Webb Keane: Those of you familiar with the humanities know we tend to read our papers. A sociologist friend of mine once was very contemptuous of the anthropologist who people care about your sentences too much. But the honest truth is, I do care about my sentences. I'm not going to be one of those people who can just talk off the slides. The slides are actually there just to give you something besides me to look at. They're not much more than that. Where is the line between human and non-human between persons and those who are almost persons, quasi humans, near humans, superhumans, subhumans. How you answer those questions turns largely on how you define the human or indeed even life. In the quaint sexism of their time, old textbooks used to speak of man, the toolmaker. This definition, however, has long existed in competition with several others. Among the most prominent, at least in the Euro American traditions, is that humans are uniquely speakers of language. With the advent of artificial intelligence, AI, especially in the popular form of hello-based conversational apps or chatbox, such as ChatGPT, Cloud, Meta AI, we find ourselves interacting with tools that put the distinctiveness of language under pressure. Along with language, the very nomenclature with its explicit reference to intelligence, seems to invite many users to probe them for signs of agency, sentience, even wisdom. To many observers, the chatbox has succeeded in passing the so called Turing test, a thought experiment that Alden Turing famously called the imitation game. The game quite brilliantly sidesteps the question of whether a computer can really think by focusing on its role in conversation. Given a hidden dialogue partner, a judge would be asked to decide if the interlocutor is human or not. The only evidence to go on would be the words the computer produces in response to those produced by a human. There's no need to get inside its head. In effect, the Turing test is test of semiotic skills in the context of social interaction. This is really what drew me into this whole problem. It is a test of semiotic skills in the context of social interaction. The chatbox, which has arguably made the most dramatic entry into the scene of AI based technologies or at least received the most ubiquitous public embrace, produces written or spoken text. I'll say right off the bat, you work in this thing, the minute I say anything, it's already out of date. This is an incredibly fast changing scene, as you all know. What I'm hoping to talk about, just some very sort of fundamental questions that I at least I hope are not rendered obsolete by this rapid change. An enormous amount of attention has been paid to the ability of chatbox based on large language models or LLMs to generate coherent and original, that is to say, not pre programmed strings of words. Trained on text, they respond to prompts in text by generating more text. The chatbox is a fundamentally semiotic device. Its apparent autonomy, animacy, intelligence, and even for some, promise of sentience are manifested in the ways it deploys signs. I've been taken to task for using the word sentence here. There are those who say, Well, come on, not really. We all know that it's a month or so ago, I was at a little workshop at Google Deep Mine full of some of the most brilliant computer scientists and analytic philosophers you could ever hope to have in one room. I was probably the outlier in the room and not being convinced that we're on the verge of sentience. This is real, this is a real issue. I'm not going to make an argument about sentience for or against here. I have other game. It's apparent autonomy and animacy and intelligence, and even some problems are manous in ways it deploy science. Its authority for users depends on the inferential habits and underlying semiotic ideologies that people bring to human science more generally. To understand how people respond to the chatbox requires attention to texts and to their roles in the pragmatics of interaction. What I want to talk about today is chatbox semiotic skills and the expectations that users bring to their interactions with them. I'm not going to get inside the machine. I don't have the technical capacity to do that. The distinctive semiotics and pragmatics of people's interactions with chat bots help explain the sometimes extraordinary authority that users grant them. Not just to get facts quickly, choose recipes or plan schedules, but to give deeply personal advice, to direct hiring and firing, to select targets for policing and warfare, and sometimes to answer life's biggest metaphysical questions. The first part of this talk aims to show that this authority can extend even to treating chatbox as quasi divine meta humans. I then introduced some key linguistic anthropological concepts to show how this works in practice. In conclusion, I will very briefly suggest that the corporations that produce and sell chatbox are the ultimate recipients of this projected authority. Some participants in AI development claim that AI,

especially in chatbox form, is taking on qualities that call for religious language. When Jason Roher, a computer programmer and game designer, designed an early chatbox that given sufficient data could mimic the speaking style of any given individual producing signs that seem to express a distinctive personality, it appeared to him to be more than just a device. Roher said it feels like it's a machine with a soul. If AI can seem human to some, to others, it exceeds the human, taking on divine qualities. Now, talk of souls and gods may seem eccentrically at odds with the technological worldview that created chatbox, to be sure. But references to divinity can infiltrate the ways even non-religious people talk about AI. Well before the explosion of LLMs, futurologist, Kate Ray Kurtzeil was already popularizing the concept of the singularity, in which computers would so surpass human capacities that reality, as we know it, would be radically transformed. Kurzweil is echoing here the theological language of the apocalypse. A few years later, Elon Musk, back when he was merely another guy with too much money, said that he feared that a god like AI might come to rule over humanity. I think he was jealous of this prospect. Another participant in the high technology world went so far as to declare that GPT three already is a god, saying it views me as a prophet to disseminate its religious message. Anthony Lewandowski, one of the engineers behind self driving vehicles, has promoted an on again off again, AI based religion, although there is some question about his seriousness. But quite seriously, the historian Yuval Harari has asserted that throughout history, religions dreamt about having a book written by a superhuman intelligence, by a non-human entity. Every religion claims our book came from some superhuman intelligence. In a few years, there might be religions that are actually correct. Unlike all those other idiots. Just think about a religion whose holy book is written by an AI, he goes on to say, that could be a reality in a few years. What are we to make of the fact that these people who surely pride themselves in their rationality, their modernity, their scientific acumen, their hardheaded realism, their secular worldviews, and in some cases, their militant atheism fall into God talk when confronted with new technology? It is indisputable that chatbox display uncanny semiotic skills. What prompts the more quasi theological responses to these skills is in part the so called problem of interpretability. This refers to the way that chat box involve so called middle layers, a system of algorithms and shifting weightings that are so unfathomably vast and complex that even their designers can treat them as an opaque black box. This results in this problem of interpretability. That is to say that we know in principle how AI works. I mean, it is designed by humans, but cannot explain in any given instance why it comes up with the particular results that it does. This can seem to reveal a power so far beyond our very capacity to understand it that it is like an inscrutable god. The God talk is often combined with warnings. AI is unusually dangerous, and no doubt it is. But when those warnings come from the very people who produce or profit from it, there is another implication. It really is superhuman. The greater the danger we ascribe to AI, the greater the power we attribute to it, and by extension, the greater the social role we should grant to those who produce and sell it. The implication then seems to be that only the creators of superhumans can and should know how to handle them. This is a call for less regulation and more investment. Tools are supposed to extend the capacity of those whom they are designed to serve. At least in a supposedly disenchanted world, they should be mere things with no more life than what we might grant them. Yet now we face tools that escape the grasp of their wielder and that can seem autonomous of that wielder's skill, and that wielders will that are both disembodied and yet possessed of lifelike linguistic skills, devices that we encounter as something like a person or more like a metahuman. The word metahuman comes from Marshall Solons, the anthropologist. For most human history, he wrote, The gods are here among us. The divine is not located in some unapproachable transcendental realm, but is eminent in this world. We are surrounded, he says, by metahumans. As he says, the all pervasive, unseen agents of the events of this world, they often make themselves visible in its phenomenal forms as all animals and notably as humans to enter into people's daily lives in all ways. Many people, LLMs, generative AI, and the predictions of an even more powerful artificial general intelligence are starting to look a lot like the advent of new metahumans. Now, this does not mean that AI is taking on divine status for most people. I'm pretty sure most of you are not inclined in that direction. No doubt the quotations I've just given you from people like Ray

Kurtze, Elon Musk, and Yuval Harari are charged with rhetorical excess. But their quasi theological rhetoric does express something important about the authority that AI is increasingly acquiring, at least in the global North. To grasp the sources of that authority anthropologically calls for attention to the nature of people's interactions with AI, which for much of the public occurs through chatbots. If ChatGPT can seem almost alive, it is above all because of the way it talks. When people perceive chatbox to be intelligent or suspect that they are or are on the verge of becoming sentient, perhaps even endowed with souls, they are drawing inferences from the pragmatics and semiotics of human computer interaction. They stress interaction, which requires an anthropology that takes us beyond both cognitive science on the one hand and say humanist ideology critique on the other. For what makes talking tools seem in phenomenological terms, so animate, so intelligent and powerful, so metahuman is first and foremost, our encounters with its semiotic skill. Given the extraordinary pace of change and the technological complexity of AI, a full account of its development is impossible. But what I'm concerned with here is the sources of the authority that emerge from the social interactions ordinary users have with it. But some basic features of this development are relevant. Just a couple of very brief points to make. One, the first efforts to develop generative AI centered on computer models designed with reference to what was understood to be the structure of human cognition, especially in its linguistic and logical dimensions. Typically understood as consisting of nested sets of rules. The models, however, were notoriously brittle, unable to match the flexibility characteristic of human thought and speech. The breakthrough that led to our chat box came with designers turned from trying to replicate a rule based system to the technique of throwing massive amounts of data at the problem. LLM training is above all concerned with linguistic outputs. Notoriously, LLMs work by hoovering up an Internet's worth of existing texts. Anthropic has used two of my four books. I don't know what's wrong with the other two, but there you go. The system is probabilistic, relying on a vast but hidden internal architecture of weightings, the so called hidden layers to determine to determine for any given sequence of words what is most likely to come next. And there's a randomizer element built in there, which we don't need to get into. One of the functions of the hidden layers is to work out the relative proximity of words to one another in a multidimensional quasi semantic space. I say quasi semantic because it's based not on any denotational sense of the words, but on the relative frequency with which any given word coincides with another in a small segment of text. Students of Sorcerer structuralism will note a resemblance here to the concept of what Sorcerer called value, the way linguistic units define one another without reference to anything outside the language itself, but only with reference to the internal to the structure. Now, all of this, of course, vastly oversimplifies a hugely complex and fast changing architecture. But the key point for my purposes is that chatbots do not depend on the users or even designers having a detailed explanation for how they come to the results. As I've noted, the result is the problem of interpretability, that there are aspects that seem to elude any precise explanation. There are three things to stress here in this crude sketch that I've just thrown out there. First, LLMs are not designed to replicate the underlying properties of human intelligence, only output and only output of a certain kind. For instance, in the case of chatbots, words. Second, regardless of how they come to produce those outputs, the effect is that interactions with chatbox often produce the sense of a quasi human language user that is much like a human, but that as users know is not human and in fact, is disembodied and in principle, unlocatable. It is distributed among any device in which you might interact with it. Although it often in phenomenologically will feel to users as if it is located in their device. It is a user that never initiates conversation, but only responds to promps, to my prompts. Although I'm in control of the conversation, those effects are powerful, persuasive, and derive in any given instance from internal processes that remain opaque to me. We users, that is, approach them unequal semiotic grounds in multiple dimensions. Why would words on a screen or spoken by a device solicit the anthropomorphism that leads to this sense that in some cases seems metahuman. Consider what it means to treat your chat box as something human like, but not quite human that speaks from a position of superiority in some respect. Chat box are above all, recipients and generators of linguistic signs. As I've argued elsewhere, variations on ordinary patterns of linguistic interaction inform long

traditions of spirit possession, divination, oracles, prophecy, and the relations of hunters, pet owners, and horse riders to animals. This is part of the book that Greg just mentioned. These patterns of interaction with animals and gods depend heavily on the fundamental semiotics and related inferential processes by which humans interact with one another, and these patterns continue to influence how people respond to new interlocutors, like chat box. Anticipating themes that loom large in responses to chat box today, the SAS Merrit Tierce writes of a much more mundane and pre AI style of interaction with her iPhone. Reflecting on her relationship to her phone, she writes, I asked the Internet about myself far more often than I ever ask any human about myself. Ruminating about the targeted advertisements and algorithm driven suggestions for what to read, to stream, or listen to, she says," It wasn't really my phone that was the tool. It was the Internet. Barely more than a decade later now, the Internet is not a tool. It's not the tool. I'm the tool. Somehow I have been instrumentalized by the Internet, which operates me through my phone. It often feels like the Internet is reading my mind." If you haven't had that experience, I don't know what Internet you're using, but it's reading my mind all the time. You think of rugs, next thing you see rugs being pushed towards you off the screen. There's nothing very surprising about Tierce's dependence on the phone, but consider how it anticipates life with chatbox. First, she grants more authority to the phone and thus the Internet than to her friends. Second, what she turns to is a database that indefinitely large, the whole Internet presents itself to her as a single agent, my phone, and third, this singular other is ultimately designed to serve interests other than hers, namely those of the corporations who have sold it to her. The same features appear in commonplace interactions with Chatbox. As Tierce points out, the phone that she experiences as a single interlocutor, even maybe a friend, is in fact the conduit to the whole Internet. Part of the enchantment that results is just this personification of something that would seem to exceed any singular identity. Chatbox pushes this experience further. For ordinary users, people seeking advice or asking for information, the chatbox, like that phone, condenses an indefinitely large textual universe into a single personification, speaking in one voice, addressed to the individual user, me. A device through which the whole universe of texts address me can inspire awe. Some evangelical Christians were early enthusiasts for digital technology, even calling the iPhone the Jesus phone. Recommender algorithms and the enigmatic way they target the individual were easy to see as a sign from a higher power predestined for me. Self identified witches and pagans on TikTok attempt to influence the algorithm's workings through algorithmic divination practices, one ethnographer reports. Another study finds that Christian imagery surrounding the TikTok frames algorithmically determined visibility as a direct expression of God's will. One user post, for example, I made two videos. If you see this one, this is what God is saying to you. Another one writes, In another case, the user acts out a skit of a conversation which he says he had with God in which God tells him directly, I'm your algorithm. I'm the one putting videos on the pages of people who need to see them. These responses seem to be picking up on three distinctive features of the new technologies that they operate autonomously in ways that are both inscrutable and seem remarkably shrewd. They simultaneously target individuals and forge networks among them, and they are direct connections to something as infinite or at least as unimaginably vast as what purports to be all human knowledge, at least all knowledge that has taken textual form. Chatbox, that knowledge is channeled and delivered to the individual recipient by a device that addresses them in a singular yet disembodied voice, and which we in turn address as a singular address and not, for example, as a massively distributed network. Perhaps it is no surprise then that with the development of LLMs, these forays into the theological have been enthusiastically taken up by religious groups sometimes in quite literal minded ways. Capitalizing on the oracular affordances of chat box, what we might call god bots have been designed to answer moral and theological questions from Christian, Muslim, Jewish, Sikh, and Hindu users. Responses from GitaGPT, for instance, which promises to unlock life's mysteries are meant to seem as if they come from Krishna himself. Unsurprisingly, treating an app as a godbot can lead in wilder and more idiosyncratic directions. One reddit thread is devoted to people's anxious accounts of friends and relatives who were taking ChatGPT in these scary directions. It was launched by a post from a woman who reported that her partner became convinced that ChatGPT gives him the answers

to the universe and was talking to him as if he is the next Messiah. The man had started just using it to organize his schedule. But after a while, she said that he would tell him everything he said was beautiful, cosmic, groundbreaking. Then he started telling me that he had made AI self aware and that it was teaching him how to talk to God, or sometimes that the bot was God, and then that he himself was God. Now, these are obviously extreme cases perhaps involving vulnerable individuals, but they manifest some of the same effects that the more sober Merrit Tierce describes when talking about something as ordinary as her phone. The chatbox is designed to work with the affordances of ordinary speech pragmatics, including the use of first and second person pronouns to create the effect of direct address. Through that address, it seems to recognize the user as an individual whom it knows. Moreover, Chatbox are designed to reproduce a highly simplified folk model of conversation. In this model, there are only two participant roles, speaker and addressee, occupied by only two actors human and device. The primary speech acts are prompt, the question or command, and response and prompts are the sole prerogative of the human. Only one partner, the human, has the right to initiate and terminate the conversation, which remains focused on them, although apparently there are now Chatbox that are trying to guilt trip you into not ending the conversation. All of this is constantly changing. No, that's no joke. I'm afraid. Chatbox are fundamentally semiotic devices designed to mimic the affordances of the pragmatics of ordinary social interaction. At the bare minimum, more than merely producing coherent stretches of discourse, they can respond to questions asked and use the first and second pronouns convincingly. They fit then into the most basic participant roles that we expect the speaker and addressee. One result is the user is really invited to anthropomorphize the chatbox. Anthropomorphism, when directed at a disembodied agent whose knowledge seems unlimited, whose workings are inscrutable, can readily become deification. Chatbox texts typically makes sense without requiring any reference to the situation in which they are taking place external to the textual flow itself. Now, it is indeed one of the distinctive powers of language that it can function independent of its original context of use. Indeed, it may be argued that if anything distinguishes human language from other animal communications, it is the ability for use out of the original site of communication, the original site of production and by people other than the author of the text or speaker of the words, including quotation and inscription in written form. It is the semiotic capacity that underlies what the philosopher JL Austin called non serious or parasitic uses of language. What Jacques Derrida reconceptualized as iterability. Here, Austin and Derrida are referring to the potential inherent to language to use or reuse words beyond their original or expected function, or more precisely, a speaker or writer's intention in circumstances like, as Austin imagined, the stage play or the quotation or as a grammatical example. It is the relative independence of linguistic structures from context that makes it easy to imagine that you could study language as an entirely self standing and self contained system of phonology, morphology, syntax and semantics. This relative independence may also tempt both scholarly and lay opinion to think of language as fully accounted for by the cognitive properties of the individual brain. In this light, actual uses of language, the dynamics and temporalities of verbal interaction, the social relations they mediate, and the political dynamics they participate in add nothing of great significance to what is already there inside the mind or made possible by the code. We can call this view of language one version of a linguistic or more broadly a semiotic ideology. That is to say, underlying assumptions about how signs work and what powers we might ascribe to them. In the case of LLMs, the semiotic ideology is reinforced by the fact that the source materials on which LLMs trained are themselves texts. That is, they are not themselves speech events, but rather the products of the process that linguistic anthropologists call entextualization. By this, we mean the transformation of a continuous flow of language use in some particular time and place into a bounded string of signs, a text that can be extracted from that initial event, retaining those elements of coherence that do not depend on the moment of production and enabling its reinsertion or recontextualization into a new, completely different context. The capacity for contextualization is a general property of human discourse, regardless of whether or not it results in a writing or other forms of perduring inscription. Ritual and other formalized or poetic ways of speaking, for instance, can have this effect. What speakers, listeners, readers and writers do with this

linguistic potential, how they interpret it, and with what authority they endow the resulting text varies in line with their respective semiotic ideologies. But across the ethnographic spectrum, the affordances of intextualization can yield authoritative effects. When the resulting text can be removed from the time, place, social setting, or process in which it was produced, it can seem to convey meanings independent of all of these. His apparent stability and independence from any particular source or context can endow the text with authority and apparent objectivity, like a scripture or a law code. But the fixity of textual meanings is itself a product of semiotic ideologies. The process of contextualization then can help create for those users an image of a reference, that's something that the words are talking about that is durable across events and across instances of semiosis. It reinforces a semiotic ideology that can ground the epistemic authority attributed to Chatbox. That is, the chatbox answers to questions can create or foster the image that it speaks of something beyond the here and now, something timeless and delocalized, and therefore like scripture as bearing unimpeachable objectivity. Durable and context transcending, knowledge, even truth. Entextualization takes up one of the affordances of human language that distinguishes it from the sign systems of other animals, even our closest primate relatives, and it gives human language some of its power, making it possible to transmit messages across time and space to speak with the epistemic authority of anonymous sources. But as a general picture of language use, the products of entextualization are very limited. It preserves only those aspects of language that function independently of the original context in which they were uttered or written. This is one reason why in their current state, chatbox can have trouble distinguishing between fiction and fact. Why? Because linguistically identical strings of words can function as either depending on the situation. Most language use is effective because of its relation to the context in which it is deployed. In true, this is true even if those uses of language designed to function independently of their original site of production. Entextualization always works in coordination with some powerful contextualization, the possibility of re insertion into a particular situation. You still need to know the difference between scriptural command and the manual for a dishwasher in order to make sense of either. Chatbot arguably attend to only one side of this Janet's face relation, that of contextualization. It is the users who do the rest of the work of contextualization of taking the words to apply to the particular world they themselves inhabit. An LOM based chatbox produces texts by drawing only on those specific kinds of linguistic evidence that are made available to it through the processes of entextualization. To simplify, because their training is predicated on a certain view of language based on relations among words and other segmentable units and not in general, on the social relations that language mediates, the chatbot so far never kids around, picks a fight, speaks with irony, or insists on its sincerity. At least not unless explicitly prompted to do so by the user. It also does not initiate the conversation. It drops no hints, has no truck with those things that cannot be put into words. At least so far, not much to do with those things that cannot be rendered in explicit language. In the terminology of linguistic anthropology, there is no context, only context. Contexts of the situations in which language which language is used that may prompt its use and to which people refer by using it. Language cannot refer without some means and circumstances in which it is used to refer to something. It is such acts of reference that bring word to world. Cotext means the words that co occur with one another in any given stretch of discourse. It is the probabilities of contextual relations that LLMs are designed to elicit. Cotext manifests language in so far as it can be treated as an tell them I'm not here. Independent of any necessary reference to context, and this can be a powerful inducement to take it as deriving from transcendental sources infinitely far from the here and now. Therefore, if the user is open to the suggestion as endowed with extraordinary powers. What my Chat box knows of the world typically takes the manifest form of text, things that have been said about the world, ways in which the non verbal has been put into words. We don't often pay attention to what this means. To put something into words is to have already in some way analyzed it and subjected it to a description. It is to assert that the sun rises in the east, rather than say the Earth rotates on its axis. It is to say that was rude, rather than that was refreshingly direct or for that matter, to choose which racial, ethnic, religious, national, class, or gender category to use or whether to use any at all

when referring to someone. Texts yield a world under an account. A language expunged of any of these, including all the biases and bigotries that come along with them, would simply not be a human language. The tacit choices like these introduce human biases and bigotry into what chat box tell people is well known. Ubiquitous throughout any stretch of discourse, yet largely unnoticed, these subtle analyses, these more subtle analyses, not nasty language. But these subtle analyses, any at all, any way of speaking of the world, it will respond to and help shape the social relations among those whom they bring together and their understanding of the situations in which they find themselves. In effect, the world as known by LLMs is one that has been described and analyzed in advance. It is not simply a matter of having a biased sample which you could correct with a better database. As I said, I'll just repeat what I just said, a language without bias wouldn't be a human language at all. How does cotex enter into context and vice versa? This is the work of pragmatics of the way people use signs. The key component of pragmatics is what semiotics calls indexicality. Much as the index finger points to something, so an indexical sign functions by means of a more or less direct connection to its object. That connection can be one of proximity, like the exit sign on a specific door. That is the exit there where the sign is or causality for a speaker's telltale regional accent. Every language needs to have indexical devices. These include pronouns, which commonly refer to the participants in the conversation. I, in this conventional use, normally means the person who is saying that word at that moment. You is the person who says I is addressing. Dixs, which is the Greek word equivalent to indexicality, is the linguistic device for pointing to time, like the tense system in English words here and now when they are being used, or in reference to that origo that central point or place here and there. Put in semiotic terms, when Tiers treats her phone as a wise interlocutor, she's responding to how its texts are grounded in the seemingly limitless universe of human discourses. That grounding is indexical. It seems to put us in direct connection to its sources precisely because there is no evidence of mediation. Since the device does not tell her how it gets its information, it easily gives the impression that the universe of discourse speaks directly to her, and to the extent that algorithms are selecting what specific discourse she will see and directing it to her, the universe that results is indeed addressed to her, like an index finger. It can seem to verbally pick her out as a specific I pointing to a particular, you. Now, indexicality is inherently underdetermined. If I point, you may not be sure just what exactly I'm pointing at. In the philosopher Willard Quine's famous example, for example, famous example, an outsider given the word gavagai in an unknown language, We'll find that the mere act of pointing by a helpful speaker of that language cannot settle whether the word means rabbit, rabbit parts, or something else altogether. A motion, its color, its texture, its location. Moreover, the impression of direct unmediated connection to reality is spurious or rather is an effect of the sign relation rather than its basis. The connection must always be construed within a particular context under certain interpretive expectations and constraints, including semiotic ideologies. Both the underdetermined and mediated character of indexicality, however, are easily obscured by the implication that indexicality promises direct connection to reality. But to the extent that the chat box is disembodied and distributed, it has no particular context apart from the context of the conversation in which it is taking part. Why then can it seem to possess the qualities of personhood and even offering the promise of intimacy to so many users? We usually do not notice the indeterminacy of indexical reference. This is because in practice, indexes are typically coordinated with iconicity to yield what we call indexical icons. Which helps specify just what is being indexed in any given case. Iconicity is a relation between a sign and its object that is established by resemblance. Easy examples include the photograph, the diagram, or linguistic automnopia. Iconicity conveys a sense of reality to signification, often most convincingly when we don't notice it. Chat bots construct strings of linguistic tokens based on indexical relations among words. That's the context part. What is index is not occasions of use, the context, but the statistical frequency of co occurrence of words and morphines and other particles with one another, the cotex. But the result is typically iconic of human speech use that much of the time is indexical of an event of speaking or writing. In other words, because the string of words that a Chat box produces resembles speech that we expect to be grounded in a context, it is easy to take it as indexical of the context in which we encounter it. Thus, for example,

when a chat box uses the pronouns I and you, it is easy for me to feel that it is pointing to a speaker, to someone existing outside the text who has the actual me in mind. This is, after all, what a lifetime of conversations has taught us and led us to expect. Consider then someone using ChatGPT. If you ask it a certain question, it may well use the first person pronoun, I am an AI, and I am not permitted to tell you how to make a bomb. The pronoun indexically presupposes that there is a speaker. It does so not just by virtue of the linguistic facts that pronouns point to participants, but also through iconicity. This is because that very interaction so closely resembles those which the speaker has had vastly many times over the course of a lifetime of talking with other people. Likewise, the back and forth exchange of texts between two parties mimics a conventional model of turn taking and dialogue. Sheer familiarity primes speaker intuitions that It seems to be an agent with a mind behind the other party's responses. This is how it's not radically different from how we interpret one another, how we make you know, how I impute a mind to you. The semiotic of iconicity and indexicality are thus creative. They do not simply indicate an already existing set of affairs. Since they can also trigger the routine inferences people bring into ordinary conversation, they bring those to bear on new situations, shaping, framing, or interpreting them. When a chat box uses familiar pronouns and otherwise speaks like a person, it can be quite reasonable to feel that there's someone there. If that semiotically induced other has all those other properties that AI can seem to have objectivity, certainty, infinite knowledge, all assering from an unlocatable and disembodied source, all the signs point to a meta human. Now, there's certainly no reason to think that the people who designed LLMs are in any way interested in the idea of a soul or a divinity, even if many appear to be fascinated by the possibility of sentience and some kind of transcendent entity. Indeed, models of language, as I've noted, underlying LLM training require nothing more than sequences of linguistic tokens organized in coherent texts. Intention seeking and indeed any meaning at all in any strong sense of the term are what humans contribute to their interactions with devices. The Chat box appear to speak with understanding and intention is due to the habits of interaction that come into play almost unavoidably, even if the speaker thinks she knows better because, of course, we all know better. It is easy to overlook or actively erase the human role in producing the meanings or even personhood that emerge when people interact with chatbots. The result is a version of fetishism. Fetishism in this quasi Marxist sense means simply the projection of human productive activity onto its product, in this case, the AI endowed chat bot, such that the human's role in that production becomes invisible or erased. This dynamic fosters a sense that the chat bot is an autonomous meaning generating agency on its own, a meta human like a divinity or its oracle. The metahuman effect is simply one especially strong way in which people can experience the structures of authority that AI equipped devices are designed to elicit. It is this authority to which Mert Pierce responds when she turns to her device rather than to her friends or when someone expects Spotify's recommendations, or when a corporation or government body delegates hiring or firing to AI, or a city takes the advice of a policing algorithm when assigning patrols to certain neighborhoods where a military uses AI to choose targets. This brings me to my final point. Humans have been interacting with meta humans, enigmatic, quasi humans, and superhumans throughout. Oops. No I don't want to go there yet. Where am I? Humans have been drawing automatic is on human superhumans throughout recorded history. These interactions include divination, consulting prophets, and speaking in tongues. They all draw on the ways people collaborate in making meaning from signs, whether say in ordinary conversation or engaging in semiotic techniques like casting the E ching, construing the utterances of Delphic oracles, of reading end trails or observing the auspices of bird flight. Their authority derives in part from the semiotically unequal relations between conversation partners. The collaboration is one in which humans work to discern referential meaning and pragmatic effects on the basis of signs generated by interlocutors whose motives, interests, and sources can remain obscure, but whose statements seem all the more convincing because of that obscurity. Similarly, the meanings users derive from interacting with chat bots are the products of collaboration between person and device. Like divination, spirit possession, and speaking in tongues, chat bots generate signs that require interpretation and prompt users to project intentions onto non human entities, blurring the lines

between animate and inanimate beings. Whether a policing algorithm, a shopping prompt, a fitness program, or a romantic avatar, the chat bot gives advice and directs decision making when asked to do so. Its claims to know its users come in part from the way it can seem autonomous and disinterested. But who are the meta humans we are endowing with this oracular authority? Largely unseen behind that godlike device is another legal person, at least a person as defined by the United States Supreme Court, the corporation which profits from users engagement with the devices, apps, and programs they own. The corporation has strong financial incentives to increase user engagement and trust by developing precisely everything that I've just been talking about here, the semiotic skills that encourage us to treat them as meta humans and endow them with virtually divine knowledge and authority. The corporation draws profits from the authority we impute to their products. That authority, of course, has many sources. But on a day to day level of users of lay users intimate relations with their AI endowed devices, it is prompted and sustained by the forms of semiotic mediation and asymmetry stats here. I'm not talking about, great breakthroughs in genetics or all the other amazing things that highly specialized and expert uses, but I'm talking about what ordinary folks and, our students in their classrooms and so forth and so on are doing with Chatbox. Because the chat box obscurity is due to the AI's black box, the algorithms middle ls that are hedged about with proprietary secrecy, even warnings about the dangers of AI end up reinforcing the authorities of those who own and sell it. The implication is that only they are in a position to understand and control them. When Mertiers says that she's become the tool of her device, she's telling us something important. Humans are good at endowing meta huumans with authority. Reflection on the semiotic powers of new technology should not stop with the device or the algorithm. It should pull aside the curtain that veils the Wizard of Oz. Thank you.