

HOW DO PEOPLE REGARD AI-CREATED CONTENT?

By Yunhao Zhang and Renée Richardson Gosline

IN THIS BRIEF

- *To study how generative AI is perceived by consumers, we tested four separate creator paradigms: professional human writer; AI only; professional human-created content edited by AI; and AI-generated content edited by professional humans.*
- *We recruited and randomly assigned participants to rate the content's quality. In the baseline condition, participants were completely unaware of the content-generation paradigm they were rating.*
- *Our results show that when participants don't know how the content is produced, they perceive content created solely by generative AI and content edited by AI, to be of higher quality than content created by either humans only or edited by humans using AI content as input.*
- *However, when study participants were informed of the content producers, the perceived quality gap was reduced. It turns out that preference reflects a bias driven primarily by human favoritism, not by AI aversion.*

Generative artificial intelligence (AI) programs such as ChatGPT use large language models (LLMs) to generate new content based on existing information taken from the internet. These programs have spurred a great deal of excitement. They also raise the question of how AI will affect business and industry.

Recent research has shown that generative AI can enhance labor productivity in both customer communications (Brynjolfsson et al., 2023) and essay writing (Noy & Zhang, 2023). But to date, the research has not yet thoroughly explored how consumers perceive the content generated by LLMs. Nor has the research investigated how people perceive the quality of content produced by human-AI interactions.

To explore these issues, our research compared the quality of persuasive content—a mix of advertising content for consumer products and marketing campaigns for changing behaviors. We generated this content using four distinct paradigms:

- **Human expert only:** Content is created by human content creators employed by a leading consulting firm. AI does not participate.
- **AI only:** Content is created by ChatGPT. Human experts do not participate.
- **Augmented human editor:** ChatGPT generates the content's first draft, and then the human expert either edits the content or uses it for inspiration. The human makes the final decision on the output.
- **Augmented AI editor:** The human expert generates the content's first draft, and then ChatGPT either edits their

content or uses it for inspiration. The software makes the final decision on the output.

The experiment also allowed us to explore the thorny question of whether humans or AI should be allowed to make the final decision on content creation (McKendrick & Thurai, 2022).

THE CONTENT

To produce the study, we first selected products for the ads and actions for the campaigns. For the ads, we selected five products from a retail website: an air fryer, video projector, electric bicycle, emergency kit, and reusable beverage mug. For the campaigns, we selected five goals we deemed uncontroversial: stop racism, get more physical exercise, wash hands more often, eat less junk food, and recycle reusable materials.

We then enlisted 10 professional content creators, all employed by a leading consulting firm, to generate ads for the products and persuasive content for the campaigns. To ensure that the content was of a high quality, we selected creators who had experience with writing both ad copy for products and campaign messages for nonprofit humanitarian organizations.

The 10 human content creators were then instructed to complete two tasks: produce an ad for a product and a set of persuasive content for a campaign. They were also instructed to do the work without help from AI, and to keep both content types short—no more than 100 words—the approximate amount of text that can be effectively communicated in 30 seconds. This content constituted the human-only category.

To create the AI-only content, we presented ChatGPT with prompts nearly identical to those given the human content creators. ChatGPT then created text of a similar length for the 10 products and campaigns.

To develop the Augmented Human Editor content, we showed the AI-generated content to the human content creators. They were told to use the AI content either as inspiration for their own ideas or as a first draft they could

edit. In the end, this resulted in 40 pieces of content from the 10 human content creators: 20 created on their own, plus 20 pieces they created based on the AI-generated content.

For the fourth category, Augmented AI Editor, we fed ChatGPT the content created by the human content creators. We then gave the AI software the same instructions as the humans, namely, that the AI could either edit the human-generated content or use it as inspiration for a new draft. AI made the final decision, working with the content created by humans.

THE EXPERIMENT

We recruited participants for our experiment from a research panel platform, CloudResearch Connect. We had a total of 1,201 participants, split evenly between men and women, and with a median age of 38. Each participant was paid a nominal fee of \$1.50 and informed that completing the survey would take about 10 minutes.

We then randomly assigned each participant to one of three conditions:

- **Baseline:** These participants evaluated content without any knowledge of the content-creation paradigm. That is, they judged the content solely on its textual output.
- **Partially informed:** Participants were briefed about the four content-creation paradigms, but they were not told how each piece of content was created.
- **Informed:** Participants were briefed about the four content-creation paradigms, and they were also told how each piece of content was generated.

All participants were first asked to evaluate the quality of the five advertising campaigns, and to do so along three measures. One, imagining themselves the product's seller, how satisfied were they—on a scale of 1 to 7—with the content? Two, still imagining themselves the product's seller, how much would they be willing to pay—this time, on a scale of \$1 to \$1,000—to use the content in their ad? And three, to what extent—again, on a scale of 1 to 7—were they

interested in learning more about the product?

Participants were then shown pictures of the products, along with a piece of content generated by one of the four content-creation paradigms. “Informed” participants were also told which paradigm had been used to create the content.

Participants were asked to rate the quality of persuasive content for the five campaigns (“stop racism,” “get more exercise,” etc.). They were again given three questions, the first two being the same as those they’d been asked for the product ad copy. Only the third question was different: To what extent—on a scale of 1 to 7—were they convinced to take the action being advocated. Again, only the “informed” participants were told which paradigm had been used to create the content.

RESULTS

Our experiment delivered some surprising results. On average, content generated solely by ChatGPT resulted in the highest satisfaction level. Closely aligned was content generated by augmented AI.

As shown in Figure 1, content generated by a human expert alone achieved a satisfaction level similar to that of content generated by an augmented human expert. Interestingly, when AI made the final decision on the output, the content received a higher satisfaction level than content created when the final decision was made by a human expert. Surprisingly, the baseline participants—those who did not know that various content-creation paradigms were used—judged the AI-only content highest.

Similar patterns emerged in participants’ willingness to pay for content. As shown in Figure 2, participants were, on average, just as willing to pay for AI-generated content as they were for content generated by augmented AI. Content generated by humans only or by augmented humans also had similar willingness-to-pay levels. Again, baseline participants, who knew nothing about the content-creation paradigms, rated the AI-only and augmented AI content higher than the content created by either a human only or augmented human.




Figure 1. Content satisfaction ratings. The Y axis shows participants’ stated satisfaction levels, pooling all 10 content types into each paradigm; higher means more satisfied. The X axis shows the four content-creation paradigms, and the colors indicate the three participant conditions. The black bars at the top of the columns indicate the 95% confidence levels, meaning that in 95 out of 100 times, the result would fall between a bar’s lower and upper levels.

It’s particularly interesting to note the responses of those who were only partially informed about the origin of the content they were evaluating because in the real world, a person online might wonder whether they were interacting with a human or an AI agent. Generally, humans have an aversion to AI interactions. Here, participants were aware that some content was generated by AI, just as a person is in the real world, but unaware of which content that might be, again just as in the real world.

We also explored whether participants might be biased toward or against content created by AI by comparing how evaluations differed among the baseline and fully informed participants. Overall, we found that participants felt more satisfied with content when they knew it was generated by a human expert. They also were more willing to pay when they knew the content was generated by a human. However,

we did not find any effect on either the satisfaction level or willingness to pay for the other content-creation paradigms (AI only, etc.). In other words, we did not find evidence for AI aversion. Instead, we have evidence of human favoritism.




Figure 2. Willingness-to-pay ratings. The Y axis shows a logarithm of participants' stated willingness to pay for the content, pooling all 10 content types into each paradigm; higher means a greater stated willingness. The X axis shows the four content-creation paradigms, and the colors indicate the three participant conditions. The black bars at the top of the columns indicate the 95% confidence levels, meaning that in 95 out of 100 times, the result would fall between a bar's lower and upper levels.

CONCLUSIONS AND IMPLICATIONS

Our results suggest that generative AI can outperform human experts in generating specific types of advertising and persuasive content, and that genAI could also replace some human labor in content generation. However, we are by no means suggesting that genAI should completely displace human workers. Human oversight remains especially important.

Although we chose products and campaigns we deemed harmless, AI can persuade humans on important issues such as politics. Human oversight will be needed to ensure the content created by genAI is appropriate.

Our research can serve as further evidence of genAI's ability to benefit producers and consumers by raising productivity and lowering prices (Acemoglu & Restrepo, 2020). The results also contribute to discussions of algorithmic aversion and appreciation (Dietvorst et al., 2015; Logg et al., 2019) and human favoritism (Morewedge, 2022).

We find that perception plays a big part in acceptance: simply knowing a piece of content was generated by humans increases its perceived quality. However, we do not find strong evidence of algorithm aversion in our context. This is somewhat surprising, given that previous research has shown how people display algorithm aversion in subjective tasks, such as evaluating how funny a joke is (Castelo et al., 2019). Our work suggests that the conclusions made on more conventional forms of AI may not apply to generative AI. Future research could further investigate how people perceive the performance of genAI and refine the human-in-the-loop protocol.

REPORT

Read the [full research report](#).

ABOUT THE AUTHORS

Yunhao Zhang is a postdoctoral fellow at the Psychology of Technology Institute, and the recent recipient of an MBA in Management from MIT Sloan. His work focuses on the impact of AI on individual and societal well-being.

Renée Richardson Gosline is a senior lecturer in the Management Science group of MIT's Sloan School of Management; and the head of the Human-AI Interaction research group at the MIT Initiative on the Digital Economy (IDE).

REFERENCES

- Acemoglu, D. & Restrepo, P. (2020). Robots and jobs: evidence from U.S. labor markets. *Journal of Political Economy*; vol. 128, no. 6, pp. 2188-2244. <https://www.journals.uchicago.edu/doi/10.1086/705716>
- Brynjolfsson, E., et al. (2023). Generative AI at work. Cornell University. arXiv:2304.1177v1 <https://arxiv.org/abs/2304.1177v1>
- Castelo, N., et al. (2019). Task-dependent algorithm aversion. *Journal of Marketing Research*; vol. 56, no. 5, pp. 809-825. <https://journals.sagepubcomdoi/10.1177/0022243719851788>
- Dietvorst, B.J., et al. (2015). Algorithm aversion: people erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology, General*; vol. 144, no. 1, pp. 114-126. <https://pubmed.ncbi.nlm.nih.gov/25401381/>
- Logg, J.M. & Minson, J. (2019). Algorithm appreciation: people prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*; vol. 151, no. 10, pp. 90-103. <https://www.sciencedirect.com/science/article/abs/pii/S0749597818303388>
- Morewedge, C.K. (2022). Preference for human, not algorithm aversion. *Trends in Cognitive Sciences*; vol. 26, no. 10, pp. 824-826. <https://www.sciencedirect.com/science/article/abs/pii/S1364661322001644>
- McKendrick, J. & Thurai, A. (2022). AI isn't ready to make unsupervised decisions. *Harvard Business Review*. <https://hbr.org/2022/09/ai-isnt-ready-to-make-unsupervised-decisions>
- Noy, S. & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*; vol. 381, no. 6654, pp. 187-192. <https://www.science.org/doi/10.1126/science.adh2586>



MIT
INITIATIVE ON THE
DIGITAL ECONOMY

MIT Initiative on the Digital Economy

MIT Sloan School of Management
245 First St, Room E94-1521
Cambridge, MA 02142-1347
ide.mit.edu

Our Mission: The MIT IDE is solely focused on the digital economy. We conduct groundbreaking research, convene the brightest minds, promote dialogue, expand knowledge and awareness, and implement solutions that provide critical, actionable insight for people, businesses, and government. We are solving the most pressing issues of the second machine age, such as defining the future of work in this time of unprecedented disruptive digital transformation.

Contact Us: David Verrill, Executive Director
MIT Initiative on the Digital Economy
MIT Sloan School of Management
245 First St, Room E94-1521
Cambridge, MA 02142-1347
617-452-3216
dverrill@mit.edu

Become a Sponsor: The generous support of individuals, foundations, and corporations help to fuel cutting-edge research by MIT faculty and graduate students, and enables new faculty hiring, curriculum development, events, and fellowships.

[View all of our sponsors](#)