

Tessler, M.H., Goodman, N.D., & Frank, M.C. (2017). Avoiding frostbite: It helps to learn from others. Commentary on Lake et al., Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, E279. doi:10.1017/S0140525X17000280

Abstract

Machines that learn and think like people must be able to learn from others. Social learning speeds up the learning process and – in combination with language – is a gateway to abstract and unobservable information. Social learning also facilitates the accumulation of knowledge across generations, helping people and AIs learn things that no individual could learn in a lifetime.

Main text

Causality, compositionality, and learning-to-learn – the future goals for artificial intelligence articulated by Lake et al. – are central for human learning. But these abilities alone would not be enough to avoid frostbite on King William Island in the Arctic Archipelago. You need to know how to hunt seals, make skin clothing, and manage dog sleds, and these skills are not easy to acquire from the environment alone. But if the Netsilik Inuit people taught them to you, your chances of surviving a winter would be dramatically improved (Lambert, 2011). Similar to a human explorer, an AI learning to play videogames like Frostbite should take advantage of the rich knowledge available from other people. Access to this knowledge requires the capacity for social learning, both a critical prerequisite for language use and a gateway in itself to cumulative cultural knowledge.

Learning from other people helps you learn with less data. In particular, humans learn effectively even from “small data” because the social context surrounding the data is itself informative. Dramatically different inferences can result from what is ostensibly the same data in distinct social contexts or even with alternative assumptions about the same context (Shafto, Goodman, and Frank, 2012). The flexibility of the social inference machinery in humans turns small signals into weighty observations: even for young children, ambiguous word-learning events become informative through social reasoning (Frank & Goodman, 2014), non-obvious causal action sequences become “the way you do it” when presented pedagogically (Buchsbaum et al., 2011); and complex machines can become single-function tools when a learner is taught just one function (Bonawitz et al., 2011).

Learning from others comes in many forms. An expert may tolerate onlookers, a demonstrator may slow down when completing a particularly challenging part of the task, and a teacher may actively provide pedagogical examples and describe them with language (Kline, 2014; Csibra & Gergely, 2009). Informative demonstrations may be particularly useful for procedural learning (e.g., hunting seals, learning to play Frostbite). Language, however, is uniquely powerful in its ability to convey information that is abstract or difficult to observe, or information that otherwise does not have a way of being safely acquired (e.g., learning that certain plants are poisonous, how to avoid frostbite; Gelman, 2009). Studying social learning is an important part of studying language learning (Goodman & Frank, 2016); both should be top priorities for making AIs learn like people.

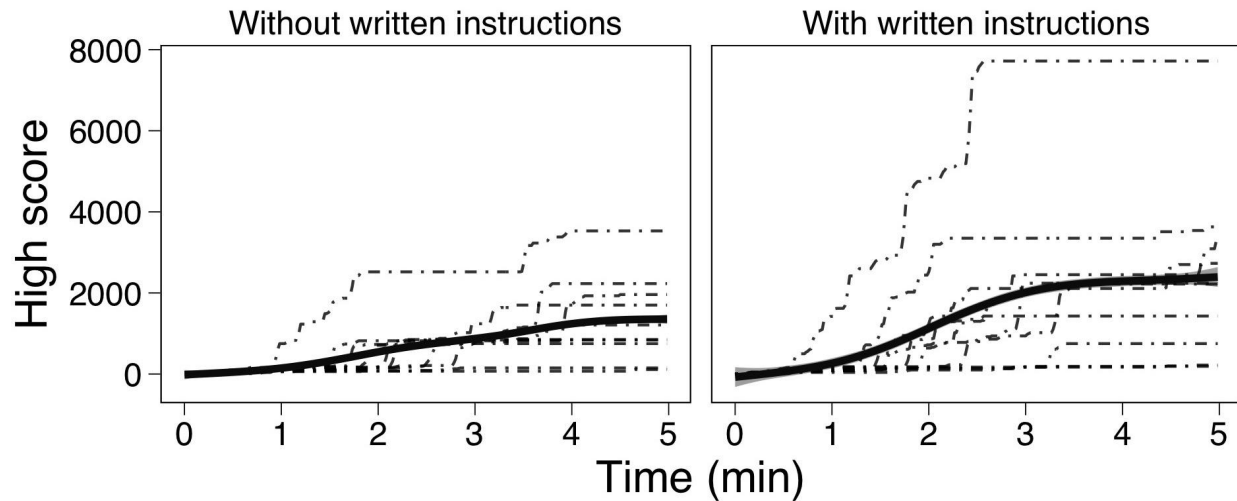


Figure 1. Score trajectories for players in the game Frostbite over time. The two panels depict results with and without instructions about the abstract structure of the game.

Focusing on Lake et al.’s key example, you can even learn the game Frostbite with less data when you learn it from other people. We recruited 20 participants from Amazon’s Mechanical Turk to play Frostbite for 5 minutes. Half of the participants were given written instructions about the abstract content of the game, adapted directly from the caption of Figure 2 in the target article. The other half were not given this information. (Everybody was told that you move the agent using the arrow keys.) Learners who were told about the abstract structure of the game learned to play the game more quickly and achieved higher overall scores ($M=2440$) than the group without written instructions ($M=1333$; Figure 1). The highest score for those without linguistic instructions was 3530 points, achieved after about 4 minutes of play. By comparison, the highest score achieved with linguistic instructions was 7720 points, achieved after 2 minutes of play. Indeed, another group (including some authors of the target article) recently found a similar pattern of increased performance in Frostbite due to social guidance (Tsvidis et al., 2017).

Learning from others also does more than simply “speed up” learning about the world. Human knowledge seems to accumulate across generations, thus permitting progeny to learn in one lifetime what no generation before them could learn (Tomasello, 1999; Boyd et al., 2011). We hypothesize that language – and its flexibility to refer to abstract concepts – is key to faithful transmission of knowledge, between individuals and through generations. Human intelligence is so difficult to match because we stand on the shoulders of giants. AIs need to “ratchet” up their own learning, by communicating knowledge efficiently within and across generations. Rather than be subject to a top-down hive mind, intelligent agents should retain their individual intellectual autonomy and innovate new solutions to problems based on their own experience and what they’ve learned from others. The important discoveries of a single AI could then be shared, and we believe language is the key to this kind of cultural transmission. Cultural knowledge could then accumulate within both AI and human networks.

In sum, learning from other people should be a high priority for AI researchers. Lake et al. hope to set priorities for future research in AI but fail to acknowledge the importance of learning from language and social cognition. This is a mistake: The more complex the task is, the more learning to perform like a human involves learning from other people.

References

- Bonawitz, E., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition*, *120*(3), 322–330. <http://doi.org/10.1016/j.cognition.2010.10.001>
- Boyd, R., Richerson, P. J., & Henrich, J. (2011). The cultural niche: why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences of the United States of America*, 10918–10925.
- Buchsbaum, D., Gopnik, A., Griffiths, T. L., & Shafto, P. (2011). Children's imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition*, *120*(3), 331–340. <http://doi.org/10.1016/j.cognition.2010.12.001>
- Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, *13*(4), 148–153.
- Frank, M. C., & Goodman, N. D. (2014). Inferring word meanings by assuming that speakers are informative. *Cognitive Psychology*, *75*, 80–96.
- Gelman, S. A. (2009). Learning from others: Children's construction of concepts. *Annual Review of Psychology*, *60*, 115–140.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*.
- Kline, M. A. (2014). How to learn about teaching: An evolutionary framework for the study of teaching behavior in humans and other animals. *Behavioral and Brain Sciences*, 1–70.
- Lambert A (2011) *The Gates of Hell: Sir John Franklin's Tragic Quest for the North West Passage* (Yale Univ Press, New Haven, CT).
- Shafto, P., Goodman, N. D., & Frank, M. C. (2012). Learning From Others: The Consequences of Psychological Reasoning for Human Learning. *Perspectives on Psychological Science*, *7*(4), 341–351.
- Tomasello, M. (1999). *The cultural origins of human cognition*. Cambridge, MA: Harvard University Press.
- Tsividis, P.A., Pouncy, T., Xu, J.L., Tenenbaum, J.B., & Gershman, S.J. (2017). Human learning in Atari. *AAAI Spring Symposium on Science of Intelligence: Computational Principles of Natural and Artificial Intelligence*.