Probabilistic representations should be defined by their computational role: reply to Rahnev et al. 2021

Raphael T. Gerraty^{*a,b}, Samuel Lippl^{*a}, John Morrison^c, and Nikolaus Kriegeskorte^{a,d,e,f}

^{*}equal contribution

^aZuckerman Mind Brain Behavior Institute, Columbia University

^bCenter for Science and Society, Columbia University

^cDepartment of Philosophy, Barnard College, Columbia University

^dDepartment of Psychology, Columbia University

^eDepartment of Neuroscience, Columbia University

^fDepartment of Electrical Engineering, Columbia University

Abstract: *Probabilistic representations promise to shed light on how humans and other organisms cope with uncertainty, but debates about them have been hindered by unclear definitions and a lack of testable predictions. We argue that neither camp in this GAC has offered a plausible way out of this situation. We describe problems with their attempts to define probabilistic representations and sketch an alternative view, in which the defining feature of a probabilistic representation is how it is used.*

We agree with the authors (Rahnev et al. 2021) about the need for a clear definition of probabilistic representation that lets us test claims that specific brain regions contain them. Such a definition would enable us to distinguish computations involving probabilistic representations from computations relying on non-probabilistic heuristics for handling uncertainty. But there is an inherent tension. Defined too loosely, the concept of probabilistic representation encompasses every neural code since we can always infer its content with a Bayesian decoder that expresses it as a probability distribution. Defined too strictly, the concept excludes paradigmatic examples of probabilistic representation. This tension is clear in the accounts offered by both teams of authors. We will first describe our disagreements and then lay out our own alternative view. In brief, our view is that a definition of probabilistic representations should focus on how representations of uncertainty are *used* in downstream computations.

To avoid giving a definition that's too loose, the pro team, Denison and Jehee, require probabilistic representations to "reflect the true amount of information available about the variable of interest" (Rahnev et al. 2021, p. 8). For example, a probabilistic representation of orientation might be based on an accurate model defining (1) how the sensory data was generated by the stimulus (likelihood) and (2) the frequency of stimuli with different orientations in the environment (prior). How accurate does the model need to be? They elsewhere clarify that it needn't be "fully accurate" and can be "incomplete" (Rahnev et al. 2021, p. 6). But that doesn't tell us where we should draw the line between probabilistic and non-probabilistic representations. More importantly, a completely wrong model can still form the basis of probabilistic inferences. In practice, this is often the case and we use probabilistic representations precisely because our model is inaccurate or incomplete. There is a useful analogy with logical inference. Just as we can use logical representations to infer a false

conclusion from false premises, we can use probabilistic representations to do the same. Requiring probabilistic representations to be true or accurate imposes an arbitrary constraint.

Denison and Jehee also assume that more than two moments of a probabilistic representation must be accessible to downstream areas. The underlying thought seems to be that we can't distinguish representations of simple distributions, such as Gaussians, from heuristic representations of uncertainty, and we should therefore restrict probabilistic representations to those where more than two moments define the density (which is not the case for Gaussian distributions). But excluding the paradigmatic example of the Gaussian distribution is too high a price. There are many important examples of latents whose (exact or approximate) probabilistic encoding could take the form of a Gaussian. Once again, the analogy of logical representations is helpful. It would be arbitrary to deny that the syllogism "p, if p then q, therefore q" is a logical inference on the grounds that logical inferences must involve more than three propositions.

The con team, Block and Rahnev, distinguishes between 'weak' and 'strong' definitions. According to weak definitions, any effect of information on behavior beyond a point estimate involves a probabilistic representation. As they point out, this would imply that probabilistic representations are ubiquitous. According to a strong definition, only representations that conform to the Kolmogorov axioms are probabilistic. This is an odd definition, because 1) there are other axiomatizations of probability, 2) it is not clear what it would mean for the brain to follow the Kolmogorov axioms (as Rahnev and Block point out, as well), and 3) it seems unlikely that a biological system evolved to ground its perception in the abstract formalism of set theory. This definition therefore appears arbitrary, lacks clarity, and is unlikely to apply to the brain. It would be better to look for definitions in the ample space between Block and Rahnev's weak and strong definitions.

We will end by briefly sketching our own view. We think that the most important criterion for a representation of a probability distribution is whether this is required by the *best explanation for how it is used in downstream computations*. An important and underappreciated problem is that, for any downstream computation, there is almost always some set of assumptions that would enable us to interpret it as a computation over probability distributions. For example, computing the mean of a set of observations is indistinguishable from performing MAP estimation with a Gaussian likelihood and a uniform prior. We call this the indeterminacy of the mental model (or "model indeterminacy"). Without a solution to model indeterminacy, we cannot distinguish probabilistic computations from simpler, non-probabilistic heuristics. One solution is to leave it up to the experimenter to choose the relevant assumptions (Block and Rahnev's "experimenter's perspective"), but we would rather ground the assumptions in the computations themselves.

Our proposal is to focus on tasks involving *probabilistic transfer*: the generalization of a computation over a probability distribution to a new setting. For some computations, such as marginalization and change-of-variables, transfer is hard to explain without invoking computations over a probability distribution. For example, research using the framework of Bayesian transfer has shown that subjects are able to take their motor uncertainty into account

on a new task even if they have not been trained on this task at different levels of uncertainty (Maloney 2002; Trommershäuser et al. 2003; Maloney and Mamassian 2009). While it is perfectly clear how probabilistic representations would enable subjects to do this, it is hard to explain using a non-probabilistic heuristic. Bayesian transfer is an example of an experimental observation that we believe is best explained by probabilistic representations. We propose broadening the notion of Bayesian transfer to include more general transfer of probabilistic computations and defining probabilistic representations as representations that are capable of such probabilistic transfer. Seen through this lens, the search for probabilistic representations should focus on finding representations of uncertainty and demonstrating that probabilistic computations over these representations generalize across tasks.

References

- Maloney, L. (2002). Statistical decision theory and biological vision. In *Perception and the physical world: Psychological and philosophical issues in perception* (pp. 145-189). Wiley.
- Maloney, L. T., & Mamassian, P. (2009). Bayesian decision theory as a model of human visual perception: Testing Bayesian transfer. *Visual neuroscience*, 26(1), 147-155.
- Rahnev, D., Block, N., Denison, R. N., & Jehee, J. (2021). Is perception probabilistic? Clarifying the definitions. <https://doi.org/10.31234/osf.io/f8v5r>
- Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003). Statistical decision theory and the selection of rapid, goal-directed movements. *JOSA A*, 20(7), 1419-1433.