Owl's behavior and neural representation predicted by Bayesian inference

Brian J. Fischer[*], Jose Luis Peña

[*]To whom correspondence should be addressed. Email: brian.fischer@ens.fr

Supplementary Information

Supplementary Discussion

Population vector approximation to Bayesian estimator

Here we show that as the size of a neural population grows, the population-vector estimate of direction will converge to the Bayesian estimate if the tuning curves are proportional to the likelihood function and the preferred directions are drawn independently from the prior distribution. A similar analysis of the center-of-mass estimator was developed previously[1].

The inference problem

The structure of the inference problem faced by the neural system can be described by the Markov chain

$\theta - Y - R$

where θ is the unknown direction (e.g., direction of a sound source), Y is the received sensory information (e.g., ITD), and R is a vector of neural responses.

We propose that the population vector decoding method should produce an estimate of θ from the neural response R that matches a Bayesian estimate of θ obtained from the posterior distribution $p(\theta \mid y)$.

Definition of population vector: The estimate of θ from the neural response R is taken to be the direction of the population vector, which is defined as

$$PV(y) = \frac{1}{N}\sum_{n=1}^{N} r_n(y) u(\theta_n),$$

where $u(\theta)$ is a unit vector pointing in direction θ, $r_n(y)$ is the response of the n^{th} neuron to the input y, and the vector of neural responses is given by
$R(y) = [r_1(y), r_2(y), \ldots, r_N(y)]^T$.

Definition of neural tuning curves: We define the neural tuning curves to be the expected value of the neural response, taken over the distribution of neural noise:

$$a_n(y) = E[r_n(y)],$$

where $E[]$ represents the expected value.

Definition of Bayes vector: The Bayesian estimate of stimulus direction θ from the sensory input Y is given by the mean of θ under the posterior distribution $p(\theta \mid y)$. The mean direction is found as the direction of the Bayes vector, which is defined as

$$BV(y) = \int_{-\pi}^{\pi} u(\theta) p(\theta \mid y) d\theta \propto \int_{-\pi}^{\pi} u(\theta) p(y \mid \theta) p(\theta) d\theta.$$

Proposition: Consider the neural inference problem defined above with neural tuning curves that have the form $a_n(y) = a(y, \theta_n)$, where for each y, $\int_{-\pi}^{\pi} p(\theta) a(y,\theta) u(\theta) d\theta$ is finite. If the preferred directions $\theta_1, \theta_2, \ldots, \theta_n$ are independent and identically distributed according to the prior $p(\theta)$, and the neural tuning curves are proportional to the likelihood function, then the expected value of the population vector, taken over the distribution of neural noise, will converge almost surely to a vector that points in the same direction as the Bayes vector as the number of neurons grows to infinity.

Proof:

The expected value of the population vector, taken over the distribution of neural noise, is given by

$$E[PV(y)] = \frac{1}{N} \sum_{n=1}^{N} a(y, \theta_n) u(\theta_n).$$

Since for each y, $\int_{-\pi}^{\pi} p(\theta) a(y,\theta) u(\theta) d\theta$ is finite, the strong law of large numbers says that

$$E[PV(y)] = \frac{1}{N} \sum_{n=1}^{N} a(y, \theta_n) u(\theta_n) \xrightarrow{a.s.} \int_{-\pi}^{\pi} p(\theta) a(y,\theta) u(\theta) d\theta.$$

If the neural tuning curves are proportional to the likelihood function, then

$$\int_{-\pi}^{\pi} p(\theta) a(y,\theta) u(\theta) d\theta \propto \int_{-\pi}^{\pi} p(\theta) p(y \mid \theta) u(\theta) d\theta \propto BV(y). \qquad \square$$

Note that this result does not depend on the structure of the correlated variability of the neural noise.

Alternative neural decoders: place code and probabilistic population code

The results indicate that a strict place-code estimate will not reproduce the owl's behavior. The shape of the tuning curves of neurons in the owl's OT are well described by the likelihood function (**Fig. 6**, main text). Therefore, estimating the sound source direction by the location in the auditory space map with maximal activity would be equivalent to performing a maximum likelihood estimate. The maximum likelihood estimate fails to capture the owl's behavior (**Fig. 4a**, main text), indicating that the maximally activated region of the network alone does not describe the owl's localization behavior.

To test decoding systems other than a population vector, we used a *probabilistic population code*[2]. The Bayesian estimate of stimulus direction θ from the population neural response $\mathbf{r}$ is given by the mean of θ under the posterior distribution $p(\theta \mid \mathbf{r})$. The estimate is computed as described for the Bayesian estimate from the posterior distribution conditioned on ITD. Specifically, the estimate from the probabilistic population code is given by the direction of the vector that points in the mean direction:

$PPCV = \int_{-\pi}^{\pi} u(\theta) p(\theta \mid \mathbf{r}) d\theta \propto \int_{-\pi}^{\pi} u(\theta) p(\mathbf{r} \mid \theta) p(\theta) d\theta$, where $u(\theta)$ is a unit vector pointing in direction θ.

Estimates of sound source direction using this decoder did not reproduce the owl's behavior when using a network matching the experimental measurements. We examined the direction estimates obtained by computing the mean of the posterior probability conditioned on the neural response vector $p(\theta \mid \mathbf{r})$. This estimator utilizes the full probability distribution of neural responses, unlike the population vector. Using the same network model described for the population vector (**Fig. 6**, main text), the Bayesian estimate from $p(\theta \mid \mathbf{r})$ greatly underestimated the source direction (**Supplementary Fig. 2a,b**). For this estimator to match the owl's behavior, the distribution of preferred directions must be wider than the distribution of preferred directions measured in OT and the tuning curves must be much sharper than the measured tuning curves (**Supplementary Fig. 2c-f**). Specifically, the Bayesian estimate from $p(\theta \mid \mathbf{r})$ will match the Bayesian estimate from $p(\theta \mid \mathrm{ITD})$ when the likelihood function based on the neural response vector $p(\mathbf{r} \mid \theta) = \int p(\mathbf{r} \mid \mathrm{ITD}) p(\mathrm{ITD} \mid \theta) d\mathrm{ITD}$ approximates the likelihood function $p(\mathrm{ITD} \mid \theta)$. The likelihoods will be similar when the probability $p(\mathbf{r} \mid \mathrm{ITD})$ is sharply peaked around the ITD that the neurons encode, i.e, when the encoded ITD is uniquely determined from the neural response. We found that this occurs only when the distribution of preferred directions is wide and the tuning curves are narrow. Even after reducing the standard deviation of the ITD-based likelihood $p(\mathrm{ITD} \mid \theta)$ from 41.2 μs to 5 μs, the tuning-curve widths required to match the owl’s behavior were sharper than the lower bound of the observed curve widths (not shown).

Supplementary Figures

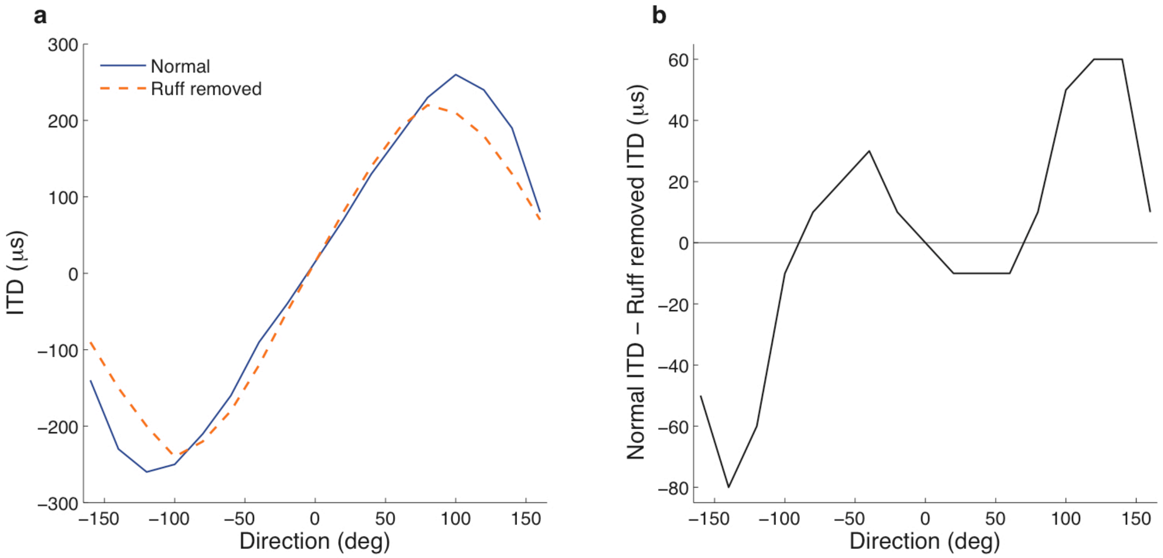

Supplementary Figure 1: Direction-dependence of ITD. (a) Measured relationship between direction and interaural time difference (ITD) (solid blue) under normal conditions[3], along with the measured relationship between direction and ITD under ruff-removed conditions (dashed orange). Note that ruff removal causes a decrease in the range of ITD and an increase in the frequency of the sinusoidal relationship between direction and ITD. (b) The difference between the normal and ruff-removed relationships between direction and ITD.

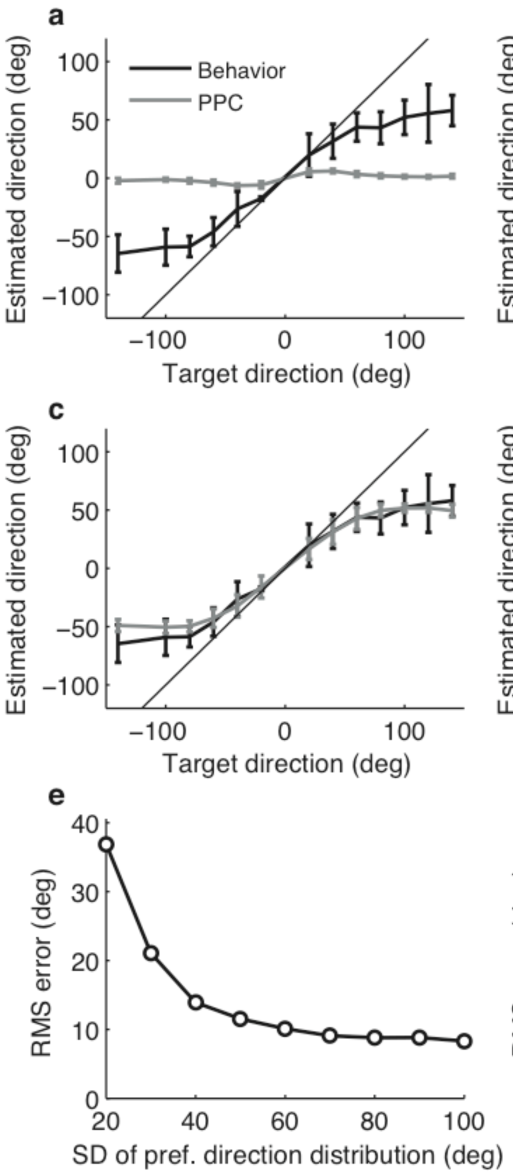

Supplementary Figure 2: Performance of the probabilistic population code (PPC). (a,b) Owl's behavior (bold black) and the estimate from the PPC, i.e., the Bayesian estimate from the neural response-conditioned posterior distribution $p(\theta\,|\,\mathbf{r})$ (gray) in the normal (a) and ruff removed conditions (b) using the network described in **Figure 6**, main text. The thin black line is the identity. (c,d) Owl's behavior (bold black) and Bayesian estimate from $p(\theta\,|\,\mathbf{r})$ using sharper tuning curves and a broader range of preferred directions (gray) in the normal (c) and ruff removed conditions (d). (e) RMS error between the owl's normal behavior and the Bayesian estimate from $p(\theta\,|\,\mathbf{r})$ as a function of the standard deviation of the Gaussian-shaped distribution from which preferred directions were drawn. (f) RMS error between the owl's normal behavior and the Bayesian estimate from $p(\theta\,|\,\mathbf{r})$ as a function of the width of the neural tuning curve of a neuron with preferred direction 0 deg.

Supplementary References

1. Fischer B.J. Bayesian estimates from heterogeneous population codes. *The 2010 Int'l. Joint Conf. on Neural Networks* (*IJCNN*). doi: 10.1109/IJCNN.2010.5596687, 2010.
2. Ma, W.J., Beck, J.M., Latham, P.E. & Pouget, A. Bayesian inference with probabilistic population codes. *Nature Neurosci.* **9**, 1432–1438 (2006).
3. Hausmann, L., von Campenhausen, M., Endler, F., Singheiser, M. & Wagner H. Improvements of sound localization abilities by the facial ruff of the barn owl (Tyto alba) as demonstrated by virtual ruff removal. *PLoS ONE* **4**(11): e7721. doi:10.1371/journal.pone.0007721 (2009).