Systems/Circuits

Accounting for Bias in the Estimation of r^2 between Two Sets of Noisy Neural Responses

Dean A. Pospisil[1,2] **and Wyeth Bair**[2,3,4]

[1]Princeton Neuroscience Institute, Princeton University, Princeton, New Jersey 08540, [2]Department of Biological Structure, University of Washington, Seattle 98195, [3]Computational Neuroscience Center, University of Washington, Seattle 98195, and [4]Washington National Primate Research Center, University of Washington, Seattle 98195

The Pearson correlation coefficient squared, r^2, is an important tool used in the analysis of neural data to quantify the similarity between neural tuning curves. Yet this metric is biased by trial-to-trial variability; as trial-to-trial variability increases, measured correlation decreases. Major lines of research are confounded by this bias, including those involving the study of invariance of neural tuning across conditions and the analysis of the similarity of tuning across neurons. To address this, we extend an estimator, $\hat{r}^2_{ER}$, that was recently developed for estimating model-to-neuron correlation, in which a noisy signal is compared with a noise-free prediction, to the case of neuron-to-neuron correlation, in which two noisy signals are compared with each other. We compare the performance of our novel estimator to a prior method developed by Spearman, commonly used in other fields but widely overlooked in neuroscience, and find that our method has less bias. We then apply our estimator to demonstrate how it avoids drastic confounds introduced by trial-to-trial variability using data collected in two prior studies (macaque, both sexes) that examined two different forms of invariance in the neural encoding of visual inputs—translation invariance and fill-outline invariance. Our results quantify for the first time the gradual falloff with spatial offset of translation-invariant shape selectivity within visual cortical neuronal receptive fields and offer a principled method to compare invariance in noisy biological systems to that in noise-free models.

Key words: **confound; correlation; invariance; sensory; statistics; V4**

Significance Statement

Quantifying the similarity between two sets of averaged neural responses is fundamental to the analysis of neural data. A ubiquitous metric of similarity, the correlation coefficient, is attenuated by trial-to-trial variability that arises from many irrelevant factors. Spearman recognized this problem and proposed corrected methods that have been extended over a century. We show this method has large asymptotic biases that can be overcome using a novel estimator. Despite the frequent use of the correlation coefficient in neuroscience, consensus on how to address this fundamental statistical issue has not been reached. We provide an accurate estimator of the correlation coefficient and apply it to gain insight into visual invariance.

Introduction

The measurement of correlation between noisy datasets is ubiquitous and plays a critical role in sensory neuroscience. The r^2 between two sets of mean neural responses is fundamentally relevant to lines of research that compare tuning curves across neurons to understand functional organization and population encoding (Gawne and Richmond, 1993; Gawne et al., 1996; Cohen and Kohn, 2011; Power et al., 2011; Kiani et al., 2015) and to studies that compare tuning curves within the same neuron across stimulus transformations to understand invariance in cortical representations (Nandy et al., 2013; El-Shamayleh and Pasupathy, 2016; Kell and McDermott, 2019; Popovkina et al., 2019). Despite the importance of estimating the correlation across tuning curves, the typical estimator, $\hat{r}^2$, is fundamentally confounded by the trial-to-trial variability of neural responses. Even if two neurons have identical tuning, or a single neuron perfectly maintains its selectivity under some stimulus transformation, the measured correlation in each case will decrease as trial-to-trial variability increases. Investigators have approached this problem by averaging over repeated trials of the same stimulus, but added data collection is expensive and never wholly removes the confounding influence of noise.

Received Jan. 24, 2022; revised Sep. 23, 2022; accepted Oct. 20, 2022.

Author contributions: D.A.P. and W.B. designed research; D.A.P. performed research; D.A.P. analyzed data; D.A.P. and W.B. wrote the paper.

This work was supported by National Science Foundation Graduate Research Fellowship DGE-1256082 (D.A.P.) and National Institutes of Health–National Eye Institute T32 Vision Training Grants EY07031 (D.A.P.), R01-EY029997 (W.B.), and R01-EY027023 (W.B.). We thank Dina Popovkina, Yasmine El-Shamayleh, and Anitha Pasupathy for sharing data and Anitha Pasupathy, Greg Horwitz, and Yen-Chi Chen for providing comments.

The authors declare no competing financial interests.

Correspondence should be addressed to Dean A. Pospisil at dp4846@princeton.edu.

https://doi.org/10.1523/JNEUROSCI.0198-22.2022

A more principled approach has been to account for trial-to-trial variability in the estimation of correlation. Correction for the attenuation of the correlation coefficient by measurement error has received considerable attention from fields outside neuroscience (Thouless, 1939; Beaton et al., 1979; Rosner and Willett, 1988; Adolph and Hardin, 2007; Saccenti et al., 2020) and has recently been adopted to estimate invariance of neural activity (Kell and McDermott, 2019). These studies followed a line of work begun by Spearman (1904), who provided a correction in the form of a multiplicative scaling of the Pearson correlation coefficient by its estimated attenuation. Here, we follow a different approach developed in the neuroscience literature (Sahani and Linden, 2003; Haefner and Cumming, 2008; Pospisil and Bair, 2021b) that sought to account for trial-to-trial variability in the estimation of the fit of a model to noisy estimated tuning curves. This line of work developed separate unbiased estimates of the numerator and denominator of the Pearson correlation coefficient. Here, we show how to extend that work to the case of measuring the correlation between two noisy, estimated tuning curves, and we call the estimator $\hat{r}^2_{ER}$ because it estimates the r^2 between the expected responses (ERs) of two neurons. We show in simulation that $\hat{r}^2_{ER}$ is less biased for small samples and less biased asymptotically than Spearman's approach.

To demonstrate the usefulness of our novel estimator, we apply it to the problem of invariant cortical representation, where single-neuron selectivity is considered to be invariant if it is maintained across a transformation of the sensory input. Transformations of visual stimuli can be geometric, such as translation (Nandy et al., 2013) or scaling (El-Shamayleh and Pasupathy, 2016), or they can modify appearance by changing surface properties, for example, color or luminance (Bushnell et al., 2011; Popovkina et al., 2019). We show in neural data how $\hat{r}^2_{ER}$ provides insights into two forms of invariance in area V4. First, for translation invariance, we find that the naive $\hat{r}^2$ causes V4 neuronal selectivity to appear to vary with small shifts in position, whereas our noise-corrected $\hat{r}^2_{ER}$ reveals that selectivity is robust to small translations. Second, we study fill-outline invariance, which examines how shape selectivity for boundaries depends on surface properties of simple objects. We find that among neurons driven strongly enough to accurately infer invariance, our estimator shows an approximately uniform distribution of invariance from zero to one across V4, indicating a striking diversity in the influence of surface properties on form selectivity. Finally, to demonstrate the more general significance of our estimator, we document how studies of place field remapping in hippocampus are particularly susceptible to the confounding effect of noise on the estimation of tuning curve correlation.

Materials and Methods

Simulation procedure. To demonstrate the bias that noise imparts on neuron-to-neuron correlation and the ability of our methods to remove this bias, we simulate the responses of neurons to m stimuli across n repeats. We generate values, $r_{i,j}$, that represent the square root (for reasons described below) of single-neuron responses to the jth repeat of the ith stimulus from a normal distribution, $r_{i,j} \sim N(\mu_i, \sigma^2)$, where the mean, μ_i, varies across stimuli while the variance, σ^2, is constant. The degree of variation across the μ_i reflects how well the stimuli modulate the neural response; we call this the dynamic range and quantify it as follows:

$$d^2 = \sum_{i=1}^{m} (\mu_i - \bar{\mu})^2, \tag{1}$$

where $\bar{\mu}$ is the average of the μ_i across stimuli. Trial-to-trial variability can be different across cells, fixing n and m; the signal-to-noise ratio

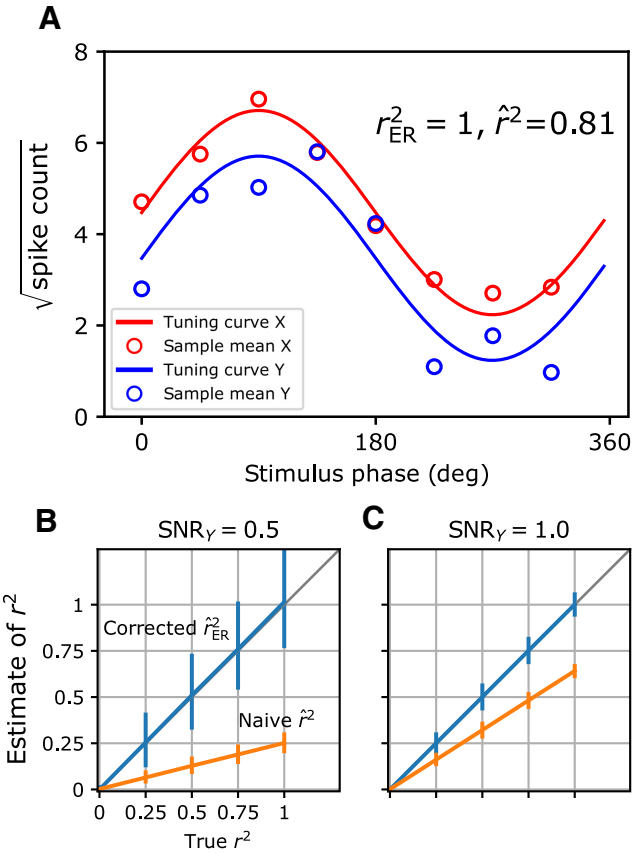

Figure 1. Simulation model of neuron-to-neuron fits and validation of $\hat{r}^2_{ER}$. ***A***, Here, $r^2_{ER} = 1$ because tuning curves (solid trace) are identical up to a shift and scaling. The estimate of correlation, $\hat{r}^2$, from trial averages (open circles) is lower than one. ***B***, Simulation of $\hat{r}^2$ (orange) and $\hat{r}^2_{ER}$ (blue) for estimating the correlation of two noisy tuning curves at varying levels of true r^2_{ER} in the case where the second neuron has a lower SNR, SNR$_X$ = 1, SNR$_Y$ = 0.5. Number of trials are n = 4, stimuli m = 371, and trial-to-trial variability is σ^2 = 0.25. Vertical bars are 95% quantiles. ***C***, The same simulation as ***B***, but both neurons have SNR = 1.

(SNR) determines the reliability with which correlation can be estimated across neuron pairs as follows:

$$\text{SNR} = \frac{d^2/m}{\sigma^2}. \tag{2}$$

This will be a critical parameter in both our simulations and our analysis of neural recordings.

For our simulations, the tuning curves for a pair of neurons, X and Y, in response to a set of m stimuli are defined to be the two sets of mean responses, μ_i and ν_i, respectively (i = 1 ... m), and these are modeled as sinusoids (Fig. 1*A*, lines), as in the following:

$$\mu_i = a_X + d_X \frac{\sin(\theta_i)}{\sqrt{\sum_{j=1}^{m} \sin(\theta_j)^2}}, \tag{3}$$

where the denominator normalizes the length of the tuning vector as follows:

$$\theta_i = \frac{(i-1)}{m} 2\pi. \tag{4}$$

The tuning curve for neuron Y is simply phase shifted to achieve the desired value of r^2_{ER}, as follows:

$$\nu_i = a_Y + d_Y \frac{\sin(\theta_i + \arccos(r_{ER}))}{\sqrt{\sum_{j=1}^{m} \sin\left(\theta_j + \arccos(r_{ER})\right)^2}}. \quad (5)$$

To validate potential estimators for r^2_{ER}, we draw responses independently from this model for n repeats for each of the m stimuli. Averaging these observations across repeats produces observed tuning curves for the two simulated neurons (Fig. 1*A*, open circles).

Assumptions for the derivation of unbiased estimators. To simplify the derivation of unbiased estimators for r^2_{ER}, we assume that neural responses have undergone a variance-stabilizing transform so that trial-to-trial variability is the same across all stimuli for a given neuron. For example, if the neural responses are Poisson distributed or, more generally, have a linear mean-to-variance relationship, the square root is a variance stabilizing transform. In all results where we apply our estimator to neural data, we take the square root of spike counts. We furthermore assume that averaged responses are normally distributed, justified by invoking the central limit theorem guaranteeing convergence to normality as the number of repeats increases. The Box–Cox transform applied to the spike counts can further ensure these assumptions of equal variance and normality are met (Box and Cox, 1964).

Unbiasing r^2. In the case where both X and Y are equal variance stochastic responses, $X_i \sim N\left(\mu_i, \frac{\sigma^2}{n}\right)$ and $Y_i \sim N\left(\nu_i, \frac{\sigma^2}{n}\right)$, we aim to unbias the naive correlation coefficient defined as follows:

$$\hat{r}^2 = \frac{\left(\sum_{i=1}^{m}(X_i - \bar{X})(Y_i - \bar{Y})\right)^2}{\sum_{i=1}^{m}(X_i - \bar{X})^2 \sum_{i=1}^{m}(Y_i - \bar{Y})^2} \quad (6)$$

to ideally achieve a corrected estimator, $\hat{r}^2_{ER}$, as follows:

$$\mathbb{E}[\hat{r}^2_{ER}] = r^2_{ER} = \frac{\left(\sum_{i=1}^{m}(\mu_i - \bar{\mu})(\nu_i - \bar{\nu})\right)^2}{\sum_{i=1}^{m}(\mu_i - \bar{\mu})^2 \sum_{i=1}^{m}(\nu_i - \bar{\nu})^2}. \quad (7)$$

In our approach, we will unbias the numerator and denominator of Equation 6 separately. The expected value of the numerator is the following:

$$\mathbb{E}\left[\left(\sum_{i=1}^{m}(X_i - \bar{X})(Y_i - \bar{Y})\right)^2\right] = \left(\sum_{i=1}^{m}(\mu_i - \bar{\mu})(\nu_i - \bar{\nu})\right)^2 + \frac{\sigma^2}{n}\left(\sum_{i=1}^{m}(\mu_i - \bar{\mu})^2 + \sum_{i=1}^{m}(\nu_i - \bar{\nu})^2 + (m-1)\frac{\sigma^2}{n}\right), \quad (8)$$

in which the first term is the desired numerator (that of Eq. 7), and the term on the second line is the bias. To remove this bias, an unbiased estimate of it can be subtracted from the numerator of Equation 6.

The denominator of Equation 6 consists of two factors that are each scaled noncentral chi-squared distributions as follows:

$$\sum_{i=1}^{m}(X_i - \bar{X})^2 \sim \frac{\sigma^2}{n}\chi^2_{m-1}\left(\frac{n}{\sigma^2}\sum_{i=1}^{m}(\mu_i - \bar{\mu})^2\right) \quad (9)$$

and

$$\sum_{i=1}^{m}(Y_i - \bar{Y})^2 \sim \frac{\sigma^2}{n}\chi^2_{m-1}\left(\frac{n}{\sigma^2}\sum_{i=1}^{m}(\nu_i - \bar{\nu})^2\right). \quad (10)$$

Thus, the expected values of these factors are as follows:

$$\mathbb{E}\left[\sum_{i=1}^{m}(X_i - \bar{X})^2\right] = \mathbb{E}\left[\frac{\sigma^2}{n}\chi^2_{m-1}\left(\frac{n}{\sigma^2}\sum_{i=1}^{m}(\mu_i - \bar{\mu})^2\right)\right] = \sum_{i=1}^{m}(\mu_i - \bar{\mu})^2 + (m-1)\frac{\sigma^2}{n} \quad (11)$$

and

$$\mathbb{E}\left[\sum_{i=1}^{m}(Y_i - \bar{Y})^2\right] = \mathbb{E}\left[\frac{\sigma^2}{n}\chi^2_{m-1}\left(\frac{n}{\sigma^2}\sum_{i=1}^{m}(\nu_i - \bar{\nu})^2\right)\right] = \sum_{i=1}^{m}(\nu_i - \bar{\nu})^2 + (m-1)\frac{\sigma^2}{n}. \quad (12)$$

Because these two random variables are independent, the expected value of the denominator, which is their product, is the product of these expected values as follows:

$$\mathbb{E}\left[\sum_{i=1}^{m}(X_i - \bar{X})^2 \sum_{i=1}^{m}(Y_i - \bar{Y})^2\right] = \left(\sum_{i=1}^{m}(\mu_i - \bar{\mu})^2 + (m-1)\frac{\sigma^2}{n}\right)\left(\sum_{i=1}^{m}(\nu_i - \bar{\nu})^2 + (m-1)\frac{\sigma^2}{n}\right) = \sum_{i=1}^{m}(\mu_i - \bar{\mu})^2 \sum_{i=1}^{m}(\nu_i - \bar{\nu})^2 + (m-1)\frac{\sigma^2}{n}\left(\sum_{i=1}^{m}(\mu_i - \bar{\mu})^2 + \sum_{i=1}^{m}(\nu_i - \bar{\nu})^2 + (m-1)\frac{\sigma^2}{n}\right). \quad (13)$$

The first term (second line of this equation) is the desired denominator (that of Eq. 7), whereas the second term (third line) is the bias. To remove this bias, an unbiased estimate of it can be subtracted from the denominator of Equation 6.

Estimators of bias terms. To compute the bias terms to be subtracted from the numerator and denominator, three unknown quantities need to be estimated, that is, the two following dynamic range values:

$$d^2_X = \sum_{i=1}^{m}(\mu_i - \bar{\mu})^2 \quad (14)$$

$$d^2_Y = \sum_{i=1}^{m}(\nu_i - \bar{\nu})^2, \quad (15)$$

and the trial-to-trial variability, σ^2. Below we provide unbiased estimators for these quantities.

In the case of the dynamic range values, the following naive sample estimator:

$$\hat{d}^2_X = \sum_{i=1}^{m}(X_i - \bar{X})^2 \quad (16)$$

has a bias, as revealed by taking its expected value as follows:

$$\mathbb{E}\left[\sum_{i=1}^{m}(X_i - \bar{X})^2\right] = \mathbb{E}\left[\frac{\sigma^2}{n}\chi^2_{m-1}\left(\frac{n}{\sigma^2}\sum_{i=1}^{m}(\mu_i - \bar{\mu})^2\right)\right] = \sum_{i=1}^{m}(\mu_i - \bar{\mu})^2 + (m-1)\frac{\sigma^2}{n}. \quad (17)$$

Thus, an unbiased estimator is the following:

$$\hat{d}^2_{X_{(\mathrm{ER})}} = \sum_{i=1}^{m} (X_i - \bar{X})^2 - (m-1)\frac{\sigma^2}{n}. \tag{18}$$

For the case of the sample variance, if we have n repeated trials, we can calculate sample variance over those trials, then because the variance is the same across stimuli and neurons (assuming a variance stabilizing transformation), we can average them for a global estimate as follows:

$$\hat{\sigma}^2 = \left(\frac{1}{m}\sum_{i=1}^{m} s^2_{i,X} + \frac{1}{m}\sum_{i=1}^{m} s^2_{i,Y}\right) \Big/ 2, \tag{19}$$

where $s^2_{i,X}$ and $s^2_{i,Y}$ are the sample variance estimates calculated across n repeats for the ith stimulus. Using these unbiased quantities to form the terms to be subtracted from the numerator and denominator of the naive $\hat{r}^2$ (Eq. 6), we finally arrive at our corrected estimator for r^2_{ER} in the following:

$$\hat{r}^2_{\mathrm{ER}} = \frac{\left(\sum_{i=1}^{m}(X_i - \bar{X})(Y_i - \bar{Y})\right)^2 - \frac{\hat{\sigma}^2}{n}\left(\hat{d}^2_{X_{(\mathrm{ER})}} + \hat{d}^2_{Y_{(\mathrm{ER})}} + (m-1)\frac{\hat{\sigma}^2}{n}\right)}{\sum_{i=1}^{m}(X_i - \bar{X})^2 \sum_{i=1}^{m}(Y_i - \bar{Y})^2 - (m-1)\frac{\hat{\sigma}^2}{n}\left(\hat{d}^2_{X_{(\mathrm{ER})}} + \hat{d}^2_{Y_{(\mathrm{ER})}} + (m-1)\frac{\hat{\sigma}^2}{n}\right)}. \tag{20}$$

Confidence intervals for $\hat{r}^2_{ER}$. To quantify uncertainty about the estimate $\hat{r}^2_{\mathrm{ER}}$ relative to the true parameter r^2_{ER}, we follow the same approach laid out in Pospisil and Bair (2021b) for estimating the correlation between a fixed model and neural tuning estimated from noisy neural data. There we found that several commonly used generic bootstrap methods did not provide the desired coverage probability in simulation. Of those methods, the parametric BCa (bias-corrected and accelerated) was the closest to the desired probability (Efron and Tibshirani, 1994). We developed an alternative confidence interval (CI) using a Bayesian approach that achieved the desired coverage probability, and we called it the estimate centered credible interval (ECCI).

ECCI finds confidence intervals centered around the estimate of $\hat{r}^2_{ER}$, which under its posterior sampling distribution, given the observed data, provides the appropriate α coverage level. To do this, we assumed uninformative priors on the dynamic range and trial-to-trial variability of the two multivariate normal distributions representing samples from the two tuning curves. Using the Metropolis–Hastings algorithm, we then conditioned on estimates of these parameters to draw from the posterior sampling distribution of $\hat{r}^2_{\mathrm{ER}}|r^2_{\mathrm{ER}}$ for an arbitrary r^2_{ER}. We then find those values of r^2_{ER} that would result in the observed $\hat{r}^2_{\mathrm{ER}}$ with probability $1 - \alpha/2$ and $\alpha/2$ for the upper and lower ends of the confidence interval. Pospisil and Bair (2021b) provide more details. We apply this method to calculate confidence intervals for estimates of $\hat{r}^2_{\mathrm{ER}}$ throughout the rest of the article.

Spearman's correction for attenuation. Spearman (1904) noted that in the case of measurements of two quantities with some underlying true correlation but with additive independent measurement error, the measured correlation would tend to be less than the true correlation. He provided the analytic expression for the attenuation, A (using the notation of Saccenti et al., 2020), under a bivariate normal model as follows:

$$A = \frac{1}{\sqrt{\left(1 + \frac{\sigma^2_{w,X}}{\sigma^2_{a,X}}\right)\left(1 + \frac{\sigma^2_{w,Y}}{\sigma^2_{a,Y}}\right)}}, \tag{21}$$

so that the observed correlation, ρ, was a scaling of the true correlation, ρ_0, as follows:

$$\rho = A\rho_0. \tag{22}$$

Here, $\sigma^2_{w,X}$ and $\sigma^2_{w,Y}$ are the variance of the additive measurement error (within-condition variance), and $\sigma^2_{a,X}$ and $\sigma^2_{a,Y}$ are the variance of the underlying quantities (across-condition variance). This is slightly different from our derivation above. Here, both the correlated and noisy components are random, whereas in the previous section our correlated component was fixed, and only the noise component was random. Thus we use $\sigma^2_{a,X}$ instead of d^2_X because the former is the variance of a normal distribution, whereas the latter is the squared deviation of a set of means.

The method to correct the attenuation is straightforward. Find an estimate for A and multiply the estimated correlation, ρ_0, by the inverse of this estimate. Spearman does not specify estimators for the unknowns in Equation 21; however, Adolph and Hardin (2007) use sample variance to estimate both the within-condition variance ($\sigma^2_{w,X}$ and $\sigma^2_{w,Y}$) and across-condition variance ($\sigma^2_{a,X}$ and $\sigma^2_{a,Y}$). Thus, using the estimators defined in the previous section and assuming equal noise variance ($\sigma^2_{w,X} = \sigma^2_{w,Y}$), Spearman's corrected estimator takes the following form:

$$\hat{\rho}_0 = \sqrt{\left(1 + \frac{\hat{\sigma}^2/n}{\hat{d}^2_X/(m-1)}\right)\left(1 + \frac{\hat{\sigma}^2/n}{\hat{d}^2_Y/(m-1)}\right)} \cdot \hat{r}. \tag{23}$$

We will square this estimator to compare it to ours (see below, Results). Because $\hat{d}^2_X$ and $\hat{d}^2_Y$ are biased estimators (Eq. 17), we also compute for comparison a version of this estimator that uses the unbiased parameter estimates shown in Equation 18, and we refer to this as $\hat{\rho}^2_{0_{(\mathrm{ER})}}$.

Electrophysiological data. To demonstrate the use of our estimator, we reanalyzed data from three previous single-unit, extracellular studies of parafoveal V4 neurons in the awake, fixating rhesus monkey (*Macaca mulatta*) of both sexes. Data from the first study, (Pasupathy and Connor, 2001), consists of the responses of 109 V4 neurons to a set of 362 shapes. There were typically three to five repeats of each stimulus, and we examined only the 96 cells that had at least four repeats for all stimuli. We computed the spike count for each trial during the 500 ms stimulus presentation. From a second study (El-Shamayleh and Pasupathy, 2016), we reanalyzed responses of 80 neurons tested for translation invariance using shapes like those of the first study but where the position of the stimuli within the receptive field (RF) was also varied. Each neuron was tested with up to 56 shapes (some of which are rotations of others) presented at three to five positions within the RF. Each unique combination of stimulus and RF position was presented for 5–16 repeats, and spike counts were averaged over the 300 ms stimulus presentation. Data from the third study (Popovkina et al., 2019), consists of the responses of 42 V4 neurons using shapes like those of the first study except in two conditions—fill and outline. In the fill condition, the interior of the shape was the same color as its outline. In the outline condition, the interior of the shape was the same color as the background. Stimuli were presented for 300 ms with a 200 ms blank interval preceding each. For each neuron, stimuli were presented for at least three repeats. Experimental protocols for all studies are described in detail in the original publications.

Experimental design and statistical analysis. The experimental design is described above in Electrophysiological data, and a full description is provided in the original studies (Pasupathy and Connor, 2001; El-Shamayleh and Pasupathy, 2016; Popovkina et al., 2019). The statistical analysis of data are described in detail above in the following sections: Unbiasing r^2, Estimators of bias terms, Confidence intervals for $\hat{r}^2_{\mathrm{ER}}$, and Spearman's correction for attenuation.

Data availability. Code for calculating estimates of $\hat{r}^2_{\mathrm{ER}}$ and associated confidence intervals are available at https://github.com/deanpospisil/er_est.

Results

Our results are organized as follows. We first describe the source of the bias in $\hat{r}^2$ using idealized neuronal tuning curves and show how our corrected estimator, $\hat{r}^2_{\mathrm{ER}}$, greatly reduces this bias. We next validate our method in simulation, comparing it to alternative methods, and we validate it using neural data in a split-trial

comparison. Finally, we demonstrate how our estimator avoids significant confounds in two applications, the measurement of translation invariance and fill-outline invariance in single-unit data collected in previous studies of shape selectivity in visual cortical area V4.

Consider a typical scenario in sensory neuroscience where the response of two neurons to n repeated presentations of m stimuli have been collected, and the average of these responses are compared (Fig. 1*A*). Even if the underlying neuronal tuning curves (Fig. 1*A*, blue and red lines) were perfectly correlated, the m sample averages, Y_i and X_i (blue and red open circles), will deviate from the expected value because of trial-to-trial variability, σ^2, scaled by $1/n$. Thus, the observed $\hat{r}^2$ can be appreciably <1, although the r^2 between the expected values of the neuronal responses is 1 (Fig. 1*A*, red and blue lines are identical up to a shift and scaling).

The quantity we attempt to estimate here is the correlation coefficient r^2 between the expected values, μ_i and ν_i, of the responses of the two neurons, that is, between the tuning curves in the absence of noise. We will call this quantity r^2_{ER}, the fraction of variance shared by the expected responses (ER) of the two neurons defined as follows:

$$r^2_{ER} = \frac{\left(\sum_{i=1}^{m}(\mu_i - \bar{\mu})(\nu_i - \bar{\nu})\right)^2}{\sum_{i=1}^{m}(\mu_i - \bar{\mu})^2 \sum_{i=1}^{m}(\nu_i - \bar{\nu})^2}. \quad (24)$$

It is tempting to estimate this quantity using the naive sample estimator as follows:

$$\hat{r}^2 = \frac{\left(\sum_{i=1}^{m}(X_i - \bar{X})(Y_i - \bar{Y})\right)^2}{\sum_{i=1}^{m}(X_i - \bar{X})^2 \sum_{i=1}^{m}(Y_i - \bar{Y})^2}, \quad (25)$$

however, the expected values of the numerator and denominator of $\hat{r}^2$ have a bias with respect to those of r^2_{ER}, as shown in Equations 8 and 13, respectively (see above, Materials and Methods). Each has a bias proportional to σ^2/n, the amount of variability remaining after averaging over stimulus repeats. The bias in the denominator, which results in the attenuation of $\hat{r}^2$, remains even as $m \to \infty$ because the effect of trial-to-trial variability is accumulated across stimuli in the calculation of total variance in the denominator. This is reflected in the scaling of the bias of the denominator by $(m - 1)$ (Eq. 13, third line). The bias does go to zero as $n \to \infty$ because averaging across repeated presentations of stimuli reduces the effect of trial-to-trial variability. This is reflected by the trial-to-trial variability factors in the bias of both Equations 8 and 13 being scaled by $1/n$.

To solve this problem, we take the straightforward strategy of finding unbiased estimators of these noise terms and subtracting them from the numerator and denominator of $\hat{r}^2$ to give our estimator $\hat{r}^2_{ER}$ (see above, Materials and Methods leading up to Eq. 20). Yet, the expectation of the ratio, $\hat{r}^2_{ER}$, is not necessarily the ratio of the expectations of the two unbiased estimates of the numerator and denominator; thus the method is approximately unbiased, although consistent. We find in simulation that for ranges of parameters typical of neural recordings, $\hat{r}^2_{ER}$ has very little bias and is the least biased of available estimators.

Validation of estimator by simulation

To test the ability of $\hat{r}^2_{ER}$ to accurately estimate true neuron-to-neuron tuning curve correlation in the face of noise, we simulated the responses of two neurons with the same trial-to-trial variability (σ^2), the same number of trials (n), but potentially different SNR (Eq. 2; because noise is fixed, SNR increases with tuning curve amplitude). Both simulated neurons had sinusoidal tuning curves (as in Fig. 1*A*), and their relative phase offset was varied to set the true tuning correlation (see above, Materials and Methods, Eq. 3). We first consider the case where neuron X has a high SNR ($SNR_X = 1.0$), whereas neuron Y has a low SNR ($SNR_Y = 0.5$). In our simulation we set $m = 371$ and $n = 4$ to represent a typical V4 experiment with many shape stimuli and few repeats of each. This parameter regime, where $m \gg n$, is typical in neuroscience because often many experimental conditions (e.g., stimuli) are needed to characterize tuning (see below, Nonvision example of correcting correlation between tuning curves for hippocampus example). In this case, the naive estimate of correlation, $\hat{r}^2$ (Fig. 1*B*, orange), is on average only one-quarter of the value of the true correlation between the simulated tuning curves. For example, $\hat{r}^2 = 0.25$ when $r^2_{ER} = 1$. On the other hand, our $\hat{r}^2_{ER}$ on average lies very close to the true correlation (Fig. 1*B*, blue line lies on diagonal). However, although $\hat{r}^2_{ER}$ is clearly the less biased estimator of r^2_{ER}, it has become more variable; the 95% quantile error bars are longer than those for the naive estimator. When both neurons have a high SNR (Fig. 1*C*), the difference between the average of $\hat{r}^2_{ER}$ and $\hat{r}^2$ shrinks commensurate with the increased SNR, as does the difference in the error bars. We note that $\hat{r}^2_{ER}$ can exceed $r^2 = 1$ or go below $r^2 = 0$; this is necessary to obtain an unbiased estimator (e.g., if all probability mass is below 1, the expected value cannot be 1). To account for these impossible values of r^2, when analyzing results for single neurons, we recommend truncating the estimator to lie in [0, 1]; this will minimally add to the bias and aid interpretation.

Comparison to Spearman's correction

To demonstrate how our estimator, $\hat{r}^2_{ER}$, performs relative to past attempts to correct the downward bias of correlation caused by noise, we compare it to that of Spearman as formulated by Adolph and Hardin (2007), $\hat{\rho}^2_0$, and to an extension, $\hat{\rho}^2_{0(ER)}$, where we substituted unbiased estimates of d^2_X and d^2_Y (see above, Materials and Methods, Spearman's correction for attenuation). For reference, we also compare these estimators to the naive estimate $\hat{r}^2$.

We compare the bias of these four estimators using simulated data in which we vary the number of stimuli m and the SNR for both neurons while holding the number of repeats constant at $n = 4$ and the true correlation fixed at $r^2_{ER} = 1$ (Fig. 2, top row). Our corrected estimates ($\hat{\rho}^2_{0(ER)}$ and $\hat{r}^2_{ER}$) can become unstable for low SNR because significant probability mass accumulates near zero in their denominators; thus, here we truncate these estimators between [−1, 2], thereby reducing the effect of extreme values. This allows our simulations to be reproducible while still reflecting the low bias and high variance of the raw corrected estimators for low SNR. In the case of a small number of stimuli ($m = 20$; Fig. 2*A*), the naive $\hat{r}^2$ (orange trace) is well below one and remains below for all levels of SNR. This substantial downward bias remains even as m is increased to 50 (Fig. 2*B*, orange) and 1000 (Fig. 2*C*, orange). The estimator proposed by Adolph and Hardin, $\hat{\rho}^2_0$, is less biased (red trace) but still consistently underestimates the true r^2 and does not improve with m. The

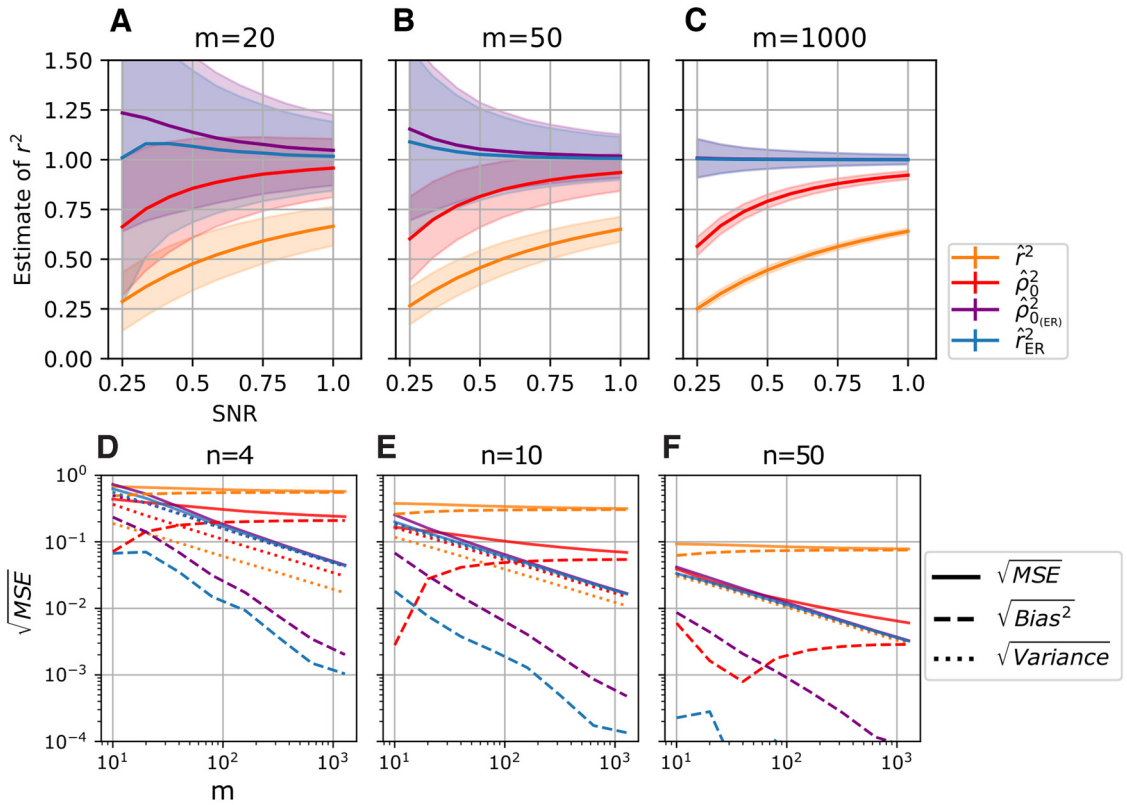

Figure 2. Comparison of $\hat{r}_{ER}^2$ to alternative estimators in simulation. ***A–C***, The mean and SD (shaded area) of four r^2 estimators are plotted as a function of SNR (equal in both neurons) and number of stimuli, m, in simulations of two sets of neuronal responses, each with $n = 4$ repeats, and true correlation is $r^2 = 1$ (see above, Materials and Methods, Simulation procedure) for 50,000 simulations. In some cases, for low SNR, the SD (shaded area) exceeds the range of the plot because the estimators become unstable so that even the 95% CI of the mean (vertical bars) remains high. In orange is the naive estimator Pearson's r^2, in red is the estimator of Spearman (1904) calculated according to the methods of Adolph and Hardin (2007), in purple is Spearman's estimator computed using an unbiased estimator of dynamic range (Eq. 18), and in blue is the estimator we use throughout the article, $\hat{r}_{ER}^2$. ***A***, Low number of stimuli (m). ***B***, Intermediate number of stimuli. ***C***, High number of stimuli. ***D–F***, The square root of the MSE (solid line), bias^2 (dashed line), and variance (dotted line) are plotted for each estimator. Here, SNR is fixed (SNR = 0.5), whereas m is varied in each plot, and n is varied across plots. ***D***, Low number of repeats. ***E***, Intermediate number of repeats. ***F***, High number of repeats.

same estimator computed using an unbiased estimate of the dynamic range, $\hat{\rho}_{0_{(ER)}}^2$ (Fig. 2, purple trace), shows an upward but, typically, smaller bias relative to the downward bias of $\hat{r}^2$ and $\hat{\rho}_0^2$, and it converges to the true value ($r^2 = 1$) as m increases. Finally, our $\hat{r}_{ER}^2$ (blue trace) shows the least bias of the four estimators over a wide range of SNR and m, although in the case of the lowest SNR values, it becomes highly variable (Fig. 2*A*, large blue SD at left). For higher values of m (Fig. 2*B*,*C*), this is not an issue. For very high m, our correction to Spearman's method $\hat{\rho}_{0_{(ER)}}^2$ and our estimator $\hat{r}_{ER}^2$ are essentially identical (Fig. 2*C*, blue and purple lines overlap).

Although we have found our corrected estimators can effectively reduce bias, they also typically have higher variability than the prior estimators, which can lead to overall higher mean squared error (MSE, a common measure of overall estimator accuracy). Here, we explore the trade-off between bias and variance across different numbers of repeats (n) and stimuli (m; Fig. 2, bottom row). When we set the number of repeats to be $n = 4$ (Fig. 2*D–F*), we found that for very low m, the MSE was highest for our corrected estimators and the naive estimator. Specifically, $\hat{\rho}_{0_{(ER)}}^2$ (solid purple trace) was highest, then $\hat{r}^2$ (orange trace), and then $\hat{r}_{ER}^2$ (blue). The majority of the MSE for the corrected estimators resulted from the variance of the estimators (dotted purple and blue lines, just below the solid lines of the same colors), and only a small amount was attributable to the bias of the estimators (dashed purple and blue lines). As m increases above 10, the bias of the corrected estimators drops below those of the prior estimators, but the MSE remains higher than Spearman's $\hat{\rho}_0^2$ until $m = 40$. As m increases beyond 40, the MSE of the corrected estimators drops below those of the prior estimators (blue, purple below red, orange solid traces). Thus, at this point our corrected estimators are more accurate than the uncorrected estimators despite their variance remaining higher. Although the variance of all estimators decreases with m (dotted lines decreasing), the bias of the naive estimator ($\hat{r}^2$) remains constant, whereas the corrected estimators continue to improve in terms of both MSE and bias (blue, purple decreasing while orange remains largely flat). This is because the corrected estimator estimates of trial-to-trial variability are improving with m, thus improving the accuracy of the correction. Interestingly, the MSEs of the corrected estimators are nearly identical, but the bias of $\hat{\rho}_{0_{(ER)}}^2$ is higher than $\hat{r}_{ER}^2$ throughout our simulations (blue, purple solid traces overlapped but blue dashed below purple dashed traces). When the number of repeats increases ($n = 10$; Fig. 2*E*), the performance of all estimators improves, but similarly to Fig. 2*D*, the MSE of the naive estimator $\hat{r}^2$ remains unimproved by increasing the number of stimuli. The rate at which the MSE of $\hat{\rho}_0^2$ improves with m increases (solid red line is steeper in Fig. 2*E* than in Fig. 2*D*), this is because the overestimation of dynamic range because of trial-to-trial variability (Eq. 17) is reduced by averaging across more repeats, and as m increases, the true dynamic range dominates. Finally, for a very high number of repeats ($n = 50$; Fig. 2*F*), the bias of all estimators (dashed lines) is low because the effect of trial-to-

trial variability has largely been averaged out. Overall, with a high number of repeats, the difference between the estimators becomes small, but when the number of repeats is low, the corrected estimators ($\hat{\rho}^2_{0_{(ER)}}$ and $\hat{r}^2_{ER}$) overcome their higher variability to outperform the prior estimators ($\hat{r}^2$ and $\hat{\rho}^2_0$) in terms of both MSE and bias over a wide range of m.

Split trial validation for V4 data

We have shown that our estimator, $\hat{r}^2_{ER}$, works well under simulations that follow the assumptions of its derivation. Because neuronal responses are not guaranteed to be well approximated by these assumptions, it is important to test the method on real neuronal data. In most cases, we do not know r^2_{ER} and thus cannot determine whether the estimator is working. One case in which we do know r^2_{ER} is across independent trials of responses of a neuron to the same set of stimuli. Theoretically, the expected response values are the same for two tuning curves computed from alternating stimulus repeats recorded from a single neuron, thus $r^2_{ER} = 1$. Here, we correlate the mean responses from the odd trials to those from the even trials for a set of single neurons with the hope that our method correctly estimates $r^2_{ER} = 1$. Our neuronal dataset consists of single-unit recordings from 96 V4 neurons with $n = 4$ repeats and $m = 371$ stimuli (see above, Materials and Methods, Electrophysiological data). We evaluate our estimators on the first and third trial correlated to the second and fourth trial.

We first computed the naive $\hat{r}^2$ for all neurons (Fig. 3*A*, orange points) and plotted these values as a function of the estimated SNR ($\hat{d}^2/m\hat{\sigma}^2$, across all trials). Despite the theoretical value of $r^2_{ER} = 1$, the unit with the highest SNR (rightmost orange point) achieves only $\hat{r}^2 = 0.70$. Thus, for real neural data with low n, the naive $\hat{r}^2$ clearly underestimates r^2_{ER}. In contrast, the scatter of $\hat{r}^2_{ER}$ values across the population of cells is approximately centered around one (Fig. 3*A*, blue points), with an increasing upward bias for low SNR. The highest SNR units (rightmost blue points) have estimates that are close to one. For neurons with lower SNR, some of the $\hat{r}^2_{ER}$ values become unstable, taking on values far >2 (Fig. 3*A*, small filled blue points, top left) or <0 (small filled blue points, bottom left) consistent with our simulations (Fig. 2*A*, shaded blue SD on left). The population average ($\hat{r}^2_{ER} = 3.04$) is made unstable by these low SNR neurons (maximum value $\hat{r}^2_{ER} = 213.3$). The median is a more appropriate method to aggregate results across a population that includes low SNR neurons with few repeats (median $\hat{r}^2_{ER} = 1.07$).

For two example neurons, plots of the raw tuning curve values on even versus odd trials (Fig. 3) provide deeper intuition into the metric. Each point in the scatter plots represents the mean response to a particular shape on the even trials versus that on the odd trials. In both plots, by the split-half construction, all the residual variance is attributable to trial-to-trial variability. The estimator $\hat{r}^2_{ER}$, being ~1.0 in both cases, is appropriately factoring out the trial-to-trial variability and is in essence predicting that with more trials, one should expect these points to settle onto a line.

Overall, our validation on simulated data demonstrates that our corrected $\hat{r}^2_{ER}$ avoids the large bias of the naive estimator, $\hat{r}^2$, and outperforms a previously proposed metric. Furthermore, our metric appears able to achieve an approximately veridical result for a population of real neuronal data in a split-half test that had very low n (only two repeats) and thus substantial noise.

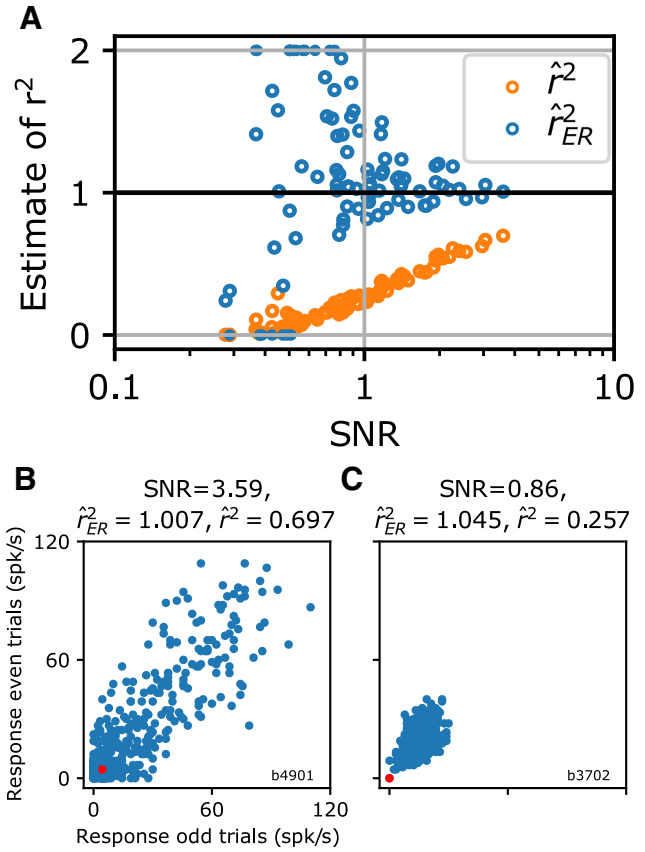

Figure 3. Motivation and validation of $\hat{r}^2_{ER}$ on split-half correlation for V4 neural responses from Pasupathy and Connor (2001). ***A***, Estimated correlation between the odd and even trials of the responses of each neuron to a set of $m = 371$ stimulus conditions, 96 neurons total, is plotted as a function of the SNR of each neuron calculated across all trials. In blue are the corrected $\hat{r}^2_{ER}$ and in orange the naive $\hat{r}^2$. Theoretically, $r^2_{ER} = 1$ (bold black horizontal line). Corrected estimates that went above two are set to two, or below zero are set to zero. ***B***, Example cell with high SNR. Each point is the average of the even (*x*-axis) and odd (*y*-axis) spike rate for a single stimulus. In red is the baseline firing rate. SNR is high because of a large dynamic range in spike count. ***C***, Example cell with smaller SNR. Despite $\hat{r}^2$ being lower in this case, our estimator predicts that similar to ***B***, the neural responses on even and odd trials have the same pattern of means.

We found that low SNR neurons tended to have unstable estimates. This is because for low SNR, the distribution of denominator estimates shifts toward zero. When estimating population level correlation this can be ameliorated by calculating the median (see below, Measuring population translation invariance). When an analysis focuses on correlation at the single-neuron level, truncation of the estimate of $\hat{r}^2_{ER}$ to lie in [0, 1] can reduce variability and improve interpretability, but the variability may be so high that these estimates provide little information. Use of confidence intervals is a principled approach to determining whether variability is in fact too high (see below, Measuring invariance of shape tuning to a change in a surface property). Next, we apply our method to estimate relationships in neuronal tuning where no ground truth exists.

Measuring population translation invariance

Translation invariance of neuronal selectivity is the degree to which a neuron maintains the same pattern of responses to a stimulus set regardless of where the stimulus set is presented in its RF. This can be quantified by measuring the correlation between the responses at one position to those at another. The naive $\hat{r}^2$ will suffer from the confounds described above. To avoid this, we use $\hat{r}^2_{ER}$ to accurately assess translation invariance

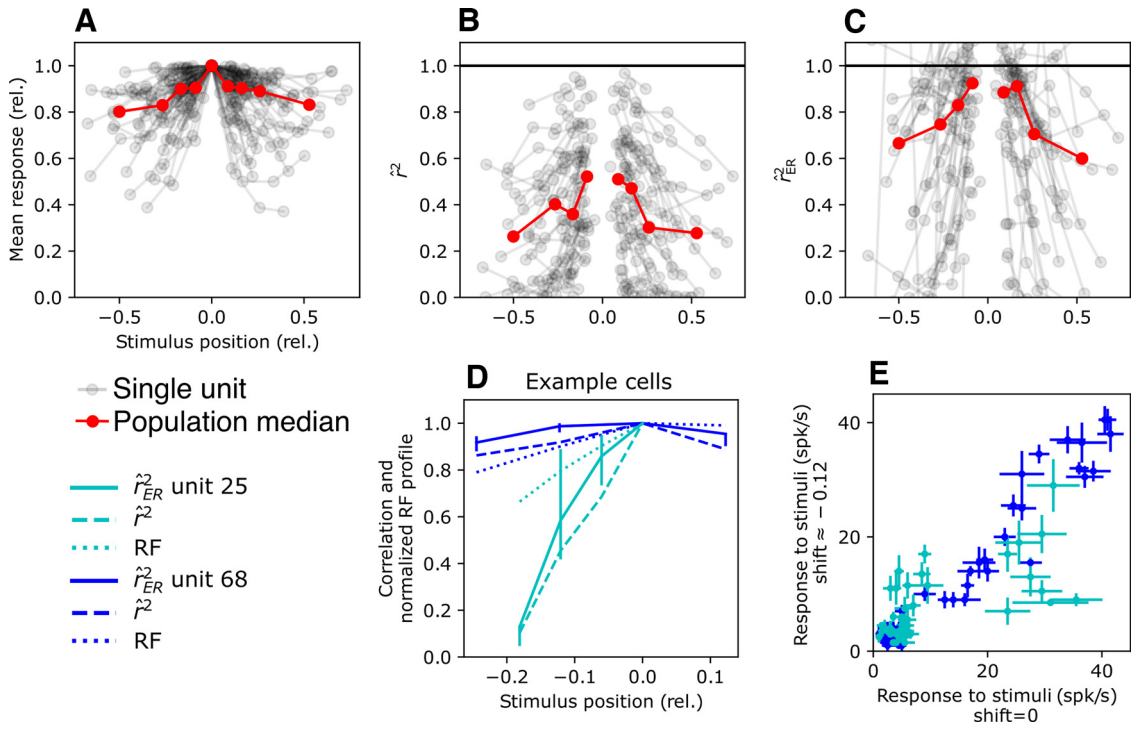

Figure 4. RFs and translation invariance across 80 V4 cells from El-Shamayleh and Pasupathy (2016). ***A***, Average response (across all stimuli) at each RF position tested (one gray line per neuron). Each trace is normalized to the peak average response defined to be the center of RF, and the *x*-axis is the position relative to the RF center, scaled by the estimated RF diameter. Red dots show the median across the population in nine nonoverlapping bins containing ~34 samples each except at the center of the RF, points are at the center of bins. ***B***, The $\hat{r}^2$ between responses at the RF center versus those at each offset position (1 trace per neuron; red shows median with same binning as ***A***). Center correlation is by definition 1. ***C***, Similar to ***B***, but correlation is computed with $\hat{r}^2_{ER}$. ***D***, Comparison of example cells chosen for low and high TI (both with high SNR) and with measurements at similar positions in the RF. The cell with high TI (dark blue) maintains a high $\hat{r}^2$ (dash line) across the RF (dotted line), and $\hat{r}^2_{ER}$ remains near 1 (solid line), reflecting near perfect invariance. For a second example cell (cyan), $\hat{r}^2$ drops off quickly and $\hat{r}^2_{ER}$ is similar; thus, the drop is not the result of noisy responses. ***E***, For the same two example cells, responses to each shape stimulus (1 point per shape) presented at shift −0.12 is plotted against the responses at RF center, along with SE bars. Unit 68 (dark blue) shows a strong linear relationship reflecting its high $\hat{r}^2$, whereas for unit 25 (cyan), one subset of stimuli evokes a higher response at the shifted position (higher points toward left) than at the center, and another subset (cyan, lower right) evokes a lower response at the shifted position; thus, tuning is clearly changing with position, consistent with the lower $\hat{r}^2_{ER}$ at −0.12 for this neuron in ***D***.

in V4. A population level estimate of invariance found by calculating the median across single-neuron estimates will inherit the low bias of the individual samples and thus will provide a low bias estimate of typical translation invariance in V4.

We reanalyzed data from El-Shamayleh and Pasupathy (2016), which consisted of recordings of 80 V4 neurons responding to a set of simple shapes presented at up to four positions within the RF (see above, Materials and Methods, Electrophysiological data). One potential concern is that the RFs of the neurons are diverse (Fig. 4*A*, gray traces); for some neurons the average response (across shapes) falls off quickly over space, whereas for others the mean remains high. Half of the neurons are at less than ~80% of mean response for the farthest stimuli (Fig. 2*A*, red trace). Given the influence of SNR on $\hat{r}^2$ (Fig. 2*A*–*C*, top row, orange traces), it is quite possible that the naive $\hat{r}^2$ would display a falloff in correlation simply because the SNR was lower farther from the center of the RF. Indeed, $\hat{r}^2$ drops off sharply from the center of the RF (Fig. 4*B*), falling from a value of 1 (by definition) at the center to a median of 0.52 at the position immediately to the left of center and to 0.51 at the point immediately to the right of center. The median $\hat{r}^2$ value continues to drop further as the stimulus is shifted to the edges of the range tested (0.23 and 0.27 at left and rightmost points, respectively). Knowing that this naive metric is biased downward by noise and depends on SNR raises the distinct possibility that such a sharp drop would not be observed if more trials were collected. To overcome this confound, we used our corrected metric $\hat{r}^2_{ER}$. The resulting single-neuron traces and their medians are shifted upward (Fig. 4*C*). This reveals that for the smallest shifts away from the center, tuning remains quite similar (e.g., median $\hat{r}^2_{ER}$ = 0.92 left of center). Thus, the initial steep dropoff in correlation (Fig. 4*B*) is attributable to trial-to-trial variability. On the other hand, a substantial dropoff in tuning correlation remains for the largest offsets tested.

By focusing on neurons with a high SNR, we can find units that are truly invariant, and we can also identify units that have tuning that changes reliably with position in the RF. For example, consider a high-SNR neuron (neuron 68, which had high d^2; Eqs. 1, 2; thus was strongly tuned for shapes) that had high translation invariance (Fig. 4*D*, solid blue trace is relatively flat) versus another (neuron 25) with tuning that is more sensitive to position (solid cyan trace falls off sharply). For the same fraction estimated RF shift (−0.12), one cell has near-perfect invariance, despite the RF sensitivity dropping off by ~11% (blue dotted trace), whereas the second example cell shares only half the variance with the tuning at the center of the RF (cyan trace, RF sensitivity drops ~20%). The important contribution of the $\hat{r}^2_{ER}$ estimator is that it removes ambiguity about whether such differences are caused by changes in firing rate over position. Examining the responses at positions 0 and −0.12 across individual shapes (Fig. 4*E*, each point corresponds to one shape), it can be seen that there is a strong linear relationship for neuron 68 (blue). Intuitively, any deviation from a line is attributable to

trial-to-trial variability because a line drawn through these points could intersect the majority of SE bars. For neuron 25 (cyan), however, there is a substantial change in tuning; many shapes with points previously in the cluster near the origin have more than doubled their responses at the shifted position, whereas for points to the right of 20 spikes/s, the average response to these shapes has been roughly halved. SE bars on these points indicate that it is unlikely their deviation from a line of best fit is solely because of trial-to-trial variability.

In summary, $\hat{r}^2_{ER}$ can distinguish with confidence between units that are sensitive versus insensitive to stimulus position and can provide accurate estimates of translation invariance across a population without needing to resort to removing all but the most well-tuned neurons.

Measuring invariance of shape tuning to a change in a surface property

We now demonstrate how our metric along with its associated CIs) can be used to accurately and systematically quantify another important visual invariance, one that relates to the representation of boundary versus surface features. The neuronal data comes from a study (Popovkina et al., 2019) in which V4 neurons were tested with a set of shapes presented as either outlines (shape interior matched the background) or fills (interior was painted the same as the boundary) with the goal of determining whether responses were dictated by boundary shape alone or were dependent on interior fill. It was hypothesized that neuronal responses would have a high correlation across the transformation from fill to outline—high fill-outline invariance—if neuronal shape tuning was largely based on boundary information, but the authors found that the r values tended to be low, suggesting that most neurons did not show strong fill-outline (FO) invariance. This is an intriguing observation because many models differ strongly from V4 in this respect; for example, an H-Max model (Cadieu et al., 2007) fit to V4 shape tuning was found to have very high FO invariance (Popovkina et al., 2019). But to be able to accurately compare noisy neuronal data with noise-free models, we need a way to correct for noise.

Our $\hat{r}^2_{ER}$ estimator can address this problem, but it is important to consider that the added variability in the corrected estimate can be confounding as it will tend to smear out the population distribution of invariance. To this end, we can make use of confidence intervals to select only those cells with reliable $\hat{r}^2_{ER}$ values. For this subset, we can be confident that the empirical distribution of invariance is accurate and not degraded by noise. This criterion naturally removes neurons that were not significantly tuned for one or both sets of stimuli; such neurons with measurable tuning for only one condition can be considered a special case of noninvariance.

We computed fill-outline $\hat{r}^2_{ER}$ values and associated confidence intervals for 42 V4 neurons and found that nearly 40% of the cells had CIs so wide they spanned the full [0, 1] range of possible values for r^2 (Fig. 5*A*, rightmost bar). These neurons correspond to the observation in the original study that many neurons failed to show tuning for one or the other of the fill or the outline stimulus sets, and this lack of tuning means that the SNR is too low to allow for accurate estimation of the correlation value.

The cumulative distribution of the naive $\hat{r}^2$ across all neurons (Fig. 5*B*, orange trace) climbs quickly along the left side of the plot, suggesting that high fill-outline invariance is rare (median $\hat{r}^2 = 0.02$). In comparison, the distribution of $\hat{r}^2_{ER}$ values for the neurons with narrow confidence intervals (Fig. 5*B*, blue line, CI < 0.5) is shifted substantially rightward (median $\hat{r}^2_{ER} = 0.38$) and becomes broader (slope of blue line is shallower). In this distribution, neurons that reliably change their pattern of selectivity across the fill and outline conditions are indicated by low values, whereas neurons for which selectivity for either fill or outline stimuli is absent or unmeasurable are now excluded. To allow a direct comparison between the naive and corrected correlation values for the same set of units, the green line (Fig. 5*B*) shows the distribution of the naive estimator for just those units with narrow CIs. From this comparison, we conclude that units with measurable tuning for both fill and outline stimuli tended to sit on the higher end of the original naive distribution (green curve lies to the right of orange curve), and that our corrected values are substantially larger than the naive values (blue lies to the right of green).

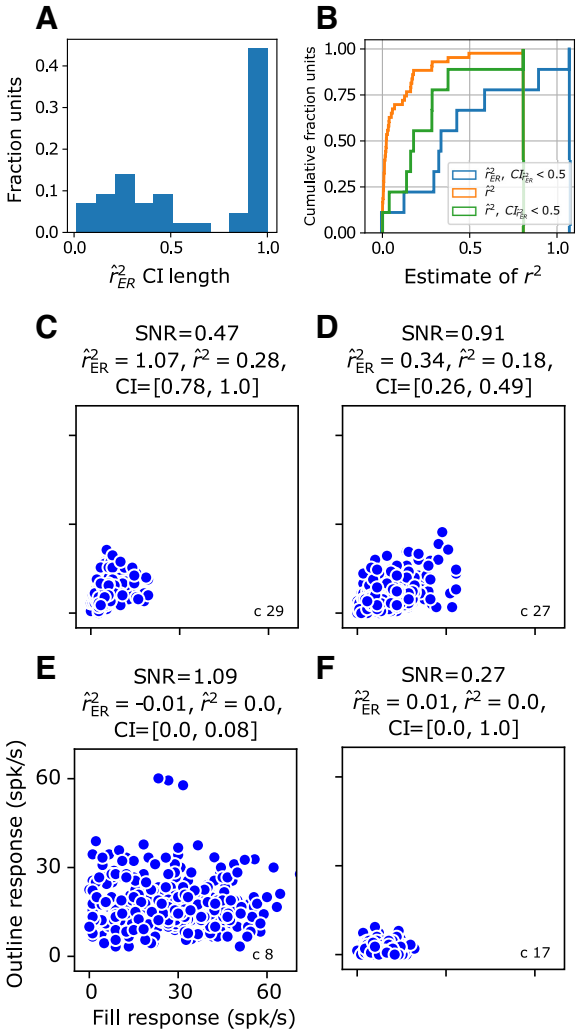

Figure 5. Comparison of $\hat{r}^2$ and $\hat{r}^2_{ER}$ for determining strength of correlation between responses to fill and outline shapes in V4 neurons. ***A***, The distribution of length of 90% confidence intervals. ***B***, Cumulative distribution of $\hat{r}^2$ (orange), of $\hat{r}^2_{ER}$ (blue) when its associated confidence interval length was <0.5 ($n = 17$ of the original 42 neurons), and of $\hat{r}^2$ for the same 17 neurons (green). ***C***, Example neuron with lower SNR but high fill-outline invariance. On the basis of $\hat{r}^2 = 0.28$, it might be assumed that fill-outline invariance is low, but given the high noise level and low dynamic range, this amount of correlation is quite high, and confidence intervals suggest the true invariance is upward of 0.7. ***D***, Example cell with high SNR but middling fill-outline invariance. ***E***, Example neuron with higher SNR but little fill-outline invariance. ***F***, Example neuron with low SNR and confidence interval of length 1. Typical of units with long CIs, this neuron evokes significant modulation for only one set of stimuli (in this case filled shapes).

Example neurons can provide some intuition into the corrected values. Figure 5*C* plots outline versus fill responses for a cell with middling $\hat{r}^2$ but high $\hat{r}^2_{ER}$ (values shown at top). The latter value indicates that the observed variation in the scatter plot is consistent with what would be expected given the trial-to-trial noise for this neuron, assuming it had identical fill and outline tuning. Cells with middling to low $\hat{r}^2_{ER}$ (Fig. 5*D*,*E*) are those where some of the variance, but not all, is attributable to noise. Finally, a cell with low SNR (Fig. 5*F*) had a confidence interval covering [0, 1] consistent with the observation that it responded significantly only to filled shapes and not to outlines.

In summary, by focusing on $\hat{r}^2_{ER}$ estimates with small confidence intervals, we were able to show that the population of neurons with significant tuning for both fill and outline stimuli displays a broad spectrum of fill-outline invariance and has a higher average invariance than would have been predicted based on a naive analysis of all neurons.

Nonvision example of correcting correlation between tuning curves

Our correction is important not only for studies of invariance in sensory areas of vision, audition, somatosensation, and olfaction (for references, see below, Discussion) but is highly relevant in any area of neuroscience where the similarity of noisy neural responses across experimental conditions or across neurons is examined. To provide a clear, concrete example from another field, we explain below how the study of hippocampal place field remapping is particularly susceptible to the confounds of noise.

Place fields are characterized by calculating firing rate tuning curves for neurons in the hippocampus with respect to the position of an animal in an enclosure. Much interest has focused on how these place fields change depending on the spatial and sensory context of the animals; the change has been termed remapping. To quantify the degree of remapping, the naive Pearson correlation coefficient has been used for more than three decades (Muller et al., 1987; Muller and Kubie, 1987; Bostock et al., 1991; Quirk et al., 1992; Skaggs and McNaughton, 1998; Anderson and Jeffery, 2003; Lee et al., 2004; Leutgeb et al., 2005, 2007; Fenton et al., 2008; Alme et al., 2014; Rueckemann et al., 2016; Fetterhoff et al., 2021), despite the fact that these experiments are especially susceptible to the confounds of trial-to-trial variability, which we have discussed above. The experimental conditions in this case are the discretized regions of the space that the animal explores, and this can easily be quite high (e.g., $m \approx 450$ for an enclosure with a diameter of 76 cm Muller et al., 1987). At the same time, the number of repeats (n) is determined by the number of times an animal crosses through the same region of space, thus putting a strong limit on how much averaging can be performed. Furthermore, place fields are sparse, only a few regions in space will evoke responses; thus SNR is suppressed by the large regions over which the neurons are silent. Commensurate with these structural difficulties, studies find that correlation of rate maps in the same experimental context (same enclosure in the same location) can be very low.

Fenton et al. (2008) performed a seminal study of remapping by estimating place field maps in a large square-shaped enclosure and smaller cylinder within the square enclosure. The authors found that the correlation of firing rate maps was lower between the cylinder and square than it was between rate maps recorded in the cylinder at different times. Yet they found that this latter quantity, the correlation of a place field map to itself at a later time, was on average only $r \approx 0.34$ ($r^2 \approx 0.11$; Fenton et al. 2008, their Fig. 2*B*). Assuming that place fields remain the same over short time periods (60 min in this case), a common assumption (Muller et al., 1987; Latuske et al., 2017), it follows that the vast majority of variance in these tuning curves is noise, and thus the SNR is very low. Furthermore, the distribution of these rate map self-correlation values across neurons is broad, being uniformly distributed between 0 and 0.8 (Fenton et al. 2008, their Fig. 2*B*, inset). At the same time, the distribution of map correlations between the cylinder and square are also broad with a substantial fraction exceeding $r = 0.34$. The diversity in these correlation distributions could be explained solely by differences in SNR, or there may be subsets of neurons whose rate maps change rapidly over time and others that remain the same between enclosures. Reanalysis of this dataset using our metric could answer this intriguing question.

In general, we find that most studies, including place-field remapping studies, account for the confound of noise by either minimizing the effect of noise (e.g., using longer recording times, fewer stimuli, and collecting many repeats) or by restricting their conclusion to relative statements (e.g., tuning curves or maps are more similar in this condition than that condition). In the former case, an inordinate amount of repeats are collected when those trials could have been used to increase the number of experimental conditions by applying our estimator. In the latter case, results often remain ambiguous because of the possibility of changes in SNR across conditions, comparisons across studies are difficult, and there is no absolute measure of similarity between tuning curves, restricting the granularity of results.

Discussion

We have developed a new estimator, $\hat{r}^2_{ER}$, for the fraction of variance shared between the expected value of two neural tuning curves. This estimator is broadly applicable to studies of neural invariance in fields beyond vision, including somatosensation (DiCarlo and Johnson, 1999; Weber et al., 2013), audition (Billimoria et al., 2008; Kell and McDermott, 2019), and olfaction (Cleland et al., 2007; Wilson et al., 2017). Additionally, it is relevant to studies of signal correlation (Gawne and Richmond, 1993; Cohen and Kohn, 2011; Pospisil and Bair, 2021a) and the stability of neuronal representations over time (Rokni et al., 2007; Margolis et al., 2012; Lütcke et al., 2013; Clopath et al., 2017; Deitch et al., 2021; Schoonover et al., 2021). Our estimator has significantly less bias than a previous, widely used (but underused in neuroscience) method of accounting for noise-attenuated correlation. We demonstrated how it avoids ambiguity and confounds that can meaningfully change conclusions in two neural datasets, particularly in comparison with the commonly used, naive $\hat{r}^2$. Our analysis, which provides the first noise-corrected estimates of translation invariance in V4, showed that invariance was typically grossly underestimated by $\hat{r}^2$. For fill-outline invariance, we demonstrated how confidence intervals on $\hat{r}^2_{ER}$ provide a criterion for including data points only when estimates are reliable. Finally, we document how our estimator could qualitatively change results on place field remapping, thus demonstrating its broader significance. Our metric with associated confidence intervals provides a novel tool for future studies that seek to avoid confounds of noise when quantifying the similarity between neural tuning curves.

Interpretation of estimator

We have developed an estimator of correlation between the true tuning curves of two neurons. Conceptually, this is the

correlation between two estimated tuning curves if the experimentalist were able to collect an unlimited number of repeats, average them, and there were no changes in experimental conditions as the experiment stretched on into infinity. In practice, however, experimental stability is difficult to maintain for long periods (e.g., electrode position may drift). Thus a key benefit of our estimator is that fewer repeats can be collected, shortening the experiment and allowing correlation to be estimated under tighter experimental control.

Extensions of the estimator

The correlation coefficient is the basis of several multivariate data analysis methods, including principal component analysis (when the variance is factored out), canonical correlation analysis (for review see Zhuang et al., 2020), and representational similarity analysis (for review see Kriegeskorte et al., 2008). Future work should examine the effect that the downward bias of noise on correlation has on the estimation of these quantities and whether applying a corrected estimator like the one developed here can improve inference.

Here, our derivations have assumed that all responses are independent (e.g., no correlation between responses to different stimuli). In some cases this assumption is incorrect. For example, studies have used correlation to assess the similarity of spike trains (de la Rocha et al., 2007), sometimes for invariance (Metzen et al., 2016). In this case, the neural response is not a set of values over stimuli, but a set of values over time. Neighboring responses in time are most likely correlated because of temporal correlation in the response of the system. Our method can be adapted to handle this type of comparison, and in a forthcoming publication we will extend it to arbitrary trial-to-trial correlations.

Translation invariance

We quantified translation invariance for a V4 population as the correlation between shape tuning at the RF center versus that at offset positions. Despite translation invariance being a key component of ventral stream models (Fukushima, 1980; Olshausen et al., 1993; Salinas and Abbott, 1997; Riesenhuber and Poggio, 1999; Ullman and Soloviev, 1999), including models specific to V4 (Pasupathy and Connor, 2001; David et al., 2006; Cadieu et al., 2007; Sharpee et al., 2013), only two studies have systematically mapped V4 RFs and presented large stimulus sets at multiple positions (Nandy et al., 2013; El-Shamayleh and Pasupathy, 2016). The number of stimulus conditions required in such experiments calls for a limited number of repeats in awake preparations, causing trial-to-trial variability to dramatically obscure tuning invariance. Correction for noise is thus crucial in studies requiring a large number of unique stimuli. Using our estimator, future studies can more accurately quantify translation invariance across the ventral stream by more densely sampling the RF with more diverse stimuli at the expense of fewer repeats.

Fill-outline invariance

We also applied our estimator to more accurately assess the invariance of V4 shape selectivity across a change in a surface property. We quantified fill-outline invariance as the correlation in the tuning curves of a neuron to a set of reference shapes with filled interiors versus the same set of shapes drawn only with outlines. Similarly to translation invariance, and consistent with the conclusions of the original study, fill-outline invariance varied substantially across neurons in V4. We found that an analysis based on the naive $\hat{r}^2$ would give the impression that very few units had strong fill-outline invariance, but our metric makes clear that V4 units uniformly span a range of invariance, from cases where the tuning curves for fills and outlines can be orthogonal to cases where shape selectivity remains unchanged. Having an estimate of the distribution of fill-outline invariance is important for making comparisons against models of form processing. In particular, the H-Max model (Cadieu et al., 2007) has nearly perfect fill-outline invariance (Popovkina et al., 2019), and the deep convolutional neural network known as AlexNet (Krizhevsky et al., 2017) has very high fill-outline invariance (Bair et al., 2019). To argue that the invariance of shape selectivity to surface features is much weaker in a critical stage of the ventral pathway than it is in networks considered to be state-of-the-art form processing models, it is critical to factor out the possibility that any lack of invariance in noisy neural data are not simply the result of noise. This is what $\hat{r}^2_{ER}$ allows us to do. Further research is needed to understand how a variety of invariances in form representation at different stages along the ventral stream compare with those in artificial visual systems.

More generally, understanding how invariance in any cortical representation differs from that in models is an important step in validating models, and our metric is critical to factor out low correlation values caused by noise. It allows one to distinguish neurons that have high invariance from those that have large changes in tuning, even in the face of of noise.

Other methods of quantifying neural invariance

Prior studies have used metrics of invariance in addition to $\hat{r}^2$. For example, the separability index (Peña and Konishi, 2001; Hinkle and Connor, 2002; Mazer et al., 2002) is a common alternative for assessing invariance (Brincat and Connor, 2004; Janssen et al., 2008; Li et al., 2009; Rust and DiCarlo, 2012; El-Shamayleh and Pasupathy, 2016). In the context of invariance, this metric is the fraction of neural tuning variance explained by a fixed tuning profile that changes only in amplitude across stimulus transformations (e.g., translation). This metric usefully provides a summary of the invariance across many stimulus transformations, but it is unclear how to correct it for noise. Furthermore, for noiseless responses, it is biased to report higher invariance for sparser RFs or selectivity (Pospisil et al., 2018). As an alternative, we proposed a metric that is essentially the average correlation of responses across position, weighted by the magnitude of responses (Eq. 6; Pospisil et al., 2018). In that study, we developed an ad hoc method for correcting for the effect of noise, but that estimator is amenable to the more principled techniques of removing bias used here.

The separability index and $\hat{r}^2_{ER}$ both quantify invariance by the strength of the linear relationship between tuning curves. Other authors have taken a more normative approach by quantifying invariance in terms of the decodability of stimulus parameters from neural activity despite transformations of the stimuli (Hung et al., 2005; Li et al., 2009; Rust and DiCarlo, 2010; Carlson et al., 2011; Rabinowitz et al., 2013). This approach doesn't account for the influence of noise, for example, a population of neurons could have poor decoding performance because they are noisy or because their tuning changes. It also depends on the choice of decoder. Another approach by Tovee et al. (1994) avoids choosing a decoder by estimating the amount of information about the position versus the identity of stimuli contained in the responses of inferior temporal cortex (IT) neurons. Less information about position indicates

more invariance. Despite the appealing generality of information as a measure of invariance (i.e., it makes no assumptions about the decoding of neural responses), this estimator, like $\hat{r}^2$, is biased when the number of repeats is low (Fig. 1; Tovée et al., 1993). Thus, future work is needed to address the effects of sampling noise in metrics of invariance that go beyond the classic r^2.

References

Adolph SC, Hardin JS (2007) Estimating phenotypic correlations: correcting for bias due to intraindividual variability. Funct Ecology 21:178–184.

Alme CB, Miao C, Jezek K, Treves A, Moser EI, Moser MB (2014) Place cells in the hippocampus: eleven maps for eleven rooms. Proc Natl Acad Sci U S A 111:18428–18435.

Anderson MI, Jeffery KJ (2003) Heterogeneous modulation of place cell firing by changes in context. J Neurosci 23:8827–8835.

Bair W, Bigelow AW, Popovkina DV, Pospisil DA (2019) Comparing simple invariances in V4 to those in a deep neural network. Society for Neuroscience, Abstracts 488.27. Available at https://nprcresearch.org/primate/sfn.php.

Beaton GH, Milner J, Corey P, McGuire V, Cousins M, Stewart E, de Ramos M, Hewitt D, Grambsch PV, Kassim N, Little JA (1979) Sources of variance in 24-hour dietary recall data: implications for nutrition study design and interpretation. Am J Clin Nutr 32:2546–2559.

Billimoria CP, Kraus BJ, Narayan R, Maddox RK, Sen K (2008) Invariance and sensitivity to intensity in neural discrimination of natural sounds. J Neurosci 28:6304–6308.

Bostock E, Muller RU, Kubie JL (1991) Experience-dependent modifications of hippocampal place cell firing. Hippocampus 1:193–205.

Box GEP, Cox DR (1964) An analysis of transformations. J R Stat Soc Series B Stat Methodol 26:211–243.

Brincat SL, Connor CE (2004) Underlying principles of visual shape selectivity in posterior inferotemporal cortex. Nat Neurosci 7:880–886.

Bushnell BN, Harding PJ, Kosai Y, Bair W, Pasupathy A (2011) Equiluminance cells in visual cortical area V4. J Neurosci 31:12398–12412.

Cadieu C, Kouh M, Pasupathy A, Connor CE, Riesenhuber M, Poggio T (2007) A model of V4 shape selectivity and invariance. J Neurophysiol 98:1733–1750.

Carlson T, Hogendoorn H, Fonteijn H, Verstraten FAJ (2011) Spatial coding and invariance in object-selective cortex. Cortex 47:14–22.

Cleland TA, Johnson BA, Leon M, Linster C (2007) Relational representation in the olfactory system. Proc Natl Acad Sci U S A 104:1953–1958.

Clopath C, Bonhoeffer T, Hübener M, Rose T (2017) Variance and invariance of neuronal long-term representations. Phil Trans R Soc B 372:20160161.

Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. Nat Neurosci 14:811–819.

David SV, Hayden BY, Gallant JL (2006) Spectral receptive field properties explain shape selectivity in area V4. J Neurophysiol 96:3492–3505.

Deitch D, Rubin A, Ziv Y (2021) Representational drift in the mouse visual cortex. Curr Biol 31:4327–4339.e6.

DiCarlo JJ, Johnson KO (1999) Velocity invariance of receptive field structure in somatosensory cortical area 3b of the alert monkey. J Neurosci 19:401–419.

Efron B, Tibshirani RJ (1994) An Introduction to the Bootstrap. CRC Press.

El-Shamayleh Y, Pasupathy A (2016) Contour curvature as an invariant code for objects in visual area V4. J Neurosci 36:5532–5543.

Fenton AA, Kao HY, Neymotin SA, Olypher A, Vayntrub Y, Lytton WW, Ludvig N (2008) Unmasking the CA1 ensemble place code by exposures to small and large environments: more place cells and multiple, irregularly arranged, and expanded place fields in the larger space. J Neurosci 28:11250–11262.

Fetterhoff D, Sobolev A, Leibold C (2021) Graded remapping of hippocampal ensembles under sensory conflicts. Cell Rep 36:109661.

Fukushima K (1980) Neocognitron: a self organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. Biol Cybern 36:193–202.

Gawne TJ, Kjaer TW, Hertz JA, Richmond BJ (1996) Adjacent visual cortical complex cells share about 20% of their stimulus-related information. Cereb Cortex 6:482–489.

Gawne TJ, Richmond BJ (1993) How independent are the messages carried by adjacent inferior temporal cortical neurons? J Neurosci 13:2758–2771.

Haefner RM, Cumming BG (2008) An improved estimator of variance explained in the presence of noise. Adv Neural Inf Process Syst 2008:585–592.

Hinkle DA, Connor CE (2002) Three-dimensional orientation tuning in macaque area V4. Nat Neurosci 5:665–670.

Hung CP, Kreiman G, Poggio T, DiCarlo JJ (2005) Fast readout of object identity from macaque inferior temporal cortex. Science 310:863–866.

Janssen P, Srivastava S, Ombelet S, Orban GA (2008) Coding of shape and position in macaque lateral intraparietal area. J Neurosci 28:6679–6690.

Kell AJE, McDermott JH (2019) Invariance to background noise as a signature of non-primary auditory cortex. Nat Commun 10:3958.

Kiani R, Cueva CJ, Reppas JB, Peixoto D, Ryu SI, Newsome WT (2015) Natural grouping of neural responses reveals spatially segregated clusters in prearcuate cortex. Neuron 85:1359–1373.

Kriegeskorte N, Mur M, Bandettini P (2008) Representational similarity analysis—connecting the branches of systems neuroscience. Front Syst Neurosci 2:4.

Krizhevsky A, Sutskever I, Hinton GE (2017) ImageNet classification with deep convolutional neural networks. Commun ACM 60:84–90.

Latuske P, Kornienko O, Kohler L, Allen K (2017) Hippocampal remapping and its entorhinal origin. Front Behav Neurosci 11:253.

Lee I, Yoganarasimha D, Rao G, Knierim JJ (2004) Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3. Nature 430:456–459.

Leutgeb JK, Leutgeb S, Moser MB, Moser EI (2007) Pattern separation in the dentate gyrus and CA3 of the hippocampus. Science 315:961–966.

Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser MB (2005) Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. Science 309:619–623.

Li N, Cox DD, Zoccolan D, DiCarlo JJ (2009) What response properties do individual neurons need to underlie position and clutter "invariant" object recognition? J Neurophysiol 102:360–376.

Lütcke H, Margolis DJ, Helmchen F (2013) Steady or changing? Long-term monitoring of neuronal population activity. Trends Neurosci 36:375–384.

Margolis DJ, Lütcke H, Schulz K, Haiss F, Weber B, Kügler S, Hasan MT, Helmchen F (2012) Reorganization of cortical population activity imaged throughout long-term sensory deprivation. Nat Neurosci 15:1539–1546.

Mazer JA, Vinje WE, McDermott J, Schiller PH, Gallant JL (2002) Spatial frequency and orientation tuning dynamics in area V1. Proc Natl Acad Sci U S A 99:1645–1650.

Metzen MG, Hofmann V, Chacron MJ (2016) Neural correlations enable invariant coding and perception of natural stimuli in weakly electric fish. Elife 5:e12993.

Muller RU, Kubie JL (1987) The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. J Neurosci 7:1951–1968.

Muller RU, Kubie JL, Ranck JB (1987) Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. J Neurosci 7:1935–1950.

Nandy AS, Sharpee TO, Reynolds JH, Mitchell JF (2013) The fine structure of shape tuning in area V4. Neuron 78:1102–1115.

Olshausen BA, Anderson CH, Essen DV (1993) A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. J Neurosci 13:4700–4719.

Pasupathy A, Connor CE (2001) Shape representation in area V4: position-specific tuning for boundary conformation. J Neurophysiol 86:2505–2519.

Peña JL, Konishi M (2001) Auditory spatial receptive fields created by multiplication. Science 292:249–252.

Popovkina DV, Bair W, Pasupathy A (2019) Modeling diverse responses to filled and outline shapes in macaque V4. J Neurophysiol 121:1059–1077.

Pospisil DA, Bair W (2021a) Accounting for biases in the estimation of neuronal signal correlation. J Neurosci 41:5638–5651.

Pospisil DA, Bair W (2021b) The unbiased estimation of the fraction of variance explained by a model. PLOS Comput Biol 17:e1009212.

Pospisil DA, Pasupathy A, Bair W (2018) "Artiphysiology" reveals V4-like shape tuning in a deep network trained for image classification. Elife 7:e38242.

Power JD, Cohen AL, Nelson SM, Wig GS, Barnes KA, Church JA, Vogel AC, Laumann TO, Miezin FM, Schlaggar BL, Petersen SE (2011)

Functional network organization of the human brain. Neuron 72:665–678.
Quirk GJ, Muller RU, Kubie JL, Ranck JB (1992) The positional firing properties of medial entorhinal neurons: description and comparison with hippocampal place cells. J Neurosci 12:1945–1963.
Rabinowitz NC, Willmore BDB, King AJ, Schnupp JWH (2013) Constructing noise-invariant representations of sound in the auditory pathway. PLoS Biol 11:e1001710.
Riesenhuber M, Poggio T (1999) Hierarchical models of object recognition in cortex. Nat Neurosci 2:1019–1025.
de la Rocha J, Doiron B, Shea-Brown E, Josić K, Reyes A (2007) Correlation between neural spike trains increases with firing rate. Nature 448:802–806.
Rokni U, Richardson AG, Bizzi E, Seung HS (2007) Motor learning with unstable neural representations. Neuron 54:653–666.
Rosner B, Willett WC (1988) Interval estimates for correlation coefficients corrected for within-person variation: implications for study design and hypothesis testing. Am J Epidemiol 127:377–386.
Rueckemann J, DiMauro A, Rangel L, Han X, Boyden E, Eichenbaum H (2016) Transient optogenetic inactivation of the medial entorhinal cortex biases the active population of hippocampal neurons. Hippocampus 26:246–260.
Rust NC, DiCarlo JJ (2010) Selectivity and tolerance ("invariance") both increase as visual information propagates from cortical area V4 to IT. J Neurosci 30:12978–12995.
Rust NC, DiCarlo JJ (2012) Balanced increases in selectivity and tolerance produce constant sparseness along the ventral visual stream. J Neurosci 32:10170–10182.
Saccenti E, Hendriks MHWB, Smilde AK (2020) Corruption of the Pearson correlation coefficient by measurement error and its estimation, bias, and correction under different error models. Sci Rep 10:438.
Sahani M, Linden JFBCq (2003) How linear are auditory cortical responses? In: Advances in neural information processing systems (Becker S, Thrun S, Obermayer K, eds), pp 125–132. Cambridge, MA: MIT.
Salinas E, Abbott LF (1997) Attentional gain modulation as a basis for translation invariance In Computational neuroscience: trends in research (Bower JM, ed), pp 807–812. Boston: Springer.
Schoonover CE, Ohashi SN, Axel R, Fink AJP (2021) Representational drift in primary olfactory cortex. Nature 594:541–546.
Sharpee TO, Kouh M, Reynolds JH (2013) Trade-off between curvature tuning and position invariance in visual area V4. Proc Natl Acad Sci U S A 110:11618–11623.
Skaggs WE, McNaughton BL (1998) Spatial firing properties of hippocampal CA1 populations in an environment containing two visually identical regions. J Neurosci 18:8455–8466.
Spearman C (1904) The proof and measurement of association between two things. Am J Psychol 15:72–101.
Thouless RH (1939) The effects of errors of measurement on correlation coefficients. Br J Psychol Gen Sect 29:383–403.
Tovée MJ, Rolls ET, Treves A, Bellis RP (1993) Information encoding and the responses of single neurons in the primate temporal visual cortex. J Neurophysiol 70:640–654.
Tovee MJ, Rolls ET, Azzopardi P (1994) Translation invariance in the responses to faces of single neurons in the temporal visual cortical areas of the alert macaque. J Neurophysiol 72:1049–1060.
Ullman S, Soloviev S (1999) Computation of pattern invariance in brain-like structures. Neural Networks 12:1021–1036.
Weber AI, Saal HP, Lieber JD, Cheng JW, Manfredi LR, Dammann JF, Bensmaia SJ (2013) Spatial and temporal codes mediate the tactile perception of natural textures. Proc Natl Acad Sci U S A 110:17107–17112.
Wilson CD, Serrano GO, Koulakov AA, Rinberg D (2017) A primacy code for odor identity. Nat Commun 8:1477.
Zhuang X, Yang Z, Cordes D (2020) A technical review of canonical correlation analysis for neuroscience applications. Hum Brain Mapp 41:3807–3833.