THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Neurons That Update Representations of the Future

Citation for published version:

Seriès, P 2018, 'Neurons That Update Representations of the Future', *Trends in Cognitive Sciences*.
<https://doi.org/10.1016/j.tics.2018.05.012>

Digital Object Identifier (DOI):

[10.1016/j.tics.2018.05.012](https://doi.org/10.1016/j.tics.2018.05.012)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Trends in Cognitive Sciences

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Neurons that update representations of the future

Peggy Seriès^{1,*}

University of Edinburgh, Informatics, Edinburgh Eh8 9AB.

pseries@inf.ed.ac.uk

Abstract: 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.

Keywords: Bayesian Brain, predictive coding, neural encoding, motor cortex.

“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 very unclear. In particular, there is still 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 of the current hand position, with a slight clockwise bias. Additionally, the farther the hand position was from the center 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 primary motor cortex (M1) and dorsal premotor cortex (PMd). Neurons in 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 prior to 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 prior to target presentation. They find that the planned reaches decoded prior to target onset were usually approximately to the position opposite of 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 long-standing debate about whether the brain uses probabilistic population codes (PPC) [2] vs sampling codes, where only a few hypotheses would be represented with frequency proportional to their probability, either across time or across the population of neurons [3,4]. At present Glaser *et al*’s data seems compatible with both explanations. In theory, PPC 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 reach direction (the posterior)? 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 sub-optimal behaviour [7].

This work might also pave the way to new neural theories of how the brain can build complex representations on fast timescales in more cognitive domains. Similar problems exist in speech processing for e.g., 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. Indeed, 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], either related 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, 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.

References

- [1] Glaser, J. I. *et al.* (2018), Population coding of conditional probability distributions in dorsal premotor cortex. *Nat. Commun.*, 3;9(1).
- [2] Pouget, A. *et al* (2013), Probabilistic brains: knowns and unknowns. *Nat. Neurosci.* 16(9):1170-8.
- [3] Fiser, J *et al* (2010), Statistically optimal perception and learning: from behavior to neural representations. *TICS*, 14:3, 119-130.
- [4] Sanborn A. N. and Chater, N. (2016), Bayesian brains without probabilities. *TICS*, 20:12.
- [5] Ma, W. and Jazayeri, M. (2014), Neural coding of uncertainty and probability. *Annu Rev. Neurosci.* 37:205–20.
- [6] Ting, C. *et al* (2015), Neural mechanisms for integrating prior knowledge and likelihood in value-based probabilistic inference. *J Neurosci*, 35(4), 1792-1805.
- [7] Rahnev D. and Denison R.N. (2018), Suboptimality in perceptual decision-making. *Behav Brain Sci*, 27:1-107.
- [8] Palmer C. J. *et al* (2017) Bayesian approaches to autism: Towards volatility, action, and behavior. *J. Psychol. Bull.*, 143(5):521-542.
- [9] Valton V. *et al*, Comprehensive review (2017): Computational modelling of schizophrenia. *Neurosci. Biobehav. Rev.* 83:631-646.
- [10] Huys Q. *et al* (2015), Decision-Theoretic Psychiatry. *Clin. Psychol. Sci.*, 3(3):400-421.