Journal Club

Editor's Note: These short, critical reviews of recent papers in the *Journal*, written exclusively by graduate students or postdoctoral fellows, are intended to summarize the important findings of the paper and provide additional insight and commentary. For more information on the format and purpose of the Journal Club, please see http://www.jneurosci.org/misc/ifa_features.shtml.

Time-Dependent Influence of Prior Probability: A Problem for the Drift-Diffusion Model?

Jill X. O'Reilly

FMRIB Centre, Nuffield Department of Clinical Neurosciences, Oxford University, Oxford OX3 9DU, United Kingdom
Review of Hanks et al.

How does the brain use prior expectations to optimize the processing of incoming sensory information? This important question is a specific instance of the problem addressed by Bayes theorem: what is the optimal way to combine newly acquired information (evidence) with that which is already believed to be true (the prior)?

Consider a classic and well studied model for sensory evidence accumulation: the moving-dots task (Britten et al., 1992). In this task, the subject sees a field of (mainly) randomly moving dots, within which a small number of dots move coherently toward the left or right; the subject's task is to indicate the direction of coherent motion. If in fact the direction of coherent motion is rightward 80% of the time and leftward only 20%, this prior information could be used to facilitate performance—either to decrease reaction times (because only a little evidence is required to confirm the prior expectation that dots are moving rightwards) or to increase accuracy (because if the subject is not sure which is the correct direction, he can increase his chances of guessing correctly by guessing in favor of the prior).

When integrating prior expectations and sensory evidence, the optimal solution according to Bayesian theory is to weight each according to its reliability—so if sensory evidence is weak, an observer should rely more upon the prior and vice versa. To achieve the optimum weighting, an estimate of reliability is needed for both sensory evidence and prior. This is particularly problematic in the case of sensory evidence, since evidence quality may vary unpredictably—in the moving-dots task, the experimenter controls the quality of sensory evidence by manipulating the proportion of coherently moving dots, but from the subject's point of view there is no way to know beforehand whether each trial will be high or low coherence.

In a recent paper, Hanks et al. (2011) offered the interesting and parsimonious suggestion that the time taken to arrive at a decision may in fact be used as a measure of sensory evidence quality. Their reasoning follows from the dominant mathematical model of evidence accumulation: the drift-diffusion model [see Hanks et al. (2011), their Fig. 3]. In the drift-diffusion model, the accumulated evidence in favor of two opposing options (such as leftward or rightward motion) is represented by a decision variable, $v(t)$. At the start of evidence accumulation, the value of $v(t)$ is zero, and a decision in favor of leftward or rightward motion will be made when $v(t)$ reaches some bound, $-A$ and $+A$ respectively. In each time step δt , $v(t)$ increases slightly if evidence during δt is in favor of rightward motion, and decreases slightly

if the evidence is in favor of leftward motion. Thus, the value of $v(t)$ drifts over time toward one of the decision bounds, with the rate of drift determined by the weight of evidence in favor of rightward versus leftward motion. High-coherence (strong-evidence) stimuli cause a rapid drift toward one or other bound and therefore lead to short decision times. In contrast, if there is little evidence in favor of one or other motion direction, the accumulator may take a very long time to reach a bound. Consequently, if the evidence accumulation process continues for a long time, this in itself indicates that the quality of sensory evidence is low.

The suggestion that long decision times are associated with low accuracy is borne out by behavioral evidence: for moving-dot stimuli with equal priors, Hanks et al. (2011) plotted the probability of a correct decision as a function of decision time (their Fig. 4). The probability of a correct decision decreased as decision time increased. They then proposed a Bayes-optimal model in which the prior is manifested as a time-dependent bias signal, $\lambda(t)$, that is proportional to the ratio of the logarithm of the odds of being correct given the before the logarithm of the posterior odds of being correct based on the evidence, given the elapsed time t [see Hanks et al. (2011), their Eq. 2]. This means that the longer the evidence accumulation process goes on, the greater the effect of the prior on the eventual decision.

Two lines of evidence support their model. First, the model incorporating the time-dependent prior was a better predic-

Received June 26, 2011; revised July 26, 2011; accepted July 26, 2011.

This research was funded by UK Medical Research Council Fellowship G0802459 to J.O.R. Thanks to Matthew Rushworth, Chris Summerfield, Nils Kolling, and Max Kleiman-Weiner for discussion in our own journal club; and to Rogier Mars for comments on the manuscript.

Correspondence should be addressed to Jill O'Reilly, FMRIB Centre, Nuffield Department of Clinical Neurosciences, John Radcliffe Hospital, Headington, Oxford OX3 9DU, UK. E-mail: joreilly@fmrrib.ox.ac.uk.

DOI:10.1523/JNEUROSCI.3253-11.2011

Copyright © 2011 the authors 0270-6474/11/3113697-02\$15.00/0

tor of behavioral performance than either a baseline-shift model (in which the prior has a static influence), or a combination of a baseline-shift and an urgency-signal model, which requires less sensory evidence strength at longer decision times, but in which there is no interaction between decision time and influence of the prior. Second, the activity of lateral intraparietal area (LIP) neurons, the putative neural correlate of evidence accumulation, ramped up toward the prior more rapidly at longer decision times; the rate of build-up, but not prestimulus baseline shifts in activity, predicted behavioral performance on a trial-to-trial basis.

How can the time-dependent bias be incorporated into current models of sensory evidence accumulation?

The use of decision time as a measure of sensory evidence quality is appealing because it gives a mechanistic explanation of how evidence quality can be assessed without recourse to an external, meta-cognitive measure of certainty (although, interestingly, monkeys do seem to be able to use an estimate of their own certainty, to “opt out” of trials in which they are unsure of the correct answer) (Kiani and Shadlen 2009). However, if the influence of prior expectations varies over time, where does this time-dependent signal come from?

The drift-diffusion model, and its putative neural instantiation, are appealingly parsimonious, but cannot account for the time-varying influence of the prior. Current understanding is that neurons in visual motion area MT are sensitive to the moment-to-moment evidence in favor of one direction of motion (Britten et al. 1992). This evidence is fed into saccade-selective neurons in LIP, which accumulate evidence in favor of one or other direction of motion (and hence, one or other saccadic response), ramping up their activity if the saccadic target in their receptive field represents the correct response. The activity of LIP neurons closely resembles the behavior of the drift-diffusion model.

Previously, the incorporation of prior probabilities has been modeled as a shift in one of the bounds of the diffusion process, or in the initial or baseline value of $v(t)$ (Simen et al. 2009; Forstmann et al. 2010). Neurally, a baseline shift could be

instantiated as a change in the baseline firing rate of the neuron. However, baseline-shift models imply a constant effect of the prior, not a decision-time dependent effect as observed by Hanks et al. (2011). More fundamentally, baseline-shift models cannot adjust the influence of the prior based on sensory evidence quality, because sensory evidence quality is not known at the start of the trial when the baseline shift would be implemented; therefore, in the context of unpredictable sensory evidence quality, any such model is necessarily nonoptimal.

In contrast, Hanks et al. (2011) modeled the effect of the prior as a temporally dependent bias signal that is added to the drift-diffusion process. This is the Bayes-optimal solution and fits the data, but a time-dependent bias signal cannot be modeled in terms of the generative processes for a drift-diffusion model (simply, the integration over time of the moment-to-moment evidence in favor of each option). This leaves two possibilities. Either the time-dependent bias must be understood in terms of some dynamics in the nervous system that are not captured by the drift-diffusion model but could perhaps be captured by a more complex neural network model, or there is a time-dependent signal being fed into LIP from elsewhere in the brain.

If future work is to address the problem, three points may be germane.

First, using the time-dependent accuracy function to compute a time-dependent bias signal is the Bayes-optimal solution, in the sense that sensory evidence and prior are weighted according to their estimated reliabilities. However, the amount of time spent in evidence accumulation for ecologically optimal behavior must also depend on the relative cost of making an error as opposed to the benefit of making a correct decision. Therefore time-dependent urgency signals might also be expected to play a role in decision-related ramping activity (Cisek et al., 2009). Hanks et al. (2011) demonstrate that the time-dependent accuracy function varies when the speed-accuracy trade-off is manipulated; to what extent are urgency and time-dependent accuracy linked?

Second, time-dependent variation in LIP activity has been observed outside the context of sensory evidence accumulation. Janssen and Shadlen (2005) manip-

ulated the probability that a saccadic target would appear as a function of time and showed the LIP activity tracked this temporally-dependent probability, even when the hazard function was bimodal, so that activity rose to a peak, fell back, and rose to a second peak. Any model for a time-dependent bias signal should be sufficiently flexible to explain these results as well as the simpler monotonic effect reported by Hanks et al. (2011).

Finally, it may be possible to separate the duration of evidence accumulation from uncertainty. In the moving-dots task, there is no way to estimate sensory evidence quality other than by the rate of evidence accumulation, but this is not necessarily true of all tasks. For example, Yang and Shadlen (2007) showed that LIP neurons accumulated evidence in favor of a leftward or rightward saccade when a series of shapes, each indicating probability that the left or right target will be rewarded, are presented sequentially. If each shape represented a probability distribution, some with more variance than others, then sensory evidence quality could be dissociated from the rate of change of sensory evidence.

References

- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12:4745–4765.
- Cisek P, Puskas GA, El-Murr S (2009) Decisions in changing conditions: the urgency-gating model. *J Neurosci* 29:11560–11571.
- Forstmann BU, Brown S, Dutilh G, Neumann J, Wagenmakers EJ (2010) The neural substrate of prior information in perceptual decision making: a model-based analysis. *Front Hum Neurosci* 4:40.
- Hanks TD, Mazurek ME, Kiani R, Hopp E, Shadlen MN (2011) Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J Neurosci* 31:6339–6352.
- Janssen P, Shadlen MN (2005) A representation of the hazard rate of elapsed time in macaque area LIP. *Nat Neurosci* 8:234–241.
- Kiani R, Shadlen MN (2009) Representation of confidence associated with a decision by neurons in the parietal cortex. *Science* 324:759–764.
- Simen P, Contreras D, Buck C, Hu P, Holmes P, Cohen JD (2009) Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. *J Exp Psychol Hum Percept Perform* 35:1865–1897.
- Yang T, Shadlen MN (2007) Probabilistic reasoning by neurons. *Nature* 447:1075–1080.