Machine learning culture and the value of old souls

BY DOMINIK DEFFNER

Technology evolves. While most periods of human (pre-)history were characterised by remarkable technological stasis, with both the Oldowan and Acheulian stone tool technologies spanning over a million years, recent decades have witnessed an explosion in technological innovations and ever-accelerating rates of technological change. Science itself is both a driver and consequence of technology. Current innovations not only determine which scientific questions we ask and deem answerable, they also inspire our collective thinking in a way that shapes conceptual frameworks and the direction of entire fields. In line with this, one leading theme of the second annual Early-carer Social Learning Researchers (ESLR) workshop in St Andrews was the potential use of machine learning techniques, AI and "big data" for social learning research. In this short piece, I will sketch how social learning research already benefits from these recent technological advances, what the potential avenues are for the future, and what could be shortcomings of jumping on the technological bandwagon.

Firstly, and unsurprisingly, there are existing scientific questions that suddenly become answerable as soon as new technology arrives. Bayesian inference, for example, had long been regarded as philosophically sound but practically useless for scientific purposes, until dramatically increasing computational power and new Markov Chain Monte Carlo algorithms revitalized the interest in Bayesian methods in the early 90s and lead to the boom we see today. As discussed at the workshop, social learning research has already greatly benefitted from new statistical modelling techniques, such as experience-weighted attraction models (EWA) and network-based diffusion analysis (NBDA), that allow researchers to study the transmission pathways and strategic learning choices underlying cultural evolution. Social learning research deals with phenomena that require integrated and dynamic explanations at different levels of analysis, ranging from (intra-)individual cognitive processes (strategic learning) to large-scale population patterns. As such, our field is also particularly likely to benefit from future advances in machine learning and statistics.

Modern technology can not only improve the methods we use to study social learning; it can also provide us with novel phenomena and new data sources. In a recent study, Miu, Gulley, Laland and Rendell (2018) used data from 14 years of online programming competitions to explore the dynamics of cumulative culture in a system exhibiting real-world complexity. They report that, within each contest population, performance increased over time through a combination of many gradual modifications and rarer innovations. Such longitudinal online data sources and data from social networks provide the opportunity to study cultural transmission on a much larger scale and with higher resolution than was previously possible. Combining these new data sources with novel modelling techniques will most likely advance our field in the future. However, it remains to be explored whether these potentially unrepresentative study systems really tap into the same cognitive machinery that underpins cultural learning in other domains.

Probably most profoundly, technological advances can influence how researchers model and, therefore, conceptualize social learning processes. Traditionally, social and individual learning have been regarded as separate processes that compete for explanatory power in any given situation.

Exemplifying this is the ever-lasting debate about associative learning accounts of social learning. New computational methods will allow researchers to model learning in cognitively more realistic ways that are grounded in both modern neuroscience and evolutionary theory. Karl Friston and others, for instance, have proposed a free energy principle (FEP) as a unified brain theory that accounts for action, perception and learning. According to this view, which draws on thermodynamics, information theory, and machine learning, brains (and other models of the world) function to minimize their free energy or surprise. Learning, therefore, is the process that optimises the connection strengths in hierarchical models of the sensory input the brain receives; organisms actively generate predictions of the variables that cause their sensory input and learn by minimizing the error between their predictions and the data they receive. Using the FEP or other Bayesian views of the brain, researchers can start to model social and individual learning as stemming from the same inferential process instead of being two separate mechanisms. That way, one can start to investigate which mechanistic factors really make social learning social and which brain systems might be involved in learning from specifically social sources.

Even though social learning research can be expected to profit immensely by implementing modern AI technology, such computational approaches cannot be our one and only way of theorizing about the evolution of social learning. But why bother with algebra if we could also just simulate? First, algebraic equations talk in a way that simulations do not; they can actually be understood. Mathematical models provide direct expressions for the dynamics of a system and can provide proofs for why a system behaves the way it does, whereas simulations provide a number of examples, which we then use to infer what algebraic expressions can directly tell. Second, the typically few recursions in mathematical models are more transparent and easier to verify and communicate compared to the numerous lines of code in simulation models. Probably most importantly, the flexibility of computational approaches often tempts modelers to include every variable of interest leading to a complicated and virtually uninterpretable model of an already baffling world (see McElreath & Boyd, 2008, for more details). At the end of the day, analytical and computational approaches complement each other and can be used for slightly different purposes. While simple analytical models serve as proof of concepts and can test the soundness of our verbal reasoning (Roger's famous and not so paradoxical model is a great example), computational approaches can extend analytical findings to more complex and realistic scenarios and, therefore, make them more amenable to empirical investigations.

To sum up, keeping up with the zeitgeist and implementing modern machine learning techniques will greatly benefit social learning research in the years to come. However, it is not always necessary nor advisable to jump from pencil to paper or chalk to board straight to the keyboard. ❐

REFERENCES

Miu, E., Gulley, N., Laland, K. N., & Rendell, L. (2018). Innovation and cumulative culture through tweaks and leaps in online programming contests. Nature communications, 9.

McElreath, R., & Boyd, R. (2008). Mathematical models of social evolution: A guide for the perplexed. University of Chicago Press.

Dominik Deffner is a first-year PhD student with Richard McElreath and Anne Kandler at the Max-Planck Institute for Evolutionary Anthropology in Leipzig, Germany. After receiving a BSc in Psychology and a BA in Cultural Anthropology from Marburg University, he moved to Scotland and worked with Kevin Laland for his Masters degree. His current research focusses on the integrated role of culture and life history in human adaptation.