



Trends in Cognitive Sciences

Figure 1. Exploring the Rich Dataset of Hartshorne, Tenenbaum, and Pinker [4] Using Item Response Theory (IRT) Uncovers Many Rich Facets of the Data. (A) A histogram of latent ability scores based on four-parameter IRT model fits. Color shows the language status of the participants. (B) Latent ability scores for monolinguals, plotted by participant age and formal education. Dot size represents the log of the number of observations; curves show LOESS functions.

In the data of HTP, even within the monolingual speakers, a higher level of education was related to faster development and a greater eventual level of latent ability (Figure 1B). Understanding the interactions between morphosyntactic testing and formal education is another important theoretical challenge for future work in this area.

In sum, through creative methodology and major effort, HTP have gathered a dataset that will keep theoreticians and modelers working for years to come. Understanding how these data reflect on the critical period hypothesis will require further formal work to develop models that link observed behavior to theoretical constructs. Psychometric theory may be a useful starting point for this development; other new tools will almost certainly be needed as well. One clear lesson from the work of HTP, however, is that bigger data provide an opportunity to construct better theories.

Acknowledgments

Thank you to Joshua Hartshorne, Joshua Tenenbaum, and members of the Stanford Language and Cognition Laboratory for valuable feedback. Data for Hartshorne, Tenenbaum, and Pinker (2018) are available at <https://osf.io/pyb8s/>; the replication code for the analyses reported here is available at https://github.com/mcfrank/critical_periods.

¹Department of Psychology, Stanford University, Stanford, CA 94305, USA

*Correspondence: mcfrank@stanford.edu (M.C. Frank).
<https://doi.org/10.1016/j.tics.2018.05.011>

References

- Johnson, J.S. and Newport, E.L. (1989) Critical period effects in second language learning: the influence of maturational state on the acquisition of English as a second language. *Cognit. Psychol.* 21, 60–99
- Werker, J.F. and Hensch, T.K. (2015) Critical periods in speech perception: new directions. *Annu. Rev. Psychol.* 66, 173–196
- Hakuta, K. et al. (2003) Critical evidence: a test of the critical-period hypothesis for second-language acquisition. *Psychol. Sci.* 14, 31–38
- Hartshorne, J.K. et al. (2018) A critical period for second language acquisition: evidence from 2/3 million English speakers. *Cognition* 177, 263–277
- Yarkoni, T. and Westfall, J. (2017) Choosing prediction over explanation in psychology: lessons from machine learning. *Perspect. Psychol. Sci.* 12, 1100–1122

- Tauber, S. et al. (2017) Bayesian models of cognition revisited: setting optimality aside and letting data drive psychological theory. *Psychol. Rev.* 124, 410
- Borsboom, D. (2006) The attack of the psychometricians. *Psychometrika* 71, 425
- Chalmers, R.P. (2012) mirt: a multidimensional item response theory package for the R environment. *J. Stat. Softw.* 48, 1–29
- McMurray, B. (2007) Defusing the childhood vocabulary explosion. *Science* 317, 631–631
- Paikeday, T.M. (1985) *The Native Speaker Is Dead! An Informal Discussion of a Linguistic Myth with Noam Chomsky and Other Linguists, Philosophers, Psychologists, and Lexicographers*, PPI

Spotlight

Neurons That Update Representations of the Future

Peggy Seriès^{1,*}

A recent article shows that the brain automatically estimates the probabilities of possible future

actions before it has even received all the information necessary to decide what to do next.

'The future depends on what we do in the present', as Gandhi said. Each action we take defines and constrains our possible future. This is true for political action, but also for everyday movements. If we are running and our leg is fully stretched, whatever we do next, the set of possible motions is constrained: they need to involve a flexion of the knee. An efficient prediction system should take this into account dynamically. It should continuously update a representation of the possible future actions before they happen, along with the associated uncertainty. Do brains do this? Glaser *et al.* [1] present evidence that indeed they do.

Predicting the future is often thought to be what brains have evolved to do. An efficient way to make predictions involving uncertainty is to represent knowledge with probability distributions and to acquire new knowledge by following the rules of probabilistic inference. It has thus become popular to think that the brain performs (an approximation of) probabilistic (a.k.a. 'Bayesian') reasoning. This idea has already had a profound impact in cognitive science and is consistent with a large body of work in human and animal behaviour [2,3].

However, the details of this hypothesis are unclear. In particular, there remains a large gap between the behavioural studies supporting the Bayesian hypothesis and uncovering the underlying neural substrate. Is the brain truly representing probability distributions? Where would those distributions live? How would they be represented? How flexibly are those representations updated, in particular when they should be dynamically changing? Glaser *et al.* [1] shed light on such issues.

In their experiment, three monkeys were trained to reach for four targets on each trial, one after the other, using their hand. On each trial, the position of the next target was conditioned on the current hand position: targets were more likely to appear approximately opposite the current hand position, with a slight clockwise bias. Additionally, the farther the hand position was from the centre of the workspace, the more likely the upcoming target was to be in the opposite direction. The authors first measured whether monkeys learned these probability distributions by looking at their behavioural performance. They found that indeed their initial reaches were biased by expectations about the target and their uncertainty.

The monkeys were implanted with electrode arrays in the primary motor cortex (M1) and dorsal premotor cortex (PMd). Neurons in the PMd are known to be active during the preparation for the reach and also during the reach itself. They are broadly tuned, responding best to one direction of reach. Glaser and colleagues find that a small population of PMd neurons, which they call 'potential response' (PR) neurons, are modulated before target presentation, based on the anticipated possible target locations. Moreover, the preferred directions of these neurons were distributed approximately in proportion to how likely upcoming movement directions were. The authors also could decode the movement that the PR neural population was planning in the 100 ms before target presentation. They find that the planned reaches decoded before target onset were usually approximately to the position opposite to the current hand position. This representation contained information about the uncertainty of the future positions, supporting the idea that it is really a probability distribution that is represented on single reaches, across the population of neurons. Such representation was not found in M1.

This line of work is important as it helps bridge the gap between neural representations and probabilistic computations. It also raises a number of questions. If PMd neurons represent the probability distribution of upcoming possible reaches, is this encoded as a continuous function or as samples of this distribution? At a theoretical level, there has been a longstanding debate about whether the brain uses probabilistic population codes (PPCs) [2] versus sampling codes, where only a few hypotheses would be represented with frequencies proportional to their probabilities, either across time or across the population of neurons [3,4]. At present Glaser *et al.*'s data seem compatible with both explanations. In theory, PPCs and sampling make different predictions, particularly about how the representation of uncertainty depends on the number of neurons involved in representation or how it would evolve in time. However, teasing them apart is proving difficult [5]. By recording more neurons, systematically decoding the neural activity using different codes, and comparing the predictions to behavioural performance, extensions of this study could possibly start answering such questions.

Other questions could be asked as well. How and where is this 'prior' distribution about likely future motion directions integrated with the information provided when the target appears (the 'likelihood') and 'read out' to lead to the actual decision? This would address how Bayes' rule is implemented, a question that has started to be investigated in various other domains [6]. Of particular importance will be to understand the nature of the necessary approximations used in these computations and how they can explain suboptimal behaviour [7].

This work might also pave the way to new neural theories of how the brain can build complex representations on fast time-scales in more cognitive domains. Similar

problems exist in speech processing; for example, where, when hearing streams of words, our brain needs to represent the syntactic and semantic structure of the sentence on the fly, anticipating future words. Cognitive flexibility may also be related to how fluidly the brain can represent likely future actions, contexts, or thoughts.

Ultimately, looking at individual differences in the flexibility of this representation could have implications in the clinical domain. It is often thought that mental disorder, in particular autism and schizophrenia, could be described as a failure mode of the predictive system [8,9], related either to the brain using wrong or incompletely learned beliefs or to failures in how neural networks implement approximate ‘Bayesian’ computations [10]. The neural substrate underlying this prediction system and the factors involved in its fluidity or its possible impairments, as well as the precise nature of the ‘code’, are still largely to be discovered.

¹University of Edinburgh, Informatics, Edinburgh EH8 9AB, UK

*Correspondence: pseries@inf.ed.ac.uk (P. Seriès),
<https://doi.org/10.1016/j.tics.2018.05.012>

References

1. Glaser, J.I. *et al.* (2018) Population coding of conditional probability distributions in dorsal premotor cortex. *Nat. Commun.* 9, 1788
2. Pouget, A. *et al.* (2013) Probabilistic brains: knowns and unknowns. *Nat. Neurosci.* 16, 1170–1178
3. Fiser, J. *et al.* (2010) Statistically optimal perception and learning: from behavior to neural representations. *Trends Cogn. Sci.* 14, 119–130
4. Sanborn, A.N. and Chater, N. (2016) Bayesian brains without probabilities. *Trends Cogn. Sci.* 20, 12
5. Ma, W. and Jazayeri, M. (2014) Neural coding of uncertainty and probability. *Annu. Rev. Neurosci.* 37, 205–220
6. Ting, C. *et al.* (2015) Neural mechanisms for integrating prior knowledge and likelihood in value-based probabilistic inference. *J. Neurosci.* 35, 1792–1805
7. Rahnev, D. and Denison, R.N. (2018) Suboptimality in perceptual decision-making. *Behav. Brain Sci.* Published online February 27, 2018. <http://dx.doi.org/10.1017/S0140525X18000936>
8. Palmer, C.J. *et al.* (2017) Bayesian approaches to autism: towards volatility, action, and behavior. *J. Psychol. Bull.* 143, 521–542
9. Valton, V. *et al.* (2017) Comprehensive review: computational modelling of schizophrenia. *Neurosci. Biobehav. Rev.* 83, 631–646
10. Huys, Q. *et al.* (2015) Decision-theoretic psychiatry. *Clin. Psychol. Sci.* 3, 400–421

Spotlight Shining Light on Social Learning Circuits

Steve W.C. Chang^{1,2,3,*} and Olga Dal Monte¹

Learning from others powerfully shapes our lives, yet the circuit-specific mechanisms underlying social learning in the brain remain unclear. A recent study in mice provides evidence that direct neuronal projections from the anterior cingulate cortex (ACC) to the basolateral amygdala (BLA) play a critical role in observational fear learning.

Social learning is arguably one of the most important and powerful learning strategies available to us. We learn about the world from our parents, family members, and close others, a process that starts at birth and continues throughout our lifespan. Our schools and universities are designed for social learning; we gain new skills and shape our worldview from knowledge gleaned from our teachers, friends, and peers. This universal human characteristic of social learning is also found in several other animals.

So, how do we learn from others? A large amount of research – encompassing behavioral and neuroscientific work done in humans and other animals – has been asking this question for long enough [1,2] that the breadth of social learning research has become too wide and complex to be easily summarized in a few sentences. We will therefore narrow our scope to focus on

one particular aspect of social learning: observational learning. Observational learning, or vicarious learning, is the learning that occurs through the observation of a model. It has been suggested that vicarious reinforcement plays a critical role in observational learning. Observing the reinforcement (either positive or negative in valence) of certain actions and outcomes in another individual results in changes in our own behavior [3,4]. Consistent with the role of reinforcement in social learning, vicariously rewarding events (e.g., seeing another individual consuming a reward) drive neuronal activity in brain regions involved in value-guided decision-making, including the ACC and the BLA [5,6]. However, the central question of how such brain regions acquire reinforcing information during vicarious learning remains unanswered.

A recent study by Allsop and colleagues [7] sheds light on this central question. This *tour de force* investigation spanning behavioral, electrophysiological, and optogenetic techniques helps delineate the circuit-specific mechanisms of observational learning. Capitalizing on the observational fear conditioning paradigm in mice (Figure 1A) – incidentally, also an effective method for studying social learning in humans [8] – the researchers uncover a direction-specific interaction between the ACC and the BLA. In the study, an observer mouse vicariously experiences a shock delivered to a demonstrator mouse paired with a predictive sensory cue. To better understand the behavioral contingencies necessary for observational fear learning, several distinct types of observer mice were tested to confirm that (i) prior experience of the shock, (ii) observation of the demonstrator mouse receiving a shock, and (iii) the vicarious shock being paired with a cue are all necessary components for observational fear learning. Next, based on electrophysiological recordings, the authors found that both ACC and BLA