



# The Profession of IT Can Machines Be in Language?

*Large language models brought language to machines.  
Machines are not up to the challenge.*

**I**N LATE 2022, large language models (LLMs) erupted into the public spotlight. Pundits were quick to claim LLMs as the next step in the path to artificial general intelligence (AGI) and even the Singularity.

LLMs are artificial neural networks (ANN) created by a complex process. First, the core ANN is trained on billions of words of text from the Internet to respond to a prompt with a list of most probable next words after the prompt. Second, the core ANN is then “fine-tuned” by a complex process called “tweaking” to make the core ANN outputs more satisfactory to humans. A large team of humans scores the quality of responses of the core ANN to a large number of sample queries. A second ANN is trained from these data to predict the most likely score a human would assign to a prompt. Third, the second ANN is then used in a mode of reinforcement learning to adjust the internal weights in the core ANN so that its outputs are even more likely to satisfy humans. Fourth, in some cases, data from user responses to LLM responses is fed back to fine-tune, further adjusting internal weights for still better results.

When all this is said and done, the basic fact remains: The LLM is an ANN that makes statistical inferences of the most likely text in response to a prompt. The likelihoods are set during initial training and tweaking, and can be changed after by further tweaking.



Although a great many benefits can come from this technology, there has also been much discussion of the dangers of LLMs running out of control and subjugating or even decimating humanity. This concern has arisen in part because these machines have finally “entered into language” and some can now pass the Turing test. People are starting to reassess possibilities and dangers based on the impressive linguistic displays of LLMs.

To find a footing in this conversation, we must have some sense of what

it means to “be in language.” Language is not simply grammar and words, it is a milieu of expression, coordination, culture, customs, interpretation, and history that fundamentally fashions our way of being in the world. Despite their surprising capacity to participate in human-like conversations, LLMs do not share other human abilities conferred by language. It is more accurate, as some are saying, to see LLMs as manifesting an “alien intelligence.” LLMs cannot match the ways we humans shape and are shaped by language.

In this column, we examine several interrelated ways language shapes human life so we can better ground speculations about the nature of this emerging alien intelligence and perhaps shine a light forward for designers of AI.

### Care

Care is one of the most fundamental aspects of being human. It is deeply intertwined with our being in language. It is not simply a feeling of affection. Care distinguishes between what matters and what does not. What we care about solicits our attention and action. We are not drawn to be in service of unimportant matters. Our sense of what is important is opened by our being in language.

We use language to articulate what we care about and bring our concerns into focus. Language enables our communities to be guided by discriminations of right and wrong, nobility and baseness, good and evil, humanity and inhumanity. Without language, no agent could be sensitive to or motivated by such standards.

We demonstrate our care by taking stands and sustaining them, with others, in word and deed over an extended time. Being in language enables us to commit our lives and coordinate with others in service to large human concerns such justice, progress, revolution, conservation, romantic love, artistic creativity, and much more. Language gives us the means to care.

One important matter that humans cannot help caring about is the honesty and sincerity of others. We care about getting things right. We care about truth. Machines can do none of this. They do not and cannot care.

### Shared Spaces of Concerns

Language enables us to explicitly share matters of common concern and to coordinate our actions to take care of them. Even seemingly trivial moments of chit-chat, such as discussing the weather, acknowledge and generate shared spaces of concerns. We often call these shared spaces “worlds” and perceive them as realities. We co-create our worlds through our conversations and interactions. We pass on our beliefs, values, and norms to our children through our conversations with

## Predicting what words come next is radically different from making the commitments expressed in those words.

them in our worlds.

This feature of language gives us a sense of belonging to a larger whole, the “we” who share commitments and norms of proper behavior. Our ability to commit and coordinate also enables us to develop lifelong friendships and to let mentors shape our lives. It enables us to clean up distrust, banish resentments, and open new futures together.

Through language, we share convictions, assessments, and opinions. We imitate and influence each other, often without even realizing it. We change each other’s minds. Language lets us socialize and adapt in the shared space, enabling us to shape each other’s ways of thinking, acting, and being.

Language can also be a tool of power and domination, a medium to establish, maintain, and contest social hierarchies. Machines do not have concerns and are incapable of forming social spaces of shared concerns.

### Commitments

Language enables humans to make and deliver on commitments. Our commitments structure our worlds.

Commitments are always social. We make commitments to other people. We hold each other responsible for how we live up to our commitments. Consistent success at fulfilling commitments generates a rapport of trust that circulates in our communities. And consistent failure generates distrust, anger, and resentment. Trust in turn shapes the interactions others are willing to have with us.

Commitments are essential for coordinating actions and co-creating a

future. Two main linguistic vehicles for coordinating actions are requests and promises—a request asks for something missing and a promise provides it. People make promises and requests because they have concerns about how things are going. But promises and requests cannot be reduced to mere formulaic sequences of words. They are events in the relationships between people. They convey commitment and generate expectations of future actions. Making a promise and accepting a request both involve setting a stake in the future and guiding our subsequent actions to take care of the underlying concern. When these expectations are not met, or when the concern has been misinterpreted, a breakdown in the relationship often happens, and further conversations are needed to repair it.

A statistical model of language can potentially track conversations, and it can keep records of agreements and convey them to the responsible parties. But without the common sense to discern the vagaries of human concern, without an embodied presence in the world to carry out future actions, and without an emotional susceptibility to breakdowns in relationships that can happen when promises and requests are broken, LLMs cannot yet participate in this all-important dance of human language.

Predicting what words come next is radically different from making the commitments expressed in those words. LLMs cannot make any commitments at all. We already know this intuitively: if an LLM fails or causes damage, we do not blame the machine or hold it responsible, we hold the designers responsible.

### Moods and Emotions

Language permeates our emotional life. Moods and emotions are among the most important ways we experience the world together. Moods are embodied dispositions that shape the possibilities we can see. Emotions are embodied reactions to events. Both are closely linked to our ability to make assessments in language. An emotion is a reactive assessment of a current event; a mood gives us assessments about the future, shaping what actions are possible for us. By examining and

## INTERACTIONS

ACM's *Interactions* magazine explores critical relationships between people and technology, showcasing emerging innovations and industry leaders from around the world across important applications of design thinking and the broadening field of interaction design.

Our readers represent a growing community of practice that is of increasing and vital global importance.



To learn more about us, visit our award-winning website <http://interactions.acm.org>

Follow us on Facebook and Twitter  

To subscribe: <http://www.acm.org/subscribe>

Association for  
Computing Machinery



sharing these assessments, we have the capability to be aware of, to anticipate, and to explore our own and other people's moods and emotions.

Moreover, language itself is permeated by emotional resonance. We can be insulted or flattered simply by the way someone talks to us. We would find it jarring if a waiter at an upscale restaurant spoke in the same way as a cashier at a late-night fast-food joint. Language enables us to bring such experiences into focus and make sense of them relative to a larger context. Not only do individuals experience moods, communities experience collective moods, such as anxieties about pandemics, joy when a sports team wins a match, or distrust of government institutions. Our actions provoke emotional and mood responses in ourselves and others. For example, someone who routinely lies or makes insincere promises evokes anger, resentment, or indignation. Competent leaders read and flow with moods, avoiding making requests when people are not in receptive moods, and making requests when they are.

Language enables us to bring moods and emotions into focus, allowing us to reflect on why we have them and deepen our sensibilities to them. Only an agent immersed in language can do likewise. Only an agent capable of moods and emotions, and the standards and concerns they express, can enter the common space of language as we live it.

Machines such as LLMs can generate text strings that signify emotions and moods. But these are statistical constructions. Having no concerns and no bodies, machines have no emotions and no moods, and no means to develop sensibilities for them.

### The Background

Our language depends upon and conveys the ripples of conversations passed down through years and centuries from prior generations. Our beliefs, customs, mannerisms, practices, and values are inherited from the conversations of our forebears, combining with the conversations we live in. We think, speak, and act against this historical background of presuppositions and prejudices without being aware of it. This background is boundless, with

no definite beginning or end, extending beyond every horizon.

We have the remarkable ability to sense and reveal what is in the background, to make what is tacit explicit, to “make sense” of current issues. Often, we do this in a process of exploration, asking each other why we said or did something. Such exploration in conversation brings forth new meanings and emotions. Poets do this professionally, by revealing our shared background and transforming our sense of it.

Paradoxically, we often react to a revelation of something hidden in the background with “that’s obvious.” It is obvious because it fits the background even though it was not obvious the moment before it was revealed. What we call “common sense” is all that goes without saying in this tacit “background of obviousness” that nevertheless makes sense when revealed and brought into conversation.

In the 1980s, failures of expert systems were attributed to missing “common sense facts” that are obvious to us, but not to the machine. Expert system designers sought compendia of common-sense facts that the machine could use. Perhaps the most famous of these efforts was the Cyc project of Douglas Lenat, which after 40 years had accumulated 25 million common-sense facts. Yet even that treasury could not add up to a background of common sense and make expert systems smart enough to be experts.

Now that we have LLMs, it is reasonable to ask whether these machines can infer background context statistically. Since all these machines can do is infer from already written texts, and since people are generally unaware of their tacit knowledge and cannot write

**LLMs can generate text strings that signify emotions and moods, but do not have either.**

about it, it seems unlikely these machines can infer text that has not been written or recorded.

Imagination is another human ability that flows from our tacit background. It is a capacity to conceive possibilities that do not exist and can become incorporated into our shared background once articulated. Although LLMs have generated some surprisingly imaginative poetry, it is more likely that these are unexpected statistical inferences rather than genuine creations relative to the background. This question deserves more exploration.

### Embodied Action Beyond Language

Language orients to what is important in our actions, but our abilities to act exceed our linguistic powers. Much of what we “know” is in the form of embodied practices rather than descriptions and rules—knowing how rather than knowing that. Even if we can linguistically describe a practice, reading the description does not impart the skill of performing the practice. Michael Polanyi, a philosopher, captured the paradox in his famous saying, “We know more than we can tell.”

Descriptions of actions can be represented as bits and stored in a machine database. However, performance skill can only be demonstrated but not decomposed to bits. Performance knowledge, what psychologists call “procedural memory” (memory of how to do things), is deeply ingrained into our embodied brains, nervous systems, and muscles. This intuitive, embodied sense of relevance resists being objectively measured, recorded, or described.

Language for discussing performance skill was invented by the brothers Stuart and Hubert Dreyfus in 1980. They defined levels of performance in a domain, which they called beginner, advanced beginner, competent, proficient, and expert. The beginner has no embodied skill and can only perform by explicitly following decontextualized rules. The expert has a fully embodied familiarity with typical situations and acts without following rules. Criteria for performance at the various intermediate levels are defined by increasing embodiment and decreasing reliance on rules. This

## How much of the language we use every day can be modeled by statistical inference?

emerging know-how can only be acquired through practice, often with the help of coaches and mentors who already have the skill. The Dreyfuses argued machines cannot attain the skills of experts because experts do not rely on rules and machines have no biological bodies.

Machines store knowledge given them by rules, algorithms, and data. This applies to traditional logic machines, which are programmed, and modern neural networks, which are trained over given data. The statistical inferences performed by LLMs are computed by the algorithms defining the operation of the neural network. Because tacit knowledge cannot be recorded, it seems unlikely statistical inference from recorded data can reveal it. In contrast, human bodies live and interact in their vast and intangible interpretative structures constantly shaped by tacit knowledge. It is impossible to distinguish the bodies of experts from their expertise. The ability to continuously adjust interpretations is far beyond the capabilities of any known or anticipated machines.

In fact, this is the reason we design and build machines. They can muster calculation speeds or marshal kinetic forces well beyond human capabilities. Machines with an exogenous “body” of hardware can get their gears, levels, hydraulics, and circuits to do tasks on a scale that is impossible for embodied humans. That is what makes machines valuable to us.

### Conclusion


LLMs reveal striking new regularities in our use of language and have harnessed these to imitate human conversation in a deeply impressive way. In doing so, they have exposed some seri-

ous gaps in our understanding of language. How much of the language we use every day can be modeled by statistical inference? Can the mathematics of inference (Bayesian methods) make inferences that no one has ever considered? Are those inferences “creations” or just “revelations” of what was hidden in the data?

These questions point to a fundamental mismatch between statistical inference and human responsibility. Inference is a third-person phenomenon susceptible to mathematical formalization. Humans make first-person commitments by taking stands, accepting risks and responsibility, staking out futures, and being open to the assessments and emotions of others. None of these is susceptible to mathematical formalization.

The best strategy is to acknowledge that humans and machines each have powers that the other lacks. Then focus on designs of machines and their interfaces that augment human powers with machine powers.

The claims that LLMs will soon include all human knowledge are nonsense at face value. Many humans live in circumstances where they cannot express themselves. Their knowledge is not recorded in the texts of humanity. Human tacit knowledge is not recorded in the texts of humanity. Moreover, the texts on the Internet are contaminated with problematic biases and are now being polluted with a large amount of synthetic text generated by LLMs.

Humans live in language. Machines are outside of language. If machines develop an intelligence, it will seem very alien to us and we might regret our achievement. 

**Peter J. Denning** (pjd@nps.edu) is Distinguished Professor of Computer Science at the Naval Postgraduate School in Monterey, CA, USA, is Editor of ACM *Ubiquity*, and is a past president of ACM. His most recent book is *Computational Thinking* (with Matti Tedre, MIT Press, 2019). The author’s views expressed here are not necessarily those of his employer or the U.S. federal government.

**B. Scot Rousse** (bsrousse@pluralisticnetworks.com) is the Director of Research at Pluralistic Networks, and a visiting scholar in the Department of Philosophy at UC Berkeley, CA, USA. His recent work has appeared in *Continental Philosophy Review* and *European Journal of Philosophy*.

The authors thank David Dohan, Stuart Dreyfus, Fernando Flores, Milan Griffes, Pentti Kanerva, and Patrick McClure for conversations as we were preparing this column.

© 2024 Copyright held by the owner/author(s).