

Prior and prejudice

Emilio Salinas

To best interpret new sensory information, populations of sensory neurons must represent the lessons of past experience. How do they do this? The same solution to this problem is now reported in two very different sensory systems, providing a classic example of computational convergence.

It is particularly incumbent on those who never change their opinion, to be secure of judging properly at first. —Jane Austen, Pride and Prejudice

A young woman rebuffs a lad. She probably does not like him—or perhaps she likes him too much. Jane Austen exploits this ambiguity between interpreting words and actions according to past experiences—prejudice—or to their literal meaning. But the tug-of-war between expectations and evidence is a fundamental problem that our brain encounters at all levels, from social interactions to the most basic perceptual judgments. Two studies in *Nature Neuroscience* now investigate how neural circuits combine the knowledge accumulated from previous encounters with sensory scenes, technically known as a prior distribution, with new stimuli. Although they analyze very different sensory computations, the determination of visual edge orientation in primates¹ and the localization of sounds in owls², they reach an identical conclusion about how populations of neurons may adapt their response properties to incorporate knowledge about the statistics of the world, and the solution is elegant.

The problem of how to combine prior expectations and current sensory information in an optimal way is addressed through the principles of Bayesian inference, which provide a mathematical recipe for evaluating their relative importance. The generality of this problem may be illustrated by sports in which players continuously update the prior describing what the opponent is likely to do. In tennis, for example, the server can direct his serve either to the middle or to the side of the court, and typically chooses whichever is hardest

for the opponent to return. However, if the serve becomes predictable, then the returner can prepare accordingly and produce a winning shot. For the returner the key trade-off is this: if the serves are slow enough (low noise in the sensory input), then he can simply see where the ball is going and choose without any bias whether to hit a forehand or a backhand, but if the serves are fast (high noise), then he must guess and commit to a particular motion early, else he has little chance of returning the serve. The Bayesian recipe finds the best probabilistic strategy between these two extremes, one that is biased toward the prior and another that is not.

A growing body of evidence indicates that human subjects often behave in such a statistically optimal way in a wide array of perceptual and motor tasks^{3–5}, and that those probabilistic calculations may also determine fundamental properties of single neurons⁶. Many such studies have specifically shown that, in making perceptual judgments, individuals indeed take into account prior distributions, whether they arise naturally^{7,8} or are artificially imposed by the experimental design^{4,5,8,9}. How then are such prior distributions represented by neural circuits, and how are they accessed?

Consider how populations of sensory neurons encode a given stimulus feature x . Typically, neuron j becomes maximally activated when x takes a particular value x_j , and the response decreases as x differs from this preferred x_j (Fig. 1a). Neurons across the population have different preferences, and their response curves as functions of x , or ‘tuning curves’, overlap to cover the full range of x . Although this type of representation has been studied thoroughly^{6,10–15}, a lingering

question is how to relate it to a prior distribution: how do those sensory neurons take into account the fact that some values of x are more likely than others?

There are two basic ways. One is to allocate more preferred values x_j to the values of x that occur more frequently (Fig. 1b) and the other is to vary the width of the tuning curves as a function of x_j (Fig. 1c). Both strategies encode some values of x with higher variability than others (high variability meaning high spread, low precision) and also create biases (high bias meaning high systematic deviation, low accuracy) with respect to the true values (Fig. 1b,c). The studies by Girshick *et al.*¹ and Fischer and Peña² dissect two examples in which these two strategies are combined (Fig. 1d). Although the neural circuits in each case are very different, the resulting computations are extremely similar.

In the study by Girshick *et al.*¹, the goal was to explain two sets of observations indicating that there is an asymmetry in the representation of orientation (x in this case) in the visual systems of mammals. First, performance in orientation discrimination tasks is consistently better at the cardinal orientations (horizontal and vertical), a phenomenon known as the oblique effect. Second, neurophysiological measurements indicate that the preferred orientations of neurons in the primary visual cortex (V1) are not distributed uniformly; rather, the cardinals are over-represented (as in Fig. 1b). From optimality arguments, a reasonable explanation for both phenomena is that visual scenes naturally contain more edges that are oriented vertically or horizontally. These ideas were not new, but Girshick *et al.*¹ tested them rigorously.

First, the authors carefully measured the actual distribution of orientations in a large collection of photographs, thus determining the prior for orientation from natural scenes. Horizontal and vertical indeed proved to be

Emilio Salinas is in the Department of Neurobiology and Anatomy, Wake Forest School of Medicine, Winston-Salem, North Carolina, USA.
e-mail: esalinas@wakehealth.edu

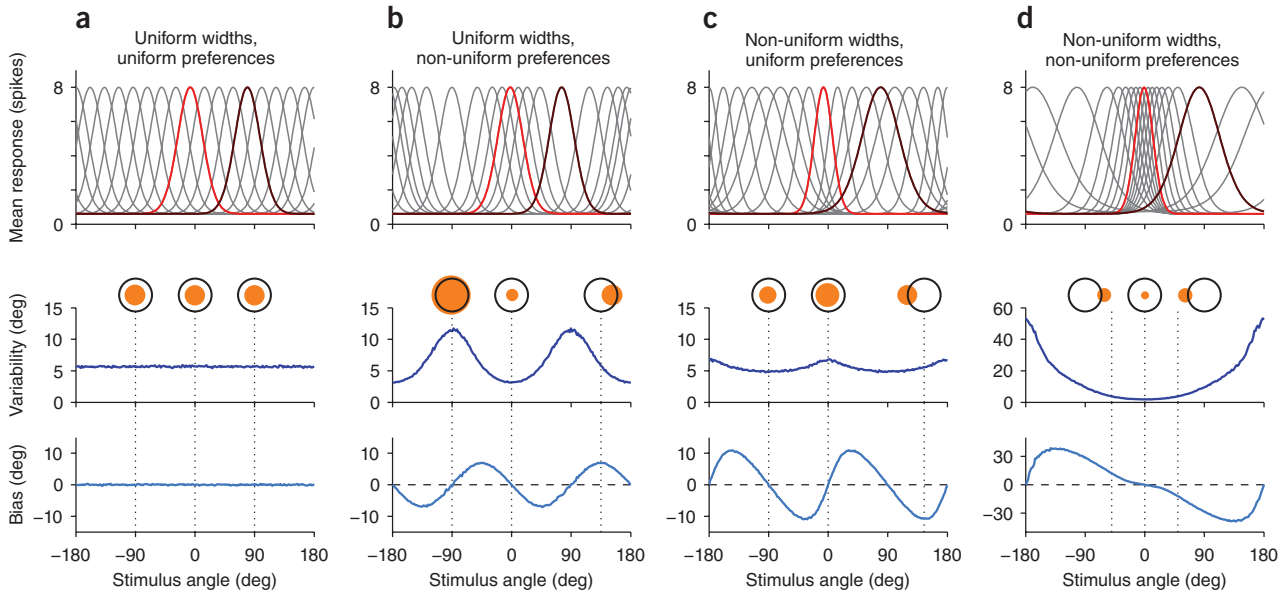


Figure 1 Encoding of a stimulus by a neuronal population. The angle on the x axis represents either the horizontal direction of a sound or the orientation of a visual stimulus (with the range of x rescaled by a factor of 2). **(a)** Top, a standard array of tuning curves with identical tuning width and uniformly distributed preferred angles. Bottom, variability (s.d.) and bias (mean) of the angle decoded over multiple trials from the responses of the population in **a**. Black circles (diameter, 10°) and orange spots depict variability (spot size) and bias (spot offset) at three stimulus angles. **(b)** Data presented as in **a**, but with variable density of preferred angles, highest at 0° and $\pm 180^\circ$ and lowest at $\pm 90^\circ$. **(c)** Data presented as in **a**, but with variable tuning-curve widths. Narrowest curves peak at 0° and $\pm 180^\circ$ and widest ones at $\pm 90^\circ$. **(d)** An array of tuning curves that approximates those found in the optic tectum of owls. All populations consisted of 50 model neurons with Poisson responses and had the same mean tuning-curve width. Encoded angles were found using the vector method.

more common. Second, they tested the discrimination capacity of people in a new task that allowed them to parametrically vary the noise or uncertainty of each oriented stimulus. This was crucial because perceived orientation varies with the amount of noise in the stimulus (that is, with its visibility): noisy stimuli appear more horizontal or more vertical. This is exactly as expected: the larger the uncertainty in the evidence, the stronger the reliance on the prior. Third, using a Bayesian model of the task applied to these psychophysical data, the authors inferred the internal prior used by the subjects and found that, on average, this prior was nearly identical to the prior obtained from natural images. This means that the neural representation of orientation in the brain is biased in a way that precisely matches the actual asymmetry found in nature. This match is a strong indicator of computational efficiency in the visual system.

Finally, Girshick *et al.*¹ simulated the responses of a population of orientation-sensitive neurons with distributions of widths and preferred orientations based on reported data from neurophysiological experiments. Their model essentially applied the dependencies shown in **Figure 1b,c** simultaneously to the same population. The objective was to investigate whether the Bayesian operations needed to combine the sensory evidence and the prior could be implemented with the types

of computations that can be easily performed by neurons (for example, weighted sums of inputs). To infer the stimulus angle encoded by the model responses at any given time, they used a simple readout scheme known as the vector method¹³. Neuron j casts a vote in favor of a vector pointing at the angle x_j such that the strength of the vote is equal to the response of the cell r_j . Then all the weighted vectors are added and the angle of the resulting vector is considered to be the angle encoded by the population. This way, the most active neurons contribute more to the final answer. With this readout or decoding method, the performance of the model population in the orientation discrimination task was indistinguishable from that of human subjects. The asymmetry in the neuronal representation fully accounted for the bias in behavior.

Fischer and Peña² studied the auditory system of owls, so the details are very different, but they adopted a remarkably similar approach and conceptual framework. Their starting point was also a notorious asymmetry in behavior. Owls can locate sounds along the horizontal plane accurately near the center of gaze, but they typically underestimate those originating further into the periphery. This central bias is substantial: on average, stimuli at $\pm 45^\circ$ elicit responses to $\pm 33^\circ$ or less². Fischer and Peña² accounted for these behavioral results with a Bayesian model with two

elements: a prior that favored sound sources near the center of gaze and a function that generated a noisy estimate of interaural time difference (ITD) for any given stimulus direction.

The ITD, which is the difference in the time of arrival of a sound to the two ears, is a crucial intermediate variable here because early auditory neurons are tuned for ITD and the horizontal angle of a sound is actually computed from it by specialized circuitry downstream. Thus, any uncertainty in ITD is carried over as uncertainty in source direction. Now, because for any sound direction the ITD that reaches the tympanic membrane is known from experimental measurements, the Bayesian model had only two free parameters: the amount of noise in the ITD estimation and the width of the prior. By adjusting these two parameters, the model accounted for the original behavioral data and for the behavior observed under two additional experimental conditions, one that altered the relationship between ITD and sound direction and another that increased the amount of noise in the owl's perception of ITD.

Next, Fischer and Peña² developed a population model describing the encoding of horizontal sound direction in the optic tectum of the owl. Again, the objective was to figure out how the neurons could implement the probabilistic operations of the Bayesian model. For this, they generated arrays of neuronal tuning

curves as functions of sound direction and compared the model responses to the behavioral data. This required two ingredients. First, they needed a read-out to infer the source angle encoded by the population's responses, and they used the very same vector method as Girshick *et al.*² Furthermore, they obtained an important theoretical result describing the mathematical conditions under which the vector method is equivalent to the Bayesian model². Second, to fit the behavioral data, they had to adjust the distribution of preferred locations across the population. Their resulting model is qualitatively similar to that shown in **Figure 1c**, except that the owl's tuning curves are not perfectly symmetric. Finally, they showed that the distribution of preferred locations in the best-fitting model matched the actual distribution measured experimentally, providing further proof of consistency between the behavioral, computational and neurophysiological results.

Both these studies create convincing links between psychophysical performance and neuronal representations using the formalism of Bayesian inference. There is a noteworthy difference between them, though. For edge orientation, the prior corresponds exactly to the frequencies with which horizontal, vertical or other orientations are encountered in a visual scene. Thus, the statistics of natural images can fully account for the asymmetries in width and density in the V1 orientation tuning curves (**Fig. 1b,c**). For the owl, in contrast, the prior does not represent the distribution of sound sources in the environment;

presumably, sounds in a forest may come from any direction. Rather, the prior function represents the relevance of the various sound directions.

Such an 'importance coefficient' of each direction may depend on many factors besides the associated frequency of occurrence. For instance, sounds coming from the back of the owl may be irrelevant because large orienting movements may alert the potential prey or require too much time or energy. In fact, the underestimation of sound directions has been reported in many species². If, for whatever reason, there is no point in responding to a particular direction, then detecting sounds from it is unnecessary; it just wastes resources¹⁴. In general, asymmetries in the distributions of preferences and widths in a population can be used to assign different weights to different stimulus values because of their frequency, their potential for higher reward, motor constraints¹⁴, and so on. In the tennis analogy, a player may ignore balls coming to his backhand side either because they are too infrequent, because he cannot see well in that direction, or because he is hurt and cannot hit backhands. As a consequence, behavioral asymmetries may have multiple causes, and resolving them may require careful analyses such as those in carried out by Girshick *et al.*¹ and Fischer *et al.*², and behavioral or neuronal responses that appear suboptimal under one prior may be optimal under another.

In a wider context, the goal is not just to identify the factors that determine the distributions of widths and preferred values of a

neuronal population, but also to understand why different neurons have tuning curves of different shapes^{14,15}. What makes a 'good' shape? What makes an optimal mixture of shapes for a population encoding a particular sensory feature? The answers will certainly depend on the organism's lifestyle and its interactions with the environment, but there is hope that general principles will emerge^{12,14,15}. The new studies have peeled a layer of mystery from this fundamental issue in computational neuroscience.

COMPETING FINANCIAL INTERESTS

The author declares no competing financial interests.

1. Girshick, A.R., Landy, M.S. & Simoncelli, E.P. *Nat. Neurosci.* **14**, 926–932 (2011).
2. Fischer, B. & Peña, J.L. *Nat. Neurosci.* **14**, 1061–1066 (2011).
3. Ernst, M.O. & Banks, M.S. *Nature* **415**, 429–433 (2002).
4. Körding, K.P. & Wolpert, D.M. *Nature* **427**, 244–247 (2004).
5. Trommershäuser, J., Maloney, L.T. & Landy, M.S. *Trends Cogn. Sci.* **12**, 291–297 (2008).
6. Ma, W.J., Beck, J.M., Latham, P.E. & Pouget, A. *Nat. Neurosci.* **9**, 1432–1438 (2006).
7. Weiss, Y., Simoncelli, E.P. & Adelson, E.H. *Nat. Neurosci.* **5**, 598–604 (2002).
8. Ashourian, P. & Loewenstein, Y. *PLoS ONE* **6**, e19551 (2011).
9. Miyazaki, M., Yamamoto, S., Uchida, S. & Kitazawa, S. *Nat. Neurosci.* **9**, 875–877 (2006).
10. Pouget, A., Dayan, P. & Zemel, R. *Nat. Rev. Neurosci.* **1**, 125–132 (2000).
11. Paradiso, M.A. *Biol. Cybern.* **58**, 35–49 (1988).
12. Berens, P., Ecker, A.S., Gerwin, S., Tolias, A.S. & Bethge, M. *Proc. Natl. Acad. Sci. USA* **108**, 4423–4428 (2011).
13. Salinas, E. & Abbott, L.F. *J. Comput. Neurosci.* **1**, 89–107 (1994).
14. Salinas, E. *PLoS Biol.* **4**, e387 (2006).
15. Bonnasse-Gahot, L. & Nadal, J.P. *J. Comput. Neurosci.* **25**, 169–187 (2008).

Anti-TANKyrase weapons promote myelination

Patrizia Casaccia

A study identifies mechanisms responsible for the inability to form new myelin after neonatal hypoxia. It identifies Axin2 as a potential therapeutic target for reversing the 'differentiation block' of oligodendrocyte-lineage cells.

Cerebral palsy and cognitive deficits represent the devastating consequences of preterm births and of perinatal hypoxic or ischemic injury of full-term infants. At a cellular level, disease severity correlates with the degree of white matter injury and is characterized by the inability of cells in the oligodendrocyte lineage to differentiate into myelin-forming cells. There are no therapies to overcome this

differentiation block. A similar deficit in the ability to form new myelin can be detected in the adult brain after demyelination in people with multiple sclerosis and is associated with lack of repair. In this issue of *Nature Neuroscience*, Fancy and colleagues¹ identify Axin2, an inhibitor of the Wnt pathway, as a promising new therapeutic target for drug development directed at favoring new myelin formation in the neonatal and adult brain.

Wnt proteins comprise a family of secreted ligands crucial for stem cell biology and embryonic development. Inappropriate regulation of Wnt signaling occurs in several types of cancer, including colon, liver and brain tumors

of neuroectodermal origin, and involves a downstream molecule called β -catenin. In the absence of Wnt, Axin1 cooperates with glycogen synthase kinase 3 (GSK3) and phosphorylates β -catenin, thereby signaling its degradation. In the presence of Wnt, β -catenin is not phosphorylated and accumulates in the cell and modulates gene expression. Active Wnt has been shown to impair oligodendrocyte progenitor differentiation and repair of demyelination^{2–5}.

Fancy and colleagues¹ identified the protein Axin2, also known as Axil (in rat) and Conductin (in mouse), as a negative regulator of β -catenin stability (**Fig. 1**), even in the

Patrizia Casaccia is in the Department of Neuroscience and Friedman Brain Institute, Mount Sinai School of Medicine, New York, New York, USA.
e-mail: patrizia.casaccia@mssm.edu