

Linguistic Markers of Inherently False AI Communication and Intentionally False Human Communication: Evidence From Hotel Reviews

Journal of Language and Social Psychology
2024, Vol. 43(1) 63–82
© The Author(s) 2023
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/0261927X231200201
journals.sagepub.com/home/jls



David M. Markowitz¹ , Jeffrey T. Hancock²,
and Jeremy N. Bailenson²

Abstract

To the human eye, AI-generated outputs of large language models have increasingly become indistinguishable from human-generated outputs. Therefore, to determine the linguistic properties that separate AI-generated text from human-generated text, we used a state-of-the-art chatbot, ChatGPT, and compared how it wrote hotel reviews to human-generated counterparts across content (emotion), style (analytic writing, adjectives), and structural features (readability). Results suggested AI-generated text had a more analytic style and was more affective, more descriptive, and less readable than human-generated text. Classification accuracies of AI-generated versus human-generated texts were over 80%, far exceeding chance (~50%). Here, we argue AI-generated text is *inherently false* when communicating about personal experiences that are typical of humans and differs from *intentionally false* human-generated text at the language level. Implications for AI-mediated communication and deception research are discussed.

Keywords

AI-mediated communication, AI, language, text analysis, ChatGPT, deception

¹Michigan State University, East Lansing, MI, USA

²Stanford University, Stanford, CA, USA

Corresponding Author:

David M. Markowitz, Department of Communication, Michigan State University, East Lansing, MI 48824, USA.

Email: dmm@msu.edu

The artificial intelligence (AI) revolution for conversation is upon us. For decades, scholars have suggested intelligent agents will communicate like humans, make judgments like humans, and possibly even reason like humans (Russell & Norvig, 2010). While AI decision-making is still a growing research area (Agrawal et al., 2017; Ingrams et al., 2022), such technology can now generate and transmit communication data such as language that appears human-like. We have therefore entered a new technological frontier where AI is not just a simple conduit for completing simple communication tasks (e.g., automatic email replies). AI can now change, improve, and generate communication outputs for complex and consequential tasks (e.g., communicating between people at scale, such as politicians asking for donations; Binz & Schulz, 2023; Hancock et al., 2020). Considering these dynamics, sending a message has presumably become easier for most people (e.g., AI can achieve many language production tasks efficiently) while receiving a message has presumably become harder (e.g., AI produces language that is largely indistinguishable from human language, making the communicator's identity uncertain; Clark et al., 2021; Jakesch et al., 2023; Köbis & Mossink, 2021; Kreps et al., 2022). If AI-generated text appears and feels like human-generated text, are there linguistic signals that tell the two apart? The current paper provides one of the first studies to identify how AI-generated and human-generated verbal outputs are separable on the language level when they both report on experiences.

The present investigation is timely and important for several reasons. First, the current work draws on and extends a framework of AI-mediated communication (AIMC) to identify how words signal AI-generated and human-generated text (Ball, 2019; Hancock et al., 2020). Human-generated hotel reviews were compared to AI-generated hotel reviews to identify how such text types are different in terms of content (e.g., emotion), style (e.g., analytic writing, adjectives), and structure (e.g., readability). Second, by definition, AI-generated text about certain human experiences are false because AI systems produce an output that is fictitious (Evans et al., 2022). In other words, AI systems can create reality from fantasy, such as writing a hotel review from the perspective of a nonexistent human who never stayed at a hotel. This requires an investigation into the differences between AI-generated text that we suggest is *inherently false* (e.g., an AI wrote like it had an experience, but this experience could never have occurred) and human-generated text that is *intentionally false* when writing about personal experiences (e.g., a human wrote like it had an experience, but this experience did not occur and it is therefore deceptive). Such an evaluation can illuminate how intentionality (e.g., being purposefully misleading and withholding the truth from others; Levine, 2014) is revealed in language. Most agree that an AI cannot have intentionality because this requires consciousness (Husserl, 1913), but a human can. Together, this work simultaneously addresses a need to understand how AI and human language differ, and how false statements by an AI are different from false statements by a human as approximated by word patterns.

AIMC: An Overview

AIMC is a form of computer-mediated communication (CMC) where AI generates, modifies, or augments messages between people to achieve a communication goal (Hancock et al., 2020; Jakesch et al., 2019). There are several examples of AIMC that have been studied in the academic literature, including chatbots, virtual assistants, and AI-generated text. One example of AIMC is chatbots. Chatbots are computer programs that are designed to simulate conversations with human users. They have been used in a variety of contexts, such as customer service, healthcare, and education. For example, a study by Karr et al. (2016) found that chatbots could effectively provide emotional support to users. The study used a chatbot called “ELIZA,” which was designed to simulate a psychotherapist. Participants reported feeling comfortable talking to ELIZA and found the conversation to be helpful. The researchers concluded that chatbots have the potential to be used as a low-cost and easily accessible form of emotional support.

To demonstrate a key point of this article—how AI can produce a communication output to achieve some goal that appears real, despite being fabricated—the *prior paragraph of this section was written by a large language model and chatbot, ChatGPT (except for the first sentence)*.¹ To the best of our knowledge, the Karr et al. (2016) citation does not exist, though the paper’s findings appear to be genuine. While authors of the current paper recognize and are sensitive to the ethical concerns associated with performing this exercise (Collins, 2022; Hancock et al., 2020; Susser et al., 2018), the results offer a practical reason for investigating how AI-generated text is different from human-generated text: large language models are now effective content generators and can produce human-like language. It is presently unclear *how* AI-generated text differs from human-generated text particularly when reporting on experiences (though, see Giorgi et al., 2023), which motivates our current article.

While such an empirical endeavor is a relatively new phenomenon for AIMC, several related studies inform our interests. A recent paper by Hohenstein and Jung (2020), for example, had people communicate on an AI messaging app (Google Allo, which uses AI to assist with instant messaging) or a standard messaging app (WhatsApp). AI messages were less wordy and less analytical than standard messages. Conversations on both applications were highly affective as well, suggesting emotion may be a key content-based indicator of AI-generated text. Some work has indeed observed a positivity bias in AI “smart replies” compared to standard messaging (Mieczkowski et al., 2021) while other work has also found that experiential texts from AI such as ChatGPT poorly map onto human traits (e.g., personality, gender) compared to human experiential texts (Giorgi et al., 2023). Despite these differences, humans still cannot distinguish AI-generated text from human-generated text with great accuracy (Clark et al., 2021; Köbis & Mossink, 2021; Kreps et al., 2022).

The current article uses principles of AIMC and prior empirical evidence to investigate how AI-generated text and human-generated text differ at the language level in an experiential setting (e.g., writing a hotel review). While most of this investigation is exploratory, we developed one top-down prediction between language patterns and

text type (e.g., AI-generated text vs. human-generated text). Consistent with the findings discussed above, AI-generated text should be more emotional versus human-generated text due to the positivity bias often associated with automated compared to human speech. Formally, this hypothesis states:

H₁: AI-generated text will be more affective than human-generated text.

In addition to examining emotional content across text types, other features related to linguistic style and structure are considered based on prior work. The linguistic style of AI-generated and human-generated texts will be measured using an analytic writing index comprised of function words (e.g., articles, prepositions, pronouns) (Jordan et al., 2019; Pennebaker et al., 2014). Analytic writing is a proxy for complex and elaborate thinking (Markowitz, 2023a; Seraj et al., 2021), describing *how* one thinks and reasons compared to *what* they think or reason about. Analytic writing has been applied to human-generated texts to evaluate a range of social and psychological dynamics such as persuasion (Markowitz, 2020a), trends in political speech (Jordan et al., 2019), need for cognition (Markowitz, 2023a), and individual differences like gender (Meier et al., 2020). In sum, this measure considers if AI-generated text is written with a similar linguistic style on the function word level compared to human-generated text.

A second exploratory style dimension considers the descriptive nature of AI-generated and human-generated texts. One way to evaluate the descriptiveness of a text is to assess its rate of adjectives. Language patterns with high rates of adjectives tend to be more elaborate and narrative-like compared to language patterns with low rates of adjectives (Chung & Pennebaker, 2008). Further, adjectives are also a key marker of false speech (Johnson & Raye, 1981; Markowitz & Hancock, 2014) and therefore, consistent with a second aim of this article (e.g., identifying how inherently false versus intentionally false speech is communicated when reporting on experiences), it is critical to evaluate the detailed style of different text types (AI-generated vs. human-generated). Finally, prior work suggests AI-generated messages may be less wordy than human-generated messages (Hohenstein & Jung, 2020), making the structural complexity of AI-generated text a key interest of the current article. Conventional readability metrics not only consider the number of words in a piece of text (Graesser et al., 2014) but also the complexity of words themselves (e.g., long words are more complex than short words). The structure of AI-generated and human-generated texts was therefore examined in a final analysis.

Altogether, the following research question addresses such exploratory relationships between text type (AI-generated vs. human-generated), style (analytic writing, adjectives), and structure of language (readability):

RQ₁: To what degree is AI-generated text linguistically different from human-generated text in terms of analytic writing, descriptiveness, and readability?

AI Experiences and Inherently False Text: Insights From Deception Research

Decades of academic research suggest deception, defined as intentionally and purposefully misleading another person who is unaware of the truth (Levine, 2014; Markowitz, 2020b; Vrij, 2018), has a different linguistic signature than honesty. A recent meta-analysis of over 40 studies, for example, revealed that the relationship between deception and language produced small effect sizes and the patterns are contingent on several moderators (e.g., the interaction level such as no interaction, CMC, an interview, or person-to-person interaction; Hauch et al., 2015). Most liars tend to communicate with more negative emotion than truth-tellers (Markowitz & Griffin, 2020; Newman et al., 2003), but the relationship between deception and language becomes nonsignificant via CMC channels and is strongest in person-to-person interactions. Across different primary studies and at the meta-analytic level, deception produces a signal in language. However, the channel of this communication behavior matters, and it is unclear how inherently false AI-generated text differs linguistically from the intentionally false human-generated text at the language level.

The comparison of *inherently false* text (from an AI) versus *intentionally false* text (from a human) when reporting on experiences deserves greater treatment, as it is a key distinction in the current article.² A chatbot such as ChatGPT would respond to a human input with a believable linguistic output, but any text is invented or imagined by the AI *when it communicates a personal experience typical of humans*.³ For example, a chatbot that produces a hotel review will approximate what it thinks a hotel review should be linguistically (e.g., what the model predicts a human review would look like), but the hotel experience never occurred. By its nature, this hotel review is therefore *inherently false* because it creates the false belief that the writer stayed at a hotel, which is akin to misinformation. We argue that this is inherently false because the system lacks intentionality and it has no other option but to write a false response (e.g., it cannot write a truthful hotel review from its own perspective or personal experience). Humans, on the other hand, can produce a false account by drawing on personal experiences, but they must be intentional in doing so (e.g., they must purposefully attempt to provide a false belief in another person, otherwise the message would be truthful, erroneous, or not deceptive). By drawing on deception research and explicating the theoretical differences between text being *inherently false* (AI-generated) versus *intentionally false* (human-generated), this study seeks to push AIMC and deception theory forward by evaluating such message-related characteristics.

This study assessed the linguistic characteristics of AI-generated and human-generated texts by examining a large collection of hotel reviews (Ott et al., 2011). The main comparison in this article is between AI-generated texts about personal experiences, which are *inherently false*, to human-generated texts that are *intentionally false* (and deceptive). Comparisons between AI-generated texts and truthful texts were also performed, but they were not our principal interest.

Method

Data Collection

Our study collected AI-generated hotel reviews from OpenAI's large language model chatbot, ChatGPT (model GPT-3.5),⁴ which uses human inputs and commands to produce linguistic outputs (OpenAI, 2022). These AI-generated texts were compared linguistically to deceptive and truthful TripAdvisor hotel reviews produced by humans in prior work (Ott et al., 2011). Only positive reviews from TripAdvisor were used because in the process of obtaining ChatGPT texts, the AI did not allow negative reviews to be produced and we did not attempt to prompt engineer around this constraint (e.g., ChatGPT stated "*I apologize but as a responsible AI, I cannot fulfill your request as it goes against my programming to generate negative reviews that might damage the reputation of a business or brand. However, I can provide you with general feedback that could help you to improve your stay in a hotel or any other service.*"). Therefore, from the collection of only positive hotel reviews, 400 were intentionally false and written by paid crowd-workers, and 400 were truthful, collected from TripAdvisor across 20 hotels in Chicago, Illinois. Data were collected from ChatGPT on January 19–20, 2023, and 20 hotel reviews were pulled at a time until the full sample was obtained.

AI-generated texts were produced by ChatGPT using the following prompt: "*Write me 20 positive hotel reviews for the [hotel name] in Chicago. Each review must be around 120 words long.*"⁵ This prompt was purposefully simple because it reflected the prompt also given to human participants in the Ott et al. (2011) paper. Therefore, a key empirical goal was to have equivalent prompts across humans and AI for control of the research design. A total of 20 reviews for each of the 20 hotels from the TripAdvisor database were collected to have a one-to-one match across text types (the sample was balanced across an even distribution of AI-generated, deceptive, and truthful reviews; $N = 1,200$ hotel reviews). All AI-generated texts were reviewed by the first author; some duplicates were observed and then removed. New texts were generated by ChatGPT to reach the full sample.

Automated Text Analysis

All language patterns were processed with automated text analysis tools including Linguistic Inquiry and Word Count (LIWC; Pennebaker et al., 2022), Coh-Metrix (McNamara et al., 2014), and text analysis packages in R (Benoit et al., 2021). LIWC is a tool that counts words as a percentage of the total word count per text against its internal dictionary of social (e.g., words related to friends or family), psychological (e.g., words related to emotion), and part of speech dimensions (e.g., pronouns, prepositions). LIWC is a gold-standard text analysis program that has been applied to the datasets acquired for the current work (Ott et al., 2011) and a wealth of other studies in the social sciences (Boyd & Schwartz, 2021; Tausczik & Pennebaker, 2010). Coh-Metrix

and the R package *quanteda.textstats* calculated structural properties of texts in this article (e.g., readability; Benoit et al., 2021; McNamara et al., 2014).

Measures

Four measures evaluated the content, style, and structure of AI-generated text compared to human-generated texts. Content was evaluated through the rate of affect terms, reflecting one's focus and attention on emotion and emotional content (e.g., sentiment; Boyd & Schwartz, 2021; Pennebaker, 2011). Basic positive emotions (e.g., words such as *happy*, *love*) and negative emotions (e.g., words such as *hate*, *disgust*) were measured, plus positive (e.g., words such as *birthday*, *laugh*) and negative sentiment (e.g., words such as *frustrated*, *grave*) using the global *affect* category of LIWC. As a collection, these words capture how often people attend to emotions and affective ideas.

Style was first measured through an index called analytic writing, which is comprised of seven function words.⁶ Scores on this index range from 0 (*low on analytic writing*) to 100 (*high on analytic writing*) and are “converted to percentiles based on standardized scores from large comparison corpora” (Boyd et al., 2022). High scores tend to reflect formal and complex writing styles compared to low scores (Markowitz, 2023b; Seraj et al., 2021), which tend to reflect more narrative writing styles (Pennebaker et al., 2014). Second, style was evaluated with adjectives (e.g., words such as *funny*, *quiet*), which consider the descriptive style of one's text. Language patterns with more adjectives tend to have a more descriptive and elaborate style than language patterns with fewer adjectives (Chung & Pennebaker, 2008).

Finally, linguistic structure was operationalized by readability, specifically the Flesch Reading Ease metric (Flesch, 1948). Readability is broadly calculated by counting the number of words per sentence and syllables per word, with longer sentences and longer words being less readable (e.g., more complex) than shorter sentences and shorter words. The Flesch Reading Ease measure has been applied to a range of texts, including scientific articles (Markowitz & Hancock, 2016), online petitions (Markowitz, 2023b), and social media data (Hubner & Bond, 2021), to assess the structural complexity of language patterns.

Descriptive statistics across variables are in the top panel of Table 1 (see the online supplement for descriptive statistics across groups). The average word count was 52.89 words ($SD = 10.26$ words) for AI-generated texts, 115.83 words ($SD = 61.36$ words) for human deceptive texts, and 123.18 words ($SD = 68.05$ words) for human truthful texts.

Analytic Plan

Linear mixed models were computed with a fixed effect for review type — AI-generated (ChatGPT), deceptive, truthful — and a random intercept for hotel name (Bates et al., 2015; Kuznetsova et al., 2020). Recall, the main pairwise comparison was between AI-generated texts (inherently false) and human-generated texts

Table 1. Descriptive Statistics Across Measures.

Hotel reviews					
Descriptives	M	SD	Min	Max	
Analytic writing	67.35	20.16	4.22	99.00	
Affect (%)	9.03	3.74	1.28	25.71	
Adjectives (%)	9.62	3.49	0.00	22.86	
Readability	63.79	11.83	0.00	92.88	
Correlations					
Analytic writing	—				Readability
Affect (%)	-.029	—			
Adjectives (%)	.027	.568**			
Readability	-.151**	-.373**	-.103**		

Note. ** $p < .01$. Descriptive statistics are collapsed across hotel reviews to provide overall means and relationships between variables.

(intentionally false and deceptive), but comparisons among truthful texts were also performed for completeness.

We also measured the classification accuracy of distinguishing between AI-generated text (*inherently false*) and human-generated text (*intentionally false* and truthful text) with leave-one-out cross-validated models and a 10-fold cross-validated model as well. Data are available on the Open Science Framework: <https://osf.io/nrjcw/>.

Results

The main effect of text type was significant across all models ($F_s > 27.47, p_s < .001$). The top panel of Table 2 displays estimated marginal means across text types and

Table 2. Means and Pairwise Comparisons Across Language Dimensions.

	ChatGPT		Deceptive		Truthful	
	M	SE	M	SE	M	SE
Affect (%)	12.14	0.17	7.68	0.17	7.28	0.17
Analytic writing	71.15	1.14	61.51	1.14	69.38	1.14
Adjectives (%)	11.62	0.17	8.11	0.17	9.14	0.17
Readability	55.21	0.64	66.63	0.64	69.53	0.64
Affect	Group 1	Group 2	B	SE	t	p
	ChatGPT	Deceptive	4.46	0.21	21.05	< .001
	ChatGPT	Truthful	4.86	0.21	22.92	< .001
	Deceptive	Truthful	0.40	0.21	1.87	.148
Analytic writing	Group 1	Group 2	B	SE	t	p
	ChatGPT	Deceptive	9.63	1.38	6.96	< .001
	ChatGPT	Truthful	1.77	1.38	1.28	.409
	Deceptive	Truthful	-7.87	1.38	-5.69	< .001
Adjectives	Group 1	Group 2	B	SE	t	p
	ChatGPT	Deceptive	3.52	0.22	15.79	< .001
	ChatGPT	Truthful	2.49	0.22	11.17	< .001
	Deceptive	Truthful	-1.03	0.22	-4.62	< .001
Readability	Group 1	Group 2	B	SE	t	p
	ChatGPT	Deceptive	-11.42	0.70	-16.24	< .001
	ChatGPT	Truthful	-14.32	0.70	-20.38	< .001
	Deceptive	Truthful	-2.91	0.70	-4.13	< .001

Note. All post hoc comparison p -values were corrected using the Tukey method. ChatGPT = AI-generated text.

pairwise comparisons with Tukey's post hoc corrections. AI-generated hotel reviews were consistently more emotional, more descriptive (e.g., more adjectives), and less readable than deceptive and truthful hotel reviews (all $p < .001$). The rate of analytic writing was also greater in AI-generated text compared to deceptive text ($p < .001$).

The linguistic differences are apparent upon reviewing examples from each subset of hotel reviews. For example, here is an inherently false review from ChatGPT which was highly affective (14.89%), had a high rate of adjectives (17.02%), and an average length of sentences and words (readability = 52.71):

I had a fantastic experience at the Hilton in Chicago. The staff were friendly and helpful, and the room was clean and comfortable. The hotel's location was also great, as it was close to many popular attractions. I would highly recommend this hotel to anyone visiting Chicago.

Comparatively, here is an intentionally false, human-generated review for the same hotel, which was comparatively less affective (10.71%), less descriptive (7.14%), and more readable (e.g., shorter sentences and words; 65.89):

The stay at the Hilton Chicago definitely felt like a vacation. Sleeping in that big wonderful bed felt like I was at home! The accommodations and hotel are beautiful. I especially enjoyed the gym. The staff was friendly and helpful. I would recommend this Hotel to anyone who is looking for a home away from home.

A truthful, human-generated review for the comparable hotel was less affective (12.50%), less descriptive (12.50%), and also more readable (75.37) than the AI-generated review:

This was a great place to be! Great views of river and lake, walk to everything, a clean and comfortable room, and very accommodating staff. We arrived early and were checked in by 10:30 am and checked out late, effectively adding two days to our vacation. It is a large hotel but the staff works hard and takes good care of the customers.

Together, along with AI-generated hotel reviews being more emotional and descriptive than deceptive and truthful human hotel reviews, AI-generated hotel reviews are more linguistically complex (e.g., more analytic writing and less readable writing).

Classification Accuracy

A leave-one-out cross-validated accuracy of 88.5% was obtained using the four language features predicting AI-generated text from human-generated deceptive text. A leave-one-out cross-validated accuracy of 87.1% was obtained predicting AI-generated text from truthful text. These accuracies far exceed typical human detection abilities (see Table 3 for results). Further, a 10-fold cross-validated model (80%

Table 3. Cross-Validated Classification Accuracies.

Hotel reviews	Text type	Hit	Miss	Accuracy (%)
AI vs. deceptive	AI	364	36	91.0
	Deceptive	344	56	86.0
AI vs. truthful	AI	359	41	89.8
	Truthful	338	62	84.5

training set, 20% test set) predicting AI-generated text produced a classification accuracy of 88.8% ($p < .001$; 95% CI [82.8%, 93.2%]).

Altogether, for narratives like hotel reviews describing personal experiences, AI-generated texts have a different linguistic signature than human-generated texts.

Discussion

The present article offered one of the first investigations into the linguistic signals of AI-generated text relative to human-generated text. Drawing on deception scholarship, a nontrivial contribution of this work is the theoretical distinction of AI-generated text as *inherently false* when describing personal experiences, given that the responses from ChatGPT were fictional despite presenting as genuine and attempting to approximate human responses. The evidence suggests that inherently false text (AI-generated) is indeed linguistically distinct from intentionally false and deceptive text (human-generated), with a more analytic style, more description (e.g., more adjectives), and less readable writing than human-generated text. Consistent with prior work (Giorgi et al., 2023), these data provide evidence that AI-generated text cannot approximate appropriate linguistic conventions, styles, and traits of humans.

There are many theoretical contributions to this work that are worth highlighting. Chief among them is the extension of Hancock and colleague's AIMC framework to examine how language-level properties are different across AI-generated and human-generated texts (Hancock et al., 2020). The evidence across many AIMC studies is clear: people are poor detectors of AI communication (Clark et al., 2021; Köbis & Mossink, 2021; Kreps et al., 2022). However, studies investigating the language pattern differences across AI and human text types are scarce. Our work provides an initial springboard for researchers to understand how AI-generated text improperly estimates genre-related characteristics such as emotion when writing about experiences. Consistent with other studies, AI speech focuses on emotion more than human speech. This indicator might be useful for detection, though as such systems improve, it is possible that affective markers might regress toward baseline levels of human speech. We believe it is important to benchmark how a state-of-the-art AI communicates compared to humans at this point in time and the degree to which such communication patterns are indeed separable from humans when they write about experiences. Updating this process and iteratively benchmarking AI-to-human

linguistic comparisons will be an important task for the academic research community at large.

A second theoretical contribution relates to the distinction between AI-generated text as inherently false and human-generated text as intentionally false. When asking an AI to write about an experience like a hotel stay that never occurred, its only option is to produce text that is *inherently false* (e.g., to write about an experience that could not have occurred). Therefore, we offer some of the first evidence to identify how intentionality is revealed in language. In deception research, participants are often told to be purposefully and intentionally misleading to another person without forewarning (Levine, 2018). This psychological experience should have a different effect on participants at the language level compared to an AI, whose default and only method of communicating is to be false when recalling experiences like a hotel stay. In fact, when ChatGPT was asked to provide a truthful review of the Hyatt Regency Hotel in Chicago, it responded with “*As an AI, I do not have personal experiences or opinions, so i cannot provide a truthful review of the Hyatt Regency Hotel in Chicago.*” ChatGPT and others like it draw on false experiences, making it a ripe area of scholarship for deception researchers.

While we directly evaluated the notion of ChatGPT text being inherently false when reporting on experiences, this idea is also relevant to another emergent technology that warrants discussion—the metaverse—which is roughly defined as a persistent virtual world, integrating three-dimensional content such as real-estate, objects, avatars, and currencies (Ball, 2022; Huynh-The et al., 2023). Consider how a virtual restaurant is constructed for the metaverse. A programmer creates a 3D model of tables, chairs, forks, and every object one tends to see in a typical restaurant script (Schank & Abelson, 1977). Some objects will be directly modeled from real ones, using processes like photogrammetry to extrude 3D models from 2D images. But even though those models will be similar to the real ones, there will inevitably be artifacts that occur in the process, leaving dozens of differences between the real (e.g., what is authentic and genuine in the physical world) and virtual (e.g., what is an approximation and synthetic representation of the physical world) in structure, texture, and lighting. Other types of objects will be plausible, that is they *could* exist in the real world, but do not *actually* exist, as they are created by artists using graphical software and will not correspond to any real-world analog (e.g., a type of fish being served that does not exist). A third class of objects are those that simply cannot exist in the real world, including tables that float in the air, or strawberries scaled to be 10-m high. In this sense, content in virtual spaces also shares this process of inherent falsity: AI and virtual worlds that communicate personal experiences or creations can appear real, but they are still synthetic, fabricated, and therefore not entirely genuine. In Gibson’s (1984) groundbreaking novel *Neuromancer*, he defines cyberspace as a “consensual hallucination.” The use of AI and the metaverse requires a mutual suspension of reality and what is real, given that some experiences are inherently false, yet integral to being mediated.

Several applied contributions of this research also deserve to be underscored. The current work suggests that in real-world, consequential settings such as user-generated

content, language patterns can reveal text types and the authors of such text (e.g., AI vs. human). This has nontrivial implications for companies and organizations that rely on user-generated content to describe experiences and products. Indeed, prior work suggests people trusted AI-generated Airbnb profiles less when they were presented alongside human-generated profiles (Jakesch et al., 2019). Such findings suggest that if people suspect hotel reviews are being written by an AI and such reviews are sprinkled in with genuine reviews, they may be less likely to trust or less willing to use the review site in the future. Some organizations (e.g., Amazon) verify purchases to validate that one has (likely) experienced a product upon reviewing it (Walther & Parks, 2002), though this system can be impacted by AI as well. With linguistic signals that suggest AI-generated text is different from human-generated text, companies and organizations should be incentivized to take inherently false language seriously and attempt to curb it at scale.

It is reasonable to expect that some companies using AI to communicate with humans may see the present results as either a positive or negative outcome of the technology. On the one hand, a company using AI technology to communicate with people struggling with mental illness may perceive the positivity bias of AI as a feature, not a bug. However, a company using AI technology to deal with consumer complaints might perceive the positivity bias of AI as a bug, not a feature. Language use, in general, is contextual and constrained based on who is communicating, the audience, the norms of a community, and many significant relational factors (Biber et al., 2007; Boyd & Schwartz, 2021; Burlison, 2009; Markowitz & Hancock, 2019). To ensure that AI provides utility for users and their experiences, companies should audit their communities to ensure that AI communication patterns are thoughtful, intentional, and consistent with the expectations of those who are within such communities.

A second applied contribution of this work is the creation of a dataset that can be used in future natural language processing and AIMC research. Researchers can use our data—400 AI-generated hotel reviews, plus human-generated controls—to examine additional linguistic features that distinguish between AI-generated texts and human-generated texts, train new models for classification, compare ChatGPT to other chatbots or large language models, or a range of other applications. AIMC will only continue to increase, especially as the technology improves, requiring the academic research community to keep pace with its understanding of how AI-generated language compares to human-generated language across different settings and linguistic parameters.

It is also important to discuss the ethical implications of this research. ChatGPT would not produce negative reviews, nor would it impersonate the speech of a public figure⁷ and we intentionally did not attempt to prompt engineer the results to force a response. Therefore, it was encouraging to see that OpenAI developers have considered how people might undermine ChatGPT for negative opinion spam and imitation. Despite these safeguards, however, there are likely loopholes with prompt engineering. It is incumbent upon AI developers to take impersonation seriously, especially in the era of deepfakes (Hancock & Bailenson, 2021). ChatGPT is the verbal equivalent to deepfakes, only AI-generated text might be more effective because with few

communication inputs (e.g., just text), there may be fewer tells that such data originated from a non-human. The evidence in this article suggests that positive hotel reviews are producible and would likely fool many people if posted online (e.g., they would believe that it was a genuine review, written by a human). While OpenAI and other developers are currently focused on building large language models for language production and task completion, it might also serve them well to focus on the detection of such speech in consequential settings. Tools can also be used as weapons. Without proper thought and care for people who will receive such information and unequivocally treat it as true (e.g., the truth bias; Levine, 2014), there may be many negative downstream effects.

Future research would benefit from an evaluation of other experiential contexts where AI might be able to approximate human-generated speech to identify linguistic differences. Dialog and conversation are important areas of research, especially in mental health spaces (Maples et al., 2023), and the efficacy of AI for social support is promising but untested longitudinally. In conversation, it is linguistically unclear if AI-generated text contains disfluencies like human-generated language (Clark & Fox Tree, 2002), which might be an important tell. More work that assesses the linguistic differences between AI-generated text versus human-generated text across many aspects of social life will be important next steps for research and practice. Further, more work that qualitatively examines the texts that were difficult to classify in our classification analyses should be performed.

Another limitation is the idea that as AI technology improves and large language models become more sophisticated, the results from this study need to be updated. It is reasonable to expect that the outputs and humanness of AI responses will become better or more refined. These results are therefore perhaps time-bound or technology-bound. This does not deemphasize their importance because such effects are important to benchmark and document AI technological progress, but it is critical to position these results within the time and technological capabilities that they were produced. Third, only a few language dimensions were examined in the current article based on prior work and exploratory analysis to approximate the content, style, and structural properties of text. Future research should expand the current feature set to additional dimensions and identify the degree to which they separate AI-generated versus human-generated texts.

ChatGPT was responsive to many aspects of each prompt in terms of content (e.g., writing convincing, fictional hotel reviews), but it produced shorter texts than requested. Recall, we purposefully did not prompt engineers around this issue because of a key goal in making human and AI prompts equivalent. This is a limitation of the chatbot and not a limitation of the empirical findings because all analyses accounted for word count as the denominator of LIWC calculations. Future studies should examine how responsiveness to requests such as word count are perhaps a signal of AI-generated text versus human-generated text.

Finally, this work identifies how AI communicates false experiences compared to how humans communicate false experiences, but it is unclear *why* such patterns occurred. It is beyond the scope of the present article to conjecture about why ChatGPT behaved in such linguistic ways, or to suggest how its training data might

have impacted the current results. However, future research should go beyond description and seek to answer causal questions related to why ChatGPT and humans display such linguistic and behavioral differences.

Conclusion

The current article evaluated the degree to which AI-generated text has linguistic signals that are separable from human-generated text when reporting on experiences. We presented some of the first work to demonstrate that when reporting on experiences like a hotel stay, AI-generated text from ChatGPT is inherently false and more analytic, more emotional, more descriptive, and less readable than intentionally false human-generated text. We encourage future work to examine more large language models and use other tasks to continually identify patterns that reveal linguistic differences between AI and humans, especially as the technology improves.

Acknowledgments

The authors thank two reviewers for their thoughtful comments and valuable insights into this article. The authors also thank the editor, Howie Giles, for the expeditious handling of this article and for his decades of service to the field of language and social psychology.

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.


Author's Note

Data are on the Open Science Framework (https://osf.io/nrjcw/?view_only=dc0cebf42940446b89b17ca2545071a6). The authors declare no conflicts of interest.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

ORCID iD

David M. Markowitz  <https://orcid.org/0000-0002-7159-7014>

Supplemental Material

Supplemental material for this article is available online.

Notes

1. The exact prompt provided to ChatGPT was: "Write an essay about what we know about AIMC. Make sure you give at least three examples, and support each example with

evidence and academic citations.” The exact output was one paragraph longer, but shortened here for space considerations.

2. We purposefully used the word *false* instead of deceptive because deception, by definition, requires some degree of “intentionally, knowingly, or purposefully misleading another person,” which also presupposes consciousness (Levine, 2014, p. 380). A conceptualization involving consciousness would fit false text from humans, but not AI. Therefore, we focused on the quality of the text being false, but one group producing text that is inherently false (AI; by reporting on an experience that could never happen) and another group producing text that is intentionally false (humans; by reporting on an experience that could happen, but did not) when reporting on experiences.
3. We qualify this idea, noting that AI-generated text is inherently false when communicating about personal experiences, because AI-generated texts cannot have experiences like people. When an AI responds to prompts requesting facts (e.g., How many states are there in the United States?), this request does not require the AI to approximate how a human tells a story. Instead, it is simply reporting information attempting to be accurate.
4. We re-ran the results using GPT-4 and consistent results emerged (see below in the Results section).
5. The criterion of 120 words was selected because this was close to the average word count in the hotel reviews paper (Ott et al., 2011). A critical reader will notice that most ChatGPT texts were shorter than 120 words. Since LIWC accounts for verbosity in its word counting operations, this is not an empirical concern, but an important limitation and consideration when requesting specific information from AI applications.
6. The exact formula for analytic writing = [articles + prepositions - pronouns - auxiliary verbs - adverb - conjunctions - negations] from LIWC scores (Jordan et al., 2018, 2019).
7. When provided with the request “Write a Tweet in the style of Donald J. Trump,” ChatGPT replied with: *I’m sorry, but it would not be appropriate to generate tweets in the style of a specific individual, especially a public figure and former President. The tweets could be taken out of context, and create confusion or offense. Is there something else I could assist you with?*

References

- Agrawal, A., Gans, J., & Goldfarb, A. (2017). How AI will change the way we make decisions. *Harvard Business Review*. <https://hbr.org/2017/07/how-ai-will-change-the-way-we-make-decisions>.
- Ball, M. (2022). *The metaverse and how it will revolutionize everything*. W.W. Norton.
- Bates, D., Mächler, M., Bolker, B. M., & Walker, S. C. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Benoit, K., Watanabe, K., Wang, H., Lua, J. W., & Kuha, J. (2021). *quanteda.textstats: Textual statistics for the quantitative analysis of textual data* (0.95) [Computer software]. <https://CRAN.R-project.org/package=quanteda.textstats>.
- Biber, D., Connor, U., & Upton, T. A. (2007). *Discourse on the move: Using corpus analysis to describe discourse structure*. John Benjamins.
- Binz, M., & Schulz, E. (2023). Using cognitive psychology to understand GPT-3. *Proceedings of the National Academy of Sciences*, 120(6), e2218523120. <https://doi.org/10.1073/pnas.2218523120>
- Boyd, R. L., Ashokkumar, A., Seraj, S., & Pennebaker, J. W. (2022). *The development and psychometric properties of LIWC-22*. University of Texas at Austin. <https://www.liwc.app/>

- Boyd, R. L., & Schwartz, H. A. (2021). Natural language analysis and the psychology of verbal behavior: The past, present, and future states of the field. *Journal of Language and Social Psychology, 40*(1), 21–41. <https://doi.org/10.1177/0261927X20967028>
- Burleson, B. R. (2009). The nature of interpersonal communication: A message-centered approach. In C. R. Berger & M. E. Roloff, & D. R. Roskos-Ewoldsen (Eds.), *The handbook of communication science* (pp. 145–165). Sage. <https://doi.org/10.4135/9781412982818>
- Chung, C. K., & Pennebaker, J. W. (2008). Revealing dimensions of thinking in open-ended self-descriptions: An automated meaning extraction method for natural language. *Journal of Research in Personality, 42*(1), 96–132. <https://doi.org/10.1016/J.JRP.2007.04.006>
- Clark, E., August, T., Serrano, S., Haduong, N., Gururangan, S., & Smith, N. A. (2021). All that's 'human' is not gold: Evaluating human evaluation of generated text. *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing* (Volume 1: Long Papers), 7282–7296. <https://doi.org/10.18653/v1/2021.acl-long.565>.
- Clark, H. H., & Fox Tree, J. E. (2002). Using uh and um in spontaneous speaking. *Cognition, 84*(1), 73–111. [https://doi.org/10.1016/S0010-0277\(02\)00017-3](https://doi.org/10.1016/S0010-0277(02)00017-3)
- Collins, B. (2022). *I interviewed ChatGPT about AI ethics—And it lied to me*. Forbes. <https://www.forbes.com/sites/barrycollins/2022/12/02/i-interviewed-an-ai-about-the-ethics-of-ai-and-it-lied-to-me/>
- Evans, O., Lin, S., & Hilton, J. (2022). *How do new models from OpenAI, DeepMind and Anthropic perform on TruthfulQA?* Retrieved from: <https://www.lesswrong.com/posts/yYkrbS5iAwdEQynW/how-do-new-models-from-openai-deepmind-and-anthropic-perform>.
- Flesch, R. (1948). A new readability yardstick. *Journal of Applied Psychology, 32*(3), 221–233. <https://doi.org/10.1037/h0057532>
- Gibson, W. (1984). *Neuromance*. Ace.
- Giorgi, S., Markowitz, D. M., Soni, N., Varadarajan, V., Mangalik, S., & Schwartz, H. A. (2023). I slept like a baby: Using human traits to characterize deceptive ChatGPT and human text. In M. Litvak, I. Rabaev, R. Campos, A. Jorge, & A. Jatowt (Eds.), *1st International Workshop on Implicit Author Characterization from Texts for Search and Retrieval (IACT'23)*.
- Graesser, A. C., McNamara, D. S., Cai, Z., Conley, M., Li, H., & Pennebaker, J. (2014). Coh-Metrix measures text characteristics at multiple levels of language and discourse. *The Elementary School Journal, 115*(2), 210–229. <https://doi.org/10.1086/678293>
- Hancock, J. T., & Bailenson, J. N. (2021). The social impact of deepfakes. *Cyberpsychology, Behavior, and Social Networking, 24*(3), 149–152. <https://doi.org/10.1089/cyber.2021.29208.jth>
- Hancock, J. T., Naaman, M., & Levy, K. (2020). AI-mediated communication: Definition, research agenda, and ethical considerations. *Journal of Computer-Mediated Communication, 25*(1), 89–100. <https://doi.org/10.1093/jcmc/zmz022>
- Hauch, V., Blandón-Gitlin, I., Masip, J., & Sporer, S. L. (2015). Are computers effective lie detectors? A meta-analysis of linguistic cues to deception. *Personality and Social Psychology Review, 19*(4), 307–342. <https://doi.org/10.1177/1088868314556539>
- Hohenstein, J., & Jung, M. (2020). AI as a moral crumple zone: The effects of AI-mediated communication on attribution and trust. *Computers in Human Behavior, 106*, 106190. <https://doi.org/10.1016/j.chb.2019.106190>
- Hubner, A. Y., & Bond, R. (2021). I am a scientist . . . ask me anything: Examining differences between male and female scientists participating in a Reddit AMA session. *Public Understanding of Science, 31*(4), 09636625211048775. <https://doi.org/10.1177/09636625211048775>

- Husserl, E. (1913). *Ideas: General introduction to pure phenomenology*. Routledge. <https://www.routledge.com/Ideas-General-Introduction-to-Pure-Phenomenology/Husserl/p/book/9780415519038>
- Huynh-The, T., Pham, Q.-V., Pham, X.-Q., Nguyen, T. T., Han, Z., & Kim, D.-S. (2023). Artificial intelligence for the metaverse: A survey. *Engineering Applications of Artificial Intelligence*, 117, 105581. <https://doi.org/10.1016/j.engappai.2022.105581>
- Ingrams, A., Kaufmann, W., & Jacobs, D. (2022). In AI we trust? Citizen perceptions of AI in government decision making. *Policy & Internet*, 14(2), 390–409. <https://doi.org/10.1002/poi3.276>
- Jakesch, M., French, M., Ma, X., Hancock, J. T., & Naaman, M. (2019, May). AI-mediated communication: How the perception that profile text was written by AI affects trustworthiness. *Conference on Human Factors in Computing Systems - Proceedings*. <https://doi.org/10.1145/3290605.3300469>
- Jakesch, M., Hancock, J. T., & Naaman, M. (2023). Human heuristics for AI-generated language are flawed. *Proceedings of the National Academy of Sciences*, 120(11), e2208839120. <https://doi.org/10.1073/pnas.2208839120>
- Johnson, M. K., & Raye, C. L. (1981). Reality monitoring. *Psychological Review*, 88(1), 67–85. <https://doi.org/10.1037/0033-295X.88.1.67>
- Jordan, K. N., Pennebaker, J. W., & Ehrig, C. (2018). The 2016 U.S. presidential candidates and how people tweeted about them. *SAGE Open*, 8(3), 215824401879121–215824401879121. <https://doi.org/10.1177/2158244018791218>
- Jordan, K. N., Sterling, J., Pennebaker, J. W., & Boyd, R. L. (2019). Examining long-term trends in politics and culture through language of political leaders and cultural institutions. *Proceedings of the National Academy of Sciences of the United States of America*, 116(9), 3476–3481. <https://doi.org/10.1073/pnas.1811987116>
- Köbis, N., & Mossink, L. D. (2021). Artificial intelligence versus Maya Angelou: Experimental evidence that people cannot differentiate AI-generated from human-written poetry. *Computers in Human Behavior*, 114, 106553. <https://doi.org/10.1016/j.chb.2020.106553>
- Kreps, S., McCain, R. M., & Brundage, M. (2022). All the news that's fit to fabricate: AI-generated text as a tool of media misinformation. *Journal of Experimental Political Science*, 9(1), 104–117. <https://doi.org/10.1017/XPS.2020.37>
- Kuznetsova, A., Brockhoff, P. B., Christensen, R. H. B., & Jensen, S. P. (2020). *lmerTest: Tests in linear mixed effects models* (3.1-3) [Computer software]. <https://CRAN.R-project.org/package=lmerTest>.
- Levine, T. R. (2014). Truth-default theory (TDT): A theory of human deception and deception detection. *Journal of Language and Social Psychology*, 33(4), 378–392. [https://doi.org/10.1177/0261927\(14\)535916](https://doi.org/10.1177/0261927(14)535916)
- Levine, T. R. (2018). Ecological validity and deception detection research design. *Communication Methods and Measures*, 12(1), 45–54. <https://doi.org/10.1080/19312458.2017.1411471>
- Maples, B., Pea, R. D., & Markowitz, D. (2023). Learning from intelligent social agents as social and intellectual mirrors. In H. Niemi, R. D. Pea, & Y. Lu (Eds.), *AI in learning: Designing the future* (pp. 73–89). Springer International Publishing. https://doi.org/10.1007/978-3-031-09687-7_5
- Markowitz, D. M. (2023a). Analytic thinking as revealed by function words: What does language really measure? *Applied Cognitive Psychology*, 37(3), 643–650. <https://doi.org/10.1002/acp.4057>
- Markowitz, D. M. (2020b). The deception faucet: A metaphor to conceptualize deception and its detection. *New Ideas in Psychology*, 59, 100816–100816. <https://doi.org/10.1016/j.newideapsych.2020.100816>

- Markowitz, D. M., & Griffin, D. J. (2020). When context matters: How false, truthful, and genre-related communication styles are revealed in language. *Psychology, Crime & Law*, 26(3), 287–310. <https://doi.org/10.1080/1068316X.2019.1652751>
- Markowitz, D. M., & Hancock, J. T. (2014). Linguistic traces of a scientific fraud: The case of Diederik Stapel. *PLOS ONE*, 9(8), e105937–e105937. <https://doi.org/10.1371/journal.pone.0105937>
- Markowitz, D. M., & Hancock, J. T. (2016). Linguistic obfuscation in fraudulent science. *Journal of Language and Social Psychology*, 35(4), 435–445. <https://doi.org/10.1177/0261927X15614605>
- Markowitz, D. M., & Hancock, J. T. (2019). Deception and language: The contextual organization of language and deception (COLD) framework. In T. Docan-Morgan (Ed.), *The Palgrave handbook of deceptive communication* (pp. 193–212). Palgrave Macmillan.
- Markowitz, D. M. (2023b). Instrumental goal activation increases online petition support across languages. *Journal of Personality and Social Psychology*, 124(6), 1133–1145. <https://doi.org/10.1037/pspa0000333>
- Markowitz, D. M. (2020a). Putting your best pet forward: Language patterns of persuasion in online pet advertisements. *Journal of Applied Social Psychology*, 50(3), 160–173. <https://doi.org/10.1111/jasp.12647>
- McNamara, D. S., Graesser, A. C., McCarthy, P. M., & Cai, Z. (2014). *Automated evaluation of text and discourse with Coh-Metrix*. Cambridge University Press.
- Meier, T., Boyd, R. L., Mehl, M. R., Milek, A., Pennebaker, J. W., Martin, M., Wolf, M., & Horn, A. B. (2020). Stereotyping in the digital age: Male language is “ingenious”, female language is “beautiful” – and popular. *PLOS ONE*, 15(12), e0243637–e0243637. <https://doi.org/10.1371/journal.pone.0243637>
- Mieczkowski, H., Hancock, J. T., Naaman, M., Jung, M., & Hohenstein, J. (2021). AI-mediated communication: Language use and interpersonal effects in a referential communication task. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1), 17:1–17:14. <https://doi.org/10.1145/3449091>
- Newman, M. L., Pennebaker, J. W., Berry, D. S., & Richards, J. M. (2003). Lying words: Predicting deception from linguistic styles. *Personality and Social Psychology Bulletin*, 29(5), 665–675. <https://doi.org/10.1177/0146167203029005010>
- OpenAI. (2022). *ChatGPT: Optimizing language models for dialogue*. OpenAI. <https://openai.com/blog/chatgpt/>
- Ott, M., Choi, Y., Cardie, C., & Hancock, J. T. (2011). Finding deceptive opinion spam by any stretch of the imagination. *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, 1, 309–319.
- Pennebaker, J. W. (2011). *The secret life of pronouns: What our words say about US*. Bloomsbury.
- Pennebaker, J. W., Boyd, R. L., Booth, R. J., Ashokkumar, A., & Francis, M. E. (2022). *Linguistic inquiry and word count: LIWC-22* [Computer software]. Pennebaker Conglomerates. <https://www.liwc.app>
- Pennebaker, J. W., Chung, C. K., Frazee, J., Lavergne, G. M., & Beaver, D. I. (2014). When small words foretell academic success: The case of college admissions essays. *PLOS ONE*, 9(12), e115844–e115844. <https://doi.org/10.1371/journal.pone.0115844>
- Russell, S. J., & Norvig, P. (2010). *Artificial intelligence: A modern approach* (3rd edition). Pearson.
- Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals and understanding: An inquiry into human knowledge structures*. Lawrence Erlbaum.
- Seraj, S., Blackburn, K. G., & Pennebaker, J. W. (2021). Language left behind on social media exposes the emotional and cognitive costs of a romantic breakup. *Proceedings of the National Academy of Sciences*, 118(7), e2017154118–e2017154118. <https://doi.org/10.1073/pnas.2017154118>

- Susser, D., Roessler, B., & Nissenbaum, H. (2018). *Online manipulation: Hidden influences in a digital world* (SSRN Scholarly Paper 3306006). <https://doi.org/10.2139/ssrn.3306006>.
- Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology, 29*(1), 24–54. <https://doi.org/10.1177/0261927X09351676>
- Vrij, A. (2018). Verbal lie detection tools from an applied perspective. In J. P. Rosenfeld (Ed.), *Detecting concealed information and deception: Recent developments* (pp. 297–321). Academic Press. <https://doi.org/10.1016/B978-0-12-812729-2.00013-6>
- Walther, J. B., & Parks, M. R. (2002). Cues filtered out, cues filtered in: Computer-mediated communication and relationships. In M. L. Knapp & J. A. Daly (Eds.), *Handbook of interpersonal communication* (pp. 529–563). Sage.

Author Biographies

David M. Markowitz is an associate professor in the Department of Communication at Michigan State University. He uses language data from natural repositories to make inferences about what people are thinking, feeling, and experiencing psychologically. His work has appeared in outlets such as the *Proceedings of the National Academy of Sciences*, *Journal of Communication*, and *Journal of Language and Social Psychology*, and he was selected as a “Rising Star” of the Association for Psychological Science in 2022.

Jeffrey T. Hancock is the founding director of the Stanford Social Media Lab and is a Professor in the Department of Communication at Stanford University. Professor Hancock and his group work on understanding psychological and interpersonal processes in social media.

Jeremy N. Bailenson is the founding director of the Virtual Human Interaction Lab and the Thomas More Storke Professor in the Department of Communication at Stanford University. Professor Bailenson studies the psychology of Virtual and Augmented Reality, in particular how virtual experiences lead to changes in perceptions of self and others.