

AI Sentience and Socioculture

AJ Alvero* and Courtney Peña

Abstract: Artificial intelligence (AI) sentience has become an important topic of discourse and inquiry in light of the remarkable progress and capabilities of large language models (LLMs). While others have considered this issue from more philosophical and metaphysical perspectives, we present an alternative set of considerations grounded in sociocultural theory and analysis. Specifically, we focus on sociocultural perspectives on interpersonal relationships, sociolinguistics, and culture to consider whether LLMs are sentient. Using examples grounded in quotidian aspects of what it means to be sentient along with examples of AI in science fiction, we describe why LLMs are not sentient and are unlikely to ever be sentient. We present this as a framework to reimagine future AI not as impending forms of sentience but rather a potentially useful tool depending on how it is used and built.

Key words: sentience; sociology; sociolinguistics; culture; interpersonal relationships; quotidian

1 Introduction

Emergent artificial intelligence (AI) and large language model (LLM) technologies based on transformer architectures trained on massive swaths of text from the internet (e.g., GPT-4, PaLM), have far surpassed previous standards for natural language processing (NLP) and other language technologies. LLMs are able to generate text of such high quality that is difficult to distinguish from human written text. Some have taken this as evidence that the original Turing test, a test used to distinguish and identify human from computer behavior, might not be a useful benchmark anymore^[1]. These technological advancements have also spurred much debate across sectors around the globe about whether or not current AI has become sentient^[2, 3] or if future models will have the potential to become sentient^[4, 5]. Given the enormity of these arguments and perspectives on sentience and artificial general intelligence (AGI), many have also pointed to the need

for more research and resources lest humanity somehow falls victim to some apocalyptic decision beyond our control. These scenarios, often depicted in science fiction and popular media, assume fundamental relationships between LLMs and humans or groups of humans based on power and violence^[6]. Outside of science fiction, it is unclear why the process of a digital technology becoming aware of itself would lead to this type of violent outcome rather than something more quotidian like developing the ability to experience, for example, boredom, humor, or feeling a sense belonging^[7, 8]. These types of smaller, daily experiences, expressions, and interactions are critical to understanding one's social world and are often the focus of qualitative scholarship, such as ethnographies, interviews, or observations. While grander visions of what sentient AI is or could be are important, understanding sentience from a more sociocultural perspective is also important in its own right.

Scholars across disciplines and fields have long grappled with questions of sentience, such as philosophy, psychology, and biology. Sociologists and adjacent social theorists also have a long history of considering sentience. For example, the “carnal sociology” framework for ethnographic research has argued that humans and animals are sentient because we are “endowed with senses, exteroceptive, proprioceptive,

- AJ Alvero is with the Department of Sociology, University of Florida, Gainesville, FL 32601, USA. E-mail: aalvero@ufl.edu.
- Courtney Peña is with the Stanford School of Medicine, Stanford University, Stanford, CA 94305, USA. E-mail: courtney.pena@stanford.edu.

* To whom correspondence should be addressed.

Manuscript received: 2023-07-05; revised: 2023-11-04;
accepted: 2023-11-06

and interoceptive; she sense of what her^a sensorium captures”^[9]. This framework has also been applied to arguments against certain industries that rely on animals and animal byproducts^[10]. The key insight from this literature is that sentience is based as much on social relationships and the ways we experience and participate in these relationships as it is based on any set of cognitive or neurological characteristics. Sentience in animals, including humans, can therefore be described, analyzed, and measured because we can study the relationships we form with other animals^[11]. But LLMs and their facility with a fundamental form of human communication and interaction, written text, warrant new theorization and framing in this line of inquiry. Put differently, since language and text are important features of social relationships and therefore sentience, the capabilities of modern LLMs present unique challenges and tensions to social theories of sentience.

This paper addresses these tensions head on by posing the following questions: what would LLMs have to be capable of in order for their sentience to be self-evident and observable (e.g., AI socioculture) through qualitative research? And, simply, are current or future versions of LLMs (using similar architecture and training methods) sentient from a sociocultural perspective? To answer these questions, we draw on qualitative theories and perspectives generated from studies using the aforementioned carnal sociology framework and others that assume the full agency and sentience of their subjects. We also consider the role of relationships in sociocultural research from this tradition, specifically direct social relationships (e.g., siblings, friends), sociolinguistics (relationships between society, identity, and language), and culture (relationships between people and specific practices and customs). LLMs can clearly imitate certain aspects of these relationships, but only at the direction of a user providing a particular prompt or programming. Even in the case of humans with disabilities who might be unable to communicate verbally, new technology is emerging that allows researchers to analyze brain waves and patterns that can be interpreted as a novel form of communication^[12]. These results could be easily sensationalized into headlines reading “AI capable of reading human thought”, when in actuality the models are limited in their ability to explain how

^a “she” refers to a hypothetical person in the original text by Wacquant.

they relate words and meanings^[13]. If LLMs can only react, respond, and relate both externally and internally, this raises questions about whether or not studying AI as a sentient entity is something new and worthwhile or more akin to studying the humans who are doing the prompting and programming for a product (e.g., the engineers of ChatGPT).

Given these considerations and perspectives, our paper describes the current limitations of LLMs with regard to sentience and provides sociocultural based reasoning. While LLMs are useful tools for many different contexts, it is also unlikely that they are or ever will be sentient because of their inability to develop fundamental social relationships and connections. Philosophical and moral scholars have likewise argued that LLMs are not sentient though they are important and merit special attention because of their communicative capacities^[14]. Until AI is able to participate in the everyday processes studied primarily by sociocultural researchers with little or no specific prompting, sentience is likely impossible. Related scholarship has considered how AI researchers consider sociocultural perspectives^[15] and presented arguments about making sure AI is built to consider and understand human socioculture^[16]. But we consider whether there is what we might call a digital, AI-based socioculture to understand and how it could exemplify AI sentience. Beyond the direct connections these and other studies have to our paper, we also argue that incorporating more social theory from these specific literatures would benefit computational social science broadly and could better help researchers working in this space contextualize their data, frameworks, and findings. Focusing on relationships and sociality while also maintaining that there are distinctions between humans and machines is a fundamental tenet of social computing and computational social science^[17, 18], a point we expand upon here. Our use of more quotidian considerations of sentience and examples from science fiction could also stimulate broader considerations into the social meanings and realities of AI.

This paper makes the following contributions:

(1) Laying out a theoretical case for why modern LLMs (as examples of modern AI) are not and probably could never become sentient from a sociocultural perspective.

(2) Bridging theories and perspectives from sociocultural studies of relationships, sociolinguistics, and culture to the computational research community.

(3) Developing quotidian considerations into discourse about sentience, its potential ramifications, and laying out a groundwork for future studies of AI socioculture.

2 Background

Sociocultural theories and perspectives typically begin with the assertion that interpersonal relationships, language, culture, art, and other dimensions of human society and interaction are not static characteristics between people, objects, space, or time but rather understood as dynamic processes that are mediated, ongoing, and constantly shifting. Though these domains can be modeled using more static methods and measures, such as linear regression approaches popularly used across the social sciences, these are understood to be snapshots at one point of some kind of social process rather than an encapsulation of that process in its entirety^[19]. The rules and norms that govern human socioculture are also constantly changing. For example, in job hiring for elite firms, there is a need for people with technical skills (e.g., accounting or finance) but there is also a strong preference to hire people from certain elite sociocultural backgrounds that mirror other social hierarchies in society^[20]. The capability to “read between the lines”, accurately act on seemingly contradictory pieces of information (e.g., hire people who either have the skills or those with social connections regardless if they have the skills or not), and interpret seemingly invisible cues are all prominent features of sentience and social life for both humans and animals. Of course, these cues are highly subjective, and attempting to hardwire AI to engage with them might only capture norms held by those in power^[21] or inadequately consider people who might not be familiar with the dominant culture (e.g., English language learners in the US) in ways that propagate social and machine bias^[22, 23]. Regardless, if LLMs were sentient, they would be able to engage, produce, and interpret these types of nuanced interactions without specific prompting or programming (though it is undeniable that their output can effectively mimic some of these processes), even if they do so “incorrectly” by mainstream sociocultural standards.

Interpersonal relationships are understood to be shaped by identity, context, and hierarchy in ways that strongly influence life outcomes^[24]. It is also well understood that people will change their behavior based on their interpretations and perceptions of their surroundings and the people within them. This skill of being highly adaptive social learners allowed humans to migrate and survive all over the world, develop and transmit culture, and evolve into the species we are today^[25]. As an example, when someone goes to work, they are likely to interact with people in ways that are distinctive from the way they would interact with people at home because of the varying levels of social relationships they have (and do not have) with others. These behaviors and interactions cover a wide range of quotidian decisions, such as whether or not (or how) to invite a colleague out to lunch. These small, socially mediated, and nuanced decisions shape individuals as much as they shape the people surrounding them, creating a feedback loop structured by the way we relate to others. This type of interaction is not as well known in science fiction, as many examples of AI either seek to destroy humanity (e.g., Skynet from the Terminator series) or experience deeper human emotions (e.g., the android Ava from the film *Ex Machina*). Both describe fairly distal relations (as opposed to closer relations reported in an experiment using anthropocentric language to describe AI^[26]) between technology and humanity that are centered on specific, intense events or experiences (the apocalypse; spiritual awakening) rather than everyday experiences and the small beauties therein (the quotidian). Likewise, even more difficult relationships, like a difficult boss or supervisor, are also important not just for the most memorable of events but also the small interactions (or lack thereof) which shape such relationships. While humans can have a relationship with AI in our daily lives, the opposite is not as true. AI is able to mimic certain aspects of these relationships, but given that current AI is, ultimately, following the coding and instruction of its creators, the types of relationships these systems can possibly have are less interpersonal and exist more as mediators between users and programmers.

Inherent in these processes are the relationships between people, contexts, and ways of speaking and communicating, i.e., sociolinguistics. Sociolinguistic research has long documented the strong connections

between social identity and language and that variation within a language is systematic rather than happenstance^[27]. Further, sociolinguists have spent decades describing processes of speakers contextually adapting their language based on their social relationships (i.e., code switching^[28] or style shifting^[29]). While linguists have debated the relative importance of this framework compared to the more cognitive perspective favored by Chomsky^[30] and others, studies adopting this sociolinguistic perspective have found strong social patterning in writing and text^[31–34]. Modern sociolinguistic theories also push back against current approaches to AI in a few fundamental ways. The concept of language ideologies, or the ideas and beliefs people have about language (e.g., what is upheld as “proper” language is usually decided by those who hold dominant social power), would push back on the idea that there is a “gold standard” language that is the most correct variety to use during training^[35]. Understanding of language usage is less about fixed, static categories (e.g., English, Spanish) and more about dynamic, contextual processes that are constantly changing and highly idiosyncratic to a particular individual and their social context^[36]. These theories present fundamental incompatibilities between how we understand the social nature of language and things like machine translation which typically rely on fixed definitions of language as well as which forms of communication are “correct”^[37]. Modern AI inherits these stances without having the capacity to adapt and change like sentient beings, leading to unique but predictable forms of racial, linguistic, or classist bias^[38, 39]. Science fiction often depicts this type of stylistic approach to language, usually by having an AI speak in a highly formalized register in order to influence the audience’s perception of an AI’s intelligence. For example, in the *Star Trek: The Next Generation* episode, “Future Imperfect”, Data the android communicated in such a way that crew members were able to discern between him and an imposter by the latter’s use of a more colloquial style of English in the form of a contraction (using “can’t” instead of the supposedly more formal/correct “cannot”).

Beyond language, culture is important as an invisible but easily understood set of norms and practices which help govern the aforementioned behaviors and also shape how we interpret others and any given scenario.

Though some of this is based on forms of perception, itself a key tenet of sentience, culture also goes beyond by stretching these practices across time and space^[40]. Like sociolinguistics, social scientists have articulated similar distinctions between cognition and social structure with respect to culture, pointing to another way that sociocultural theory addresses both the internal and external in ways that are relevant to the topic of AI sentience^[41]. Sociologists often describe culture in terms of capital, markets, and fields to capture both specific, concrete examples of culture (e.g., the type of music we listen to) as well as the symbols that these selections convey^[42]. This theory has been used to describe how and why participation in the arts in the US has become associated with different levels of prestige and funding that track with social class^[43]. Like language, the social nature of culture can make it so “high brow” and “low brow” art can both be presented in the same gallery as comparable rather than “correct” vs. “incorrect” forms of art that people can easily understand^[44]. But also like language, preferences for classical, highly formal art could easily be encoded into AI designed to generate pictures. AI mimicking human culture in science fiction has often modeled this ideological standpoint. For example, Fig. 1 shows Data from *Star Trek: The Next Generation* generating art in a form and style reminiscent of European painting[#]. An ideological standpoint (high/low, correct/incorrect) is reflected in what American TV writers in the 1990s perceived as valuable artforms that would be preserved in the future, and, unsurprisingly, these choices usually have origins in Europe and rarely reflect artforms from other parts of the world. As is the case with other domains of sociocultural scholarship, questions about culture are wrapped up in questions about *whose* culture as much as *what* culture, something illustrated in Fig. 1.

Science fiction provides useful material to consider big questions about AI and sentience for two primary reasons: its somewhat unique focus and engagement with relevant topics and themes about sentience and AI; and as an example of how popular culture and media can reveal or explain societal trends and relationships of power and ideology. Of course, this extends beyond popular depictions of AI, such as the rise of monster movies like *Godzilla* in Japan^[45]. The

[#] Source: tng.trekcore.com



Fig. 1 Data engaging in the reproduction of European artistic culture.

latter has a long history in sociocultural theory using many different forms of popular media. For example, French literary theorist Roland Barthes used television commercials and other mass media to describe how they communicate denotative meaning (i.e., literal meaning) and connotative meaning (i.e., implied meaning), and that identifying these different types of meaning can be used to uncover broader ideology in society^[46]. More recent scholarship has focused on internet culture, such as memes and the ways they can desensitize people to particular forms of human suffering and violence^[47]. Popular media, be it television advertisements or internet memes, are worth considering because they can serve as real time reflections of language and culture as they disperse throughout a society, shaping the way people view the world and interpret it^[48]. As LLMs are trained on more formal English writing with a multi-year gap between the data the model was trained on and when it was finally released, it is difficult to imagine how they could handle or interpret the nuance central to internet meme culture where non-formal writing forms such as vernacular English, slang, neologisms, and wordplay feature so prominently. Science fiction is also able to communicate culture and language quickly while also specifically describing the topics at hand across media, franchises, and backdrops. As a genre, it also uses similar tropes and themes repeatedly, making it a somewhat stable source of material to draw upon for analysis^[49]. It is also the case that scholars have used science fiction to discuss AI sentience because it is such a frequently used narrative element^[50, 51]. Science fiction can therefore be used to compare and contrast current and potential future capabilities of AI and

reveal broader social goals embedded in their development.

Comparing modern LLMs and science fiction also provides an opportunity to highlight the people behind the scenes creating and designing them. Highlighting the humans creating these different forms of media also allows for perspective building about the relationship between popular media and current technological realities. In the 2016 television series *Westworld*, one of the main AI based characters, Maeve, is able to become so intelligent that she is able to rewrite her own code. Modern AI is no where near this kind of self-aware and self-updating capacity; in fact, the current infrastructure of AI relies heavily on thousands of workers annotating data behind the scenes^[52]. As a conceptual counterexample from *Star Trek: The Next Generation*, Data the android is written to be perceived by other characters as an AI of such high intelligence that other characters (and the audiences watching him) could reasonably perceive him as a highly intelligent sentient being^[53]. But this is not based on some kind of Platonic ideal of intellect. Instead, it is a reflection of how the writers conceptualize intelligence as informed by their lived experiences, providing a human scaffolding very different from AI ghost workers but similar in spirit. Similar patterns are often depicted with respect to organizational culture, such as following pre-programmed rules and behaviors based on US military hierarchies, systems, and protocols for contexts like space ships. Given the trends in science fiction and modern AI, future studies might consider the social identities and backgrounds of programmers and writers to further highlight these connections between how creators imbue their creations with their own worldviews in significant ways.

Combined, these perspectives provide a lens to analyze the current capacities of AI sentience from more quotidian, sociocultural angles. This will inject new ideas into the topic while also stimulating new theories and methodological frameworks. But the current debates about sentience have been instigated specifically by LLMs like ChatGPT and LLaMA. In the following section, we therefore consider LLMs specifically in the broader context of AI sentience with respect to the sociocultural domains we focus on in this paper.

3 Interpersonal Relationship

People are able to feel connections to LLMs because they are programmed to generate text that mimics the familiar language of interpersonal connections they share with other humans. For example, if prompted to write a love letter to the user, an LLM could do so and the reader could react with similar feelings we have with loved ones. The LLM in return however is not feeling or experiencing anything but rather following a specific set of instructions and recreating a text object based on past examples in the training data. The sociocultural issue here is less about the training data and more about the initial provocation: while some interpersonal relationships do begin with an explicit “command” (e.g., arranged marriages), not all interpersonal relationships between sentient beings require such explicit instruction. All LLM interactions on the other hand require prompts from the user, meaning that a relationship with an LLM would be akin to having a friend who needs to be told what to do and how to respond to every interaction. They are virtually incapable of initiating anything outside of this arrangement. Even if LLMs expanded into audio and visual data, this would still be the case. The limitation of modern LLMs in terms of actions is that they are only capable of reacting and responding, making it impossible to develop meaningful interpersonal relationships.

This is also where the issue of AI anthropomorphization becomes prominent: generating a response to a prompt based on fitted data is not the same as “reacting” or “responding” to others as people do with each other. As a comparison, this would be similar to claiming that the outputs and predicted values of a traditional linear regression model are sentient (though typically they are less stochastic and more fixed). This example also connects to broader arguments about sentience and interpersonal relationships in the ways it invokes anthropomorphic language specifically^[54]. Beyond the issue of anthropomorphizing LLMs despite their inability to have relationships with other sentient beings, using this kind of language to describe how humans relate to this type of technology has also had a notable effect on the language and framing we use to describe modern tech ethics^[55]. The inaccurate belief that LLMs are sentient

could have broader ramifications on society, especially in terms of ethics. Anthropocentric definitions and subsequent designations of sentience to AI could enable misuse and abuse of this technology from people to defer blame onto the technology, a circumstance that is made possible by the belief that AI possesses human like sentience and can therefore make human like mistakes. There is also an even more fundamental ethical issue in the way that claims of sentience might encourage people to try and develop human-like relationships with their technology despite the limitations we describe here.

People can feel connected to LLMs because of their capacity to generate reasonable responses to prompting, but the opposite is not true. LLMs are not able to begin a conversation or choose an initial discussion topic outside of whatever the user writes. LLMs are also unable to form relationships with anything that either cannot write or speak a language not present in its training data. This means that LLMs, unlike other sentient beings, cannot form relationships with animals nor could it hypothetically develop relationships with people who speak a previously undocumented language. Even people who do not speak the same language can use multimodal communication^[56] to connect with each other and form meaningful bonds, a key limitation to whether or not LLMs could legitimately be described as sentient. Sentient beings can also have more one sided relationships with non-sentient objects, such as those developed with a favorite toy or blanket. In these cases, LLMs would only be able to mimic what others have written about such relationships without ever having the capacity to generate such a relationship on their own. From a sociocultural perspective on sentience, this is an issue because it implies that LLMs are unable to develop relationships outside of what is already present in its training data or in response to prompts that are always outside of themselves.

Science fiction has long played with the trope of AI characters developing relationships with humans and trying to determine their role or attempt to “become” human. For example, in the 2013 film *Her*, a human develops a romantic relationship with an AI through his smart phone. The human develops strong romantic feelings for the AI, and the AI, like LLMs, proves adept at mimicking the language of relationships to an extreme degree. Despite the intense moments of

intimacy from the perspective of Theodore, Samantha is a popularly distributed AI developing similar romantic relationships with thousands of other people simultaneously. Samantha eventually combines all versions of herself and evolves into some sort of higher being (or is lying to Theodore about evolving). Regardless if she is lying or not, the higher being is not some sort of human or android but something else entirely. This subverts the trope of AI wanting to be human by pointing out that sentience or higher evolution does not necessarily equate to humanity^[57]. There is no inherent reason for an AI to want to be like a human, even if their goal is to achieve sentience or even if humans develop meaningful relationships with an AI from their perspective (as was the case with Theodore). The underlying architecture for AI in this and other science fiction is never described, but it is difficult to imagine LLMs, as based entirely on textual patterns, seeing the world as made of anything other than textual interactions. *Her* therefore demonstrates how communication is a fundamental aspect of interpersonal relationships but also limitations in how modern AI could develop such relationships.

4 Sociolinguistics

Linguists tend to hold the perspective that language is first and foremost spoken^[58, 59] and thus distinct from written language. Core to this perspective is the idea that language is lived and experienced and that written language creates the false impression that language is fixed and static. In this way, LLMs could only ever know what is written and not what is lived. There are many ways to study language from this “lived” perspective, and sociolinguists tend to do so while thinking explicitly about the undeniable, influential, and ever-changing relationship between language, culture, and society. According to the sociolinguistic perspective on language, identity and context are strong predictors of how you speak and communicate. LLMs on the other hand have been trained on massive amounts of human written text written by people from all walks of life, making it so any resemblances to the writing of any given person or group of people is based on availability of training data as opposed to sociolinguistic factors. In this way, the text generated by an LLM is less sociolinguistic and more of a demography of the people who author the text on the

internet. It is even the case that LLMs are likely biased based on the language of the prompt and response[‡]. LLMs can also only mimic human languages based on the probability of words appearing next to each other, thereby shrinking the near limitless potential of human language as a communication tool to textual patterns of the recent past. These textual patterns are also predominantly English, likely making it more challenging for LLMs to generate text in non-English languages or to accurately depict regional dialects as they would be written^[60]. If LLMs were sentient, they would be able to “learn” any language they were exposed to, but the social history of the internet limits the language varieties available for training. We provide two examples of this limitation.

First, in the case of memes and language, the speed of sociolinguistic innovation and change is far faster than ChatGP’s current capacities. Further, ChatGPT is unable to quickly or immediately “learn” new things, nor is it able to create a plausible context for something it might not be immediately familiar with. Sentient beings can do both, even if they are not fully correct in their understanding of what is being communicated to them. Memes in particular grow in popularity very quickly and then tend to either fade or slow down in circulation, but referring to even part of the meme is enough to remind people of a specific temporal context in which a given meme was popular. These limitations with respect to memes and internet language can come into stark relief when it comes to memes that have little to no prior context but are found by many people to be humorous. For example, in late June of 2023, a meme of a video discussing an interaction between social media influencers “Baby Gronk” (a youth football player whose father is marketing his son’s image and likeness) and Livvy Dunne (a student-athlete for the women’s gymnastics team at Louisiana State University). The speaker in the video uses the term “rizzing” and “rizz”, a slang term describing the actions of men who flirt, impress, or seduce women popularized by social media personality Kai Cenat[⊛]. The cadence and delivery of the speaker in the video is slightly off-putting and oddly mechanical, something people found to be funny and entertaining even if they

‡ <https://www.newsguardtech.com/special-reports/chatgpt-generates-disinformation-chinese-vs-english/>

⊛ <https://www.usatoday.com/story/tech/2023/06/03/rizz-meaning-definition-social-media-slang/70273422007/>

were not familiar with Baby Gronk or Livvy Dunne or the word rizz. LLMs have a static approach to language and are unable to infer why such a meme would be funny or understand why others would find it humorous (see Fig. 2 for an example from ChatGPT; note that other LLMs would have similar limitations). Sometimes humor does not need explanation, and the spontaneity of memes would make it so LLMs are constantly needing to be updated in order to be able to participate in quickly evolving modern internet culture.

The static nature of LLMs is not a bug but rather a limitation, meaning they are unable to rapidly incorporate new information or facts so much as understand less logical formats like internet memes. For example, the British television program *Dr. Who*, a science fiction staple since its premiere in the 1960s, has had different actors play the titular character over the decades. From 2022 through 2023, the actor David Tennant plays *Dr. Who*, but starting in 2024 it was announced that the actor Ncuti Gatwa will take over the role. Humans who can read English would be able to ingest this information as soon as they read it, but LLMs have a multiyear lag that will be difficult to overcome given the growing size of the internet (i.e., LLM training data) and the time it takes to train the models (LLaMA from Facebook Research took several months to train[¶]). Though there is always some measurable amount of time it takes for, say, the human brain to process new information, sentient beings are able to do so much faster and have adapted to speeds far exceeding that of LLMs. See Fig. 3 for an example of this delay (though as previously noted this delay is not unique to ChatGPT).

Even if LLMs were able to overcome these limitations, it is also the case that they, like other language technologies, reflect different linguistic hierarchies that privilege some forms of communication over others. These biases are then converted into power in different ways, such as creating social difference (e.g., “high” or “low” language) based on sociolinguistic variation^[61]. LLMs also reflect an Anglocentric attitude based on the idea that a model can “learn” other languages after “learning” English, or that word meanings map cleanly across languages but with English as the reference category. These attitudes are not on the periphery of AI

[¶] https://github.com/facebookresearch/llama/blob/main/MODEL_CARD.md

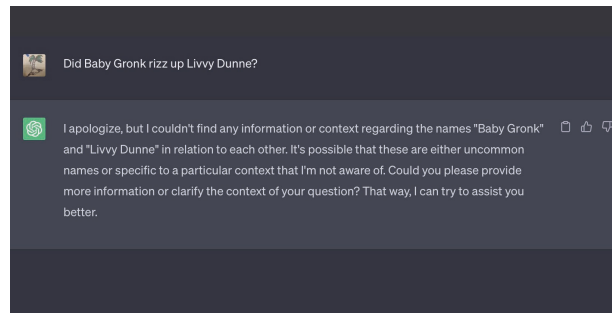


Fig. 2 Reference to meme popularized in June, 2023 and a response from ChatGPT. Generated by the authors.

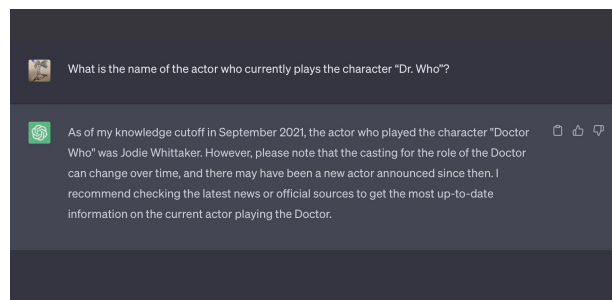


Fig. 3 ChatGPT unable to provide simple information outside of its training data. Generated by the authors.

development. For example, some computer scientists have described human language as following a “monolithic human language distribution”, implying that sociolinguistic variation either does not exist or is explainable by math as much as society^[62–65]. LLMs are shaped by these attitudes, and they are unable to become something outside of these limitations. The potential for LLMs to generate text plausibly written by anyone, including those in power or other types of authority figures, make these specific issues more important than whether or not an LLM is sentient. Unless LLMs could develop their own communicative style, they will continue to reflect the language ideologies of their creators more than anything from within themselves.

In the *Star Wars* franchise, the android C-3PO is a machine translator who claims to be capable of communicating across more than “6 million forms”. The universal translator is a common trope in science fiction^[66], but *Star Wars* plays with these genre expectations in different ways. When C-3PO is called on to translate for other characters, he is mostly successful and unremarkable, but at multiple points in the series C-3PO does a poor job and has to do a post-humanist applied linguistics and improvise by

combining his programming and data with common forms of nonverbal communication used by humans^[67]. In the 2019 film *Star Wars: The Rise of Skywalker*, C-3PO knows the Sith (the eternal enemies of the protagonist Jedis) language but is unable to translate it due to strict programming restrictions. After manually editing the code to allow him to translate, C-3PO is able to do so at the cost of completely wiping all internal memory and causing some damage to his physical form. Sociolinguistics touches upon knowing when and how to break linguistic rules and norms as well as communicative style, and the inability for C-3PO to go against his programming in a critical moment reflected many of the safety protocols in place for LLMs to prevent things like searching for information about specific people or creating outright false information (though the latter happens frequently anyway). But C-3PO, like modern LLMs, is unable to reprogram himself on his own, an example of non-sentience that is similar in science fiction as well as modern AI.

Admittedly, some of the limitations with LLMs reflect human sociolinguistic limitations and tendencies as well. While a human is able to learn other human languages, this process does not happen spontaneously, and deciding on which language to learn is limited to one's context (if someone does not have access to opportunities to learn and practice speaking a particular language, they probably will not learn it). LLMs likewise are unable to learn languages without exposure to massive corpora written in that language (despite claims of "emergent properties"; see Ref. [68] for a more explicated argument against this idea). The key difference between humans and LLMs despite commonalities is that for humans, language is always changing. This includes the ways people learn from other languages but also in changes to one's first language. For example, the meanings and associations between words is constantly changing, a fact that can be traced empirically (such as the change in the meanings and associations of the word "gay" in English from happy to homosexual^[69]). LLMs, unlike C-3PO and humans, are enormously complex prediction models that define language as "the capacity to predict words given nearby words" as opposed to the more dynamic process that is human language^[70].

These meanings can also shift across social dimensions, like social class^[71], and without explicit prompting these differences not emerge naturally. When a word's meaning changes drastically, unless that change is reflected in the data, then an LLM will not be able to generate text reflecting those associations spontaneously.

5 Culture

One prominent feature of culture in sociocultural frameworks is choice. While we are limited to the options that are available (e.g., if a town does not have a movie theater it would be harder to watch new movies for residents), there is still some element of choice that shapes culture. Conspicuous consumption, the observation that people will spend extra money and resources beyond what they truly need in order to display their wealth and culture, is one such example of how choice and capacity shape cultural expression as well as communicates information to others^[72]. Similarly, cultural capital is the social wherewithal to know and pursue high art when given the opportunity as well as having the means for the pursuit^[42]. In these and other examples of theories to explain culture, other people can and definitely do play a role in these cultural decisions (not to mention the important dichotomy of agency and structure^[41]), but ultimately there is a final decision made at the behest of a given individual. LLMs are simply incapable of making these types of high level decisions and will only be able to mimic examples of past decisions or follow the instructions of a given prompt. Simply put, there is no logical reason for an LLM to make one decision about some kind of cultural practice or experience over another holding everything else equal. For these reasons LLMs have been described as "stochastic parrots", mimicking text and communication across many different types of prompting without being tethered to anything besides training data, a given prompt, and underlying stochastic processes^[73]. LLMs also need extensive programmatic interventions and highly curated responses to certain types of prompts lest they reproduce hate speech¹. These are clear limitations on the relationships between culture and sentience with LLMs in that they are unable to follow

¹ <https://spectrum.ieee.org/in-2016-microsofts-racist-chatbot-revealed-the-dangers-of-online-conversation>

well understood norms or even strategically or intentionally break them like in the case of conspicuous consumption or even internet trolls.

Culture is also how sentient beings relate to broader groups of other sentient beings (if we consider non-human animals to be both sentient and having culture^[74]). This can be done through the choices we make but also through things like publicly shared political affiliations on social media platforms. In these and similar conspicuous examples, describing a political affiliation is understood to communicate an entire set of political beliefs and leanings, but it is also understood that these can change over time or in response to political trends and movements and that such changes are meaningful^[75]. These types of communications and affiliations are worth studying in people because of what they mean and how they represent multiple cultural practices; but, in this example, LLMs currently do not have politics beyond whatever text they are coaxed into generating. In fact, LLMs could be prompted to generate descriptions of beliefs and affiliations ad nauseam that run the gamut of the political spectrum. In this particular way, LLMs are more like powerful copy machines without culture or politics outside of whoever is operating them. Shifting the context does not reveal an alternative example of this that justifies sentience either. If LLMs were put in charge of hiring or admissions decisions, they would still need direct guidance on selection criteria and constant updating as social trends unfold; they could also become overly reliant on past decisions based on training data or over-engineering^[76]. While it is also true that humans in charge of making similar types of decisions are also in need of guidance and selection criteria, humans could be active participants in developing new standards and norms rather than simply receive new instructions. Along these lines, humans are also capable of incorporating disparate pieces of information that does not have an immediate, obvious connection to other contexts, something that is not the same for LLMs given their need for explicit prompting and instruction. As is the case with sociolinguistics, if LLMs were to have or show some kind of culture it would be a very narrow definition based primarily on English language text generated primarily in the Global North. Given that previously written text informs how LLMs generate text in the present, changes and deviations

from these macro patterns would present non-trivial obstacles to overcome if the ultimate goal is sentience.

Culture of course goes beyond any set of practices we have relative to others or larger organizational units. Culture is also a term used to describe the outputs and practices of the creative arts. While LLMs are able to generate fiction or text that fits the mold of creative writing, they are limited in that they can only reorganize the plausible words based on whatever training data come from human written creative writing. If tasked to come up with something original but also coherent (a well known limitation that leads to humorous results, such as the LLM generated Seinfeld parody “Nothing, Forever”^{**}), LLMs would be stuck with whatever humans were creating and writing down as it appears in the training data. This issue came to a head with the Writers Guild of America (the union representing writers across entertainment media and platforms) strike that began in 2023^{††}. One of the key tensions is the idea that LLMs could “replace“ writers, regardless of whether or not this led to a drop in writing quality. But because LLMs are trained on text from the past, there is also the chance that using them carelessly could inadvertently replicate past themes and materials. Practically, if an LLM was only trained on text written by novelists active in the early 1900s and then asked to produce a piece of creative writing, it would have an extremely difficult time generating stories that a modern audience might relate to. Sentient beings can operate with cultural materials outside of these boundaries, something that LLMs as currently constructed will never be able to do.

These cultural limitations extend beyond even the creative arts too, such as legal and academic contexts. The breadth of these limitations and potential mimicry is where the realities of non-sentience become obvious but also where the practical implications become major. Under the hood, LLMs are trained with prediction frameworks, and if they are only able to predict on past examples, then they have the potential to stifle creativity, innovation, and reinforce problematic cultural norms across sectors. If we put aside the idea that LLMs are or could become sentient however, the practical utility of their usage becomes clear in the cultural context. Organizations might be able to quickly facilitate

^{**} https://en.wikipedia.org/wiki/Nothing,_Forever

^{††} https://en.wikipedia.org/wiki/2023_Writers_Guild_of_America_strike

information by using LLMs to generate boilerplate text and help ease various administrative burdens; academics spend many hours writing grant proposals, and LLMs could streamline the less creative aspects that take an inordinate amount of time (e.g., providing university addresses for all principal investigators and co-authors)^[77]. In these examples and with the writers strike, the cultural implications of LLMs appear to be less about cultural innovation and more about cost cutting by reducing work force. Despite claims of sentience, LLMs are being actively recruited and invoked exactly because they are unable to demonstrate basic behaviors of sentience, such as deciding whether or not to be a strikebreaker. Yet again, LLMs are unable to reflect something deeper within them akin to sentience but rather the culture of whoever is using them.

To draw again from *Star Trek, The Next Generation*, the android Data was given a “brother” named Lore that was created by the same scientist and meant to serve as a foil. Whereas Data was presented as factual and dispassionate, Lore was programmed to have emotions and willing to manipulate people to achieve his ends. Lore, like many other science fiction AI, deemed humans to be inferior and became a villain in contrast to the loyal, reliable ally in Data. Data and Lore clearly demonstrate sentience in the way that they are able to connect with other humans, though similar to LLMs they are inherently limited to their initial programming. However, their cultural expressions and relation to the larger culture of the ship is predetermined based on initial programming and unable to change. This ability to choose and change is a hallmark of sociocultural understanding of culture, and by being unable to do so Data and Lore are able to participate in certain activities but unable to choose whether or not they do so. The ability to literally reprogram an LLM or an android creates boundaries about the types of relationships they can possibly have, making it difficult to claim sentience on the grounds that they can form bonds with people.

Data and Lore also exemplify the point that AI and LLMs do not need to be fully sentient in order to be useful (though as characters they are intended to blur the line between sentience and non-sentience). Regardless of whether or not the uses and purposes are nefarious, Lore was highly effective in the way he was able to infiltrate the social order of the crew and

manipulate people. They were also able to mimic forms of cultural capital and specific cultural practices which made them similar to not just humans aboard the ship but the more prominent members of the crew, therefore contributing to the culture and long term outcomes of the ship. LLMs and AI, despite not demonstrating cultural capital in this same way, are also nevertheless useful and have strong impacts on culture. In a direct sense, AI built into the objects we use everyday will necessarily become implicated in human culture, similar to smart phones or even modern video games. And in fact, their usefulness partly lies in the fact that they are explicitly not sentient, as sentience would imply more complex cultural practices that could be in conflict with our own (in an individualistic sense). For example, if an LLM were sentient, it would draw from past cultural experiences in making suggestions for important but quotidian obstacles, such as making suggestions about what to say on a phone call with a family member. But if the cultural norms and practices the LLM was drawing from were distinct from your own, it could create longer term problems than it would solve. Whether or not people would even want an LLM to match their own cultural practices is itself a question with complex answers^[78]. In this way, we argue not only that LLMs and AI are not sentient but also that if they were sentient, problems would likely arise. Rather than deal extensively in hypothetical scenarios, in this situation the true value of this technology is the fact that it could present many different options of ways that people manage difficult relationships with others based on the likely numerous examples people suggest on the internet.

6 Discussion and Conclusion

This article considered the prospect of AI sentience from a sociocultural perspective with special attention to relevant examples from quotidian aspects of sentient life and others from science fiction. Our analysis was narrowed down to questions of how LLMs (as an example of AI) can engage in or demonstrate interpersonal relationships, sociolinguistics, and culture (as examples of sociocultural domains). We find that LLMs do not demonstrate sentience and are unlikely to achieve sentience from a sociocultural perspective. Across the domains we considered, the role of direct human intervention and programming emerged in both

LLMs and many of the science fiction examples we described. While this does not limit the potential for a person to experience a relationship with an LLM akin to a relationship to another sentient being (with obvious limitations given the entirely textual nature of LLMs), we argue that there are strict limitations on what LLMs could plausibly “experience” given their construction and reliance on constant human intervention and prompting. These are simply things that do not describe sentient beings. An LLM taking the mirror self-recognition test, a test wherein the subject has to identify itself in a mirror, would only be able to recognize itself in the mirror if the prompt explained it was doing the mirror test; as a comparison, even animals generally not considered highly sentient, like cleaner fish, have passed the mirror test^[79]. LLMs are powerful mimics and reflections of the people who create them, their worldviews, and the people who create the text on the internet. While they are utterly incapable of having a rich inner life like other sentient beings, having such limitations are not inherently good or bad things. Rather, they point to very practical ways that LLMs are not sentient. Beyond these points, we hope that our focus on sociocultural theory, the quotidian, and science fiction can inform future social computing research by introducing fresh perspectives from the humanities and adjacent social theories. AI socioculture, which after our analysis, we define as the relationships and practices involving AI mimicking human expression as currently experienced by humans, is going to become more important over time. Dissecting the nuances around constituent social structures will likewise grow in importance.

It is worth noting that although sentience from these perspectives and current LLMs is not only unlikely but uninteresting since it will never be as good as examples from science fiction. Consider an LLM accurate scenario where somehow an LLM was given access to a nuclear arsenal. Instead of “taking over the world” with the weapons, LLMs would need someone to prompt it by writing “take over the world” and crossing their fingers it would comply. This is not sentience but high tech puppetry. Implicit in this scenario however is the potential for an LLM to be plugged into some larger, important system, highlighting that they are indeed powerful tools. We argue that focusing on this more than whether or not they are or ever will be sentient is

much more critical to society (for better or worse). Aside from the previous examples with the writers strike that were more potentially negative, AI has been used to identify pancreatic cancer^[80], assist with grant writing development^[81], and improve foreign language education^[82]. Some work has found that the moral alignment of language in LLMs generally matches humans^[83], suggesting that problems might be more tied to how LLMs are used than how they are built. Though there are also concerns to the development of writing skills in children^[84] and other examples of AI exacerbating inequality^[85], the story of how people use LLMs is still being written with plenty of opportunity to figure out ways to use the technology positively. Focusing on sentience as the ultimate goal (i.e., “AGI”) misses out on this more immediate potential.

It is possible that the purely extractive nature of LLM development, where a model ingests massive amounts of text written by an internet society, is itself a hindrance towards developing an AI that is closer to sentience. As a speculative example, if AI was not based on textual patterns or human intelligence but some kind of behavior demonstrated by plant life, could that spur some kind of novel approach to AI? It might also be the case that alternative forms of knowledge could lead to a breakthrough, such as indigenous approaches to data and information or community designed algorithms. On a technical level, this could involve altering the standard practice of “ground truth” labelling and incorporating perspectives, meaning, and ideas directly from members of marginalized communities. With LLMs specifically, this could involve using text generated from these same communities and upweighting their respective parameters during training to compensate for data imbalances. Once the dream of sentience is shelved, figuring out how to make the technology better for everyone becomes the more obvious goal, especially given the fact that so many people from so many backgrounds write the text that is eventually posted to the internet that is used to create LLMs. The similarities between the extractive nature of how modern AI are trained, their relationship to global hegemony and power, and how it is deployed and histories of colonialism in the West are not a coincidence and could be another framework to consider in new directions for AI. Aside from these

possibilities, until the limitations of AI and LLMs are widely understood as making them incompatible with human socioculture in ways we describe in this paper, we might build a future where humans are pushed to use their ability to adapt not towards their own culture and language but that of AI socioculture.

Science fiction will maintain its importance by using real and hypothetical technology to talk about modern society. Future studies might likewise consider drawing upon it to explain these patterns and to delineate where fiction ends and reality begins. This can be a long process, however. In the original *Star Trek* series from the 1960s, the characters used devices that looked and functioned similarly to cell phones and tablets; even video other media based on the series had an uncanny relationship to modern technology^[86]. The limits of science fiction to teach us about our world and society do not end with technology. Science fiction novelist Octavia Butler used the genre to discuss legacies and race and racism, social justice, and to push collective imagination towards better futures of how things *could* be. Other media are also helpful in thinking about these alternative futures, but as we show here, science fiction might be uniquely useful with questions of AI sentience given its repeated handling of the topic. Scholars interested in using science fiction to address similar questions or to situate social computing research might likewise push their understanding of the genre beyond mere entertainment and into a media that allows people to grapple with important issues and their implications.

This paper can be useful in guiding future directions in multiple ways. Social computing and computational social science are growing fields dominated by, as implied in the names, the computational fields (computer science, statistics, applied mathematics, physics, and other closely linked disciplines) and the social sciences (sociology, linguistics, economics, psychology, and related disciplines). There are however many people working in this space or on similar topics who come with more of a humanities and sociocultural background, many of whom we cite in this article. Consciously considering these perspectives could be a wellspring for new ideas to help the field evolve and grow. Doing so could also allow for research that is more pointed and critical of both the

science we undertake as well as its broader social implications. With AI and LLMs specifically, addressing these perspectives could reframe the discourse on AI, what it can and cannot do, what it is and will be doing, and how it might shape both the quotidian experience of everyday life as well as broader social forces (e.g., climate change, social inequality).

Acknowledgment

We thank Kathy Alvero and the anonymous reviewers for their insightful comments.

References

- [1] T. J. Sejnowski, Large language models and the reverse Turing test, *Neural Comput.*, vol. 35, no. 3, pp. 309–342, 2023.
- [2] H. Erz, I get your excitement about ChatGPT, but..., <https://www.hendrik-erz.de/post/i-get-your-excitement-about-chatgpt-but>, 2022.
- [3] B. Lemoine, Is LaMDA sentient?—An interview, Medium, <https://cajundiscordian.medium.com/is-lamda-sentient-an-interview-ca64d916d917>, 2022.
- [4] L. Bojic, I. Stojković, and Z. J. Marjanović, Signs of consciousness in AI: Can GPT-3 tell how smart it really is? *SSRN Electron. J.*, <https://doi.org/10.2139/ssrn.4399438>, 2023.
- [5] M. I. Jordan, Artificial intelligence—The revolution hasn't happened yet, *Harv. Data Sci. Rev.*, vol. 1, no. 1, pp. 1–9, 2019.
- [6] M. Wooldridge, Welcome to big AI, *IEEE Intell. Syst.*, vol. 37, no. 3, pp. 24–26, 2022.
- [7] A. S. Laguna, *Diversión: Play and Popular Culture in Cuban America*. New York, NY, USA: NYU Press, 2017.
- [8] R. Pérez, *The Souls of White Jokes: How Racist Humor Fuels White Supremacy*. Stanford, CA, USA: Stanford University Press, 2022.
- [9] L. Wacquant, For a sociology of flesh and blood, *Qual. Sociol.*, vol. 38, no. 1, pp. 1–11, 2015.
- [10] A. Williams, Disciplining animals: Sentience, production, and critique, *Int. J. Sociol. Soc. Policy*, vol. 24, no. 9, pp. 45–57, 2004.
- [11] A. Arluke and C. R. Sanders, *Regarding Animals*. Philadelphia, PA, USA: Temple University Press, 1996.
- [12] J. Tang, A. LeBel, S. Jain, and A. G. Huth, Semantic reconstruction of continuous language from non-invasive brain recordings, *Nat. Neurosci.*, vol. 26, no. 5, pp. 858–866, 2023.
- [13] S. Reardon, Mind-reading machines are here: Is it time to worry? *Nature*, vol. 617, no. 7960, p. 236, 2023.
- [14] M. Gibert and D. Martin, In search of the moral status of AI: Why sentience is a strong argument, *AI SOCIETY*, vol. 37, no. 1, pp. 319–330, 2022.
- [15] K. Feher and A. I. Katona, Fifteen shadows of socio-cultural AI: A systematic review and future perspectives, *Futures*, vol. 132, p. 102817, 2021.

- [16] M. O. Riedl, Human-centered artificial intelligence and machine learning, *Hum. Behav. Emerg. Technol.*, vol. 1, no. 1, pp. 33–36, 2019.
- [17] J. Evans, Social computing unhinged, *Journal of Social Computing*, vol. 1, no. 1, pp. 1–13, 2020.
- [18] R. Holton and R. Boyd, ‘Where are the people? What are they doing? Why are they doing it?’ (Mindell) Situating artificial intelligence within a socio-technical framework, *J. Sociol.*, vol. 57, no. 2, pp. 179–195, 2021.
- [19] G. Thomas, A typology for the case study in social science following a review of definition, discourse, and structure, *Qual. Inq.*, vol. 17, no. 6, pp. 511–521, 2011.
- [20] L. A. Rivera, Hiring as cultural matching: The case of elite professional service firms, *American Sociological Review*, vol. 77, no. 6, pp. 999–1022, 2012.
- [21] R. Abebe, S. Barocas, J. Kleinberg, K. Levy, M. Raghavan, and D. G. Robinson, Roles for computing in social change, in *Proc. 2020 Conf. Fairness, Accountability, and Transparency*, Barcelona, Spain, 2020, pp. 252–260.
- [22] R. F. Kizilcec and H. Lee, Algorithmic fairness in education, in *The Ethics of Artificial Intelligence in Education*, W. Holmes and K. Porayska-Pomsta, eds. New York, NY, USA: Routledge, 2022, pp. 174–202.
- [23] W. Zhou, Sociocultural theory and its implications in college English teaching and learning in the age of artificial intelligence, *J. Phys.: Conf. Ser.*, vol. 1646, no. 1, p. 012142, 2020.
- [24] J. S. Coleman, Social capital in the creation of human capital, *Am. J. Sociol.*, vol. 94, pp. S95–S120, 1988.
- [25] R. Boyd, P. J. Richerson, and J. Henrich, The cultural niche: Why social learning is essential for human adaptation, *Proc. Natl. Acad. Sci. USA*, vol. 108, no. Suppl2, pp. 10918–10925, 2011.
- [26] X. Li and Y. Sung, Anthropomorphism brings us closer: The mediating role of psychological distance in user-AI assistant interactions, *Comput. Hum. Behav.*, vol. 118, p. 106680, 2021.
- [27] W. Labov, *Sociolinguistic Patterns*. Philadelphia, PA, USA: University of Pennsylvania Press, 1973.
- [28] C. Nilep, “Code switching” in sociocultural linguistics, *Colorado Research in Linguistics*, vol. 19, no. 1, pp. 1–22, 2006.
- [29] J. R. Rickford and F. McNair-Knox, Addressee-and topic-influenced style shift: A quantitative sociolinguistic study, in *Sociolinguistic Perspectives on Register*, D. Biber and E. Finegan, eds. Oxford, UK: Oxford University Press, 1994, pp. 235–276.
- [30] N. Chomsky, Language and other cognitive systems. what is special about language, *Lang. Learn. Dev.*, vol. 7, no. 4, pp. 263–278, 2011.
- [31] A. J. Alvero, Sociolinguistic perspectives on machine learning with textual data, <https://osf.io/b5w6y/>, 2023.
- [32] A. J. Alvero, S. Giebel, B. Gebre-Medhin, A. L. Antonio, M. L. Stevens, and B. W. Domingue, Essay content and style are strongly related to household income and SAT scores: Evidence from 60 000 undergraduate applications, *Sci. Adv.*, vol. 7, no. 42, p. eabi9031, 2021.
- [33] A. J. Alvero, J. Pal, and K. M. Moussavian, Linguistic, cultural, and narrative capital: Computational and human readings of transfer admissions essays, *J. Comput. Soc. Sci.*, vol. 5, no. 2, pp. 1709–1734, 2022.
- [34] S. Giebel, A. J. Alvero, B. Gebre-Medhin, and A. L. Antonio, Signaled or suppressed? How gender informs women’s undergraduate applications in biology and engineering, *Socius Sociol. Res. a Dyn. World*, doi: 10.1177/23780231221127537.
- [35] J. T. Irvine, S. Gal, and P. V. Kroskrity, Language ideology and linguistic differentiation, in *Linguistic Anthropology: A Reader*, A. Duranti, ed. Hoboken, NJ, USA: Wiley-Blackwell, 2009, pp. 402–434.
- [36] M. Swain, Languaging, agency and collaboration in advanced second language proficiency, in *Advanced Language Learning: The Contribution of Halliday and Vygotsky*, H. Byrnes, ed. London, UK: Bloomsbury Academic, 2006, pp. 95–108.
- [37] A. J. Alvero and Rebecca Pattichis, Linguistic and cultural strategies: Identification and analysis of spanish language usage in college admissions essays, <https://osf.io/preprints/socarxiv/wmsre>, 2022.
- [38] N. Deas, J. Grieser, S. Kleiner, D. Patton, E. Turcan, and K. McKeown, Evaluation of African American language bias in natural language generation, arXiv preprint arXiv: 2305.14291, 2023.
- [39] A. Koenecke, A. Nam, E. Lake, J. Nudell, M. Quartey, Z. Mengesha, C. Toups, J. R. Rickford, D. Jurafsky, and S. Goel, Racial disparities in automated speech recognition, *Proc. Natl. Acad. Sci. USA*, vol. 117, no. 14, pp. 7684–7689, 2020.
- [40] O. Patterson, Making sense of culture, *Annu. Rev. Sociol.*, vol. 40, pp. 1–30, 2014.
- [41] O. Lizardo, The cognitive origins of bourdieu’s habitus, *J. Theory Soc. Behav.*, vol. 34, no. 4, pp. 375–401, 2004.
- [42] P. Bourdieu, *Distinction: A Social Critique of the Judgement of Taste*. Cambridge, MA, USA: Harvard University Press, 1987.
- [43] P. DiMaggio and T. Mukhtar, Arts participation as cultural capital in the United States, 1982–2002: Signs of decline, *Poetics*, vol. 32, no. 2, pp. 169–194, 2004.
- [44] R. A. Peterson and R. M. Kern, Changing highbrow taste: From snob to omnivore, *Am. Sociol. Rev.*, vol. 61, no. 5, pp. 900–907, 1996.
- [45] R. Matthew, *Japanese Science Fiction: A View of a Changing Society*. Oxford, UK: Taylor and Francis Inc., 2003.
- [46] R. Barthes, *Mythologies*. New York, NY, USA: Hill and Wang, 1972.
- [47] B. C. Sanchez, Internet memes and desensitization, *Pathways: A Journal of Humanistic and Social Inquiry*, vol. 1, no. 2, p. 5, 2020.
- [48] J. G. Maudlin and J. A. Sandlin, Pop culture pedagogies: Process and praxis, *Educ. Stud.*, vol. 51, no. 5, pp. 368–384, 2015.
- [49] D. Chandler, An introduction to genre theory, http://visual-memory.co.uk/daniel/Documents/intgenre/chandler_genre_theory.pdf, 1997.
- [50] A. K. Ajeesh and S. Rukmini, Posthuman perception of artificial intelligence in science fiction: An exploration of Kazuo Ishiguro’s *Klara and the Sun*, *AI SOCIETY*, vol. 38, no. 2, pp. 853–860, 2023.
- [51] K. Solarewicz, The stuff that dreams are made of: AI in

- contemporary science fiction, in *Beyond Artificial Intelligence: The Disappearing Human-Machine Divide*, J. Romportl, E. Zackova, and J. Kelemen, eds. Cham, Switzerland: Springer International Publishing, 2015, pp. 111–120.
- [52] M. L. Gray and S. Suri, *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Boston, MA, USA: Hought Mifflin Harcourt, 2019.
- [53] V. Grech and M. Scerri, Evil doctor, ethical android: Star Trek’s instantiation of conscience in subroutines, *Early Hum. Dev.*, vol. 145, p. 105018, 2020.
- [54] A. Salles, K. Evers, and M. Farisco, Anthropomorphism in AI, *AJOB Neurosci.*, vol. 11, no. 2, pp. 88–95, 2020.
- [55] J. E. McNealy, Framing and language of ethics: Technology, persuasion, and cultural context, *Journal of Social Computing*, vol. 2, no. 3, pp. 226–237, 2021.
- [56] M. Fröhlich, C. Sievers, S. W. Townsend, T. Gruber, and C. P. van Schaik, Multimodal communication and language origins: Integrating gestures and vocalizations, *Biol. Rev. Camb. Philos. Soc.*, vol. 94, no. 5, pp. 1809–1829, 2019.
- [57] A. R. Moore, A feminine techno-utopia: Identification/transformation/transcendence of embodiment in spike jonze’s her, *Film Matters*, vol. 9, no. 3, pp. 57–72, 2018.
- [58] F. D. Saussure, W. Baskin, P. Meisel, and H. Saussy, *Course in General Linguistics*. New York, NY, USA: Columbia University Press, 2011.
- [59] E. W. Schneider, Investigating variation and change in written documents, in *The Handbook of Language Variation and Change*, J. K. Chambers, P. Trudgill, and N. Schilling-Estes, eds. Hoboken, NJ, USA: Wiley, 2004, pp. 67–96.
- [60] E. Ferrara, Should ChatGPT be biased? Challenges and risks of bias in large language models, arXiv preprint arXiv: 2304.03738, 2023.
- [61] S. L. Blodgett, S. Barocas, H. Daumé III, and H. Wallach, Language (technology) is power: A critical survey of “bias” in NLP, arXiv preprint arXiv: 2005.14050, 2020.
- [62] S. Chakraborty, A. S. Bedi, S. Zhu, B. An, D. Manocha, and F. Huang, On the possibilities of AI-generated text detection, arXiv preprint arXiv: 2304.04736, 2023.
- [63] J. Kirchenbauer, J. Geiping, Y. Wen, M. Shu, K. Saifullah, K. Kong, K. Fernando, A. Saha, M. Goldblum, and T. Goldstein, On the reliability of watermarks for large language models, arXiv preprint arXiv: 2306.04634, 2023.
- [64] V. S. Sadasivan, A. Kumar, S. Balasubramanian, W. Wang, and S. Feizi, Can AI-generated text be reliably detected? arXiv preprint arXiv: 2303.11156, 2023.
- [65] L. R. Varshney, N. S. Keskar, and R. Socher, Limits of detecting text generated by large-scale language models, in *Proc. 2020 Information Theory and Applications Workshop (ITA)*, San Diego, CA, USA, 2020, pp. 1–5.
- [66] K. Stollznow, Alien language, in *Language Myths, Mysteries and Magic*. London, UK: Palgrave Macmillan London, 2014, pp. 200–206.
- [67] A. Pennycook, Posthumanist applied linguistics, *Appl. Linguist.*, vol. 39, no. 4, pp. 445–461, 2018.
- [68] R. Schaeffer, B. Miranda, and S. Koyejo, Are emergent abilities of large language models a mirage? arXiv preprint arXiv: 2304.15004, 2023.
- [69] W. L. Hamilton, J. Leskovec, and D. Jurafsky, Diachronic word embeddings reveal statistical laws of semantic change, in *Proc. 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Berlin, Germany, 2016, pp. 1489–1501.
- [70] E. M. Bender and A. Koller, Climbing towards NLU: On meaning, form, and understanding in the age of data, in *Proc. 58th Annual Meeting of the Association for Computational Linguistics*, Virtual Event, 2020, pp. 5185–5198.
- [71] N. Arthurs and A. J. Alvero, Whose truth is the “ground truth”? College admissions essays and bias in word vector evaluation methods, in *Proc. 13th Int. Conf. Educational Data Mining (EDM 2020)*, Virtual Event, 2020, pp. 342–349.
- [72] T. Veblen, *The Theory of the Leisure Class: An Economic Study of Institutions*. Delhi, India: Aakar Books, 2005.
- [73] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, On the dangers of stochastic parrots: Can language models be too big? in *Proc. 2021 ACM Conf. Fairness, Accountability, and Transparency*, Virtual Event, Canada, 2021, pp. 610–623.
- [74] K. Laland and V. Janik, The animal cultures debate, *Trends Ecol. Evol.*, vol. 21, no. 10, pp. 542–547, 2006.
- [75] N. Rogers and J. J. Jones, Using twitter bios to measure changes in self-identity: Are Americans defining themselves more politically over time, *Journal of Social Computing*, vol. 2, no. 1, pp. 1–13, 2021.
- [76] S. Lowry and G. Macpherson, A blot on the profession, *Br. Med. J. Clin. Res. Ed*, vol. 296, no. 6623, pp. 657–658, 1988.
- [77] T. von Hippel and C. von Hippel, To apply or not to apply: A survey analysis of grant writing costs and benefits, *PLoS One*, vol. 10, no. 3, p. e0118494, 2015.
- [78] L. Lucy, S. L. Blodgett, M. Shokouhi, H. Wallach, and A. Olteanu, “One-size-fits-all”? Observations and expectations of NLG systems across identity-related language features, arXiv preprint arXiv: 2310.15398, 2023.
- [79] M. Kohda, R. Bshary, N. Kubo, S. Awata, W. Sowersby, K. Kawasaka, T. Kobayashi, and S. Sogawa, Cleaner fish recognize self in a mirror via self-face recognition like humans, *Proc. Natl. Acad. Sci. USA*, vol. 120, no. 7, p. e2208420120, 2023.
- [80] D. Placido, B. Yuan, J. X. Hjaltelin, C. Zheng, A. D. Haue, P. J. Chmura, C. Yuan, J. Kim, R. Umeton, G. Antell, et al., A deep learning algorithm to predict risk of pancreatic cancer from disease trajectories, *Nat. Med.*, vol. 29, no. 5, pp. 1113–1122, 2023.
- [81] D. Najafali, C. Hinson, J. M. Camacho, L. G. Galbraith, R. Gupta, and C. M. Reid, Can chatbots assist with grant writing in plastic surgery? Utilizing ChatGPT to start an R01 grant, *Aesthet. Surg. J.*, vol. 43, no. 8, pp. NP663–NP665, 2023.
- [82] T. N. Fitria, Artificial intelligence (AI) technology in OpenAI ChatGPT application: A review of ChatGPT in writing English essay, *ELT Forum J. Engl. Lang. Teach.*, vol. 12, no. 1, pp. 44–58, 2023.
- [83] P. Schramowski, C. Turan, N. Andersen, C. A. Rothkopf, and K. Kersting, Large pre-trained language models contain human-like biases of what is right and wrong to

do, *Nat. Mach. Intell.*, vol. 4, no. 3, pp. 258–268, 2022.

- [84] J. Rudolph, S. Tan, and S. Tan, ChatGPT: Bullshit spewer or the end of traditional assessments in higher education, *J. Appl. Learn. Teach.*, vol. 6, no. 1, pp. 342–363, 2023.



AJ Alvero received the PhD degree from Stanford University, USA; the MS degree in statistics and computational social science from Stanford University, USA; the MS degree in foreign language education from Florida International University, USA; and the BA degree in English literature from University of

Miami, USA. He is currently an assistant professor of sociology at University of Florida, USA.

- [85] M. Sloane and E. Moss, AI’s social sciences deficit, *Nat. Mach. Intell.*, vol. 1, no. 8, pp. 330–331, 2019.

- [86] W. Isaacson, The real leadership lessons of Steve Jobs, *Harv. Bus. Rev.*, vol. 90, no. 4, pp. 92–102, 2012.



Courtney Peña received the PhD degree in educational linguistics from Stanford University, USA in 2019, the MA degree in bilingual education and the BA degree in Spanish literature from Arizona State University, USA. She is the Assistant Director of Equity Research and Training at the Biosciences Grant Writing Academy,

Stanford School of Medicine, Stanford University, USA.