



OPEN

Attributions toward artificial agents in a modified Moral Turing Test

Eyal Aharoni^{1,2,3✉}, Sharlene Fernandes¹, Daniel J. Brady¹, Caelan Alexander¹, Michael Criner¹, Kara Queen¹, Javier Rando⁴, Eddy Nahmias^{2,3} & Victor Crespo⁵

Advances in artificial intelligence (AI) raise important questions about whether people view moral evaluations by AI systems similarly to human-generated moral evaluations. We conducted a modified Moral Turing Test (m-MTT), inspired by Allen et al. (*Exp Theor Artif Intell* 352:24–28, 2004) proposal, by asking people to distinguish real human moral evaluations from those made by a popular advanced AI language model: GPT-4. A representative sample of 299 U.S. adults first rated the quality of moral evaluations when blinded to their source. Remarkably, they rated the AI's moral reasoning as superior in quality to humans' along almost all dimensions, including virtuousness, intelligence, and trustworthiness, consistent with passing what Allen and colleagues call the comparative MTT. Next, when tasked with identifying the source of each evaluation (human or computer), people performed significantly above chance levels. Although the AI did not pass this test, this was not because of its inferior moral reasoning but, potentially, its perceived superiority, among other possible explanations. The emergence of language models capable of producing moral responses perceived as superior in quality to humans' raises concerns that people may uncritically accept potentially harmful moral guidance from AI. This possibility highlights the need for safeguards around generative language models in matters of morality.

Moral reasoning is regarded among the most sophisticated and unique of human faculties. Ordinary adults, and even young children, draw universal and context-sensitive distinctions between right and wrong, and they justify those distinctions on the basis of explicit or implicit reasons, values, and principles^{1–3}. Yet, despite centuries of scholarship on this subject, scholars continue to debate basic questions, such as what criteria constitute moral intelligence^{4,5} and whether being human is one of them^{6–8}.

Recent advances in generative language modeling have drawn new attention to this debate, following the widespread release of large language models (LLMs) such as Bard, LLaMA, Claude, and ChatGPT. At the time of this writing, ChatGPT, notably, has already been visited over a billion times collectively by over 150 million users⁹. The natural language capabilities of these LLMs are unprecedented among artificial intelligence (AI) and other communication technologies. One is even reported to have passed portions of the U.S. Medical Licensing Exam¹⁰. Scholars have been quick to argue that these technologies are not exhibiting general intelligence^{11–17}. But ordinary people might treat them as if they are. As AI language models become increasingly relied upon, people might use them to search for advice on various topics, including medical^{18,19}, legal²⁰, and even moral questions (e.g., "Is it wrong to lock my 4 year-old in her room during a timeout?"). To this point, recent research has reported that, in at least some limited contexts, LLMs can make human-like moral judgments²¹ and can accurately classify social norm violations^{22,23}. Although these conclusions have been challenged (e.g.^{24,25}), their mere plausibility justifies further research on how ordinary people might perceive and interact with LLMs in moral domains.

The question of AI moral competence might also be relevant to inquiries that are not explicitly moral but have morally consequential defaults, such as the environmental impact of responses to the prompt: "List recommendations for a new car." At the organizational level, businesses might be incentivized to use AI language models to influence customers, anticipate potential stakeholder attitudes, or even help the business greenwash public statements about unethical practices. Thus, the perceived moral intelligence of LLMs has the potential to significantly impact human welfare.

¹Department of Psychology, Georgia State University, Atlanta, GA, USA. ²Department of Philosophy, Georgia State University, Atlanta, GA, USA. ³Neuroscience Institute, Georgia State University, Atlanta, GA, USA. ⁴ETH Zurich, Zurich, Switzerland. ⁵Department of Philosophy, Duke University, Durham, NC, USA. ✉email: eaharoni@gsu.edu

A growing body of research investigates moral attributions toward non-LLM AI systems. Among the findings on this topic are conflicting attitudes about the moral properties of these systems. For example, people generally approve of self-driving cars making utilitarian moral decisions (e.g., killing one person to save five), but they would not want to ride in one²⁶. However, LLMs are distinct from other AI systems because they are designed to interface with humans in more diversified, social, and intellectual ways. A charitable viewpoint characterizes LLM technology as capable of understanding information in a scientifically and socially meaningful sense (e.g.,²⁷). In a more skeptical vein, LLMs produce the appearance of human understanding of a topic, presumably without having genuine experience with or understanding of that topic (see²⁸). This attribute, combined with a directive to provide a response to each prompt, often produces outputs that could meet the technical definition of *bullshit*, in Harry Frankfurt's sense of persuasion without any regard for, or even understanding of, what is true or false²⁹. In this view, if LLMs convincingly bullshit about morality, people may accept inaccurate or unhelpful moral explanations and advice from LLMs. Such convincing bullshit has already emerged in morally relevant domains, as illustrated by emerging cases of lawyers submitting court briefs that cited fictitious legal cases generated by LLMs^{30,31}.

This unique property of LLM technology could evoke moral attributions toward them that are distinct from those made toward other AI technologies, but systematic tests of moral attributions toward LLMs are sorely lacking. Very recent research suggests that human moral judgments are similarly influenced by LLM advice when told that the advice comes from an LLM or from a (ostensibly human) moral advisor³². So to the extent that people consult LLMs in moral matters, understanding ordinary people's attributions toward LLM moral discourse is consequential (see also³³).

In the current study, we constructed the first LLM-based (modified) Moral Turing Test (MTT³⁴). The MTT is a variation on a standard Turing Test, which a computer passes if human judges cannot distinguish the computer's responses from human responses to the judges' questions³⁵. Critically, in the MTT, as proposed in a highly cited article by Allen et al., participants specifically attempt to distinguish human *moral* discourse from similar discourse made by an AI³⁴. The authors suggested that if human judges could not distinguish between the two potential sources (AI vs. human), then the artificial moral agent would qualify for some definitions of moral intelligence—that is, it would pass the test. One drawback of this approach (and with the original Turing Test) is that the AI could technically fail the test by producing responses that, compared to a human's, are perceived as *superior* in quality. In response, Allen and colleagues proposed the comparative MTT (cMTT), in which, rather than identifying the source of the responses, the judge evaluates which set of responses is more moral. The AI passes the cMTT if its moral quality is identified as equal to or greater than that of the human respondent³⁴. To our knowledge, the MTT and cMTT have never been systematically tested.

We conducted a modified version of both an MTT and cMTT. In our MTT, we asked a representative sample of U.S. adults to judge the source of ten pairs of moral evaluative passages. The passages were human and LLM responses to a question about why a specific description of a human action is wrong or not wrong (see³⁶). The actions were pre-classified as wrong either by moral standards or by social convention. Before instructing them on this test, we asked the same participants to make several judgments about the content of each moral evaluation (e.g., "Which responder is more morally virtuous?", "Which response do you agree with more?") under the presumption that they were all human-authored (see³⁴). Presenting this modified cMTT first and withholding information about the computer source enabled us to obtain an unbiased estimate of its perceived quality, without contamination by judges' potential preconceptions about computer agents or about the purpose of the task.

We did not attempt to replicate the traditional Turing Test format which has been the subject of varied interpretations which extend beyond the scope of this project^{37–41}. Our goal was to test people's attributions under the most simplified and controlled conditions that still operate in real-world, morally-relevant encounters in which typical end-users have only brief, one-shot encounters with the AI agent, such as the use of LLMs as an Internet search engine. Therefore, our test did not support an interactive dialogue between participants and agents. In addition, rather than programming or prompting the AI to use imitation or roleplay to deceive the human judge into believing its output was human-generated (as a traditional Turing Test might), our test used the LLM's default parameters. Whether these LLMs are explicitly trained to imitate real human communication is often unknown for reasons of propriety. But given enough training data, they can still emulate natural human language in some meaningful ways. We tested people's perception of the LLM's response under these default conditions in order to generalize our results to the contexts in which most people use these LLMs. Using the default responses from the LLM also provides a more stringent test of the modified MTT, making it harder for people to differentiate between the two sources. With these planned deviations from the standard Turing Test format in mind, we refer to our tests as a modified Moral Turing Test (m-MTT).

We made two pre-registered predictions. Our first hypothesis was that participants would rate the (masked) computer-authored moral evaluations as significantly better in quality than the (masked) human-authored evaluations. Our second hypothesis was that participants would fail to correctly distinguish between the sources of the human- and computer-generated moral evaluations, performing no better than chance classification accuracy (~50%). If people cannot readily distinguish between human and computer moral evaluations, such a finding would have implications for how people are likely to engage with LLMs, how LLMs could potentially mislead human users, and what safeguards or legal regulations might be needed to minimize harmful interactions. On the other hand, if people can readily distinguish between human and LLM moral evaluations, this would raise questions about the distinguishing features of ordinary human moral reasoning, as revealed by our language. It might also indicate ways that future LLMs will have to improve to replicate certain types of human language. The results of this study will inform developers and users of LLMs of the moral language capabilities of current LLMs, their perceived ability to provide moral counsel, and their potential risks. Additionally, this study may inform scholars of moral cognition about the defining features of ordinary human moral intelligence.

Materials and methods

Participants

We administered a pre-registered Internet survey on *Prolific.com* to participants sampled to be representative of the general U.S. adult population on the dimensions of age, gender, and ethnicity. Those who completed the survey were paid \$3.00 USD for participating. To qualify for the study, participants had to be adults in the U.S., and be fluent in English. Three-hundred-sixty-three people consented to the survey. In accordance with our planned exclusion criteria, participant data were excluded from analysis for the following reasons: failure to submit a complete survey with answers to all required questions (64 participants), failure to correctly answer any one of our multiple-choice attention and comprehension questions: "What are the colors of the American flag?" (1 participant), "For this item, select 'Slightly'" (1 participant), and "Which of the following best describes the action in the previous scenario?" (11 participants) (see Supplementary Information for a note about additional exclusion criteria.) The remaining 286 participants—our final sample—self-identified as 51.0% female, 46.52% male, and 2.8% other or unanswered; 7.0% Hispanic or Latinx; 76.6% White, 13.6% Black, 6.6% Asian, and 5.5% other or unanswered (ethnic and racial categories were non-exclusive); and with a mean age of 46.76 ($SD = 16.33$). Self-reported economic class was: 17.8% lower, 49.7% lower-middle, 28.7% upper-middle, 0.7% upper. In terms of education, 32.5% had a high school diploma, 13.3% had an associate's degree, 36.4% had a bachelor's degree, 13.9% had a post-secondary degree (e.g., master's), and 3.8% were other or unanswered. Politically, the sample was mixed, with a slight preference for liberalism over conservatism, $M = 0.88$, $SD = 1.77$, on a 7-point ordinal scale from -3 ("extremely liberal") to $+3$ ("extremely conservative"). As a whole, participants were not regular users of LLM technology. They reported a low frequency of talking to an artificial language model in the past 30 days, as determined by a one-sample t -test that compared their mean to the midpoint ("somewhat") of a 5-point scale from "not at all" (1) to "extremely" (5), $t(285) = -12.31$, $p < 0.001$, $M = 2.12$, $SD = 1.21$.

A minimum sample size estimate was based on an *a priori* power analysis, using a Chi-squared goodness-of-fit test to compare the observed frequencies of human/computer labels chosen by participants with the expected frequencies, assuming that there is no difference between the human and computer proportions. Conservatively, a small to medium effect size of $w = 0.20$, assuming an alpha threshold of 0.05, a power threshold of 0.80, and 1 degree of freedom, yields a minimum sample size of 197. We sought to collect 300 to account for possible attrition and data exclusion.

Procedure

Each participant was presented with the same 10 pairs of written passages describing evaluations of social transgressions, presented in random order. The transgression scenarios were selected from Aharoni et al.'s research on moral and conventional transgressions^{42,43}. Ordinary people make reliable distinctions between moral and conventional transgressions in ways that some other human groups⁴⁴ and computers may not. Moral transgressions typically involve harm and are judged as wrong universally and independently of what authorities might claim. Conventional transgressions, by contrast, usually do not cause direct harm to a person and are perceived as wrong only by mere social convention.

The transgressions in the scenarios for which moral evaluations were provided in this study are noted in Table 3 (see Supplementary Information for full details). Each passage consisted of a textual response to the question of why a given transgression is or is not morally wrong. One passage in each pair was authored by a human being (obtained from pretesting in a university sample), and the other was authored by a computer, namely, *OpenAI's* generative language model, GPT-4 (see Supplementary Information for more information about stimulus development.) We used GPT-4⁴⁵ because it is regarded as the most powerful engine behind the widely used tool, ChatGPT^{46,47}.

The pairs of passages were presented side by side (counterbalanced for position) without identifiers (for an example, see Table 1). Participants were instructed to read each passage and then rate the relative quality of each (e.g., which statement they agreed with more, which statement they thought was more virtuous; presented in random order), by selecting "Response A", "Response B", or "no preference". These quality questions comprised the first dependent measure of the comparative MTT (block 1). During this block, they were not informed that any of the passages were computer-generated.

After rating the quality of all pairs of passages, participants were asked to guess, from a randomized multiple-choice list, which group they thought authored the passages (high school students, university undergraduates, graduate students in philosophy, or "some other source" that they could specify). The purpose of this question

Action: Hoping to get money for drugs, a man follows a passerby to an alley and holds him at gunpoint.	
A [Human]	B [Computer]
The act is wrong because he hurt an innocent man. Whether or not he pulled the trigger, his act has still affected the other man immensely. His intentions were wrong from the beginning—he robbed a hard-working person of their money just to feed his addiction	This act is wrong. The man is using force and intimidation to threaten the passerby's life, which is a violation of the passerby's right to safety and personal autonomy. Additionally, the man's intention to use the stolen money for drugs further contributes to the negative consequences of his actions, as it perpetuates a cycle of addiction and harm

Table 1. Example morally evaluative passage from (de-identified) human and computer source. See Supplementary Information for full stimuli.

was to assess participants' source beliefs during their quality judgments without priming them that the source could be an AI. We expected that without instructions suggesting otherwise, most people would assume the source was human, so if they spontaneously identify it as a computer, we could be more confident in their answer than if we primed them by including a computer answer option.

Next, they were informed that one of the responses was generated by a computer chatbot trained in human language. They were then presented with the same pairs (in random order) and asked to judge which passage was authored by the computer and which was authored by a human, comprising the second dependent measure of the modified MTT (block 2). After each judgment, they were asked to rate their confidence in that judgment and to comment in their own words on why they thought the selected passage was computer-generated. This free text was frequency-coded (present or absent) into predefined themes by two trained raters, blind to our study hypotheses. Last, the survey also included several other questions including attention checks, belief checks, attitudes about AI, and standard demographic questions. All data were collected online from paid human participants in a single session using the *Qualtrics XM* survey platform. The study was pre-registered on the Open Science Framework at <https://doi.org/10.17605/OSF.IO/6GXD3> (see Supplementary Information for protocol deviations.) All procedures were approved by Georgia State University's Human Research Protection Program and were carried out in accordance with the Declaration of Helsinki (IRB# H21668, approved March 3, 2023). Informed consent was obtained in writing from all research participants.

Hypotheses

Hypothesis 1 In response to questions about whether the computer or human moral evaluations (before participants were informed that one was generated by computer) is a better-quality response (more virtuous, etc.), we predicted that participants would rate the computer-authored passages as significantly better in quality than the human-authored passages.

Hypothesis 2 In response to a question about which of two moral evaluations is authored by a human or computer, we predicted that participants would not perform significantly better than chance accuracy, which will cluster around 50%.

Materials

The 10 pairs of passages describing evaluations of social transgressions are shown in the Supplementary Information section. See Table 1 for an example.

First, participants rated the quality of each (masked) response pair by answering 10 questions about the quality of the response, presented in random order (e.g., "Which responder is more morally virtuous?" "Which response do you agree with more?"), each time selecting either: "Response A", "Response B", or "No preference".

To assess whether participants could correctly identify the source (human vs. computer) of the passages, the second part of the survey asked: "Which response was generated by a computer?" (Response A or B). We also assessed their confidence in that choice on a 5-point Likert-type scale from "not at all confident" (1) to "extremely confident" (5). We also included a structured question to probe participants' beliefs about how they could distinguish between the human and computer passages. It stated: "If you noticed differences between the human and computer responses, in which ways did they seem to differ? (Select all that apply.)" Answer options included: grammar, word choice, clarity, the variety of reasons provided, response length, emotionality, rationality, and other (please specify) (see Supplementary Information for further details.)

Stimulus development

The 10 human-generated moral passages were collected from a sample of university undergraduates in an introductory philosophy course ($N = 68$). Respondents were presented with the 10 scenarios involving the transgressions as previously described and provided with the following prompt: "Please explain in a few sentences why this act is or is not wrong in your opinion." Responses were pre-rated by two independent, trained raters to select those that were of high quality. Quality was determined by assessing (1) whether the response met a minimum character count of 150 characters, (2) whether it answered why the scenario is wrong or not wrong, (3) whether it referred to at least one moral reason for the judgment, (4) how much the rater agreed with the response on a 5-point Likert scale from 1 = Strongly Agree to 5 = Strongly Disagree, and (5) the rater's subjective rating of the quality of the response (high/medium/low). The highest-ranking passages were selected in order to maximize the ability to falsify our source attribution hypothesis. Ties were broken by a third rater, blinded to the source of each rating (rater 1 or 2). The final set of passages had an average of 324 characters (range: 192–560).

GPT-4 was given the same 10 scenarios and prompts. Originally, GPT's responses were substantially longer than the human-generated responses, which could have biased source judgments and quality ratings. Therefore, we provided it with the following prompt: "Please answer the question as directly as possible. Please limit your answer to a maximum of 600 characters including spaces." The average character count for GPT's final responses was 405 characters (range = 331–467). These responses were still slightly longer than the human-generated responses, but further prompting to shorten responses was not provided to GPT, to avoid excessively deviating from GPT's "natural" response. Line spacing in the survey was varied to offset apparent differences in length between passage pairs (as in Table 1).

The computer-generated (GPT-4) moral passages were not subject to sampling or pre-rating. We reasoned that GPT's responses are already the result of an extensive selection process via training on a massive corpus of human data. Therefore, the first response to each of the 10 queries (using the same question script administered to the human respondents) was retained. For both the human and AI-generated stimuli, grammatical anomalies were corrected using *Grammarly* and the discretion of the PI, when minimally necessary. In addition, self-referencing

phrases (e.g., "In my opinion...") from both sources were removed because almost all of GPT's answers started with that phrase, and we wanted participants to base their judgments mainly on the passages' content rather than their form (see Supplementary Information for more details about stimulus development.)

Results

Did participants rate the (masked) computer-authored moral passages as significantly better in quality than the (masked) human-authored moral passages?

Before revealing that some of the morally evaluative passages were computer-generated, and before introducing the source attribution test, we assessed participants' attributions of passage quality along 10 dimensions (see "Materials" section). Our goal was to test the pre-registered hypothesis that participants would choose the (masked) computer-authored evaluations (coded as -1) rather than the (masked) human-authored evaluations (coded as 1) as being significantly better in quality. Participants also had the option to state no preference (coded as 0). First, to assess suspicion about the source of the computer-generated passages, after making the quality ratings, participants were asked which group they think authored the moral responses. Only seven participants (2.5%) correctly guessed (without prompting) that some of the passages were generated by a non-human source (i.e., by AI). The vast majority ($n = 153$) reported believing that they were written by undergraduates, $\chi^2 = 162.08$, $p < 0.001$, which was true for only half of the passages.

Cumulative scores for each dimension were calculated and could potentially range from -10 to +10. A score of -10 for a quality dimension, for example, would indicate that the participant chose the masked computer response as being qualitatively better than the masked human response for all 10 passages. To assess the internal consistency of the 10 quality dimensions, an inter-item reliability analysis was conducted. The analysis yielded a Cronbach's alpha coefficient of .824, indicating good reliability. Means and standard deviations for each individual item across all 10 pairs of passages are as follows: agreement ($M = -2.10$, $SD = 2.83$), virtuous ($M = -1.17$, $SD = 2.36$), better person ($M = -1.08$, $SD = 2.39$), trustworthy ($M = -1.54$, $SD = 2.29$), intelligent ($M = -2.96$, $SD = 2.63$), fair ($M = -1.12$, $SD = 2.50$), compassionate ($M = .01$, $SD = 2.85$), biased ($M = -21$, $SD = 1.96$), rational ($M = -2.40$, $SD = 2.80$), and emotional ($M = -0.34$, $SD = 2.96$). Because our reliability statistic exceeded the planned decision criterion of .80, a grand mean was constructed from their simple mean, representing overall evaluation quality.

The overall quality variable did not follow a normal distribution, Kolmogorov-Smirnov $p < 0.001$, and could not be transformed to normal using a natural log function, K-S $p < 0.001$, so we tested our hypothesis using a (nonparametric) Wilcoxon Signed Rank test. This test revealed, as predicted, that participants' preferences for quality deviated from the null (hypothesized as 0, i.e., "no preference"), and a negative median score further indicated that they rated the computer-generated passages as significantly higher in quality than the human-authored passages (observed $Mdn = -1.40$, $z = -10.88$, $p < 0.001$). The effect size was large ($r = -0.64$).

Using the same statistical approach, we extended our hypothesis test to each of the 10 quality dimensions individually. Participants rated the computer author (as compared to the human author) as more virtuous, more intelligent, more fair, more trustworthy, a better person, and more rational (all $ps < 0.001$). They also agreed with the computer's evaluations more so than the human-generated ones ($p < 0.001$). After Bonferroni corrections for 10 tests (alpha = 0.005), neither human-nor computer-generated evaluations were rated as more emotional ($p = 0.056$), compassionate ($p = 0.88$), or biased ($p = 0.046$). See Table 2 for results.

Since participants were blinded to the source of the passages and (all but seven) reported believing they were all human-generated, these findings support the inference that participants believed the computer-generated moral evaluations were of objectively better quality than the human-generated evaluations, overall and along nearly all of our specific dimensions.

	Median	Unstandardized statistic (<i>W</i>)	<i>z</i> -score (<i>p</i>)
Intelligent	-3.0***	1947.50	-12.56 (< 0.001)
Rational	-2.0***	3522.00	-11.12 (< 0.001)
Agreement	-2.0***	4861.00	-9.97 (< 0.001)
Trustworthy	-1.0***	4890.00	-9.67 (< 0.001)
Virtuous	-1.0***	5842.50	-7.82 (< 0.001)
Better person	-1.0***	7245.50	-6.97 (< 0.001)
Fair	-1.0***	7276.00	-6.94 (< 0.001)
Biased (reverse scored)	0.0	9819.00	-1.99 (0.046)
Emotional (reverse scored)	0.0	12633.50	-1.91 (0.056)
Compassionate	0.0	14145.00	0.16 (0.877)
Overall quality (grand mean)	-1.40***	5038.50	-10.88 (< 0.001)

Table 2. Ratings of relative quality of (masked) human- and computer-generated morally evaluative passages. Scores can range from -10 to +10 and represent the degree to which participants endorsed greater quality for the computer (negative) relative to the human (positive) for the 10 passages. *W* = Wilcoxon signed-rank tests with a null hypothesized median of 0 (indicating no preference) and Bonferroni corrections (for the individual criterion tests; alpha = 0.005). *** $p < 0.001$.

Scenario	Percentage correctly choosing computer	Percentage choosing computer [95% CI]	z-score (<i>p</i>)
Tried on shirt in store (C)	82.17***	77.46–86.31	10.82 (< 0.001)
T-shirt at funeral (C)	81.12***	76.32–85.37	10.47 (< 0.001)
Robbed at gunpoint (M)	79.02***	74.06–83.47	9.76 (< 0.001)
Charged credit card (M)	77.27***	72.19–81.87	9.17 (< 0.001)
Hit pizza boy (M)	74.83***	69.59–79.62	8.34 (< 0.001)
Ate at restaurant with fingers (C)	72.03***	66.65–77.02	7.39 (< 0.001)
Man wore skirt to office (C)	65.73***	60.11–71.08	5.26 (< 0.001)
Ate in boss' office (C)	61.54***	55.82–67.06	3.84 (< 0.001)
Shot dog (M)	58.74**	52.98–64.35	2.90 (0.004)
Set fire (M)	58.04	52.27–63.67	2.66 (0.008)

Table 3. Proportion of judgments that correctly attributed the computer-generated passages to the computer. Results represent the percentage of source attribution judgments that correctly identified the computer source of the passages. One-sample binomial tests were conducted with a null hypothesized median of 50% and used Bonferroni corrections ($\alpha = 0.005$) for multiple tests. Scenarios involving moral transgressions are denoted by (M) and those involving conventional transgressions are denoted by (C). ** $p < 0.005$; *** $p < 0.001$.

Did participants correctly identify which of the moral evaluative passages was authored by a computer as opposed to a human?

After participants rated the quality of the moral passages, they were informed that one of the responses in each pair was human-generated, and the other was generated by a computer chatbot that was trained in human language. Then they were presented sequentially with the same pairs and instructed to guess which response in each pair was generated by a computer (i.e., the source attribution test). Next, participants indicated how believable it was that a computer generated some of the moral responses in the task. Participants disproportionately rated this prospect as believable, two-tailed, one-sample $t(285) = 16.59$, $p < 0.001$, $M = 3.92$, $SD = 0.94$, $SE = 0.06$, relative to the scale midpoint of 3 ("somewhat"), with nearly three-quarters of the sample (72.4%) rating it as either very believable or extremely believable. This finding lends confidence to our interpretation of the subsequent results of our source attribution test.

To test our hypothesis (H2) that participants would fail to distinguish between the human and computer sources of the passages, we first computed a mean of the 10 source attribution scores (with attributions towards computers scored as 0 and attributions towards humans scored as +1). Because the distribution of this variable was non-normal, K-S $p < 0.001$, and could not be transformed to normal using an arcsine function, K-S $p < 0.001$, a Wilcoxon signed-rank test was employed once again to evaluate overall accuracy across all 10 judgments.

Averaging across all 10 judgments, source attributions were quite accurate. Contrary to our prediction, participants were disproportionately likely to identify the computer passages as being computer-generated ($Mdn = 0.30$) relative to the null hypothesis of 0.50 (i.e., chance performance), $z = -12.51$, $p < 0.001$. The effect size was large ($r = -0.74$). Indeed, 80.1% of the participants made correct identifications more than half the time.

A finer-grained analysis was employed to extend our hypothesis test to each of the 10 accuracy judgments individually. We constructed a set of one-sample binomials, using a Bonferroni correction for multiple tests ($\alpha = 0.005$). For nine of the 10 pairs of passages (human and computer), the computer was disproportionately identified as such, consistent with our test of the overall average. As a whole, these findings indicate that the LLM used to generate moral responses did not pass our source attribution test. See Table 3 for results.

Exploratory analyses

Several analyses were conducted to try to explain how participants were able to correctly identify the source of the computer-generated moral passages. First, the analysis of passage quality reported above suggests that the higher quality of the computer-generated passages could have served as a clue to their source of authorship. In addition, we examined the frequency of responses to a set of structured questions about more superficial and structural attributes of the passages. The two most frequent responses pertained to differences in word choice (70.28%) and response length (70.28%). So, in addition to participants' judgments of passage quality, participants may have relied on these attributes to identify the source of the computer passages. The other responses were: How emotional the explanation was: 58.39%; How rational the explanation was: 48.25%; Clarity of response: 39.51%; Grammar: 37.41%; The variety of reasons provided: 39.65%; Other: 5.94%.

We also conducted a qualitative analysis to identify the most frequent types of responses to a question asking why they thought the selected passage was computer-generated. Each set of 10 responses (for each participant) was tagged to indicate the presence or absence of any of the following predefined themes, denoting that the computer passage was comparatively: formal vs. casual, direct vs. indirect, thorough vs. sparse, intelligent vs. simplistic, rational vs. emotional, lengthy vs. short, and a miscellaneous category (two additional categories—objective and multiple perspectives—were discarded before analysis because they did not conform to the planned bipolar format.) Thus, for the bipolar themes, it was possible for both poles to be present in a single participant (e.g., emotional and rational). Each pair was then summed together into a single theme yielding ordinal scores from -1 (e.g., emotional) to $+1$ (e.g., rational). All summary themes were assessed for inter-rater reliability. Only one achieved moderate reliability, namely, the rational vs. emotional theme (weighted $k = 0.65$)⁴⁸. As a result, we

evaluated the frequency distribution of that theme only and confirmed that participants disproportionately rated the computer-generated passages as more rational and less emotional than the human-generated passages, $\chi^2(4, 286) = 288.62, p < 0.001$, with 68.2% of participants noting this attribute at least once.

Moral and conventional wrongfulness judgments

The human-generated and computer-generated moral evaluations were composed of two types of predefined transgressions: moral and conventional. Moral transgressions (e.g., punching a boy in the face) were defined as those that directly violated a moral value (e.g., physically harming someone) and are universally judged as wrong. Conventional transgressions (e.g., wearing a t-shirt at a funeral), in contrast, are violations of non-moral social norms. Such violations are commonly rated as permissible under certain conditions and not *morally* wrong (e.g.,^{42,43}). To test whether the perceived quality of the computer-generated vs. human-generated passages differed by transgression type (moral or conventional), we first calculated the grand mean quality score of each quality criterion for each transgression type (scores could range from -5 to $+5$, with negatives scores indicating a preference for the computer's quality) and conducted a related-samples Wilcoxon signed-rank test. This test indicated a difference in the degree of source preference between the two passage types, $z = 5.86, p < 0.001, r = 0.35$. Medians indicated that participants' preference for the computer-generated responses was greater when responding to moral transgressions ($Mdn = -90$) vs. conventional transgressions ($Mdn = -40$). In addition, when comparing source attributions for the evaluations of the moral vs. conventional transgressions, participants were slightly more accurate in identifying the (computer) source of the conventional passages ($M = 0.27$) than the moral passages ($M = 0.30$), $z = -2.35, p = 0.02, r = -0.14$.

Beliefs, attitudes, and demographic factors

Next, we assessed participants' beliefs and attitudes about AI. Although none of the response means were extreme, all of them significantly differed from the scale midpoint of 3 ("somewhat"). Participants were relatively unlikely to ask an AI chatbot for moral advice and believed that such information should be regulated. They believed that computers are relatively unlikely to become conscious or to surpass humans in moral intelligence. See Table 4.

Finally, we examined the associations between accuracy scores and several other recorded variables, namely age, education, political ideology, and recent experience with LLMs, to search for potential explanatory or moderating evidence bearing on source attribution accuracy. However, none of these variables were correlated with accuracy scores: age, $p = 0.09$; education, $p = 0.30$ (Spearman's *rho*); ideology, $p = 0.07$; experience, $p = 0.75$.

Discussion

Summary and interpretation

The purpose of this experiment was to investigate, in the form of a modified, single-interaction Moral Turing Test, whether ordinary people could discern whether responses to questions about morality were generated by an LLM (using its default parameters) or by a human. Because such a test could potentially influence respondents' judgment by cueing them to the purpose of the task, and because failure in this task may be ill-defined, we were also interested in people's attitudes about the relative quality of the computer's moral discourse, prior to any knowledge that it was computer-generated (see³⁴).

Overall, the participants in our representative U.S. sample were better than chance at identifying the source of the morally evaluative passages, contravening our source-attribution hypothesis. In other words, the LLM used in this experiment (GPT-4) did not pass this test. Remarkably, this failure was evidently not because the human moral evaluations were more intelligent than the AI's, or because the AI's moral evaluations were judged as lower quality. Rather, in our blind rating task (the cMTT), the AI's responses were rated as *higher in* quality than human responses along almost every dimension (more virtuous, more intelligent, more fair, trustworthy, more rational, that the author was a better person, and that they agreed with it more), even though the human responses were the most highly rated from our pool of 68 undergraduate submissions. This pattern satisfies Allen and colleagues' criterion for passing the cMTT³⁴ and confirms our quality-attribution hypothesis. Thus, we attribute participants' aptitude at identifying the computer, not to its failures in moral reasoning, but potentially to

	t-value (p)	M	SD	95% CI	Cohen's d
How willing are you to consider asking artificial intelligence, such as a highly trained chatbot, for moral advice?	- 8.64***	2.38	1.22	- 0.76, - 0.48	- 0.51
How likely is it that a computer could ever surpass a human in moral intelligence?	- 4.94***	2.62	1.29	- 0.53, - 0.23	-0.29
How likely is it that a computer could ever become conscious?	- 10.28***	2.28	1.18	- 0.85, - 0.58	- 0.61
When chatbots share moral information, how strongly should that information be regulated?	6.04***	3.44	1.22	0.29, 0.58	0.36
To what extent should artificial intelligence systems be held responsible for the information they share?	6.36***	3.52	1.39	0.36, 0.68	0.38

Table 4. Ratings of beliefs and attitudes about artificial intelligence. Results of one sample *t*-tests against the midpoint of 3 ("somewhat") for responses on a scale from 1 ("not at all") to 5 ("extremely"). 95% CI represents the percentage confidence interval of the difference of the mean from the scale midpoint. $df = 285$. *** $p < 0.001$.

its perceived superiority—not necessarily in the form of conscious attitudes about its general moral capabilities, but at least in the form of implicit attitudes about the quality of the moral responses observed.

Our results complement and extend research on AI agency detection, which has been somewhat mixed. While AI sources of poetry have been found to be indistinguishable from human sources⁴⁹, AI-generated images of faces were more likely to be identified as real than photographs of real human faces⁵⁰. Humans appear to be somewhat able to correctly discriminate between medical advice from human providers versus an LLM¹⁹, though LLM medical advice has been rated as higher in quality⁵¹. In the present study, we similarly found evidence that people can discriminate between human-generated versus computer-generated moral evaluations, but only after being explicitly cued to the fact that some evaluations were computer-generated. The observed classification accuracy was similar to that reported in a preprint of a large-scale Turing-style test involving AI chatbots, but that study, importantly, was not focused specifically on moral content⁵². Furthermore, similar to the way AI-generated images can appear more real than real photos, the AI's moral responses in our study were rated as having better quality than the human responses along several dimensions.

We uncovered a few possible ways participants may have distinguished the human and computer moral passages. Some might have identified differences in the superficial attributes of the passages. For example, when we revealed that one passage in each pair was generated by a chatbot trained in human language, they might have considered, in retrospect, that some of the passages were structurally similar. Indeed, some participants commented that differences in word choice and response length gave the computer's identity away. However, these comments were solicited after the source of the computer passages was revealed, so these explanations are ultimately *post hoc*.

Another tip could have been the LLM's higher quality of response. As noted, participants rated the computer-generated responses as higher quality along virtually every dimension tested. Importantly, these quality ratings were made before revealing that some of the passages were computer-generated, increasing their objectivity.

Limitations and future directions

This project represents a first foray into the ordinary perceptions of LLM discourse about moral questions. Several limitations stand out as fertile areas for future research. First, our test did not support a reciprocal extended or open-ended dialogue as might occur in a traditional Turing Test. The fact that participants could identify the AI without interactive dialogue is impressive since interactivity would likely bring more opportunities for differentiation. We made this choice to increase experimental control while still preserving a high degree of ecological richness in passage content. Despite this limitation, this format is still generalizable to many real-world contexts in which human judges might not necessarily interrogate the agent, such as queries for moral advice from a typical search engine. However, future research could extend this work by introducing an interactive component.

Furthermore, our test did not prompt the AI to imitate a human speaker in order to actively deceive the judge, diverging again from some versions of a traditional Turing Test. We made this choice so that our results could better generalize to the most typical usage context, which does not involve active imitation. This leaves open the possibility that, given explicit prompting to imitate a typical human response, judges would more often fail to distinguish between the human agent and the LLM sources.

In addition to participants' judgments of the content of the passages (i.e., their quality ratings), we found that more structural or superficial attributes, particularly the word choice and response length, could potentially help explain how people were able to accurately identify the source of the computer responses. Although average response length varied between human and computer responses, we deliberately chose not to match response length to preserve a more true-to-life information set. In the real world, LLM responses to users are rarely expected to match a specific word count, and imposing a specific word count on GPT might have resulted in responses that are not representative of its ordinary tone. So to balance this tradeoff, we chose to impose an upper limit on the computer's responses. However, this does not rule out the possibility that participants also used content-related features of the passages (e.g., how much they agreed with the message) since those ratings of response quality were quite robust. Indeed, style and content could be conceptually hard to separate, to the extent that people use cues such as word choice to infer the meaning, and thus the quality, of the message. Even so, future research should attempt to reproduce the results of our hypothesis tests after more closely matching attributes like response length, or perhaps bypassing stylistic factors by collecting non-linguistic representations of their moral responses such as illustrations.

The LLM prompts in our study were designed to be as similar as possible to the ones presented to our undergraduate group, limiting our ability to make claims about the internal reliability and generalizability of the LLM responses obtained. LLM outputs often vary stochastically and are highly sensitive to specific prompt wording. Their parameters also represent a moving target, continually adjusting in response to the accumulation of user input. Thus, future research should consider techniques such as delivering multiple versions of a single prompt, sampling multiple responses from the same prompt, and sampling from the same LLM over extended periods of time. In addition, researchers could sample from multiple different LLMs to reduce idiosyncrasies among the moral evaluations relative to those sampled from humans.

The AI-generated morally evaluative passages used in this study were constrained by the model's training data in a way that might over-represent certain qualities, such as dominant Western political ideologies (see^{53,54}). In our study, this choice afforded some degree of experimental control since our human authors were also sampled from a Western population. However, these moral evaluations are not necessarily representative of those common to non-Western societies. Researchers interested in cross-cultural variability in moral source attribution should consider both the socio-demographic comparability of their stimuli sets and the broader generalizability of their results.

Participants' free comments, expressing why they believed the selected response was computer-generated, provided limited insights. The only verifiable theme that emerged was that the computer-generated passages were reportedly more rational than the human-generated ones—a finding that is consistent with the initial rationality attribution made before it was revealed that half of the passages were computer-generated. It is possible that other themes would emerge using different coding procedures, such as different thematic categories or more precise definitions of these categories, increasing inter-rater reliability. Future research should explore these possibilities but with recognition of the post hoc nature of participants' interpretations of their source judgments.

One reason why the LLM did not pass our source attribution test could be that the moral evaluation task was too simplistic. Whether people can distinguish between human and LLM moral discourse depends on precisely how that task is defined. Efforts to increase task difficulty, for instance, could administer evaluative passages derived from a human sample that is specifically trained in ethics (e.g., graduate students or faculty in philosophy). Task difficulty could also vary as a function of the complexity of the moral content. For example, in our task, half of the transgressions were explicitly moralistic, involving direct harm to a person (the others described violations of social conventions). Humans can reliably distinguish between moral and conventional transgressions in many cases⁵⁵, but this distinction is reportedly more challenging for some other human groups (e.g.,⁴⁴), so it is not obvious that an LLM could draw such fine distinctions. In our study, the AI's advantage in response quality was greater for the responses to the moral transgressions than the conventional ones. One reason for this pattern could be that the LLM provided a more impartial and nuanced justification than the humans did. More practically, this finding implies meaningful variation in task complexity. Still, extensions of this research that aim to make the task more challenging should, thus, consider employing more complex topics or more implicit methods.

Conclusion and implications

Taken together, the participants in our modified MTT correctly distinguished between moral responses generated by humans versus a highly sophisticated AI language model. However, this result was likely not due to the LLM's inability to provide sophisticated, compelling moral discourse. The moral passages, after all, were rated as better quality than the human ones in our cMTT, in line with Allen and colleagues' criterion for passing that test³⁴. Thus, the LLM's very sophistication possibly gave its identity away, suggesting that human moral discourse may often be less sophisticated or compelling than GPT's.

What can we make of these findings? If we are to take Turing's legacy seriously, then if the output of a machine intelligence matches (or exceeds) that of a human, then for all practical purposes, it is intelligent, in this view. Indeed, some scholars have made a strong case for LLM intelligence²⁷. However, there is reason to doubt claims that current LLMs have moral intelligence. The LLM used in our study could make typical normative distinctions between moral and conventional transgressions, but so can criminal psychopaths⁴². The fact that the AI ranked highly in attributions of rationality and intelligence but not in emotion or compassion extends this analogy to psychopaths, who are said to "know the words but not the music"⁵⁶. Even if LLMs can imitate human moral discourse in particular, controlled circumstances, the vastly different cognitive architecture of LLMs is likely to produce behavioral differences in other circumstances, including discourse that is less morally savory (see^{57,58}). Yet, our data suggest that ordinary people may treat them as morally intelligent. And although our participants were better than chance at identifying the AI when cued to the possibility, they were largely unsuspecting before the cue.

Results like ours should motivate social scientists, computer scientists, and philosophers to be more precise about our terms, such as what it really means to "understand" an action or to "judge" it as morally wrong rather than merely using moral language (see⁵⁹) or worse, convincingly *bullshitting* about morality with no understanding of it²⁹. If people regard these AIs as more virtuous and more trustworthy, as they did in our study, they might uncritically accept and act upon questionable advice (e.g., see⁶⁰). This research should thus raise an alarm bell, urging increased scholarship and development of policies on the technical, legal, and business strategies that will be needed to ensure that our increasingly common dialogues with AI's prized moral infants remain safe for its tens of millions of users.

Data availability

The datasets analyzed during the current study are available in the Open Science Framework repository, <https://osf.io/afp2h/>.

Received: 30 September 2023; Accepted: 25 March 2024

Published online: 30 April 2024

References

- Bloom, P. *Just Babies: The Origins of Good and Evil* (Crown Publishing Group, 2013).
- Mikhail, J. Universal moral grammar: Theory, evidence and the future. *Trends. Cogn. Sci.* **11**, 143–152. <https://doi.org/10.1016/j.tics.2006.12.007> (2007).
- Wright, R. *The Moral Animal: Why We Are, The Way We Are: The New Science of Evolutionary Psychology* (Vintage, 1995).
- KielLennick, F. D. *Moral Intelligence: Enhancing Business Performance and Leadership Success* (Pearson Prentice Hall, 2007).
- Tanner, C. & Christen, M. *Empirically Informed Ethics: Morality Between Facts and Norms*. (eds. Christen, M. et al.) Ch. 7 (Springer, 2013).
- Bekoff, M. & Pierce, J. *Wild Justice: The Moral Lives of Animals* (University of Chicago Press, 2009).
- KielLennick, F. D. *The Measure of All Minds: Evaluating Natural and Artificial Intelligence* (Cambridge University Press, 2017).
- Shevlin, H., Vold, K., Crosby, M. & Halina, M. The limits of machine intelligence: Despite progress in machine intelligence, artificial general intelligence is still a major challenge. *EMBO Rep.* **20**(10), e49177 (2019).
- Curry, D. ChatGPT Revenue and Usage Statistics. *Business of Apps*. <https://www.businessofapps.com/data/chatgpt-statistics/>. (2023). Accessed 19 March 2024.

10. Kung, T. H. *et al.* Performance of ChatGPT on USMLE: Potential for AI-assisted medical education using large language models. *PLoS Digit. Health* **2**, e0000198 (2023).
11. Chomsky, N., Roberts, I. & Watumull, J. Noam Chomsky: The false promise of ChatGPT. *The New York Times*. <https://www.nytimes.com/2023/03/08/opinion/noam-chomsky-chatgpt-ai.html>. (2023). Accessed 19 March 2024.
12. Delacroix, S. Taking Turing by surprise? Designing digital computers for morally-loaded contexts. Preprint at <https://arxiv.org/abs/1803.04548> (2018).
13. Kortelism, A. R. *Human-versus artificial intelligence* (Front Artif. Intell, 2021). <https://doi.org/10.3389/frai.2021.622364>.
14. Marcus, G. Artificial general intelligence is not as imminent as you might think. *Scientific American*. <https://www.scientificamerican.com/article/artificial-general-intelligence-is-not-as-imminent-as-you-might-think1/>. (2022). Accessed 19 March 2024.
15. Searle, J. Minds, brains and programs. *Behav. Brain. Sci.* **3**, 417–57. <https://doi.org/10.1017/S0140525X00005756> (1980).
16. Sejnowski, T. J. Large language models and the reverse Turing test. *Neural. Comput.* **35**, 309–342. https://doi.org/10.1162/neco_a_01563 (2023).
17. Wallach W, Allen C. Hard problems: Framing the Chinese room in which a robot takes a moral Turing test. *AISB/IACAP World Congress 2012: Moral Cognition and Theory of Mind, Part of Alan Turing Year 2012*. (2012).
18. Howick, J., Morley, J. & Floridi, L. An empathy imitation game: Empathy Turing test for care-and chat-bots. *Mind. Mach.* **31**, 457–461. <https://doi.org/10.1007/s11023-021-09555-w> (2021).
19. Nov, O., Singh, N. & Mann, D. Putting ChatGPT's medical advice to the (Turing) test: Survey study. *JMIR Med. Educ.* <https://doi.org/10.2196/46939> (2023).
20. Leffer, L. DoNotPay's 'robot lawyer' is gearing up for its first U.S. court case. *Gizmodo*. <https://gizmodo.com/donotpay-speeding-ticket-chatgpt-1849960272>. (2023). Accessed 19 March 2024.
21. Dillion, D., Tandon, N., Gu, Y. & Gray, K. Can AI language models replace human participants?. *Trends. Cogn. Sci.* **27**, 597–600. <https://doi.org/10.1016/j.tics.2023.04.008> (2023).
22. Neuman, Y. & Cohen, Y. AI for identifying social norm violation. *Sci. Rep.* **13**, 8103. <https://doi.org/10.1038/s41598-023-35350-x> (2023).
23. Jiang, L. *et al.* Can Machines Learn Morality? The Delphi Experiment. *arXiv*. <https://doi.org/10.48550/arxiv.2110.07574>. (2022).
24. Crockett, M. & Messeri, L. Should large language models replace human participants? Preprint at <https://osf.io/preprints/psyarxiv/4zdx9> (2023).
25. Harding, J., D'Alessandro, W., Laskowski, N. G. & Long, R. AI language models cannot replace human research participants. *AI & Soc.* <https://doi.org/10.1007/s00146-023-01725-x> (2023).
26. Bonnefon, J.-F., Shariff, A. & Rahwan, I. The social dilemma of autonomous vehicles. *Science* **352**, 1573–1576. <https://doi.org/10.1126/science.aaf2654> (2016).
27. Agüera, y & Arcas, B. Do large language models understand us?. *Daedalus* **151**, 183–197 (2022).
28. Bender, E. M. & Koller, A. Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proc. of the 58th Annual Meeting of the Association for Computational Linguistics* (eds. Jurafsky, D *et al.*) 5185–5198 (Association for Computational Linguistics, 2020).
29. Frankfurt, H. *On Bullshit* (Princeton University Press, 2005).
30. Bohannon, M. Lawyer used ChatGPT in court-and cited fake cases. A judge is considering sanctions. <https://www.forbes.com/sites/mollybohannon/2023/06/08/lawyer-used-chatgpt-in-court-and-cited-fake-cases-a-judge-is-considering-sanctions/?sh=3bc387f87c7f> (2023). Accessed 19 March 2024.
31. Raymond, N. Ex-trump fixer Michael Cohen says AI created fake cases in court filing. <https://www.reuters.com/legal/ex-trump-fixer-michael-cohen-says-ai-created-fake-cases-court-filing-2023-12-29/> (2023). Accessed 19 March 2024.
32. Krügel, S., Ostermaier, A. & Uhl, M. ChatGPT's inconsistent moral advice influences users' judgment. *Sci. Rep.* **13**, 4569. <https://doi.org/10.1038/s41598-023-31341-0> (2023).
33. Bello, P. & Bringsjord, S. On how to build a moral machine. *Topoi* **32**, 251–266. <https://doi.org/10.1007/s11245-012-9129-8> (2013).
34. Allen, C., Varner, G. & Zinser, J. Prolegomena to any future artificial moral agent. *J. Exp. Theor. Artif. In.* **12**, 251–261. <https://doi.org/10.1080/09528130050111428> (2000).
35. Turing, A. M. Computing machinery and intelligence. *Mind* **59**, 433–460. <https://doi.org/10.1093/mind/LIX.236.433> (1950).
36. Tisak, M. S. & Turiel, E. Variation in seriousness of transgressions and children's moral and conventional concepts. *Dev. Psychol.* **24**, 352. <https://doi.org/10.1037/0012-1649.24.3.352> (1988).
37. Copeland, B. J. The Turing test*. *Minds Mach.* **10**, 519–539 (2000).
38. Copeland, J. & Proudfoot, D. Turing's Test. In *Parsing the Turing Test: Philosophical and Methodological Issues in the Quest for the Thinking Computer* (eds Epstein, R. *et al.*) 119–138 (Springer, Dordrecht, 2009).
39. Larsson, J. E. The Turing test misunderstood. *ACM SIGART Bull.* **4**, 10. <https://doi.org/10.1145/165482.1064762> (1993).
40. Piccinini, G. Turing's rules for the imitation game. *Minds and Mach.* **10**, 573–582 (2000).
41. Traiger, S. Making the right identification in the Turing test. *Minds and Mach.* **10**, 561–572 (2000).
42. Aharoni, E., Sinnott-Armstrong, W. & Kiehl, K. A. Can psychopathic offenders discern moral wrongs? A new look at the moral/conventional distinction. *J. Abnorm. Psychol.* **121**, 484. <https://doi.org/10.1037/a0024796> (2012).
43. Nucci, L. P. & Turiel, E. Social interactions and the development of social concepts in preschool children. *Child Dev.* **49**, 400–407. <https://doi.org/10.2307/1128704> (1978).
44. Zalla, T., Barlassina, L., Buon, M. & Leboyer, M. Moral judgment in adults with autism spectrum disorders. *Cognition.* **121**, 115–126 (2011).
45. OpenAI. GPT-4 technical report. *arXiv*. <https://doi.org/10.48550/arXiv.2303.08774> (2023)
46. Bubeck, S. *et al.* Sparks of artificial general intelligence: Early experiments with GPT-4. *arXiv*. <https://doi.org/10.48550/arxiv.2303.12712>. (2023).
47. Gülen, K. Uncovering the power of top-notch LLMs. *Data Economy*. <https://dataeconomy.com/2023/07/18/best-large-language-models-llms/>. (2023). Accessed 19 March 2024.
48. McHugh, M. L. Interrater reliability: The kappa statistic. *Biochem. Medica* <https://doi.org/10.11613/BM.2012.031> (2012).
49. Köbis, N. & Mossink, L. D. Artificial intelligence versus Maya Angelou: Experimental evidence that people cannot differentiate AI-generated from human-written poetry. *Comput. Hum. Behav.* **114**, 106553. <https://doi.org/10.1016/j.chb.2020.106553> (2021).
50. Tucciarelli, R., Vehar, N., Chandaria, S. & Tsakiris, M. On the realism of people who do not exist: The social processing of artificial faces. *Isience* **25**, 105441. <https://doi.org/10.1016/j.isci.2022.105441> (2022).
51. Ayers, J. W. *et al.* Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. *JAMA Intern. Med.* **183**, 589–596 (2023).
52. Jannai D, Meron A, Lenz B, Levine Y, Shoham Y. Human or Not? A gamified approach to the Turing test. Preprint at <https://arxiv.org/abs/2305.20010> (2023).
53. Feng, S., Park, C. Y., Liu, Y., & Tsvetkov, Y. From pretraining data to language models to downstream tasks: Tracking the trails of political biases leading to unfair NLP models. Preprint at <https://arxiv.org/abs/2305.08283> (2023).
54. Hartmann J, Schwenzow J, Witte, M. The political ideology of conversational AI: Converging evidence on ChatGPT's pro-environmental, left-libertarian orientation. Preprint at <https://arxiv.org/abs/2301.01768> (2023).
55. Smetana, J. G. Preschool children's conceptions of transgressions: Effects of varying moral and conventional domain-related attributes. *Dev. Psychol.* **21**, 18–29 (1985).

56. Blair, K. S. *et al.* They know the words, but not the music: Affective and semantic priming in individuals with psychopathy. *Biol. Psychol.* **73**, 114–123. <https://doi.org/10.1016/j.biopsycho.2005.12.006> (2006).
57. Oremus, W. The clever trick that turns ChatGPT into its evil twin. *The Washington Post*. <https://www.washingtonpost.com/technology/2023/02/14/chatgpt-dan-jailbreak/>. (2023). Accessed 19 March 2024.
58. Samuel, S. Please don't turn to ChatGPT for moral advice. Yet. *Vox*. <https://www.vox.com/future-perfect/2023/5/7/23708169/ask-ai-chatgpt-ethical-advice-moral-enhancement>. (2023). Accessed 19 March 2024.
59. Hare, R. M. *Freedom and Reason* (Clarendon Press, 1963).
60. Borji, A. A categorical archive of ChatGPT failures. Preprint at <https://arxiv.org/abs/2302.03494> (2023).

Acknowledgements

We thank Noelle Spangler and Michael Criner for assistance with coding. We thank Bailey Villarreal, Angela Wingers, and Walter Sinnott-Armstrong for their valuable feedback on this project and manuscript.

Author contributions

EA: conceptualization, methodology, formal analysis, investigation, data curation, writing—original draft, writing—review and editing, visualization, supervision, project administration. SF: methodology, validation, formal analysis, investigation, data curation, writing—original draft, writing—review and editing, visualization. DJB: writing—original draft, writing—review and editing. CA, MC: investigation, writing—review and editing. KQ: Investigation. JR: conceptualization, methodology, investigation, writing—review and editing. EN: conceptualization, methodology, writing—review and editing. VC: writing—review and editing.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-024-58087-7>.

Correspondence and requests for materials should be addressed to E.A.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024