Are People Still Smarter than Machines? If so, Why?

James L. McClelland (jlmcc@stanford.edu)
Department of Psychology, Stanford University, USA and
DeepMind Research, London, England, UK
Postal Address: Room 340-420
450 Jane Stanford Way, Stanford, CA, 94305, USA

Keywords: Computational Intelligence, Language Understanding, Cognitive Modeling, Neural Networks, Mathematical Cognition

Breakthroughs and Remaining Challenges

The last few years have witnessed amazing breakthroughs in machine intelligence. Google's Neural Machine Translation System, DeepMind's AlphaGo, and current 'Transformer' based Language Models (BERT and GPT-3) have amazed many with their success in addressing abilities humans possess but were previously hard to capture with machines. These models rely on neural networks of the kind we advocated in *Parallel Distributed Processing* (Rumelhart, McClelland *et al,* 1986). In these volumes, we argued that people were smarter than machines, and that human-like abilities would be easier to capture using neural networks rather than the discrete symbolic approaches in vogue at the time, and in some ways, this recent progress bears this out. Yet in this talk, I will argue, in agreement with other commentators (e.g. Lake *et al.,* 2017) that we still have a long way to go before we can say any machine has truly captured human like cognitive and learning abilities.

Unlike Lake *et al.* and some other commentators, however, I will argue that we should seek the reasons for many of the amazing achievements of human intelligence not so much in built in biases toward systematicity or special purpose start-up software, but more in a fuller appreciation of the roles of culture and experience. Within culture and experience, I will argue for a central role for culturally constructed formal systems as powerful tools that extend human abilities beyond what can be achieved without these resources. These systems may be as elementary as the counting numbers or as advanced as quantum theory. Between these extremes lie principled systems for logical reasoning and theorem proving, the formal structures of linguistic theory, and the powerful tools of modern computing languages. I will also argue for a central role of language-based instruction and explanation.

Starting Places

A starting place for this is the idea that we should view language as a system for transmitting information about construals of situations between naturally and culturally grounded learners, whether human or artificial. A situation is a set of relationships between objects and their properties, such as the one conveyed by the sentence 'the cat is on the mat' or the one conveyed by the formal statement 'for all θ, $\cos(-\theta) = \cos(\theta)$'. Both sentences implicate construals of objects and relationships that are shaped by convention and experience, though the first arose more naturalistically while the second arose from efforts to predict the locations of celestial bodies and to guide ships safely to their destinations, efforts that gradually became the foci of academic and scientific investigation. In recent work on natural language processing, we (Rabovsky *et al.,* 2018) showed how a neural network model that maps language to representations of situations could capture a wise range of experimental data on the N400, a signal we argued reflected the extent to which new incoming language input results in an update in the listener's internal representation of a situation. In McClelland *et al.* (2020), we laid out a characterization of the understanding system in the brain, construed as a widely distributed collection of brain structures that together allow information from multiple sensory modalities as well as spoken or written language to constrain the construction of situation representations.

In Mickey and McClelland (2017), we turned to a consideration of trigonometry, construing it as the codification (among other things) of facts about the relationship between positions on a circle and ratios of distances in orthogonal directions in space. Specifically, we considered trigonometric expressions like $\cos(-\theta) = \cos(\theta)$, and demonstrated that when students are shown how to treat such expressions as describing relationships between the horizontal or vertical positions of points on a unit circle they acquire a productive understanding of the expression that generalizes to other trigonometric relationships, but when they treat them as arbitrary rules, no generalization to other relationships occurs.

Natural Numbers

Cognitive developmentalists once frequently proposed that children possessed an understanding of the principles underlying the natural or counting numbers, until the paper by Gordon (2004) showed that indigenous Amazonians whose language lacked number words displayed an understanding of approximate but not exact number. Number systems vary extensively across cultures, suggesting that these systems may be culturally constructed and thus must be acquired. In pursuit of an effort to understand the learning processes that might lead to this, I have worked with several others to develop neural network learning systems that can learn to count and that display behaviors reminiscent of the counting behavior of young children. In both of the relevant efforts (Fang *et al,* 2018; Sabatiel *et al.,* 2020) we have envisioned the child as an embodied learner, capable of learning to imitate the actions of a skilled counter. In both studies, the skilled counter has 'announced' the task they are about to perform, allowing the network to learn cue-able task specific procedures such as counting all of the objects in an array, or 'giving' (picking up and placing) a specified number of objects to a target location. The first paper showed interesting hints that the ability to count or produce numbers in the range 5-9 emerged at the same time, so that a network that could count to 5 reliably could also count these larger numbers. In both papers, the networks exhibited knowledge sharing across task contexts, so that in general, the more number-related tasks the network already knew, the more quickly it could learn others. This work does not yet account for all aspects of children's early number behavior, but it points to the possibility that a cultural-specific, experience-dependent approach will provide a useful basis for understanding many aspects of children's performance as they learn to count.

Systematic Cognition

The final line of work takes up the role of instruction and explanation in human learning. One way in which we as humans exceed all non-human species is that we can use direct instruction and explanation to guide our learning and behavior. Some (Fodor & Pylyshyn, 1988; Lake *et al.,* 2017) have argued that the ability to think systematically is a hallmark of human cognition. I agree that the special powers that humans in advanced societies have exhibited through the construction of advanced reasoning systems require systematicity, but I argue that the ability to think systematically is acquired, and the ability to do so in an abstract and general way is only acquired by those with advanced education or other experience thinking in a formally structured way. Studies by both psycholinguists (Gleitman & Gleitman, 1970) and mathematics educators (Burger & Shaughnessy, 1986) support this view. These works and recent work of ours (Nam & McClelland, in preparation) support the view that the ability to rapidly acquire a new systematic reasoning skill is associated with more advanced educational achievement and with the ability to provide an explanation of the basis of one's correct performance. I will also mention approaches to interfacing language instruction and task performance (Abrahamson *et al.,* 2020; Lampinen & McClelland, 2020).

In sum, I hope to argue that one of the important reasons why humans are still smarter than machines is that they rely on language and culturally constructed systems of thought to leverage their intuitive and more basic cognitive abilities. Future work should seek ways to incorporate these features into general purpose computational models that capture these intuitive and more basic abilities.

Acknowledgments

I would like to thank the many cognitive and computer scientists who have influenced my thinking on these issues over the last several years. These include the co-authors of all of the relevant recent papers I have listed in the references, as well as many other colleagues and students.

References

Abramson, J., Ahuja, A., Barr, I., Brussee, A., Carnevale, F., Cassin, M. *et al.* (2020). Imitating interactive intelligence. *arXiv preprint arXiv:2012.05672*.

Burger, W. F., & Shaughnessy, J. M. (1986). Characterizing the van Hiele levels of development in geometry. *Journal for research in mathematics education*, *17*(1), 31-48.

Fang, M., Zhou, Z., Chen, S., & McClelland, J. L. (2018). Can a recurrent neural network learn to count things? *Proc Cog Sci 40*, 360-365. Austin, TX: Cognitive Science Society.

Gleitman, L. & Gleitman, H. (1970). *Phrase and Paraphrase.* New York: W. W. Norton & Co.

Gordon, P. (2004). Numerical cognition without words: Evidence from Amazonia. *Science*, *306*(5695), 496-499.

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, *40*.

Lampinen, A., & McClelland, J. L. (2020). Transforming task representations to perform novel tasks. *PNAS, 117*(52), 32970-32981.

McClelland, J. L., Hill, F., Rudolph, M., Baldridge, J., & Schuetze, H. (2020). Placing language in and integrated understanding system: *PNAS, 117*(42), 25966-25974.

Mickey, K. & McClelland, J. L. (2017). The unit circle as a grounded conceptual structure in pre-calculus trigonometry. In D. C. Geary, *et al.,* (Eds.), *Acquisition of Complex Arithmetic Skills and Higher-Order Mathematics Concepts.* Elsevier/Academic Press.

Nam, A. J., & McClelland, J. L. (in preparation). What underlies rapid learning and out-of-distribution transfer in humans?

Rabovsky, M. Hansen, S. S., & McClelland, J. L. (2018). Modeling the N400 brain potential as change in a probabilistic representation of meaning. *Nature Human Behaviour, 2*, 693-705.

Rumelhart, D. E., McClelland, J. L., & the PDP research group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Vols I & II.* Cambridge, MA: MIT Press.

Sabathiel, S., McClelland, J. L. & Solstad, T. (2020). Emerging Representations for Counting in a Neural Network Agent Interacting with a Multimodal Environment. *Proc Alife 2020.* (pp. 736-743). Cambridge, MA. MIT Press.