected improved reaction time before experimenting using an allegedly high refresh rate. Interestingly, the participants still believed in their improved reaction time when using a low refresh rate in reality while using the placebo improvement narrative. However, no objective improvement in reaction times was found. Our research shows that the narrative of using technologies that improve human capabilities influences the assessment of users during studies, thus having wide implications on how the HCI community must deal with the assessment of technologies in the future.




Figure 2: Description of the target selection reaction time task. (a): A button appears on which the participants must click as fast and accurately as possible. (b): A new button appears on the screen, and the old button disappears.

2 METHODOLOGY

We conducted a within-subjects study design in a reaction time task to investigate the following hypotheses (see Table 1) concerning different Frames per Second (FPS) for display refresh rates.

- H1:** Participants who will play with a true FPS of 60 Hz expect to perform better when told they will play with 240 Hz before playing the reaction time task.
- H2:** Participants who will play with a true FPS of 240 Hz expect to perform worse when told they will play with 60 Hz before playing the reaction time task.
- H3:** After playing with true 60 FPS, participants rated their performance better when told to play with 240 Hz.
- H4:** After playing with true 240 FPS, participants rated their performance worse when told to play with 60 Hz.
- H5:** Players' reaction times differ between 60 and 240 true frame rate.

2.1 Task & Experimental Setup

We describe the target selection reaction time task in the following. The reaction time task is a target selection task. For the reaction time task, we designed a circular button that, upon clicking, relocates to a new position on the screen. To ensure the new location was not entirely randomized and comparable between the conditions and the participants, we implemented criteria dictating that the button's new position should be within a specific distance from its previous location (see Figure 2). We opted for four different distances (i.e., 50px, 100px, 150px, 200px), randomly presenting the button at these distances in each round, with each distance occurring precisely five times. After clicking the displayed button, the next button appeared without delay in the defined vicinity of the previous button. Participants were required to complete the task over four rounds, each involving 20 button clicks. The participants conducted 80 trials throughout the study. Participants were primed with four narratives of a refresh rate overall before playing each round (see Table 1). The button's size remained consistent, with 100px across all rounds. The task completion time was recorded for each click. The participants were instructed to click the button as fast and as accurately as possible. We used a Lenovo Y25-30

display that can run 60 Hz and 240 Hz using a computer that can render frames with the intended refresh rate for each condition. We verified the refresh rate correctly before starting each condition. We used a Razer Viper Ultimate with a polling rate of 1000 Hz. We set the mouse to 200 dpi for each participant.

2.2 Independent Variables

Our experimental design consists of two independent variables. We define the *stated refresh rate* as the refresh rate manipulated by the experimenter through a narrative. However, the real displayed refresh rate was a different one, i.e., 60 Hz when the narrative was 240 Hz and 240 Hz when the stated refresh rate was 60 Hz. In other words, the placebo refresh rate is not the actual refresh rate but a refresh rate that acts as a placebo. Similarly, we define the *true refresh rate* with the rates 60 Hz and 240 Hz. Conditions with the same *stated refresh rate* and *true refresh rate* act as a baseline condition.

2.3 Dependent Variables

We measure the reaction times between the appearance of the target item and the time the participant takes to click on the button. We calculate the expected performance rating using a five-point Likert scale (see Table 2) before and after playing each condition.

2.4 Procedure

This study employed a full factorial within-subjects design, conducting individual tests for each participant in four conditions. Participants were seated with a computer on the table and provided with an iPad equipped with an Apple pencil. An iPad presented an informed consent form in PDF format for participants to read and sign. Following this, participants were instructed to conduct one test trial of the reaction time task, which appeared on the entire screen and began with text outlining the task, algorithm, and a brief description of refresh rates. This trial aimed to acquaint participants with the task and allow them to request clarification from the instructor. Subsequently, participants were introduced to two question forms on the iPad. The first form featured Likert scale questions (see Table 2) about performance expectations for each

ID	Question
Q1	I am using the 60 Hz refresh rate, I think I will perform faster than the 240 Hz refresh rate.
Q2	I am using the 240 Hz refresh rate, I think I will perform faster than the 60 Hz refresh rate.
Q3	I used the 60 Hz refresh rate, I think I performed faster than the 240 Hz refresh rate.
Q4	I used the 240 Hz refresh rate, I think I performed faster than the 60 Hz refresh rate.

Table 2: Likert items used to assess the user expectations before (Q1 & Q2) and after (Q3 & Q4) each reaction time task trial.

refresh rate. Participants completed questions Q1 and Q2 before the respective condition and Q3 and Q4 after conducting the reaction time task. This ensures that the participants believe the narrative before the conditions, thus manipulating their expectations before each condition. While participants were informed of the refresh rate before each condition, they were unaware if they were playing with the stated or true refresh rate. Participants were then directed to select the refresh rate announced by the instructor for the reaction time test of that condition. After selection, the screen briefly turned black to simulate the system adjusting to the refresh rate. A 'click' button appeared, initiating the reaction time test. Afterward, participants completed a second question form on the iPad, evaluating their expected performance after the condition. This process was repeated three more times, resulting in four repetitions. The conditions were counterbalanced. Upon completion of the conditions, participants provided demographic information and details about their prior experience with refresh rates. We informed the participants about the experiment's deceptions and allowed them to withdraw their data from the experiment. The institute's review board gave ethical approval.

2.5 Participants

Overall, 18 participants were recruited for the study (nine female, eight male, and one preferred not to disclose their gender). The average age of the participants was 22.83 years ($SD = 2.90$). Eight participants were aware of the functionality of display refresh rates. Five participants heard about refresh rates but have yet to learn about the exact functionality. Five participants never heard of refresh rates before. When asked about their own used refresh rates, ten participants stated that they use a 60Hz refresh rate display, and three said that they use a 120 Hz refresh rate display. Two participants stated that they use a 144 Hz display. Two participants stated that they use a 165 Hz monitor. Only one participant said they usually use a 240 Hz refresh rate monitor.

2.6 Data Analysis & Results

For H1 to H4, we tested the difference in performance assessment between the stated 60 Hz and the stated 240 Hz by the Wilcoxon signed-rank test. We corrected the alpha levels using Bonferroni correction for multiple tests. We calculated the effect size using the signed-r statistic. For H5, we calculated a regression that tested whether different stated and true refresh rate levels predict reaction times.

2.7 User Expectations

Figure 3 displays expected and evaluated performances by true and stated refresh rates. Participants who played with 60 Hz expected to

perform significantly better when told they would play with 240 Hz than with 60 Hz, $Z = 32.5$, $p < .001$, $r = 0.88$. When participants played with 240 Hz they expected to perform significantly better when they were told that they played with 240 Hz than with 60 Hz, $Z = 3.00$, $p < .001$, $r = 0.99$.

After playing with 60 Hz, participants rated their performance significantly better when they had been told to play with 240 Hz rather than 60 Hz, $Z = 42.50$, $p < .001$, $r = 0.86$. After playing with 240 Hz, participants rated their performance significantly better when they were told to play with 240 Hz than 60 Hz, $Z = 16.50$, $p < .001$, $r = 0.92$.

2.8 Reaction Times

Figure 3 displays a descriptive plot of reaction times in different true and placebo refresh rates. However, placebo and true refresh rates did not significantly affect reaction times, $F(4, 157) = 0.82$, $p = .51$.

3 DISCUSSION

We conducted a study in which participants were presented with various narratives about display refresh rates while performing a reaction time task. Without the participants' knowledge, the refresh rates they were told differed from those used in reality. Although there were noteworthy variations in participants' anticipated ratings before and after the task, there was no significant distinction in reaction times between the stated and true refresh rates. In addition, objective reaction times did not differ between 60 Hz and 240 Hz true refresh rates. Although placebo effects can improve the user experience without providing a functional system, it is still being determined which factors study narratives must include to elicit placebo effects and for how long they are functional. This has severe implications in safety-critical scenarios, where users rely on improvements through interactive systems that may not contribute to safety. We discuss the impact of our results and the need to control for placebo effects in the following.

3.1 User Expectations Manipulate Self-Assessments Prior Interaction

Our results indicate that participants anticipated performance in line with the stated refresh rate before engaging in the reaction time test. Participants who were informed that they would be using a refresh rate of 240 Hz rated their performance higher than when using a refresh rate of 60 Hz. Conversely, participants rated their performance lower when informed they would be using 60 Hz instead of 240 Hz. This pattern was consistent for both the stated and true refresh rates. The statistical significance of the ratings before

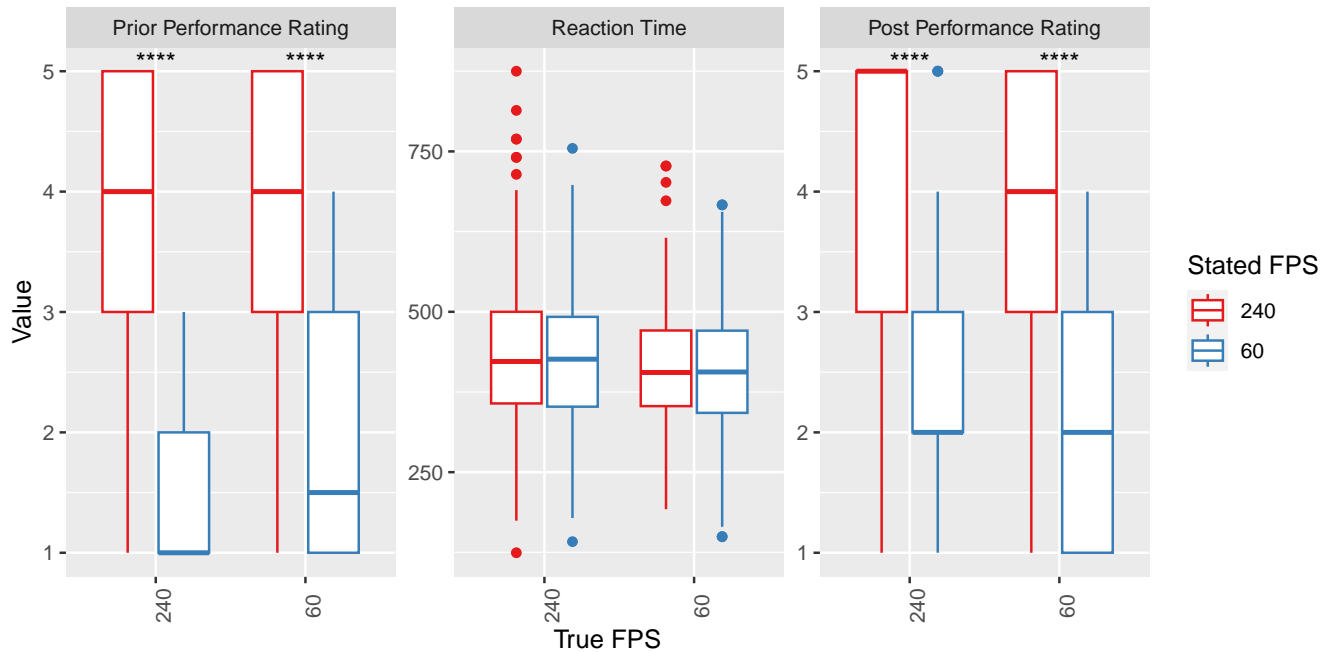


Figure 3: Aggregated user expectations before and post-interaction. Our narrative manipulated the user expectations towards the used refresh rates. Allegedly used high refresh rates elicited high user expectations before and after interaction, even when lower refresh rates were used. In contrast, a narrative of low refresh rates lead to lower performance expectations even when high refresh rates were used. No significant difference was found for objective reaction time measures.

interaction supports the confirmation of **H1** and **H2**. Our findings align with previous research examining the impact of artificial intelligence priming on anticipated user performance [12, 13] and human augmentation [24]. These studies suggest that when users are unaware of a system’s improved functionality that is not easily perceptible, their expectations tend to increase. Consequently, these heightened expectations may bias subjective scores when users are unaware of an improvement. Therefore, assessing subjective user performance ratings before interaction indicates user expectations and should be approached cautiously, considering the potential for placebo effects.

3.2 User Expectations Manipulate Self-Assessments Post Interaction

Moreover, we assessed the performance rating following the interaction with the reaction task game. Similar to the prior subjective performance ratings, the subjective performance scores exhibited significant differences even when participants engaged with a stated refresh rate that differed from the true refresh rate. This suggests that participants continued to hold onto the performance improvement or degradation post-interaction narrative. As a result, we substantiate the validity of **H3** and **H4**. This aligns with prior research demonstrating that voice assistants, whether portrayed as benevolent or malevolent when neutral, contribute to biased participant perceptions after interaction [19], indicating sustained effects post-interaction [23].

3.3 Objective User Performance

We assessed the performance in the reaction task for each condition, which was measured in time per click trial. Our statistical analysis did not reveal a significant effect between the conditions. Similar findings were reported in other studies, where no direct differences in objective performance were observed [5, 13]. Thus, we were unable to support **H5**. Interestingly, in contrast, another study identified significant differences in gaming performance when presenting phantom latencies as network latency in shooter games [9] or when using sham-AIs combined with nocebo conditions [12]. Potential placebo effects in objective game performance may be task-dependent and influenced by specific task-related designs or placebo display configurations. For example, changing the task design to include moving elements while using different refresh rates may change the objective task performance [26]. Future research will investigate the design parameters of placebo indicators and their impact on task performance.

3.4 Controlling for Placebo Effects in User-Centered Studies

Our findings highlight the significance of evaluating the user’s subjective ratings before and after a study condition to determine the susceptibility of a system with latent functionalities to placebo effects. Examining the performance expectations before and after interaction and comparing scores between stated and actual refresh

rates effectively indicate the system's vulnerability to placebo effects. This approach aligns with the recommendations from prior research [2], emphasizing the importance of assessing user expectations before and after each experimental condition to mitigate biases stemming from user expectations in each condition. Thus, studies should include a placebo condition to test from biased expectations based on study narratives.

4 CONCLUSION

Display manufacturers provide users with various display options featuring different refresh rates. Users often opt for displays with higher refresh rates based on their intended use, expecting enhanced performance in their tasks. In our study, we showed the participants' expectations to be manipulated according to the stated refresh rate, even when dealing with placebo-stated refresh rates rather than the true refresh rate employed in reality. While our experiment successfully manipulated the participant's expectations when playing with different refresh rates, the reaction time did not show a significant effect. Nevertheless, our research demonstrates an approach for investigating narratives that claim to enhance or reduce perceived system functionalities, affecting user performance ratings. The manipulation of user expectations and the active control of other subjective measures in future HCI research poses a challenge for future research.

REFERENCES

- [1] Henry K. Beecher. 1955. The powerful placebo. *Journal of the American Medical Association* 159, 17 (12 1955), 1602–1606. <https://doi.org/10.1001/jama.1955.02960340022006> arXiv:https://jamanetwork.com/journals/jama/articlepdf/303530/jama_159_17_006.pdf
- [2] Walter Boot, Daniel Simons, Cary Stothart, and Cassie Stutts Berry. 2013. The Pervasive Problem With Placebos in Psychology Why Active Control Groups Are Not Sufficient to Rule Out Placebo Effects. *Perspectives on Psychological Science* 8 (07 2013), 445–454. <https://doi.org/10.1177/1745691613491271>
- [3] Michael Correll. 2020. What Do We Actually Learn from Evaluations in the "Heroic Era" of Visualization? arXiv:2008.11250 [cs.HC]
- [4] Alena Denisova and Paul Cairns. 2015. The placebo effect in digital games: Phantom perception of adaptive artificial intelligence. In *Proceedings of the 2015 annual symposium on computer-human interaction in play*. 23–33.
- [5] Alena Denisova and Paul Cairns. 2019. Player experience and deceptive expectations of difficulty adaptation in digital games. *Entertainment Computing* 29 (2019), 56 – 68. <https://doi.org/10.1016/j.entcom.2018.12.001>
- [6] Alena Denisova and Elliott Cook. 2019. Power-Ups in Digital Games: The Rewarding Effect of Phantom Game Elements on Player Experience. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play* (Barcelona, Spain) (CHI PLAY '19). Association for Computing Machinery, New York, NY, USA, 161–168. <https://doi.org/10.1145/3311350.3347173>
- [7] Paul Enck, Ulrike Bingel, Manfred Schedlowski, and Winfried Rief. 2013. The placebo response in medicine: minimize, maximize or personalize? *Nature reviews Drug discovery* 12, 3 (2013), 191–204.
- [8] Paul Enck and Sibylle Klosterhalfen. 2019. Does sex/gender play a role in placebo and nocebo effects? Conflicting evidence from clinical trials and experimental studies. *Frontiers in neuroscience* 13 (2019), 160.
- [9] David Halbhuber, Maximilian Schlenczek, Johanna Bogon, and Niels Henze. 2022. Better be quiet about it! The Effects of Phantom Latency on Experienced First-Person Shooter Players. In *Proceedings of the 21st International Conference on Mobile and Ubiquitous Multimedia* (Lisbon, Portugal) (MUM '22). Association for Computing Machinery, New York, NY, USA, 172–181. <https://doi.org/10.1145/3568444.3568448>
- [10] Yan Jin, Jaehong Kim, and Jang Jin Yoo. 2023. Study on reaction time depending on display parameters of gaming displays. *Journal of the Society for Information Display* 31, 8 (2023), 511–521. <https://doi.org/10.1002/jsid.1198> arXiv:<https://sid.onlinelibrary.wiley.com/doi/pdf/10.1002/jsid.1198>
- [11] Ted J Kaptchuk. 1998. Powerful placebo: the dark side of the randomised controlled trial. *The lancet* 351, 9117 (1998), 1722–1725. [https://doi.org/10.1016/S0140-6736\(97\)10111-8](https://doi.org/10.1016/S0140-6736(97)10111-8)
- [12] Agnes M. Kloft, Robin Welsch, Thomas Kosch, and Steeven Villa. 2023. "AI enhances our performance, I have no doubt this one will do the same": The Placebo Effect Is Robust to Negative Descriptions of AI. <https://doi.org/10.48550/arXiv.2309.16606> arXiv:2309.16606 [cs.HC]
- [13] Thomas Kosch, Robin Welsch, Lewis Chuang, and Albrecht Schmidt. 2023. The Placebo Effect of Artificial Intelligence in Human-Computer Interaction. *ACM Transaction Computer-Human Interaction* 29, 6, Article 56 (jan 2023), 32 pages. <https://doi.org/10.1145/3529225>
- [14] Louis Lasagna, Frederick Mosteller, John M von Felsinger, and Henry K Beecher. 1954. A study of the placebo response. *The American journal of medicine* 16, 6 (1954), 770–779.
- [15] Curtis E Margo. 1999. The Placebo Effect. *Survey of Ophthalmology* 44, 1 (1999), 31–44. [https://doi.org/10.1016/S0039-6257\(99\)00060-0](https://doi.org/10.1016/S0039-6257(99)00060-0)
- [16] Guy Montgomery and Irving Kirsch. 1996. Mechanisms of placebo pain reduction: an empirical investigation. *Psychological science* 7, 3 (1996), 174–176.
- [17] Koshiro Murakami, Kazuya Miyashita, and Hideo Miyachi. 2021. A Study on the Relationship Between Refresh-Rate of Display and Reaction Time of eSports. In *Advances on P2P, Parallel, Grid, Cloud and Internet Computing: Proceedings of the 15th International Conference on P2P, Parallel, Grid, Cloud and Internet Computing (3PGCIC-2020) 15*. Springer, 339–347.
- [18] M Ojanen et al. 1994. Can the true effects of exercise on psychological variables be separated from placebo effects? *International Journal of Sport Psychology* 25, 1 (1994), 63–80.
- [19] Pat Pataranutaporn, Ruby Liu, Ed Finn, and Pattie Maes. 2023. Influencing human-AI interaction by priming beliefs about AI can increase perceived trustworthiness, empathy and effectiveness. *Nature Machine Intelligence* (2023), 1–11. <https://doi.org/10.1038/s42256-023-00720-7>
- [20] Manfred Schedlowski, Paul Enck, Winfried Rief, and Ulrike Bingel. 2015. Neurobio-behavioral mechanisms of placebo and nocebo responses: implications for clinical trials and clinical practice. *Pharmacological reviews* 67, 3 (2015), 697–730.
- [21] Steve Stewart-Williams and John Podd. 2004. The placebo effect: dissolving the expectancy versus conditioning debate. *Psychological bulletin* 130, 2 (2004), 324.
- [22] Kristen Vaccaro, Dylan Huang, Motahhare Eslami, Christian Sandvig, Kevin Hamilton, and Karrie Karahalios. 2018. The illusion of control: Placebo effects of control settings. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [23] Lucia Vicente and Helena Matute. 2023. The Inherited Bias Effect: the propagation of artificial intelligence biases to human decisions. (2023).
- [24] Steeven Villa, Thomas Kosch, Felix Grelka, Albrecht Schmidt, and Robin Welsch. 2023. The Placebo Effect of Human Augmentation: Anticipating Cognitive Augmentation Increases Risk-Taking Behavior. *Computers in Human Behavior* (2023), 107787. <https://doi.org/10.1016/j.chb.2023.107787>
- [25] Steeven Villa, Robin Welsch, Alena Denisova, and Thomas Kosch. 2024. Evaluating Interactive AI: Understanding and Controlling Placebo Effects in Human-AI Interaction. In *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). ACM, New York, NY, USA. <https://doi.org/10.1145/3613905.3636304>
- [26] Jialin Wang, Rongkai Shi, Wenxuan Zheng, Weijie Xie, Dominic Kao, and Haining Liang. 2023. Effect of Frame Rate on User Experience, Performance, and Simulator Sickness in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 29, 5 (2023), 2478–2488. <https://doi.org/10.1109/TVCG.2023.3247057>