GETTING REAL: EXPLORING THE ESSENCE OF LANGUAGE IN THE AGE OF AI

UNIVERSITY OF TORONTO
FACULTY OF LAW

PRESENTED BY
Sonali Ravi, PhD

March 1, 2024

Future of Law Lab
University of Toronto Faculty of Law

FUTURE OF LAW

About the Future of Law Lab

The Future of Law Lab is a platform for students, academics, lawyers, and other professionals to participate in collaborative initiatives exploring how the law will evolve in the future. We will dive into the intersection of law, technology, innovation, and entrepreneurship, with programing dedicated to each of these streams. As a hub of interdisciplinary activity, we are dedicated to bringing together individuals from all backgrounds to examine the changing face of the legal profession.

About the University of Toronto Faculty of Law

Established in 1887, the Faculty of Law is one of the oldest professional faculties at the University of Toronto, with a long and illustrious history. Today, it is one of the world's great law schools, a dynamic academic and social community with more than 50 full-time faculty members and up to a dozen distinguished short-term visiting professors from the world's leading law schools, as well as 600 undergraduate and graduate students.

The Faculty's rich academic programs are complemented by its many legal clinics and public interest programs, and its close links to the Faculty's more than 6,000 alumni, who enjoy rewarding careers in every sector of Canadian society and remain involved in many aspects of life at the law school.

Introduction

The use of language and its diverse expressions are arguably what define the essence of humanity. Yet, in the 21st century, we find ourselves at a crossroads where technology and artificial intelligence have not only augmented but also challenged the very essence of linguistic expression. In this contemporary linguistic landscape, Large Language Models (LLMs) are redefining the boundaries of communication. Platforms like ChatGPT have become companions of thought and conversation, seamlessly integrating into our daily lives. Whether we seek answers, advice, or creative inspiration, they are ever-willing interlocutors. Central to their capabilities is the process of tokenization, where LLMs navigate colossal amounts of data to predict the next sequence of words with precision that mimics human speech and thought. The question that arises is this: does their proficiency in predicting text truly equate to possessing the essence of language? What grants models like ChatGPT the distinction of being regarded as *language* models? And in the context of human communication, can they ever authentically emulate the subtleties of our linguistic expressions? This essay invites us to contemplate the inner workings of language, in order to answer the difficult question of whether LLMs can genuinely embody the spirit of language in the age of artificial intelligence.

What is language?

While there is no single definition of language, various attempts to define[1] it have revealed certain characteristics of human language. Natural language that humans possess is unique in that it is:

[1] Philosophers have grappled with the question of what language is. Different schools of thought exist in the Western philosophical canon, including the Cartesian rationalist school, Derrida's structuralist school and Chomsky's universal grammar theory. For more, see René Descartes, *Discourse on Method* (1637), Jacques Derrida, *Monolingualism of the other; or, the Prosthesis of origin* (1998), and Noam Chomsky, *Knowledge of Language: Its Nature, Origin and Use* (1986).

(a) Innate – According to Noam Chomsky's theory of universal grammar,[2] babies are born with innate linguistic capacity. They are not blank slates who, upon birth, acquire language through social interaction. Instead, they possess innate linguistic machinery that allows them to learn language rules and systems. According to this theory, genetic endowment and social stimuli might play a role in comparative language capacity and what language(s) a child acquires, but the ability to acquire language is pre-programmed.

(b) Social – Humans use language not only to accomplish strictly communicative, but also social and cultural functions. Language signifies identity, signals group belonging (and who is outside of it) and elicits a range of emotions.

(c) Semiotic – A system of signs can combine in an infinite set of ways to convey meanings in human languages. They are unlike animal or artificial language (like code), which have a finite number of signals. By contrast, human language has complex grammar and semiotic systems to signify meaning. These systems do not directly correspond to objects or phenomena, but nevertheless have shared meaning through signs and symbols.[3]

What makes LLMs succeed or fail as models of human language?

Diverse meanings attach to the term 'language,' which prompts an exploration into the multifaceted nature of linguistic expression. The three French words for 'language' signal the complex ways in which we use the term to refer to the concept of language. The French *langage* refers to language as an abstract concept, *langue*, to a particular linguistic system like English, and *parole*, to the spoken variety of a linguistic system. In English, we also use the term 'language' to refer to code and other computer languages that are not natural but artificial language systems. For

[2] Noam Chomsky, *Knowledge of Language: Its Nature, Origin and Use* (New York, NY: Praeger, 1986).
[3]Ferdinand de Saussure, *Course in General Linguistics* (trans. Roy Harris) (London: Duckworth, 1916).

the purposes of this essay, we are not interested in language in this technical sense. Instead, our goal is to determine what, if any characteristics, LLMs possess that make it both like and unlike natural human language.

LLMs' success as a model of human language

LLMs have the ability to predict what comes next based on data used to train them. This leads to the question of whether LLMs simply store and regurgitate pre-learned data from training sets by picking up on statistical regularities. Research suggests that LLMs like ChatGPT do more than provide probabilistic predictions. While their responses are based on conditioning by previous words, researchers have shown that LLMs have the ability to generate novel data, new combinations of words, and grammatical structures that were not originally in their training set.[4] In fact, Mahowald et al. argue that advanced LLMs like ChatGPT have formal linguistic competence akin to humans.[5]

Formal linguistic competence is not just linguistic knowledge surroundings rules and grammar in a language (like knowing that sentences in English follow the subject-verb-object structure). It is more nuanced. Specifically, it involves the ability to know and use linguistic rules flexibly. Mahowald et al. define formal linguistic competence as possessing knowledge of a language's vocabulary, rules, syntactical structures and idiosyncrasies, to produce grammatical output.[6] By being trained on large date sets, LLMs learn deeply about language structure, to succeed (and in some cases outperform) human language processing, acquisition, generation and

[4] Kyle Mahowald, Anna A Ivanova, Idan A Blank, Nancy Kanwisher, Joshua B Tenenbaum, and Evelina Fedorenko, "Dissociating language and thought in large language models: a cognitive perspective" (2023) [Mahowald et al.]
[5] Ibid at 8.
[6] Ibid at 4.

comprehension. [7] They have the ability to abstract, learn hierarchical structures, use idiosyncratic phrases, and combine various words and phrases to arrive at novel sentences not present in training data. In this sense, LLMs succeed as models of formal linguistic learning and processing.

Language limitations of LLMs

While LLMs have the ability to generate seemingly human-like responses to prompts, they nevertheless lack several characteristics unique to human language. LLM's language is:

a) Not innate – Unlike human language, LLMs are trained on pre-existing data. They do not possess innate grammar rules that they acquire and refine through social interaction. Instead, LLMs learn language patterns from large-scale data. ChatGPT also contains a feedback mechanism called Reinforcement Learning from Human Feedback (RLFH), which helps to constantly fine-tune this model.

b) Not social – Whereas human language is social in both how it is acquired and used, LLMs have more or less an *asocial* communicative function. The ends of LLM's language are communicative – when asked to write a song, it does; when questioned about the meaning of a concept, it explains it. Beyond this, it does not use language for a variety of other social purposes that humans use language for (more on this <u>below</u>). LLMs like ChatGPT have a dialogic feature, which suggests social form and intent. However, despite this dialogic nature, ChatGPT's response is only as good as the prompt that users input. Natural language, while dialogic, is *not* contingent upon input. Interlocutors often exceed the quality of dialogue of their partners, or their response might not even be related to the topic of conversation at hand. This multifaceted communicative fluidity is absent from LLMS.

[7] Researchers have noticed significantly lower performance in low-resource languages. For more, see Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie B, et al. "A Multitask, Multilingual, Multimodal Evaluation of ChatGPT on Reasoning, Hallucination, and Interactivity" (2023).

Taking together both the successes and failures of LLMs as models of human language, it is clear that large language models might not have all the features of human natural language. Yet, they possess the characteristics necessary for functional linguistic competence. As the technology currently stands, we might conclude that what LLMs possess is indeed language, although one that is different from human natural language.[8]

Can LLMs ever speak like human experts in professional contexts?

This issue consists of two prongs. First, we need to evaluate whether LLMs can speak like humans. And second, we need to examine whether they can speak like experts in professional fields like law and medicine. We will explore each prong in turn.

Speaking like humans

In many ways, LLMs do not speak like humans. Linguistic analysis conducted by Natural Language Processing (NLP) experts has revealed key differences between human language and AI-generated language.

(a) Formal vs Functional Linguistic Competence – As outlined above, GPT does not use language in human-like ways. For example, it is not good at conveying or eliciting emotion. Some researchers theorize that this is because LLMs have formal, but not functional linguistic competence. Functional competence, according to Mahowald et al, involves the "non-language-specific cognitive functions that are required when we use language in real-world circumstances."[9] It is the set of extralinguistic skills that make up human thought (like "formal reasoning, world knowledge, situation modeling, and social cognition"[10]) and

[8] Even between GPT 3.5 and 4, there is significant improvement in functional features. With constant technological progress, LLM's language capacity will likely advance, which could very well make these conclusions untenable.
[9] Mahowald et al, *supra* note 4 at 5.
[10] *Ibid* at 1.

support social language use in different contexts. These skills help us know common world facts (such as that a person is smaller than a house) and give us the social know-how to understand that we do not talk to a friend or to a parent the way we would to a boss. Such non-linguistic skills are separate from formal linguistic competence, and make language usable in the real world, competencies that LLMs arguably do not possess.

(b) Monolingualism vs Multilingualism – Humans constantly live between multiple languages and registers (formal, informal, slang, etc). In his book *Poetics of Relation*,[11] writer-philosopher Édouard Glissant calls attention to the hybridity and plurilingual nature of human existence. Using the concept of the "Tout-monde" or "Whole-Word," he posits that within the self lies the totality of the world, such that all beings and their cultures become a product of multiple peoples and civilizations crossing and mixing with one another. In this sense, we are already and always multilingual. Needless to say, we switch codes and language depending on context. Even language pedagogy involves teaching multiple language competencies (in the plural), including grammatical, socio-cultural and discourse competence. LLMs can undeniably understand and respond in multiple world languages. However, they are nevertheless monolingual in many ways. They speak in one register, without the hybridity of language that characterizes the richness of human social interactions.

(c) Neutral vs Emotive Language – Researchers found that ChatGPT's language tends towards formality. It often contains several conjunctions (firstly, secondly, in addition, by contrast, one the one/other hand, in conclusion) that signal logical flow of thought. This kind of

[11] Édouard Glissant, *Poetics of Relation* (trans. Betsy Wing), (U of Michigan P, 1997).

language use suggests an attempt towards neutrality.[12] On the other hand, human language is highly emotive. This includes the strategic use of language to elicit emotions. For example, humor, irony and antiphrasis are prevalent not only in oral, but also in written language.[13] Such language acts require decoding the actual intention of the speaker,[14] instead of taking utterances at face value, tasks that LLMs typically perform poorly. [15] Consequently, they do not tend to use humor and irony, sticking to neutral, less emotive language.

Even if LLMs have acquired considerable linguistic knowledge through training data, they are bound to the corpus they learn from, in ways human language is not. If we consider language to be not only a reflection of purely linguistic skill, but also a window into an entity's mind (including how they think, reason, and interact with the world), then LLMs do not possess the power of language the same way that humans do.

Speaking like human experts

A model that does not speak like humans, however, does not necessarily mean cannot speak like human experts. Can LLMs use language in the peculiar ways required for inclusion in professional epistemic communities? In other words, can LLMs ever speak the language of experts in professions like law, finance, psychology, and medicine? The answer is contingent upon the field and level of expertise required. Researchers found that ChatGPT can be more helpful than human counterparts in finance and psychology, while not in medicine. In the former, they found

[12] Biyang Guo, Xin Zhang, Ziyuan Wang, Minqi Jiang, Jinran Nie, Yuxuan Ding, Jianwei Yue, and Yupeng Wu, "How close is ChatGPT to human experts? comparison corpus, evaluation, and detection" (2023) [Guo et al.].
[13] Sonali Ravi. *Humoring Pathos: Comedic Discourse and its Paradoxes in Contemporary Francophone Literature and Culture*. 2022. Princeton University, PhD dissertation.
[14] Bart Holterman B, Kees van Deemter K, "Does ChatGPT have Theory of Mind?" (2023).
[15] Some have argued that LLMs are getting better at decoding human intentions and decision-making, thereby suggesting that LLMs possess a form of theory of mind. For more, see ibid.

that ChatGPT provided advice that was more concrete and approachable than responses by finance and psychology experts. In medicine, however, researchers concluded that the answers were too long to be useful, while human medical experts generally provided short and straightforward responses.[16] Therefore, the expertise and helpfulness of LLMs is a function of the discipline in question.

Extrapolating these findings to the legal field, we might ask whether LLMs can speak like and be as helpful as human lawyers and legal experts. The answer, again, is that it depends. When it comes to basic legal questions and advice, LLMs can be a useful starting point for clients facing legal issues. But for in-depth legal advice or mounting legal arguments, LLMs are likely not helpful, given their propensity to make up case law and misinterpret legal provisions. What is particularly dangerous is ChatGPT's ability to speak in ways that are peculiar to the legal community and produce realistic responses, even fooling a lawyer of 30 years.[17] In the face of such hallucinatory tendencies, this very ability to convincingly speak and sound like a legal expert poses problems, especially for non-lawyers entering into conversation with the model.

Some argue that we can teach LLMs to write and reason like lawyers do. Even the most complex legal reasoning consists of structures and patterns – stating facts, spotting issues, conducting analysis, citing case law to provide authority – all of which can potentially be taught. Curran et al. propose that technologists can program LLMs to move from one task to another using a multimodal algorithm.[18] This can teach LLMs to tokenize when necessary and quote verbatim without taking creative liberties when appropriate, which could partially combat the hallucination

[16] Guo et al., *supra* note 12 at 5.
[17] Benjamin Weiser and Nate Schweber, "The ChatGPT Lawyer Explains Himself," NYT, Jun 8, 2023.
[18] Shawn Curran, Oliver Bethell, Sam Lansley, "Hallucination is the last thing you need" (2023).

issue. However, given where the technology is now, it is unclear whether LLMs possess the professional language and knowledge base required for true legal expertise.

Conclusion

Language is more than a process used to achieve an outcome. It transcends mere functionality, embodying emotions, and signifying belonging. This kind of real-world language use requires skills that skills go beyond linguistic proficiency, and it is unclear whether they can be learned through the tokenization process. Current AI technology, particularly LLMs, possess stellar linguistic competence, with the ability to even sound like experts in a field. However, what remains doubtful is their ability to sound quintessentially human – the essence of human communication, with its nuanced emotional depth, remains elusive to LLMs.

Future of Law Lab

University of Toronto Faculty of Law