



Who is the human in the machine? Releasing the human–machine metaphor from its cultural roots can increase innovation and equity in AI

Gwyneth Sutherland¹

Received: 11 August 2023 / Accepted: 14 November 2023

This is a U.S. Government work and not under copyright protection in the US; foreign copyright protection may apply 2023

Abstract

Computer science and cognitive science have a shared past, with many intertwined goals and perspectives. The conceptual metaphor, shaping the discoveries of these fields for decades, has been *the human mind–machine*. New cross-cultural findings indicate that it is time that we interrogate the origin of the metaphor and develop a more global representation of attributes labeled *human*. This paper describes a gap in fairness research in cross-cultural bias affecting international participation in the field. It further outlines opportunities to diversify and test core concepts inspiring design and increasing equity. The proposed adaptation would shift our approach to knowledge and technology creation by (1) altering the attributes of the *human mind–machine* metaphor that define intelligence, memory, categorization, logic, inference, perception, concepts of time and space, concepts of personhood, and other cognitive terms which both fields study; (2) interrogating the universality implied by the conceptual metaphor to both machine and end-user; and (3) seizing the broadened conceptual metaphor to create new math, science, and disrupt the current paradigm scripting the inferences of research findings in computer science and cognitive science. A more globally attuned conceptual metaphor, updated to enfranchise the full membership the term *human* implies, will increase our collective ability to investigate, describe, and develop new science and technology and increase the equity of those involved in the process.

Keywords Artificial intelligence · Cross-cultural · Cognitive science · Equity · Human · Machine · Metaphor

1 Introduction

Historically, we have amended phrases such as *man's inalienable rights* to become *human rights* in an acknowledgement of a pervasive gender and racial bias affecting laws and policies at a fundamental level which require new language to provoke holistic change. The term *man* originated from its authors, reflecting the original members of a group which was later, slowly extended to describe others. Research which explored evidence-based similarities or differences coincided with specificity in terms guiding our assertions and inferences. In our technological era, we, the community of scientists, must interrogate the term *human* as it is commonly used to relate to and understand artificial intelligence

(AI) through the *human mind–machine* metaphor, because this shapes power balances in each sector of society, where AI plays a role. This term has its own origins stemming from the membership of its authors. Marvin Minsky, often referred to as the father of AI, wrote about ‘mechanical brains’ in 1961 [1] while also describing efforts to understand the human mind. The language we use to solve complex, abstract problems delineates how we think about them as Lakoff and Boroditsky have consistently described over decades. [2–5] It follows that the language we use to think about AI design has the power to shape our path to innovate. This paper outlines a gap in fairness research, describing the root of the ethical dilemma affecting global equity in AI, and an opportunity to test a potential solution by asking: who is this *human* in the machine?

There has always been an unnamed, ideal human giving shape to the conceptual metaphor, “the mind is like a computer, and vice versa,” which Crawford [6] quoting Ullman [7] writes has, “infected decades of thinking in the computer and cognitive sciences.” This universal, ideal human made

✉ Gwyneth Sutherland
gwyneth.b.sutherland.civ@ndu.edu

¹ College of Information and Cyberspace, National Defense University, Washington, DC, USA

sense several decades ago within the western hub of these sciences when the researchers and subjects in these endeavors were more homogenous—few women, participants of color, or cultures from outside the U.S. and western Europe were recorded in literature as contributing to early cognitive or computer science. More recently, our fields have increased our acknowledgement and understanding of what diversity of perspective brings to innovation. However, we have yet to update this imagined, anonymous human lurking in our research. It persists in its original concept, asserting universal qualities for all types of humans without having integrated the latest ideas on what that could mean. As this article will outline, cognitive science demonstrates more and more ways that humans do not all think alike, so why do we continue designing ‘thinking machines’ as though there is a single, model human, an artificial intelligence, that our designs should aspire to? How can current cognitive science, capturing the breadth of human thinking, once again inspire computer science and vice versa through this conceptual metaphor of the *human mind–machine*, so that we may increase the variables and dimensions of innovation alongside the equity of participation in the process?

Computer science and cognitive science have shaped one another through the powerful conceptual metaphor of the *human mind–machine* allowing us to explore the circuitry of our brains and design complex thinking machines. Baria [8] reviewed the pervasive nature of this metaphor in science and popular culture in his paper titled, “The Brain is a Computer is a Brain.” This mind–machine meld began in the 1950s, with what Miller called, “The Cognitive Revolution.” [9] It combined breakthroughs in psychology, linguistics, and computer science charting a new understanding for how we think; the terms brain and mind both entered the literature shaping the debate on distinct aspects of the human experience. This paper is concerned with how the use of such terms, in a conceptual metaphor, constrain the process of exploration to attributes authored decades ago by pre-determining what qualities a human will or will not possess. We have been laboring under a narrow definition of this invisible, implied *human*, lurking in our machines since the 1950s that invades current research guided by the *human mind–machine* metaphor; it is due further scrutiny by a wider field of scientific perspectives.

There are numerous research areas in cognitive science and computer science where this inequity plays a role. The fields overlap because the language we use to describe our inquiry has been co-developed. Due to their intertwined history, both fields perform research on memory, attention, problem-solving, logic, reasoning, cause and effect, spatial and temporal concepts, categorization, perception, vision, and more [10]. Complicating things further, the fields merge as we consider the human–machine interface, the ways in which we communicate with machines. Here, the

anonymous *human* in the metaphor stands in for both someone who will use the technology and the imagined set of cognitive processes extended to the ‘thinking machine.’ By imposing the conceptual metaphor on both the technology and the consumer, we obscure a means to critically evaluate the validity of our work. We have defined what we will certify as logical and intelligent, determining a priori the qualities of the archetypal human mind.

A frequent term associated with the *human mind–machine* metaphor is *artificial intelligence*. It evokes a distillation of the mind into something more potent, the stuff of science fiction machines surpassing humans. But why *the* human mind, singular? Do all humans really think alike? And who is this idealized human that drives our scientific ambitions and fuels our fears of a robot-rebellion? We are so accustomed to the metaphor that we have failed to interrogate the underlying assumptions and corresponding research outcomes. This paper is not concerned with the debate about whether AI can match or surpass human reasoning [11], rather it is focused on unpacking the generalized and overlooked *human* which features consistently in the discourse. Certainly, there is research describing the role of conceptual metaphors in reasoning [12]; debate advocating a shift in the lexicon of this particular metaphor [13]; discussion exposing the harm in the narrow concept of intelligence [14]; uncertainty about the anthropomorphization of machines integrating into our lives in valuable roles or driving a reductionist, gendered regression [15]; and many demonstrations of the power constructs inherent in AI which have omitted under-represented groups in design propagating bias in terms of age, gender, race, and accessibility [16, 17]. While these debates have generated research devoted to fairness, the cultural roots of the metaphor have remained unchallenged [18]. Re-imagining the concept of the *human mind–machine* to include the breadth of global cultural perspectives creates vast potential for both computer science and cognitive science to explain phenomena, solve hard problems, and increase authorship among scientists and technologists globally. This strategy presents an opportunity to re-write the bedrock of both disciplines and remake them in a new image which may include more math, diverse concepts, or even redirect research priorities.

2 A dividing metaphor

This metaphor has been pervasive, affecting scientists, technologists, and the consumers of their work. It has established two cultural groups—a *producing culture* that authored the concepts reflecting views from regions known as the west or the global north, and the *culture of use* which receives and consumes these concepts. This

dominant vs. other dichotomy is a powerful, even purposeful, barrier to scientific advancement.

This first group is the culture where cognitive science and computer technology have evolved together. Several studies [19, 20] combine to paint a picture of just how limited the knowledge of the producing culture remains about the culture of use, despite the perceived openness of knowledge sharing via internet connectivity. These studies describe instances of formal knowledge from academic journal submissions and acceptance rates to informal platforms including Wikipedia and GitHub. Knowledge about how humans behave, theories from sociology, anthropology, psychology, etc., continues to represent a small fraction of the earth's populations from which scientists from the producing culture extrapolate without sufficient evidence to support the claim of universality about what is *human* across all cultural groups. For example, it is widely recognized that, "Most studies on memory have tested individuals that come from western, educated, industrialized, rich and democratic societies [WEIRD] – all characteristics which are rather atypical when compared to those of other humans. Moreover, the languages they speak hardly represent the linguistic diversity found across the world." [21, 22] Can we then infer that the concept of memory that originally inspired computing memory is only one of many potential ways to conceptualize human memory? What could this mean for computing in terms of efficiency and processing?

'War of the Ghosts,' often described as the first cognitive science experiment, brings cultural cognitive variations, such as those linked to orality, memory, and event perception, into relief. First performed by Sir Frederic Bartlett in the 1920s, the results consistently illustrate the dynamics of one culture having authorship over another's information [23]. Participants read a translated Chinook story aloud, find it difficult to understand because it does not match their anglophone narrative schema, and for these same reasons, struggle to accurately recall the story because their memory cannot make sense of information presented in this unexpected pattern. Upon retelling, they reorganize the details into a more linear timeline, change speakers, and strip away or change details to conform to their own narrative expectations. By homogenizing the foreign narrative to fit their norms, they have reauthored the story. This is how technology designed in the producing culture defines the terms for users from other cultures, prescribing rules of cognitive engagement. As we consider perception, vision, attention, categorization, inference, memory, and other cognitive processes, how can we adapt how we talk about cognition as well as how we conceptualize it by re-imagining the metaphor of the human mind-machine?

Blasi et al. made an extensive review of cognitive science literature to argue there is:

...an emerging body of evidence that highlights how the particular characteristics of English and the linguistic habits of English speakers bias the field by both warping research programs (e.g., overemphasizing features and mechanisms present in English over others) and overgeneralizing observations from English speakers' behaviors, brains, and cognition to our entire species [24].

Creating this knowledge stovepipe has a more serious, compound effect. An extensive report in *Nature* by Park et al. [25], reviewed millions of citations across several science and technology fields over six decades. The authors assert that research is "less likely to break with the past" and there has been a "decline in disruptiveness" due to trends in citation building on previous work. The human mind-machine metaphor is reproduced through this trend, and it is worth asking if the increased frequency is due to the scientific merit or narrative ease. The pervasive reference to a mono-cultural construct throughout the literature is a profound obstacle to innovative thinking. Technical textbooks in English teach how to design and develop computing advances cementing this approach to problem-solving. Barton notes the same English-language trend for mathematics which he asserts privileges certain concepts that can be described in English. Cognitive science has demonstrated the process of problem-solving, linking cause and effect, imaging, even creativity is culturally variable, so could we unlock more scientific breakthroughs by altering this widespread conceptual metaphor, and inviting additional cultural perspectives? Isn't it worth testing if the inclusion of more ways of thinking and seeing the world will enrich how we design AI?

The second group is the *culture of use* that encompasses the myriad of populations where the knowledge and technology has spread that are increasingly distant to the producing culture. This distance is not geographic, linguistic, or an amorphous cultural divide. The distance is cognitive. It can be seen in how cultures think, their worldview, logic, decision-making, sense of right and wrong, the concepts they would imbue into the design of the technologies that capture and share their knowledge and communication. Currently, the science and technology of one culture is imposed on another through the conceptual metaphor of the *human mind-machine* which dictates the construct of thinking tools and intelligent machines.

3 What we know about the human mind

Cognition has been inextricably linked to computer design from the start. In 1987, George Lakoff described it as, "... a new field that brings together what is known about the mind

from many academic disciplines: psychology, linguistics, anthropology, philosophy, and computer science.”[26, 27] More recent work in the social science fields has challenged our understanding of how our minds work, revealing cognition to be less fixed, often shaped by culture. The new findings have not been integrated into the conceptual metaphor, and so have not yet impacted AI design. For example, Amici, F. et al.[28] identify the cultural variability of working memory depending on the direction of written language and discuss the connection to higher cognitive functions such as problem-solving and planning; Odejobi and Adegbola [29] assert that technology needs to represent African concepts and logics; Bidwell [30] observes design mismatches between local and universal (i.e., dominant) in social media interface design; Tefera and Gamlen [31] contribute to the theory of temporal logics across cultures and provide an example of locations that are understood by their pace of life rather than geographic coordinates or landmarks; Lakoff and Nunez [32] contend we use spatial concepts including motion, bodily orientation, manipulation of objects (rotation, stretching) to conceptualize math such as algebra and calculus, while research [33, 34] has shown these concepts to have high cultural variation. Blasi et al.[35] confirm an anglophone bias with their extensive review of numerous concepts that vary by culture, including those related to mathematics. Ethnomathematics and ethnocomputing [36], fields that describe cultural concepts to teach from global perspectives outside the dominant paradigm, offer a glimpse of possible additions to the current human mind–machine metaphor.

Nesbitt[37] contrasts western and east Asian concepts including logic and categories emphasizing the western reliance on rules ordered by categories to perform many cognitive processes while east Asians use contextual, relational factors to convey complexity and may consider the use of logic (as understood by westerners) to signal immature thinking. Nesbitt’s comments on the role of categorization overlap with potential topics of investigation including problem-solving, event conceptualization, and causal inference. As each of these is shaded by culture, so too are the thinking processes which employ them. Nesbitt’s observations highlight the frequent sampling gap for oral cultures and the diverse processing auditory perception has in all higher cognitive functions. Together, these authors’ findings are an extension of sociocultural shifts throughout academia that seek to reexamine theories built from distinctly homogeneous anglophone samples whose lack of diversity undermines any claims of universality in how humans think. Atari et al. [38] similarly assert from their own review of recent cognitive science research that, “it is misleading to refer to a monolithic category of “humans” when so much psychological diversity lies across human populations.”

If cognitive science has new insights, an expanded range of potential models on which computing and, critically, AI can be based, what could these look like? The current asynchrony of these previously wedded sciences has exposed a gap between the culture of production and the cultures of use whose cognitive processes are now being studied more deeply.

4 Thinking machines based on how humans think—all of us

There has been a call to arms to develop localized technologies, transforming from many *cultures of use* into producing cultures. Groups all over the world are coalescing around their own requirements which reflect unique cultural and cognitive variations, advocating a rejection of a single producing culture’s norms. The requirements include being able to capture and reflect their own concepts, logics, identity (e.g., the spectrum including collectivist and individualist), non-linear time, cause–effect relationships, alternatives for personhood that permit dual or multifaceted roles simultaneously, or agency [39] categories just to name a few [40, 41].

Personhood is particularly tricky, but foundational to capture the very core of authorship for research or technology design, and foreground notions about status, kinship, number, gender, membership or relationship, and (more ambiguously) persons vs. selves [42]. It is perhaps the most complex application of categorization. Lakoff cautioned:

Categorization is not a matter to be taken lightly. There is nothing more basic than categorization to our thought, perception, action, and speech. Every time we see something as a kind of thing, for example, a tree, we are categorizing. Whenever we reason about kinds of things—chairs, nations, illnesses, emotions, any kind of thing at all—we are employing categories. . . . Without the ability to categorize, we could not function at all, either in the physical world or in our social and intellectual lives [43].

It is Lakoff’s use of *we* that is of interest. And in the current paradigm, a narrow group of authors imply universal human characteristics modeled in their image.

Examining how cultural differences in thinking would manifest for a local AI ecosystem, Kalyanakrishnan et al. argue for, “. . . the need to plan ‘AI for India’ from the *bottom up*, by paying attention to India’s social, political, cultural, and economic configuration [44].” The authors suggest developing a context-specific suite of technologies that departs from the conventional paradigm. The authors explain to readers from the technology producing culture why considering local requirements are necessary:

Since India is significantly behind many other countries in its technological development, it is natural for technologists and policy makers to look to transplant successful ideas from other contexts into India. A growing body of literature warns of the inefficiency, even danger, of such an approach. . . . Our proposal goes in the opposite direction. . . [45].

They go on to articulate how to develop less complex, building-block technologies including user interfaces, search engines, subtitling, and all the other information and communication technologies (ICTs) that can generate linguistic richness for natural language processing (NLP), a core element of AI. The authors suggest creating a local digital ecosystem, the kind readers take for granted, but which makes the AI revolution possible in the western world. Most importantly, these researchers explain the impact, both good and bad, when incorporating culture as a variable in design and considering the application of AI to significant problems within a society such as health, governance, economy, and the environment. For these reasons, we must examine how early computing design choices, made without consideration of cultural variation in cognition, persist in maintaining a dichotomy between a culture of production and culture of use.

Hiring a local workforce is not a guarantee of localizing the design for thinking tools. Most textbooks that teach coding are in English. Even in India, a technical powerhouse, they have only just started creating manuals to teach new coding languages like Python in any of the 190 vigorous living languages [46]. (discussion board requests for Python manuals and courses observed by the author.) Authors like Abbate [47] have argued that diversification of workforces is a new economic strategy to make companies or products more competitive. Major funding often comes from producing culture giants expanding into new locations. Initially, mirroring existing technology (the transplant approach) has shifted to encourage some localization without affecting the underlying conceptual metaphor scripting what an intelligent, thinking tool should be. The research from cognitive science and linguistics explains how problem solving, memory, and imagination is formed in a second language reinforcing concepts from the producing culture, rather than exploring design representing local ways of thinking.

To ascribe bias to the data and algorithms of today is to see only the surface of the problem that is projected vividly as the image of an African American face erroneously tagged as a gorilla [48]. Lamenting the limited participation in the construction of AI, again, fails to see how long this project has been in the making. A singular worldview has been inculcated in all the technologies that have collected, aggregated, monitored, analyzed, and now learned among us. To truly understand how users outside the producing

culture have been impacted over decades, and are now critically at risk from AI technologies that seep into the social, cultural, and political decision-making roles of government and business, we must investigate the logic or mental model captured by conventional technology design and recognize the extent to which it contrasts with other user cultures' mental models, thus limiting their abilities to convey identity, intent, plan, problem-solve, share emotion, monitor health, make decisions, and represent justice on their own terms.

Culture has often been marginalized as the domain of the humanities and social sciences, as the outdated pursuit of anthropologists with notepads in a jungle, or simply reduced to the words and grammar studied by linguists. Amici et al. [28] highlight a core issue that language and culture are frequently conflated, particularly regarding cognition. Put simply, language is one element of culture. The profound effects of culture and cognition have not been adequately researched. This topic of inquiry is particularly taboo when studying cognition raising suspicions of asserting superiority of one group over another. In fact, quite the opposite is true. By recognizing that each culture has a different worldview, set of values, approach to problem-solving, collection of memories, we begin to acknowledge both commonalities and uniqueness, which scientists continue to document through experiments with more diverse participants [49].

5 Broaden the human mind–machine metaphor: a proposed empirical method

Borrowing from Grozinger et al. [50] who suggested a method for challenging the dominant metaphor in their field of cellular computing by describing clear benchmarks to identify the benefits of a potential new conceptualization, I argue the current version of the human mind–machine metaphor may be constraining future innovation by obscuring a more culturally diverse set of concepts and perspectives, while a more globally attuned version offers clear benefits. Hendricks and Boroditsky introduced foreign spatial temporal metaphors to English speakers and found, "...results suggest that learning new relational language can be a powerful tool in constructing new representations and expanding our cognitive repertoire." [51]. Adapting the metaphor will alter how we approach problem-solving, potentially opening paths for innovation.

The research community could test the merits of expanding the metaphor for both computer science and cognitive science by identifying distinct areas of inquiry that have high cognitive cultural variability and explore how they might change methods, inferences, or outcomes of our research. These areas of inquiry may include problems that computer science is struggling with that cognitive variations would provide better solutions for than current methods, i.e., new

concepts that could be operationalized as mathematical variables or formulas. Below are four suggested areas of further research.

An example of the human mind–machine metaphor in action currently with AI are neural networks. When used in natural language processing, neural networks are a type of machine learning that identifies patterns in grammar leading to programs that can anticipate the next word in a sentence. Large language models (LLMs) are a result of this approach. These are built from English grammar concepts which relies on prediction and is itself an artifact of the producing culture. Prediction as a cognitive task is more prevalent in text-based languages vs oral languages [52, 53]. Among the nearly 7000 languages, approximately one hundred have a culture of writing and literature while the vast majority remain oral. This presents a vast, largely unexplored area of research to develop LLMs or generative AI based on an alternative concept, such as one from an oral language. Orality influences the process of memory [54]. Consider, for example, trying to remember directions to a location or doing long-division without writing. Planning and problem-solving processes develop around writing. These are similarly conceptualized and encoded in AI tools. What cognitive processes could we better understand if we examined orality and thought? The field of reinforcement learning could be approached anew inspired by memory and learning concepts described in an oral culture. Blasi et al. remind us to be critical of universality demanding robust sampling and theory and, “...not to sweep variation under the rug” [55].

Abstraction is another cognitive process underpinning machine learning which, according to Zheng et al., is, “ubiquitous to humans’ ability to process vast amount[s] of information and derive general rules, and principles.” [56] What would happen if abstraction was not a universally human process? How could culturally variable approaches to processing information, such as those that function in high-context environments more effectively than current models, add to the machine learning research field?

Emotion recognition, a highly fraught area of AI research [57], could be redefined by the right conceptual metaphor that reimagines emotion in a way that formulates deeper understanding. For example, to illustrate Mesquita’s [58] cultural construction of emotion in which she asserts a novel ‘ours vs. mine’ framework, one might adapt a concept from pacific island navigation to conceptualize her idea. She asserts emotions are located externally to individuals; they are shared, dynamic, negotiated, and of course, not universal [59]; therefore, emotions could be operationalized mathematically as external, independent objects that change or move relative to the individuals who experience them, rather than being descriptive attributes of an individual, thus the emotions move as boats relative to the islands (individuals)

experiencing them. This ‘geometry of influence’ is further reflected in the language of several cultures in the pacific region who employ words meaning fear/anxiety jointly shared by the sender/receiver. [60]. This geometry can be mapped by the shifting trajectory of the emotional vessel as its path, distance, and acuteness.

Differences in moral choices are certainly not trivial. Costa et al. [61] performed a few experiments that touched on moral decision-making across cultures. These researchers asserted that morality was language-dependent by posing two hypothetical questions based on the Trolley Problem to study participants’ written responses in their native languages and in English. They found the responses, as to whether to save one life or many, varied significantly between languages for the same individual. The researchers concluded there was more emotional attachment in the primary language and more distance in the second which accounted for a more utilitarian response in the non-primary language. The current human mind–machine metaphor significantly guided the inferences to understand how emotion, memory, perception, language, and morality are linked. How can a broader cultural concept of mind yield a more robust toolkit to understand these complex and important questions? How can it improve research and development of AI that must navigate morality across cultural boundaries?

6 Future-minded machines

Again, building on Grozinger et al. [62], new inquiries should be compared to conventional methods based on both qualitative and quantitative criteria. Creating these evaluation criteria would be a valuable collaboration within the community. Initially, qualitative descriptions of observed differences such as richness, precision, and equity could be useful and could point toward viable quantitative benchmarks, such as accuracy, speed, efficiency, and repeatability.

A common critique to performing research that includes cultural variations like those described in this paper from members of the producing culture is that, while there may be cognitive variations across cultures, these are slight and would not be meaningful for computing design. First, how do we know, have we tested this hypothesis? Largely, no. It is assumed that members of cultures whose thinking approaches and worldviews might be more robustly captured also feel these variations are slight and not impactful for computing design. Therefore, a single culture, without evidence, decided not to investigate the potential solutions available across cultures and from the inclusion of wider perspectives. It is a glaring scientific and ethical omission that presents opportunities to the community to challenge and learn from. Anecdotally, when discussing this topic in various countries over the last decade, the most memorable

comment I've heard has been, "This difference is obvious. We were wondering when you would notice!"

Without expanding the conceptual metaphor of the human mind–machine, which so strongly influences the interconnected fields of cognitive science and computer science, we risk a kind of recolonization through the design and control of 'thinking tools.' By defining intelligence and outlining logical inference for both fields, a single culture is determining authorship of the qualities attributed to the idealized human mind for everyone else. What are the implications for those who do not fit this mold? As scientists, we must dig into these assumptions and consider how they have scripted global concepts of intelligence, logic, categorization, personhood, memory, time, place, morality and so many crucial elements common to our fields which may, in fact, have wide variability across cultures. How could these variations add efficiencies or elements we have not yet imagined because our slate of variables has been limited by last century's knowledge? Broadening the conceptual metaphor to include more perspectives will have a lasting effect on both cognitive science and computer science research by increasing equity and knowledge.

Declarations

Competing of interest No competing interests or disclosures.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Minsky, M.: Steps toward artificial intelligence. *Proc. IRE* **49**, 8–29 (1961)
- Lakoff, G., Núñez, R. E.: Where mathematics comes from. how the embodied mind brings mathematics into being. **39**, (Basic Books, 2000)
- Boroditsky, L.: Does language shape thought?: Mandarin and english speakers conceptions of time. *Cogn. Psychol.* **43**(1), 1–22 (2001)
- Thibodeau, P., Boroditsky, L.: Measuring effects of metaphor in a dynamic opinion landscape. *PloS One* <https://doi.org/10.1371/journal.pone.0133939> (2015)
- Chkhaidze, A., Buyruk, P., Boroditsky, L.: Linguistic metaphors shape attitudes towards immigration. *Proceedings of the Annual Meeting of the Cognitive Science Society*, **43**, (2021)
- Crawford, K.: *Atlas of AI: Power Politics and the Planetary Costs of Artificial Intelligence*. **6**, (Yale University Press, 2021)
- Ullman, E.: *Life in code: a personal history of technology*. (McDonnell Douglas Systems Integration Company. 136–137 (2017)
- Baria, A., Cross, K.: The brain is a computer is a brain: neuroscience's internal debate and the social significance of the computational metaphor. [arXiv:2107.14042](https://arxiv.org/abs/2107.14042) (2021)
- Miller, G.: The cognitive revolution: a historical perspective. *Trends Cogn. Sci.* **7**(3), 141–144 (2003)
- Pavlenko, A.: Eyewitness memory in late bilinguals: evidence for discursive relativity. *Int. J. Biling.* **7**(3), 257–281 (2003)
- Large-scale AI language systems display an emergent ability to reason by analogy. *Nat Hum Behav* **7**, 1426–1427 (2023)
- Thibodeau, P., Boroditsky, L.: Metaphors we think with: the role of metaphor in reasoning. *PloS One* (2011)
- Baria, A., Cross, K. (2021)
- Bender, E., Gebru, T., McMillian-Major, A., Mitchell, S.: On the dangers of stochastic parrots: can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FACCT '21)*. Association for Computing Machinery, 610–623 (2021)
- Noble, S. U.: Your robot isn't neutral. In: your computer is on fire. Mullaney, T., et al. (eds). (MIT Press, 2021)
- Kalluri, P.: Don't ask if artificial intelligence is good or fair, ask How it shifts power. *Nature* **583**(7815), 169–169 (2020)
- Mullaney, T.: Your Computer is on Fire. In: Your Computer is on Fire. Mullaney, T., et al. (eds.) (MIT Press, 2021)
- Brockman, J. (ed.) *Possible minds: 25 ways of looking at AI*. (Penguin Books, 2019)
- Graham, M.: Geographies of information in Africa: Wikipedia and user-generated content. *R-Link: Rwanda's official ICT magazine*. Kigali: Rwanda ICT Chamber, 40–41 (2013)
- Graham, M., De Sabata, S., Straumann, R., Ojanpera, S.: Uneven digital geographies Part 2. In: *Kollektiv Orangotango* (eds.). This is Not an Atlas: a global collection of counter cartographies. (2018)
- Amici, F., et al.: The word order of languages predicts native speakers' working memory. *Nat. Sci. Rep.* **9**(1124), 1 (2019)
- Henrich, J., Heine, S.J., Norenzayan, A.: The Weirdest people in the world? *Behav. Brain. Sci.* **33**, 61–83 (2010)
- Sutherland, G.: "War of the ghosts": How we think through information technology across cultures." In: *teaching public diplomacy and the information instruments of power in a complex media environment*. Walker, V., Finley, S. (eds). (United States Advisory Commission on Public Diplomacy, 2020)
- Blasi, D., Henrich, J., Adamou, E., Kemmerer, D., Majid, A.: Over-reliance on english hinders cognitive science. *Trends Cogn. Sci.* **26**(12), 1153–1170 (2022)
- Park, M., Leahey, E., Funk, R.J.: Papers and patents are becoming less disruptive over time. *Nature* **613**, 138–144 (2023)
- Lakoff, G.: *Women, fire, and dangerous things: What categories reveal about the mind*. (University of Chicago Press, 1987)
- Miller, G. (2003)
- Amici, F. et al. (2019)
- Odejobi, T., Adegbola, T.: Computational and engineering issues in human computer interaction systems for supporting communication in African languages. In: *handbook of research on discourse behavior and digital communication: language structures and social interaction*. O.A. Taiwo (ed.), (2010)
- Bidwell, N.: Moving the center to design social media in rural Africa. *AI & Soc.* **31**(1), 51–77 (2016)

31. Tefera, G.W., Gamlen, A.: Temporal logics in urban place-making: the case of refugee background ethiopians in Australia. *Urban Geography* (2021)
32. Lakoff, G., Nunez, R.: **28** (2000)
33. Brown, A., Gullberg, M.: Bidirectional cross-linguistic influence in event conceptualization? Expressions of path among Japanese learners of English. *Bilingualism: Language and Memory* **14**(1), 79–94 (2011)
34. Boroditsky, L.: Does language influence culture? *The Wall Street J.* (2010)
35. Blasi, D., et al. (2022)
36. Barton, B.: *The language of mathematics: telling mathematical tales.* (Springer, 2008)
37. Nisbett, R.: *The Geography of Thought: How Asians and Westerners Think Differently... and Why.* (Simon and Schuster, 2010)
38. Atari, M. et al. Which humans? *PsyArXiv*, (2023)
39. Fausey, C., Long, B., Inamori, A., Boroditsky, L.: Constructing agency: the role of language. *Front. Psychol., Sec. Cultur. Psychol.* **1**, (2010)
40. Mawere, M.: *African belief and knowledge systems: a critical perspective.* (Langaa Research & Publishing CIG, 2011)
41. Sutherlin, G.: *The Myth of the Universal user-- pursuing a cultural variable in ICT design for conflict management through quantitative analysis: implications from a Ugandan case Study.* PhD Thesis. (University of Bradford, 2014)
42. Jarvis, S., Pavlenko, A.: *crosslinguistic influence in language and cognition.* (Routledge, 2007)
43. Lakoff, G. (1987)
44. Kalyanakrishnan, S., et al. AIES proceeding: opportunities and challenges for artificial intelligence in India. New Orleans, LA, USA. February 2–3, 1–2 (2018)
45. Kalyanakrishnan, S., et al. **3**, (2018)
46. Simons, G., Fennig, C. (eds.) *Ethnologue: languages of the World*, Twenty-first edition. (SIL International, 2018)
47. Abbate, J.: Coding is not empowerment. In: *your computer is on fire.* Mullaney et al. (eds.) (MIT Press, 2021)
48. Byrne, W.: “Now is the time to act to stop bias in AI.” *Fast company.* (2018)
49. Jarvis, S. & Pavlenko, A. (2007)
50. Grozinger, L., Amos, M., Goroehowski, T.E., et al.: Pathways to cellular supremacy in biocomputing. *Nat. Commun.* **10**, 5250 (2019)
51. Hendricks, R., Boroditsky, L.: New space-time metaphors foster new nonlinguistic representations. *Top. Cogn. Sci.* **9**, 3 (2017)
52. Brouwer, S., Ozkan, D., Kuntay, A.: Verb-based prediction during language processing: the case of Dutch and Turkish. *J. Child Lang.* **46**(1), 80–97 (2019)
53. Huettig, F., Mani, N.: Is prediction necessary to understand language? probably not. *Lang. Cognit. Neurosci.* **31**(1), 19–31 (2016)
54. Ong, W.: *Orality and literacy: The technologizing of the Word.* (Routledge, 1982)
55. Blasi, D., et al. **6**, (2022)
56. Zheng, H. et al. Take a step back: evoking reasoning via abstraction in large language models, [arXiv:2310.06117](https://arxiv.org/abs/2310.06117) [cs.LG] **2** (2023)
57. Long, H., Peluso, N., Baker, C.I., et al.: A database of heterogeneous faces for studying naturalistic expressions. *Sci. Rep.* **13**, 5383 (2023)
58. Mesquita, B.: *Between us: how cultures create emotions.* (Norton, 2022)
59. Mesquita, B., De Leersnyder, J., Boiger, M.: The cultural construction of emotion. *Curr. Opin. Psychol.* **8**(31–36), 29 (2016)
60. Jarvis, S. & Pavlenko, A. **126**, (2007)
61. Costa, A., Foucart, A., Hayakawa, S., Aparici, M., Apesteguia, J., Heafner, J., Keysar, B.: Your morals depend on language. *PLoS One* **9**(4), e94842 (2014)
62. Grozinger, L., Amos, M.: Goroehowski, T.E. et al. (2019)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.