



Population coding in mouse visual cortex: response reliability and dissociability of stimulus tuning and noise correlation

Jorrit S. Montijn^{1*}, Martin Vinck¹ and Cyriel M. A. Pennartz^{1,2}

¹ Cognitive and Systems Neuroscience, Faculty of Science, Center for Neuroscience, Swammerdam Institute for Life Sciences, University of Amsterdam, Amsterdam, Netherlands

² Research Priority Program Brain and Cognition, University of Amsterdam, Amsterdam, Netherlands

Edited by:

Arnulf Graf, California Institute of Technology, USA

Reviewed by:

Rava Azeredo Da Silveira, Ecole Normale Supérieure, France
Marianna Yanike, California Institute of Technology, USA

*Correspondence:

Jorrit S. Montijn, Swammerdam Institute for Life Sciences, Universiteit van Amsterdam, Science Park 904, Amsterdam 1098XH, Noord-Holland, Netherlands
e-mail: j.s.montijn@uva.nl

The primary visual cortex is an excellent model system for investigating how neuronal populations encode information, because of well-documented relationships between stimulus characteristics and neuronal activation patterns. We used two-photon calcium imaging data to relate the performance of different methods for studying population coding (population vectors, template matching, and Bayesian decoding algorithms) to their underlying assumptions. We show that the variability of neuronal responses may hamper the decoding of population activity, and that a normalization to correct for this variability may be of critical importance for correct decoding of population activity. Second, by comparing noise correlations and stimulus tuning we find that these properties have dissociated anatomical correlates, even though noise correlations have been previously hypothesized to reflect common synaptic input. We hypothesize that noise correlations arise from large non-specific increases in spiking activity acting on many weak synapses simultaneously, while neuronal stimulus response properties are dependent on more reliable connections. Finally, this paper provides practical guidelines for further research on population coding and shows that population coding cannot be approximated by a simple summation of inputs, but is heavily influenced by factors such as input reliability and noise correlation structure.

Keywords: population coding, calcium imaging, spatial organization, mouse, visual cortex, orientation tuning, noise correlations, signal correlations

INTRODUCTION

Stimuli in an animal's environment are encoded by the pattern of action potential firing of a neuron (Hubel and Wiesel, 1974). It remains unclear however how precisely information about the outside world is encoded and processed in the brain on a network level. Using information coded by single neurons is straightforward, but when the activity of multiple neurons is used to code information, higher-order statistics such as temporal dependencies can strongly influence neural codes (Butts et al., 2007; Ainsworth et al., 2012).

To study these general principles of population coding, it is necessary to choose a suitable model system. For many years, the primary visual cortex has been serving as such a model (Tong, 2003). Many properties of V1 appear to be preserved across mammalian species, including neuronal tuning properties, but the spatial organization of these neurons differs. For example, in primates orientation tuned neurons are organized in columns (Hubel and Wiesel, 1974), while in mouse V1 neurons that are tuned to different orientations are organized in an intermingled “salt-and-pepper” way (Hübener, 2003). Despite the apparent absence of orientation columns in mouse V1, pyramidal neurons still exhibit selectivity to many visual stimulus properties, such as retinal location, orientation or motion direction, contrast, spatial and temporal frequency, and speed (Hübener, 2003; Niell and

Stryker, 2008; Andermann et al., 2011). These factors, in combination with the relative ease to work with mice and the many options for genetic modifications, make them an excellent animal model for studying population coding of visual information.

One of the biggest challenges of population coding is understanding how higher-order neuronal populations can read out their input from primary sensory populations, especially considering the relatively high noise of neural systems (Faisal et al., 2008), and the complex spatiotemporal structures this noise can assume (correlated trial-to-trial variability between neurons that occurs regardless of any similarity in the mean response of these neurons to stimuli, i.e., noise correlations) (Cohen and Kohn, 2011). Even when disregarding noise, this reading-out—or decoding—is non-trivial, since neurons are often responsive to a plethora of stimulus properties, while higher-level populations need to extract specific information from this entangled stimulus representation (Pagan et al., 2013). In addition to the already large number of stimulus dimensions that influence responses of neurons in V1, their firing patterns can be modulated by non-stimulus related processes such as attention (Motter, 1993), wakefulness (Greenberg et al., 2008), reward-related activity (Goltstein et al., 2013), or even higher-order properties of the visual scene, such as whether an object belongs to the foreground or background of a scene (Lamme, 1995).

One way to study population coding is with the use of computational algorithms that decode information from the population activity (e.g., Zhang et al., 1998; Pillow et al., 2008; Quiroga and Panzeri, 2009; Graf et al., 2011). Comparing a decoder's performance in different situations can then provide insight into how much information is present in population spiking activity. Moreover, comparing the performance of different decoders may provide insight into how population coding works in the brain and into which variables are important. Using decoding algorithms has been particularly fruitful in the investigation of the effects that specific spatiotemporal structures of noise (noise correlations) have on population coding (Pillow et al., 2008; Graf et al., 2011), although it remains unclear whether they are beneficial (e.g., Averbeck et al., 2006; Ecker et al., 2011; da Silveira and Berry, 2013) or detrimental (e.g., Mitchell et al., 2009). Moreover, it is still unknown whether the same anatomical substrates (i.e., synaptic connections) underlie both noise correlations and stimulus tuning. It has been hypothesized for some years that noise correlations are dependent on horizontal cortico-cortical connections (Ts'o et al., 1986; Smith et al., 2012), and it has more recently been shown that L2/3 orientation tuned neurons have a higher intra-laminar connection probability with similarly tuned neurons (Ko et al., 2011). This would make it plausible that these two processes indeed share similar anatomical substrates, but this hypothesis has—to our knowledge—not yet been tested.

A number of decoding algorithms has been applied to neural data recorded from several brain regions, but to our knowledge the performance of these algorithms has not been compared in mouse primary visual cortex (V1). Especially the influence of single neurons on the decoding performance of ensembles has been studied little, even though it has been known for several years that neuronal populations in V1 show sparse coding, where single attributes in a visual scene are represented by only a very small number of neurons (Vinje and Gallant, 2000). We will therefore compare the performance of a classical population vector, template matching, variability-normalizing template matching, and Bayesian maximum-likelihood algorithm as a function of neuronal population size. We also relate their performance to their mathematical assumptions to show which variables determine the applicability of the different decoders. We performed these analyses using population data obtained by *in vivo* two-photon calcium imaging of ensemble activity in V1, which directly reflects spiking activity (Kerr et al., 2005; Greenberg et al., 2008). Furthermore, we investigate the dependence of decoding accuracy on orientation selectivity, signal correlations and noise correlations, and show the relationships between these neuronal properties. Methodological descriptions of several analysis methods are provided so they can be used as guidelines by other researchers for future investigation of population coding. The results from these analyses reveal the caveats and critical factors of applying population vector and template matching algorithms to population activity of V1 neurons, including the pivotal importance of response reliability (i.e., little variation over trials in response magnitude). Secondly, our results show that there exists an anatomical intersomatic distance dependence of noise correlations, but not of the difference in preferred visual stimulus direction between pairs of cells. These results lead us to formulate

two hypotheses for which we also present testable predictions. First, we argue that a Hebbian plasticity mechanism might lead to a normalization of neuronal responses based on their variability. Secondly, we hypothesize that noise correlations are due to global non-specific synaptic barrages in local circuits, while stimulus response properties are dependent on stronger or more reliable connections.

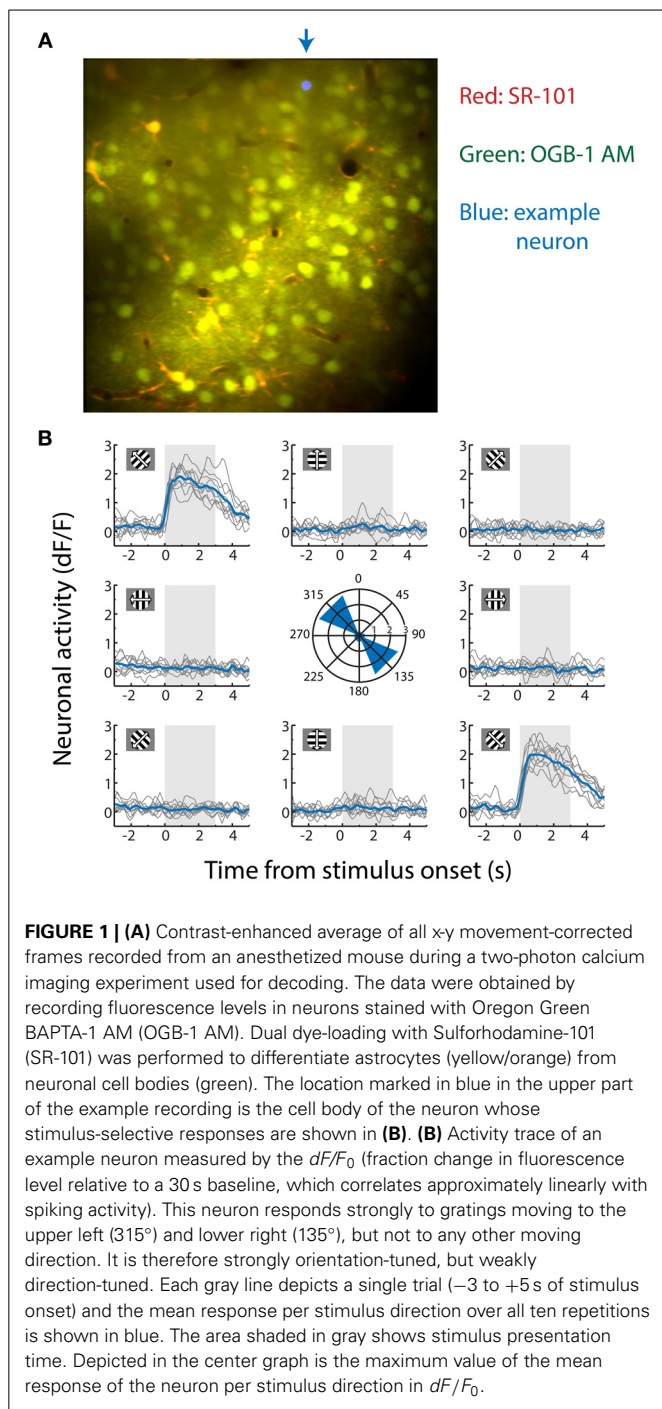
METHODS

ANIMAL PREPARATION

All experimental procedures were conducted with approval of the animal ethics committee of the University of Amsterdam and are largely similar to those described in more detail in Goltstein et al. (2013). In short, 6 adult male C57BL/6 mice (Harlan) were implanted with a titanium head bar prior to the imaging experiment and were allowed to recover for a minimum of 3 days. On the day of the two-photon calcium imaging experiment, buprenorphine (0.05 mg/kg) was injected subcutaneously 30–60 min before induction of anesthesia with isoflurane (3.0% in 100% oxygen). Intrinsic signal imaging (ISI) was performed to localize the primary visual cortex (V1) while the animal was anesthetized lightly (0.8% isoflurane). We subsequently performed a small (2 mm) craniotomy above the retinotopic area responding to visual stimulation with drifting gratings. After the craniotomy, the dura was kept wet with an artificial cerebrospinal fluid (ACSF: NaCl 125 mM, KCl 5.0 mM, MgSO₄ * 7 H₂O 2.0 mM, NaH₂PO₄ 2.0 mM, CaCl₂ * 2 H₂O 2.5 mM, glucose 10 mM) buffered with HEPES (10 mM, adjusted to pH 7.4) (Svoboda et al., 1999). After removal of the skull, multi-cell bolus loading with Oregon Green BAPTA-1 AM (OGB) and Sulforhodamine 101 (SR101) was performed 230–270 microns below the dura as previously described (Stosiek et al., 2003; Goltstein et al., 2013). After injection of the dyes, the exposed dura was covered with agarose (1.5% in ACSF) and sealed with a circular cover glass that was fixed to the skull using cyanoacrylate glue.

APPARATUS AND STIMULUS PRESENTATIONS

Dual-channel two-photon imaging recordings (filtered at 500–550 nm for OGB and 565–605 nm for SR101; see **Figure 1A**) with a 512 × 512 pixel frame size were performed at a sampling frequency of 25.4 Hz. We used an *in vivo* two-photon laser scanning microscopy setup (modified Leica SP5 confocal system) with a Spectra-Physics Mai-Tai HP laser set at a wavelength of 810 nm to simultaneously excite OGB and SR101 molecules, as previously described (Goltstein et al., 2013). Mice ($n = 6$) were kept lightly anesthetized (0.8% isoflurane) during the entirety of the two-photon calcium imaging recordings ($n = 6$, 1 data set/animal), while they were presented with 10 repetitions of 8 different directions of square-wave drifting gratings ($n = 80$ trials/recording). Visual stimulation duration was 3 s and was alternated by a 5 s blank inter-trial interval during which an isoluminant gray screen was presented. Visual drifting gratings (diameter 60 retinal degrees, spatial frequency 0.05 cycles/degree, temporal frequency 1 Hz) were presented within a circular cosine-ramped window to avoid edge effects at the border of the circular window. All visual stimulation was performed on a 15 inch TFT screen with a refresh rate of 60 Hz positioned at



16 cm from the mouse's eye, which was controlled by MATLAB using the PsychToolbox extension (Brainard, 1997; Pelli, 1997). A field-programmable gate array (FPGA, OpalKelly XEM3001) was connected to the microscope setup and interfaced with the stimulus computer to synchronize the timing of the visual stimulation with the microscope frame acquisition.

DATA PREPROCESSING

After a recording was completed small x-y drifts were corrected with an image registration algorithm (Guizar-Sicairos et al., 2008)

and the recording was manually checked for movement artifacts along the z-axis. If z-drifts occurred during a recording, it was rejected and no further analyses were performed. Regions of interest (neurons, astrocytes, and blood vessels) were determined semi-automatically using custom-made MATLAB software and subsequently dF/F_0 values for all neurons were calculated as previously described (Goltstein et al., 2013). In short, for each image frame i a single dF/F_0 value was obtained for each neuron by calculating the baseline fluorescence (F_{0i}), taken as the mean of the lowest 50% during a 30-s window preceding image frame i . dF is defined as the difference between the fluorescence for that neuron in the given frame and the sliding baseline fluorescence ($dF_i = F_i - F_{0i}$) (Goltstein et al., 2013). The mean number of simultaneously recorded neurons/session was 118 (range: 95–144).

NEURONAL RESPONSES TO VISUAL STIMULI

A neuron's response to a certain stimulus was defined as the mean dF/F_0 value of all frames recorded during that single trial's stimulus presentation (see **Figure 1B** for an example neuron). Stimulus presentation lasted 3.0 s (77 frames) and a single stimulus presentation yields a single response value R , equal to the mean dF/F_0 over all 77 stimulus frames. To obtain a measure of a neuron's orientation selectivity, we calculated the orientation selectivity index (OSI) as $(1 - \text{circular variance})$, as previously described by Ringach et al. (2002). To further parameterize neuronal orientation tuning, we also calculated each neuron's preferred direction by fitting a double von Mises distribution to the neuron's responses, where the peaks of both von Mises functions are opposite to each other (separated by 180°):

$$f(x | \theta, \kappa_1, \kappa_2, \mu_0) = \frac{e^{\kappa_1 \cos(x - \theta)}}{2\pi I_0(\kappa_1)} + \frac{e^{\kappa_2 \cos(x + \pi - \theta)}}{2\pi I_0(\kappa_2)} + \mu_0 \quad (1)$$

Here, $I_0(\kappa)$ is the modified Bessel function of order 0 and x represents the stimulus angle. As can be seen in the equation, we defined the free parameters as θ (preferred direction), κ_1 (concentration parameter at θ), κ_2 (concentration parameter at $\theta + \pi$) and μ_0 (baseline response). A neuron's preferred direction was defined as the angle with the highest concentration parameter (which could be either κ_1 or κ_2).

SIGNAL CORRELATIONS

In many cases neuronal responses in V1 to drifting gratings can be fairly accurately approximated by circular von Mises distributions. However, in order to more fully capture similarities and differences between the responses that neurons show to drifting gratings of different directions, we calculated signal correlations between all neuronal pairs in each recording. We defined each direction as a separate stimulus type and calculated a neuron's mean response vector $\bar{\mathbf{R}}$, where the elements of $\bar{\mathbf{R}}$ are the neuron's mean responses to each direction θ (\bar{R}_θ):

$$\bar{\mathbf{R}} = [\bar{R}_0, \bar{R}_{45} \dots \bar{R}_{315}] \quad (2)$$

We then calculated the pairwise signal correlation as the Pearson correlation between two neurons' (i, j) response vectors:

$$\rho_{i,j}^{\text{signal}} = \text{corr}(\bar{R}_i, \bar{R}_j) \quad (3)$$

Since the decoding algorithms used to extract information from the population response depend on the difference in response between neurons for different stimulus directions, high signal correlations between two neurons indicate that there is a large redundancy in information that these neurons provide to the decoder.

NOISE CORRELATIONS

Complementary to signal correlations, which provide an index of the similarity between pairs of neurons in their mean response to the different stimulus types, are noise correlations, which give an indication of the similarity in trial-by-trial response variability between neurons. We first calculated a response vector for each stimulus type θ , where each element in the vector is the neuron's response to a single presentation i of that stimulus direction:

$$\mathbf{R}_\theta = [R_{\theta_1} \dots R_{\theta_n}], \quad (4)$$

where n is ten, since we have ten repetitions per direction. Because we aim to compare a single noise correlation value per neuronal pair, we took the mean noise correlation over all eight stimulus directions $\theta = 0^\circ\text{--}315^\circ$ (with steps of 45°):

$$\rho_{i,j}^{\text{noise}} = \frac{\sum_{\theta=0}^{315} \text{corr}(\mathbf{R}_{i,\theta}, \mathbf{R}_{j,\theta})}{8} \quad (5)$$

The noise correlation is therefore an index of the mean shared trial-by-trial variability over all stimulus directions.

DECODING SAMPLE BOOTSTRAPPING

All decoding algorithms (described in detail in the following paragraphs) were tested on their ability to recover the presented stimulus direction from the neuronal population activity during single trials. To quantify their dependence on the number of neurons included, we took random subsamples of neurons (100 iterations) from each data set (ranging from 95 to 144 neurons simultaneously recorded neurons) for all tested sample sizes (1–90 neurons, each neuron was included only once). For each random sample we trained the decoder on the activity of the randomly included neurons for all 80 trials (8 directions, 10 repetitions) and subsequently decoded the presented direction of each trial from the population activity of only these neurons. Since the smallest data set consisted of 95 neurons, we ran our resampling procedures up to a maximum population size of 90 neurons. We performed these resamplings independently for each decoding algorithm.

POPULATION VECTOR DECODING

To decode the presented stimulus direction during a given trial, we used the population vector method (Georgopoulos et al., 1986) as follows. For each trial, the activity of each neuron was represented as a vector, where the angle (θ) represents the neuron's preferred direction and the magnitude (ρ) corresponds to its activity level (dF/F_0). A resultant population vector can then

be calculated by taking the circular mean over all neurons. When the resultant angle was within 22.5 degrees of the actual stimulus direction, it was counted as a "correct" decoding, in all other cases it was marked as "incorrect." After decoding all trials, the percentage correct decoded was then calculated over all trials, so it could be plotted as in **Figure 4**.

TEMPLATE MATCHING ALGORITHM

One step up in complexity from the population vector is the template matching algorithm (Lehky and Sejnowski, 1990; Zhang et al., 1998; van Duuren et al., 2008, 2009). While in the population vector method each neuron is represented only by a preferred direction and its activity level, template matching compares the population activity to response templates of different stimulus types. These templates are generated by taking the mean activity across trials during the presentation of each stimulus direction (θ) for each neuron, resulting in a "template population activity." The similarity of this template (\mathbf{R}^θ) to the actual population activity (\mathbf{R}^{stim}) is given by

$$\Theta_\theta = \frac{\sum_{i=1}^N \mathbf{R}_i^{\text{stim}} \cdot \mathbf{R}_i^\theta}{\|\mathbf{R}^\theta\| \cdot \|\mathbf{R}^{\text{stim}}\|}, \quad (6)$$

where i indexes the N elements (neurons) of \mathbf{R} , and \cdot indicates multiplication. The similarity index Θ is calculated for all eight directions and the decoded output is determined by taking the direction with the highest similarity to the population activity.

THE NON-SCALE SENSITIVE FANO FACTOR (rFANO)

The Fano factor provides an indication of the variability of neuronal responses. Two-photon calcium imaging yields neuronal responses that are given as dF/F_0 , where the relationship between the actual spiking rate and the dF/F_0 is approximately linear per neuron, but varies over neurons (Kerr et al., 2005; Greenberg et al., 2008). Compared to firing rates, neuronal dF/F_0 responses are therefore scaled by an arbitrary constant c by which both the mean (μ) and standard deviation (σ) of its responses over the entire recording are multiplied. The scaling sensitivity is described by the following equation:

$$\text{Fano} = \frac{\sigma^2 \cdot c^2}{\mu \cdot c} \quad (7)$$

Therefore, Fano factors calculated from calcium imaging data are only accurate when $c = 1$, as scaling (of the firing rate) has a disproportionately larger effect on the variance than the mean. For this reason, raw Fano Factors from calcium imaging cannot be simply translated to the Fano Factors of the raw firing rates. To correctly calculate the variability of calcium transient data, the variance and mean activity level per neuron must be normalized, thereby removing the constant c from the equation. While it is not possible to compute the raw Fano factors of the unknown firing rates, it is possible, however, to compute a Fano factor based on relative changes in mean and variance across stimulus dimensions. We therefore define the ratio Fano factor (r_{Fano}) as follows:

$$r_{\text{Fano}} = \frac{\sigma_p^2 / \sigma_{np}^2}{\mu_p / \mu_{np}} \quad (8)$$

where μ_p and σ_p^2 are the mean response and variance over trials to the neuron's preferred (P°) and opposite ($P^\circ + 180^\circ$) direction, and μ_{np} and σ_{np}^2 are the mean response and variance to the other six moving directions. This way, the mean and variance are described as fractional increases in activity and variability from non-preferred stimulus epochs to preferred stimulus epochs. Because c^2 appears in both variance terms and c appears in both mean response terms, the r_{Fano} factor is no longer scale sensitive.

NORMALIZED TEMPLATE MATCHING

In an attempt to improve the performance of the template matching algorithm, we performed a similar decoding operation, but now on the z-scored rather than the raw dF/F_0 activity values. The new equation becomes

$$\Theta_\theta = \frac{\sum_{i=1}^N Z_i^{\text{stim}} \cdot Z_i^\theta}{\|Z^\theta\| \cdot \|Z^{\text{stim}}\|}, \quad (9)$$

where Z is an index that describes a neuron's dF/F_0 response (R) for any time point (i) as the number of standard deviations (σ) from the mean dF/F_0 calculated over the entire recording (\bar{R}_{all}):

$$Z_i = \frac{(R_i - \bar{R}_{\text{all}})}{\sigma_{R_{\text{all}}}} \quad (10)$$

Therefore a neuron that shows a high mean activation level (\bar{R}_{all}), but an even higher variability $\sigma_{R_{\text{all}}}$ (high Fano factor) will contribute less to the template similarity than a neuron with a low activation level and an even lower variability (low Fano factor).

BAYESIAN MAXIMUM-LIKELIHOOD (ML)

The final algorithm we used was a Bayesian decoder. Bayes' rule is given by

$$P(\theta | A_{\text{pop}}) = \frac{P(A_{\text{pop}} | \theta) P(\theta)}{P(A_{\text{pop}})}, \quad (11)$$

where $P(\theta)$, the prior, is the prior probability of having direction θ ; $P(\theta | A_{\text{pop}})$, the posterior, is the probability of this trial's population activity being caused by direction θ ; $P(A_{\text{pop}} | \theta)$, the likelihood, represents the probability that stimulus direction θ will result in pattern A_{pop} ; and $P(A_{\text{pop}})$ is the probability of activation pattern A_{pop} , or model evidence—a normalization term. It has been shown before that there is little difference in performance between uniform and natural priors—at least for retinal ganglion cells (Jacobs et al., 2009)—so for the following results we assumed a uniform prior [$P(\theta)$ is equal for all directions] to simplify computational procedures. Note that $P(A_{\text{pop}})$ is identical for all stimulus directions, since the term θ is not present. Therefore, for a given direction θ , the posterior probability is proportional to the likelihood.

The likelihood distribution for a certain direction, say $\theta = 90^\circ$ degrees, for any neuron i was given by a Gaussian approximation

based on the mean (μ) and standard deviation (σ) of its response to stimuli with that direction during training trials (**Figure 2**); in our case, we included all 80 trials to determine the likelihood distributions. Note that these Gaussians are not tuning curves, but rather an approximation of the dF/F_0 response distribution for a single stimulus direction. The posterior probability for that direction can then be extracted by taking the value of the Gaussian likelihood distribution at the given neuronal activation level A . The population posterior probability distribution for that direction can be calculated by taking the product over the posterior probabilities for all n neurons:

$$P(\theta | A_{\text{pop}}) \propto \prod_{i=1}^n P(\theta | A)_i \quad (12)$$

Now the decoded direction can be read out by taking the direction with the highest probability.

JACKKNIFING PROCEDURE

One question that can be addressed with these decoding algorithms is whether certain properties of a neuronal response correlate with decoding performance. We therefore performed a bootstrapping procedure (1000 iterations) per recording to select random groups of neurons (ranging in size from 2 to 15) from the whole datasets of 95–144 neurons. Each bootstrap resampling was followed by a jackknife procedure applied to all neurons

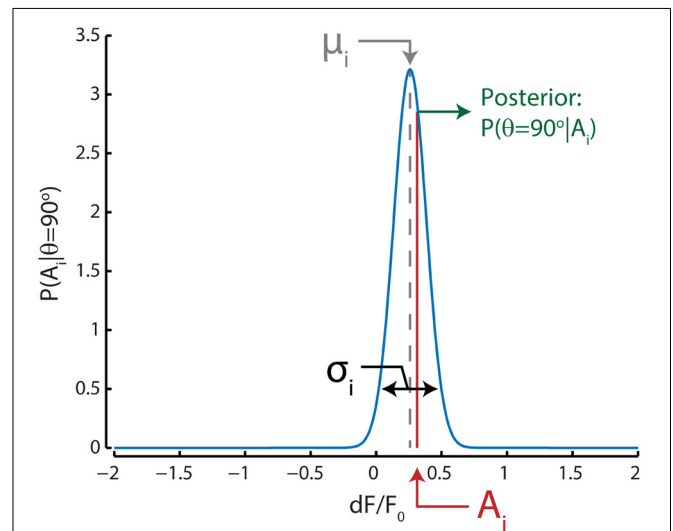


FIGURE 2 | Graphical representation of the Bayesian decoding scheme shows the process for one example neuron i for a single stimulus direction (90°) in one test trial. The probability of responses to a 90° stimulus [$P(A_i | \theta = 90^\circ)$] for this neuron is approximated by a Gaussian distribution (blue curve) with a mean (μ_i , gray) and standard deviation (σ_i , black) equal to those estimated from training trials where this stimulus direction was presented. For this example to-be-decoded test trial, the neuron's activation level (A_i , red) was close to the neuron's mean response to stimuli with a direction of 90 degrees; therefore the posterior probability [$P(\theta = 90^\circ | A_i)$, green] that the stimulus was 90 degrees is relatively high based on only this neuron's activation level and tuning. This procedure is repeated for all stimulus directions and the posterior probabilities for all neurons are multiplied. The stimulus direction with the highest probability is subsequently chosen as the maximum-likelihood (ML) read out.

in the randomly selected group to quantify the contribution of these single neurons to the decoding algorithm's performance. A jackknifing procedure provides insight into the contribution of a single neuron within a sample by taking away that neuron from the rest of the sample and comparing the properties (i.e., decoding performance) of the sample with and without this neuron. This procedure can be repeated for each neuron in the sample and yields a change in decoding performance per neuron that can be compared to other properties of the neuron. For example, jackknifing a certain neuron from a sample that has a high OSI might result in a different change in decoding performance than jackknifing a different neuron from the same sample having a low OSI. This way, the impact of several neuronal properties on decoding performance can be assessed. We normalized the decoding improvement effected by the jackknifed neuron for the number of neurons in the sample by defining the decoding improvement index D_i as

$$D_i = ND - (N - 1)D^{-i} \quad (13)$$

where N is the sample size, D the decoding accuracy of the algorithm using the entire cluster and D^{-i} is the accuracy using the entire cluster except the jackknifed neuron (i). The decoding improvement was then binned and averaged over all points per bin from all 6 data sets (see **Figure 6**).

TUNING PROPERTY INTERDEPENDENCY

Three important factors that may influence how much information on stimulus direction a neuron contributes to a neuronal cluster are the neuron's preferred orientation (PO), its signal correlation (SC) and its noise correlation (NC) with other members of the cluster. To further quantify if and how these properties show interdependencies, and vary with intersomatic distance (ID) in the imaging plane, we performed a pairwise comparison between neurons for these properties. For all neuronal pairs in all sessions we computed the SC, NC, ID and angular difference in preferred orientation (dPO). We then pooled data points from all sessions and performed a regression analysis to test if there were significant correlations between ID-SC, ID-NC, ID-dPO, SC-NC, SC-dPO and NC-dPO. Two properties were judged to be significantly interdependent when the 99% confidence interval for the regression slope did not overlap with 0. To determine the significance of a difference in slopes between dependent properties for a given independent property (either ID or SC), we z-scored the values of the dependent variable and performed another set of linear regressions. Because of this normalization, we could now quantitatively compare the relative slopes between ID-SC, ID-NC and ID-dPO, and between SC-NC and SC-dPO by assessing whether the 99% confidence intervals of these regression slopes overlapped.

REMOVING BIASES IN TUNING PROPERTY INTERDEPENDENCY

As will be further elaborated in the Results section, it is possible that the previously described analysis shows significant relationships between neuronal properties that depend on mean differences in those properties between data sets, when in fact these properties are not significantly related within data sets. To

control for such possible biases due to across-recording differences, we pooled all neuronal pairs into data bins for ID (0 to 200 with steps of 40 microns), SC (−0.7 to +0.7 with steps of 0.2) and NC (−0.175 to +0.275 with steps of 0.05). These bins were chosen so that at least 90% of all neuronal pairs per data set would be included in the analysis, with the number of points per bin still sufficiently large for robust data analysis. We then computed the minimal number of pairs for each bin for each calcium imaging data set and used this number as the resample size for a bootstrapping procedure (256, 305, and 84 pairs/bin/recording for ID, SC, and NC respectively). For each bootstrapping iteration ($n = 1000$), all bins received this number of randomly selected data points (neuronal pairs) from each data set, so that no across-recording differences in mean neuronal property values could lead to biases in the pooled resampled data set. We then performed a linear regression on each randomly sampled set and computed the 99% confidence intervals of the regression slopes to determine whether there was a significant correlation between two neuronal properties.

RESULTS

THE CLASSICAL POPULATION VECTOR FAILS AT DECODING MOVING DIRECTIONS

One of the simplest algorithms to decode neuronal population activity is the population vector method (see Methods; Georgopoulos et al., 1986). The only determinants of the contribution of a neuron's activity to the population vector are its preferred stimulus property (i.e., preferred direction) and activity level. Stimulus direction was decoded by calculating the resultant vector over all neurons (**Figure 3**). The performance of the population vector (**Figure 4**; green line), as measured by the percentage of trials that was decoded correctly, was above chance level, but this decoder clearly performed much worse than more complex algorithms. Moreover, the performance rose only slowly with the number of neurons in the sample, such that it would take a very large sample for its performance to catch up with the other algorithms.

TEMPLATE MATCHING DECODING REVEALS A STRONG SENSITIVITY TO HIGHLY ACTIVE, HIGHLY VARIABLE NEURONS

Secondly, we tested the decoding performance of a template matching algorithm (Lehky and Sejnowski, 1990; Zhang et al., 1998; van Duuren et al., 2008, 2009). Its performance (**Figure 4**; purple line) was better than the population vector method, but it suffered from a significant bias, as neurons with higher response amplitudes dominated the algorithm. If the reliability of a neuron would scale linearly with its response amplitude to the stimulus (dF/F_0), such a weighting might be beneficial, because the inherent bias incurred by high response amplitudes would favor reliable neurons. To examine whether there is indeed a bias and whether this bias poses a problem for the algorithm, a calculation of the Fano factor of the neuronal responses is in order (Fano, 1947).

The Fano factor is an index of variability and is defined as σ^2/μ (the variance of the response divided by the mean). With neuronal spiking traces, simply pooling all spikes per neuron in time bins allows the calculation of a Fano factor per neuron from

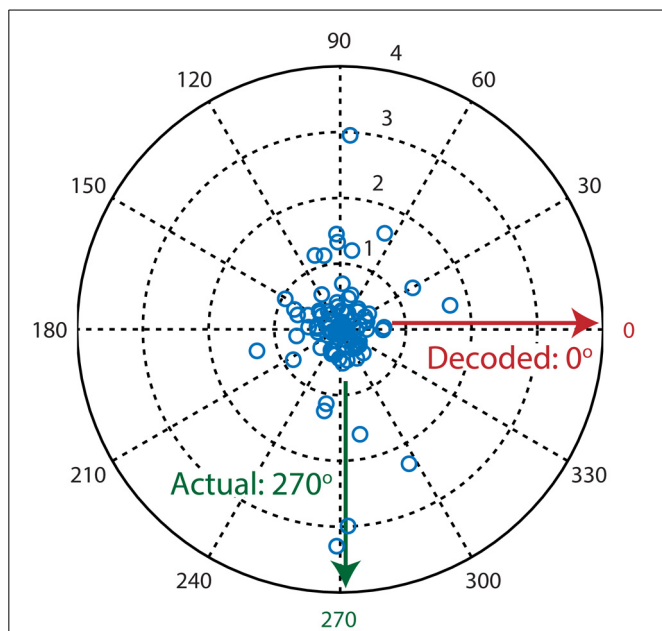


FIGURE 3 | Example decoding trial where the population vector algorithm fails to decode the presented stimulus direction. Each blue circle in the polar plot is a single neuron where its preferred direction is represented by θ (angle) and its activation level (dF/F_0) is represented by ρ (magnitude). The population vector algorithm's output is calculated by taking the circular mean over all vectors (neurons). Since most neurons are sensitive to two moving directions that are separated by 180° , calculation of the resultant vector leads to mutual cancellation of the two peaks at 90° and 270° , therefore resulting in an incorrectly decoded direction of 0° .

the mean spiking rates and variance in those bins. However, Fano factors calculated from the raw dF/F_0 trace of calcium imaging data are biased. Although there is a linear relationship between dF/F_0 and spiking activity for rates commonly observed in pyramidal cells (Kerr et al., 2005), the percentage increase in dF/F_0 due to a single action potential can differ from one neuron to the next. This means that the Fano factor has a scale sensitivity that needs to be corrected per individual neuron. We therefore developed a non-scale-sensitive ratio Fano factor (r_{Fano}) instead (see Methods).

The mean r_{Fano} calculated over all 6 recordings was 1.39 with a standard deviation of 0.20 (see Figure 5), which is close to the Fano Factor of 1 for a theoretical Poisson process and in the range of the typically observed Fano Factors for experimental spiking data (Dayan and Abbott, 2001). This value is higher than 1, which implies that neurons with higher activation levels actually have lower reliability on average. This would suggest that the template matching algorithm is indeed prone to overrepresentation of highly active and relatively variable neurons.

The dominance of the template matching algorithm by neurons with high activity levels can be accommodated by taking the mean z-score normalized activation level across trials for all neurons instead of their actual activity level (see Methods). Consider the hypothetical case that two neurons are used for decoding, where one neuron is highly active and has a high Fano factor, while the other is less active, but has a lower Fano factor. The

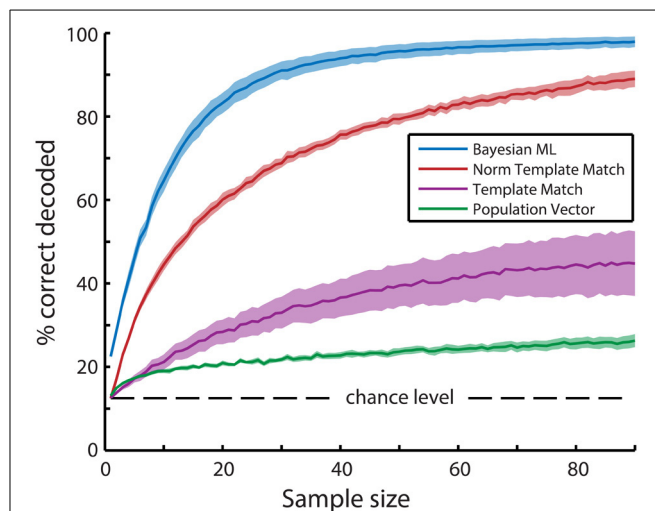
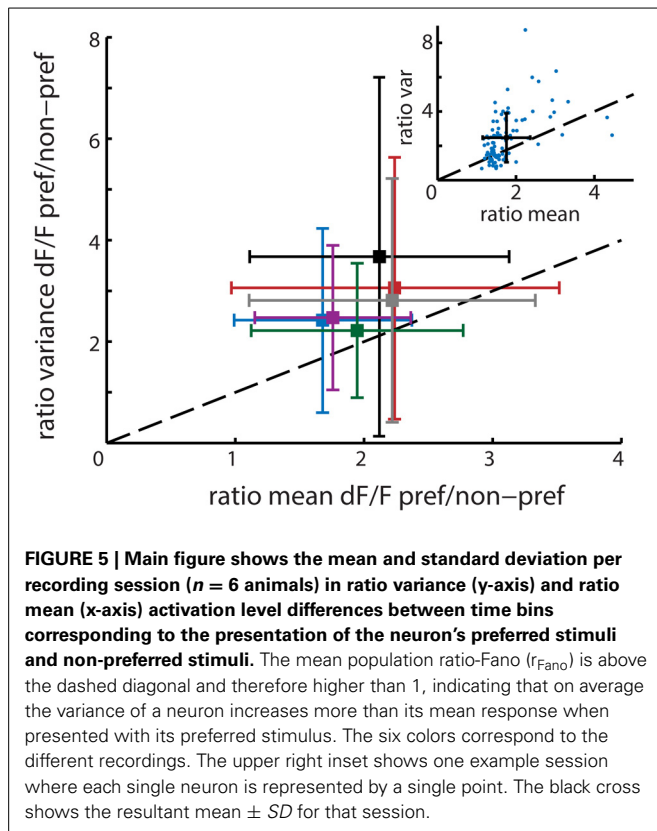


FIGURE 4 | Mean bootstrapped performance over all recordings ($N = 6$) reveals superior performance of the Bayesian maximum-likelihood (ML) direction decoding algorithm. Bootstrapping was performed with 100 random resamplings with sizes ranging from 1 to 90 neurons. Dashed lines indicate standard error of the mean bootstrapped performance. A Three-Way ANOVA (algorithm, recording, sample size) revealed significant main effects of all three variables ($p < 10^{-20}$). *Post-hoc* *t*-tests with Tukey-Kramer correction for multiple comparisons showed significant differences between all combinations of the 4 types of algorithms ($p < 10^{-10}$). In descending order of performance, the tested algorithms are Bayesian ML (blue), Normalized Template Matching (red), Non-normalized Template Matching (purple), and Population Vector (green).

former neuron's mean stimulus-driven activation level over the entire recording is $1.0 dF/F_0$ with a standard deviation of $1.0 dF/F_0$ and its response to a certain stimulus is $2.0 dF/F_0$ on average, while the latter neuron's mean activation level is $0.5 dF/F_0$ with a standard deviation of 0.25 and this neuron's response to the same stimulus is $1.5 dF/F_0$ on average. In this case, both neurons' specific stimulus-responses are $1.0 dF/F_0$ higher than their average response, so the original template matching algorithm will weigh the activation levels of these neurons equally when determining the probability of this stimulus direction. However, the former, variable neuron's response to the stimulus would on average be one standard deviation away from its mean, while the latter, reliable neuron's response to the stimulus is four standard deviations away from the mean. The more reliable neuron therefore provides much more information about stimulus identity than the former neuron. Using z-scored activation levels therefore allows the template matching algorithm to take into account the reliability of the neurons involved. The performance of the normalized template matching algorithm (Figure 4; red line) is indeed clearly superior to the non-normalized template matching algorithm.

BAYESIAN MAXIMUM-LIKELIHOOD (ML) DECODING WITH INTERMEDIATE TO LARGE SAMPLE SIZES SHOWS NEAR-PERFECT DIRECTION DECODING

For further analysis of tuning property dependencies we created a Bayesian maximum-likelihood (ML) decoder. In contrast to the



population vector and template matching algorithms we previously described, the ML decoder calculates posterior stimulus probabilities on a single neuron basis and then takes the product over probabilities contributed by all neurons per stimulus. Because the probabilities are computed on a single-neuron level and thereafter aggregated into population posterior probabilities, the ML decoder inherently corrects for differences in response reliability between neurons. The Bayesian ML decoder indeed easily outperforms all other decoding algorithms (Figure 4; blue line). We therefore performed all following decoding analyses with the ML decoder.

SINGLE NEURON CONTRIBUTIONS TO DECODING ACCURACY DEPEND CRITICALLY ON ORIENTATION SELECTIVITY INDEX, SIGNAL CORRELATIONS AND NOISE CORRELATIONS AT LARGER NEURONAL SAMPLE SIZES

One question that can be addressed with a decoding algorithm is whether certain properties of a neuronal response have any effect on decoding performance. For instance, an intuitive hypothesis holds that decoding performance depends on a neuron's orientation selectivity index (OSI; see Methods). To investigate this further, a bootstrapping procedure with 1000 iterations was performed to select random clusters of neurons (2–15) and was followed by a jackknife procedure on all neurons per cluster to quantify the contribution of single neurons to the decoder's performance. We calculated each neuron's contribution to the decoding accuracy as an index normalized for sample size, allowing us to compare the effects over different sample sizes [see

Equation (13); Figure 6A]. The results show that at small sample sizes a neuron's OSI has little effect on the decoding performance. However, at larger sample sizes, the addition of a neuron with a high OSI results in a larger improvement of the decoder's accuracy.

However, the OSI captures only a small part of the way a neuron is tuned to properties of visual stimuli. To more fully capture the similarity in visual responses between neurons we also calculated the mean pairwise signal and noise correlations between the jackknifed neuron and the other members of the sample. The results of this analysis show a clear dependence of decoding performance on the jackknifed neuron's mean signal correlation with the sample, where negative values are associated with higher decoding performance (Figure 6B). Similarly to the pairwise signal correlations, strong positive noise correlations had a deleterious effect on decoding performance (Figure 6C). For both the signal and noise correlations, the effect was stronger for larger clusters, indicating that response heterogeneity is important for efficient coding by large assemblies.

ANALYSIS OF INTERDEPENDENCE BETWEEN RESPONSE PROPERTIES REVEALS A DISSOCIATION BETWEEN STIMULUS TUNING PROPERTIES AND RANDOM FLUCTUATIONS

The previous paragraph described the dependence of decoding accuracy on similarity in stimulus tuning properties (signal correlations and OSI) and random fluctuations (noise correlations). While these neuronal response properties may exert independent effects, it is also possible that they are strongly interdependent and might in reality have a single origin that is partly reflected in all response similarity indices. Because the spatial distance between cortical somata strongly influences the probability that two neurons are interconnected (Gilbert and Wiesel, 1983; Nicoll and Blakemore, 1993), we first investigated whether neuronal response properties depend on intersomatic distance. We therefore analyzed whether the angular difference in preferred orientation (dPO), signal correlation (SC) or noise correlation (NC) depend on the intersomatic distance (ID) by pooling all neuronal pairs from all recordings and performing a linear regression to test for significant interdependence (Figure 7).

As reported before in other animal species, such as cat (Ch'ng and Reid, 2010) and macaque (Smith and Kohn, 2008), we found that noise correlations significantly vary with intersomatic distance (Figure 7B; regression slope $\neq 0$, $p < 0.01$), where pairs of neurons that are closer together show higher noise correlations. Noise correlations are assumed to reflect common input between neurons (Nienborg and Cumming, 2010), so one could hypothesize that if noise correlations are high, there should also be high similarities in the tuning properties of neurons. However, the spatial dependencies of preferred orientation difference and signal correlations were non-significant (Figures 7A,C), indicating that there might be a dissociation in origin between stimulus tuning properties and random fluctuations. It is important to note here that mice do not possess orientation tuning columns and that the spatial scale investigated is <400 microns, which is much smaller than for example the scale on which Smith and Kohn (2008) found a spatial dependency of correlations in macaque V1.

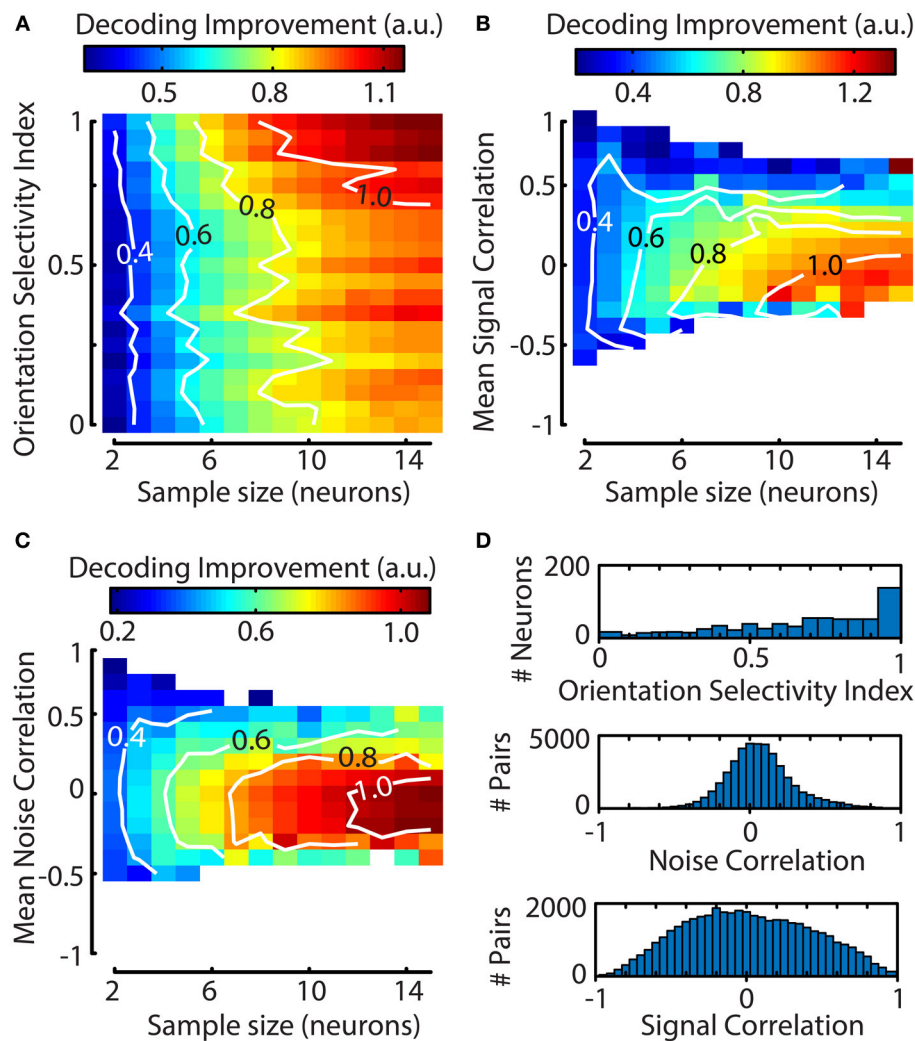


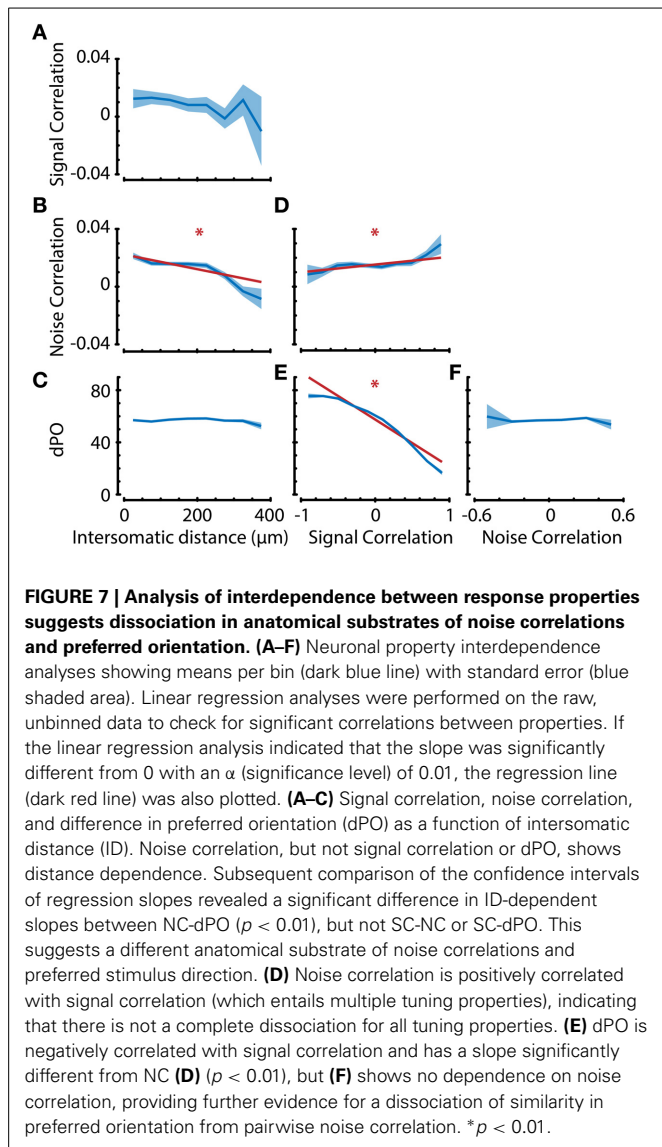
FIGURE 6 | The contribution of a single neuron to the decoding performance of a Bayesian ML algorithm depends on the neuron's orientation selectivity index (OSI) and its mean correlation with other neurons. (A–C) Graphs were created by using a bootstrapping procedure for random cluster selection and a subsequent jackknifing procedure to quantify the impact of a single neuron's addition to the cluster. Contiguous contour lines are plotted in white for increased readability. **(A)** The influence of a neuron's OSI on the decoding performance of its cluster is stronger for larger clusters. Neurons with a higher OSI improve the algorithm's performance more. **(B)** The addition

of a neuron that has a negative mean signal correlation ($\rho \approx -0.2$ – -0.1 for $N = 10$ – 15) with the rest of its cluster leads to a larger increase in decoding performance. **(C)** Slightly negative to low noise correlations ($\rho \approx -0.3$ – 0 for $N = 10$ – 15) improve decoding performance; as with **(A,B)** this effect is stronger for larger sample sizes. **(D)** Population distributions pooled over sessions. Top panel; population distribution of orientation selectivity index (OSI), Middle panel; noise correlation, Bottom panel; signal correlation. The means for the distributions of noise and signal correlations do not correspond to the values that lead to maximum decoding improvement.

We next investigated the dependence of dPO on SC and NC to see if the similarity in stimulus tuning properties is correlated with the strength of random fluctuations (**Figures 7E,F**). As expected, we found that dPO is dependent on signal correlations (**Figure 7E**; regression slope $\neq 0$, $p < 0.01$). However, dPO did not show any dependence on noise correlations (**Figure 7F**), which—in line with the distance dependence analysis—supports the conclusion that input fluctuations are dissociated from the inputs that determine a neuron's preferred orientation. The analysis of NC/SC dependence (**Figure 7D** regression slope $\neq 0$, $p < 0.01$) verified previously published observations that NCs are

positively correlated with SCs (e.g., Lee et al., 1998; Kohn and Smith, 2005; Cohen and Maunsell, 2009), suggesting that this dissociation is not complete.

To check whether the suggested dissociation still holds under more stringent statistical analysis, we tested whether the confidence intervals for the regression slopes were significantly different. After z-scoring the dPO, NC, and SC values in order to quantitatively compare the slopes for these properties, we found that the ID-dependent slopes of NC-dPO ($p < 0.01$), but not of SC-NC or SC-dPO were significantly different. Following the same approach, we also compared the signal



correlation-dependent slopes of dPO and NC, and found them to be significantly different ($p < 0.01$). Taken together, these results indicate that a neuron's preferred orientation is dissociated from the similarity in random fluctuations in spiking activity. It also indicates that there are stimulus tuning aspects (such as bandwidth, peak-trough ratio, etc.) captured by signal correlations that are different from the neuron's preferred orientation, and that these aspects might have anatomical substrates similar to those of noise correlations (Figure 7).

When comparing property-dependent differences across sessions such as described in the previous paragraphs, a corrective measure must be employed to ensure that these significant inter-property dependencies cannot be explained by across-session differences. Suppose that we compare two hypothetical recordings where the mean noise correlations are higher for recording 1 than recording 2, and the field of view also differs between recordings. In this case, recording 1 has a field-of-view resulting in a mean intersomatic distance (ID) of 100 with a maximum

ID of 200 microns and the other recording has a mean ID of 150 with a maximum ID of 300 microns. Even if there were no interdependence between noise correlation and intersomatic distance, a non-debiased analysis could detect a correlation, since all neuronal pairs with IDs more than 200 microns will come from session 2, which has a lower mean noise correlation. A regression analysis on the pooled neuronal pairs could then show a significant negative correlation where there is none. One way to correct for across-recording differences in mean or variance is to z-score the values per session and then pool the z-scored data sets. However, such a z-scoring approach does not correct for all biases. Using the same hypothetical example recordings as described above, z-scoring the intersomatic distance (ID) will create a bias rather than correct for it, because it would equate neurons with an ID of 100 in recording 1 with an ID of 150 in recording 2, since they are both equal to the recording-mean. To control for possible biases due to across-recording differences, we performed a debiased property interdependence analysis that takes a fixed and equal number of neuronal pairs from all recordings for several data bins (see Methods). The debiased analysis confirmed the results we obtained from the non-debiased regression analyses; ID-NC, SC-NC, SC-dPO showed significant linear correlations ($p < 0.01$), while ID-SC, ID-dPO, and NC-dPO did not. The neuronal property interdependencies we found are therefore not due to across-recording differences.

DISCUSSION

EXPLAINING THE WEAK PERFORMANCE OF THE POPULATION VECTOR METHOD

The population vector exhibits the weakest performance of the tested algorithms, because it relies on two important assumptions that are violated. First, the population vector method assumes a uniform distribution of preferred directions over the encoded dimension. Clearly, decoding a stimulus direction with a sample size of one neuron will not yield a uniform distribution of preferred directions. This results in a decoded direction that is biased to the direction that is overrepresented by random resampling from the entire population. Moreover, neurons' preferred directions are not uniformly distributed in the visual cortex, but have a bias toward overrepresentation of cardinal directions (Kreile et al., 2011).

The second violated assumption is that each neuron has a single preferred direction. In V1 most neurons are selective for an axis of orientation (e.g., responding to both upward moving and downward moving) rather than one specific direction (e.g., responding to only upward moving), resulting in a tuning curve with two peaks rather than one (see Figure 3 for an example trial). When the stimulus-driven population activity is divided in two peaks that are separated by 180 degrees, taking a circular mean over the whole population will reduce the size of the signal that is extracted. This reduction takes place because two vectors with opposite directions will result in an average vector that has a magnitude equal to the magnitude in the preferred direction minus the magnitude in the opposite direction. When the population responses in the preferred and opposite direction have similar strengths, the resultant vector has a small magnitude and the decoded direction is especially prone to random

imbalances in activation (noise). In essence, the bimodal nature of orientation tuning curves leads to a large decrease in signal-to-noise ratio (SNR) when decoding directions with a population vector. Therefore, although population vector decoding of orientations (Vogels, 1990) and of directions with only direction-selective and no orientation-selective neurons work well (Seung and Sompolinsky, 1993), direction decoding based on all (mixed) neurons shows very low accuracy. Since fixing or avoiding these two violations would require significant alteration of the population vector algorithm, we conclude that the classical population vector is not suitable for decoding visual stimulus directions using populations of V1 neurons.

THE CRITICAL DEPENDENCY ON INPUT RELIABILITY SUGGESTS A NEURAL MECHANISM FOR VARIABILITY NORMALIZATION

The template matching (TM) algorithm—which does not assume any particular shape of tuning curves—performed much better than the population vector method. However, our results showed that the TM decoder is highly sensitive to dominance by unreliable, highly active neurons. It has been previously reported that neuronal firing rate variability increases more than the mean spiking rate; i.e., neuronal Fano factors are usually >1.0 (Tolhurst et al., 1981; Baddeley et al., 1997; Dayan and Abbott, 2001). It has been shown that such codes could serve to maximize the spike-count entropy, and thereby the amount of information coded by neurons (Baddeley et al., 1997). However, this interpretation becomes problematic considering that the neuronal recordings in the previously mentioned studies were performed using extracellular microelectrodes, which are notoriously biased toward detecting only the most highly active neurons (Olshausen and Field, 1997; Wohrer et al., 2013). Because Fano factors >1.0 have a disproportionately large effect on highly active neurons, the increase in informational entropy might come at the cost of increased decoding complexity due to the discrepancy in variability over different neurons.

This would entail that the previously mentioned analyses of interactions between response variability and informational entropy (Tolhurst et al., 1981; Baddeley et al., 1997; Dayan and Abbott, 2001) hold for only a small subset of the entire population of neurons. The results presented in the current paper are based on two-photon calcium imaging; a non-biased recording technique. Our analyses confirm the previous observations that neuronal Fano factors are >1.0 for highly-active, but also for less active neurons. However, our current analyses also show that normalizing neuronal data by Z-scoring improves the decoder's accuracy tremendously. While this operation does not negate response non-stationarity, it does account for differences in spiking rates and Fano factors. This stimulus-invariant response normalization is somewhat reminiscent of the more general principle of population divisive normalization that has already been used to explain a host of other neuronal response properties, such as contrast-invariant neuronal responsiveness, attentional modulation, stimulus adaptation and multi-areal time-dependent attentional progression effects (Lee and Maunsell, 2009; Reynolds and Heeger, 2009; Ringach, 2010; Montijn et al., 2012; Benucci et al., 2013). Possibly, the stimulus-invariant variability normalization we applied to enhance the template matching algorithm's

performance represents a neural mechanism enabling divisive normalization at a neuronal population scale.

One way variability normalization might be neurally implemented is through a simple Hebbian plasticity rule, if neurons with a high variance in firing rates also produce unreliable spike timings. The hypothesis holds that such neurons are presynaptic to neurons in higher visual areas and that their input strength inversely correlates with their firing-rate variability. Given the observation that synaptic plasticity is strongly dependent on the exact timing of spikes, as in spike-timing dependent plasticity (STDP) (Bi and Poo, 1998), this interpretation makes two testable predictions: (1) neurons with high variances in stimulus-driven spike rates have high variances in spike timing; and (2) the synaptic strength between orientation-tuned neurons and their targets depends on the spike rate reliability over trials of the presynaptic neurons.

BAYESIAN MAXIMUM-LIKELIHOOD DECODER SHOWS ROBUST PERFORMANCE

The Bayesian maximum-likelihood (ML) decoder—with its assumption of Gaussian response distributions within stimulus types (Figure 2)—performs even better than the template matching algorithm, even though Gaussian distributions are often poor approximations of neuronal response distributions. It would seem that in our case the exact shape of the response distribution of any neuron is less important than the reliability of its response. The Bayesian ML decoder also shows robust performance despite another violated assumption. In our case, the ML decoder assumes statistical independence between neurons, because of its product-rule transformation from individual posterior probabilities per neuron to a population posterior probability. Even though we have shown that inter-neuronal noise correlations are non-zero, the decoder algorithm shows an accuracy higher than that of all other tested algorithms, and reaches $>90\%$ correct with a fairly small sample size (>28 neurons).

IN LARGER POPULATIONS NARROWLY TUNED NEURONS ENCODE MORE INFORMATION THAN BROADLY TUNED NEURONS

Results from jackknifing single neurons from a larger cluster show that orientation selectivity (OSI), signal correlations and noise correlations can be important factors influencing decoding performance. Orientation selectivity has a strong influence on a single neuron's information contribution in large clusters, but has hardly any relation to the contribution in small clusters. This observation can be related to the idea that sharply tuned neurons would contribute more information than broadly tuned neurons if the population size is sufficiently large to offer a complete coverage of the stimulus dimension (Hinton et al., 1986; Seung and Sompolinsky, 1993; Schoups et al., 2001; Kang et al., 2004; Montemurro and Panzeri, 2006). Within this framework, the relative unimportance of OSI at small sample sizes can be explained by the trade-off between orientation selectivity and having a tuning curve that encompasses a high proportion of the entire stimulus dimension ($[0-360]$ degrees). In the extreme case of a sample size of 1 neuron, a highly selective neuron that responds to only 1 stimulus direction out of 8 will provide much information about 1/8th of all stimuli, but leaves the decoder to

randomly guess in all other cases, while a more broadly tuned neuron would allow the decoder to make educated guesses in all 8 cases. With larger sample sizes and highly selective neurons, the chance decreases of having gaps in the representation of stimulus space, while these cells would still offer the benefits of high selectivity.

The idea that narrowing a neuron's tuning curve often increases its encoded information has been challenged more recently, but no definitive conclusions have yet been reached. For one, a computational study comparing two models for orientation selectivity that either included or excluded cortical sharpening of tuning curves reported that the amount of information coded by neurons is strongly dependent on the amount and structure of noise covariance (Seriès et al., 2004). While its results were striking, an important confound was that the output of the model was constrained to produce identical neuronal tuning curves. In essence, the study looked at the effect of cortical tuning curve sharpening, given the same output tuning curve of cortical neurons before and after sharpening, by changing the input tuning curves (LGN afferents). However, a more appropriate comparison would have been the comparison of coded information between pre- and post-sharpened tuning curves of cortical neurons, given the same LGN afferent tuning curves.

Another consideration regarding the optimal shape of tuning curves is that the stimulus information relies solely on the Fourier amplitude spectrum of the orientation tuning curve and not necessarily on other tuning curve properties that give a lower-order approximation for this Fourier-space tuning curve, such as bandwidth (Ringach, 2010). This would suggest that the narrow-vs.-wide debate, including some seemingly contradictory results regarding changes in bandwidth (e.g., Yang and Maunsell, 2004; Goltstein et al., 2013), might simply result from looking at parameters that describe tuning curves in a way that does not correctly capture the properties that are most important to their information coding potential. Fully capturing a neuron's information coding potential in a single—or even a handful of—tuning curve parameters might very well be impossible.

Perhaps even more importantly, the shape of neuron's tuning curve that allows maximum information coding potential is heavily dependent on the distribution of tuning curves within its population, and tuning curves within sensory cortical populations are very heterogeneous (Wohrer et al., 2013). Computational studies that look into how tuning curve properties influence the maximum amount of encoded information are important to further our understanding of stimulus coding in the visual system, but their often relatively homogeneous population of neuronal tuning curves and sometimes arbitrary parameter sets mean that one should be cautious in applying their theoretical findings too readily as a realistic model of sensory processing.

An important caveat to our observation that high OSI leads to higher decoding accuracy in large samples is that the less orientation-selective neurons we measured could have a lower signal-to-noise ratio (SNR), which would thereby convey a low orientation selectivity, as well as the measured, but not necessarily neurally coded information. To draw definitive conclusions, one would have to correct for neuronal SNR while analyzing the dependence of decoding accuracy on orientation

selectivity, which could be an interesting subject for future validation research. Despite this caveat, our results are consistent with the more classical assumption that at least under our experimental conditions, greater orientation selectivity increases the amount of encoded information for larger populations of neurons.

UNCORRELATED, HETEROGENEOUS NEURONAL RESPONSES INCREASE THE AMOUNT OF ENCODED INFORMATION

Our results also show that the effect of signal- and noise correlations on population coding is sample-size dependent. Moreover, neurons that are added to a population increase the population information coding most when their response is dissimilar from the extant population (negative signal correlations) and they are uncorrelated (noise correlations near zero). Heterogeneity of a population becomes more important as the neurons' representation of stimulus space gets crowded, because there are more neurons to represent the same space. This also means that for decoders that do not take into account the entire covariance structure of their input population, noise correlations are detrimental. To further validate if the negative effect of noise correlations on decoding we present in this paper is truly detrimental or could be remedied by taking into account the neuronal correlation structure, it would be an option to create a decoder algorithm that takes into account this covariance matrix. However, accurately estimating the covariance matrix of a group of 90 neurons that respond to 8 different stimulus types would require the fitting of eight 90-dimensional multivariate Gaussians. Since the experiments we performed have only 10 repetitions per stimulus type, this would most probably result in poor fits of the covariance matrix. If neurons can indeed take into account the entire covariance matrix of their synaptic inputs (which can consist of up to $5 \cdot 10^7$ covariance values Huttenlocher, 1979), this would mean that although our results are strongly indicative of detrimental effects of noise correlations on population coding in V1, they do not offer conclusive evidence.

While extracting a proper covariance matrix from experimental data on populations of neurons is extremely difficult, several modeling studies have addressed the effects of correlation on population coding. One approach would be to create a neuronal network model and fit the model's parameters to experimentally measured data. For example, Pillow et al. (2008) showed that a computational model can decode 20% more information from experimentally measured retinal ganglion cell activity when it does not assume independence of the cells' responses. In general, when the signal correlation and noise correlation of a neuronal pair have opposite signs (i.e., negative vs. positive correlations); encoding fidelity can be improved (Averbeck et al., 2006). However, since high signal correlations between pairs of neurons tend to often also show high noise correlations (Lee et al., 1998), the theoretical possibility of increased performance from correlations might in practice not apply to neural systems in the cortex.

To settle this debate more definitively, two matters will have to be addressed. First, future research would have to determine whether neuronal pairwise noise correlations in V1 are stationary up to a point where synaptic plasticity could keep up with

changes in the covariance matrix. If noise correlations are indeed adequately stable over time, a sufficiently large number of functionally connected neurons would then have to be measured over a relatively long time period, so that enough neurophysiological measurements can be collected to accurately estimate the covariance matrix of that group of neurons. The actual action potential firing of these neurons would then have to be compared to the theoretically predicted pattern when the neurons are decorrelating their synaptic inputs. Suffice it to say, this is unfortunately not possible with the currently used data set, but hopefully technical advances in ensemble recordings will make these kinds of measurement possible in the near future.

NOISE CORRELATIONS AND STIMULUS TUNING PROPERTIES MAY HAVE DIFFERENT ANATOMICAL SUBSTRATES

Noise correlations, but not signal correlations and angular preferred orientation difference, depend on intersomatic distance (Figure 7). This could indicate that noise correlations reflect common input (Nienborg and Cumming, 2010), in the sense that they are indicative of the temporal fluctuations in global synaptic barrages on their many dendritic arbors, but that tuning properties arise from mechanisms unrelated to overall fluctuations in spiking activity. In other words, robust visual response properties may depend on relatively reliable connections, while noise correlations are more related to non-specific global barrages conveyed by locally correlated inputs that become uncorrelated across anatomical space.

This last interpretation is in line with other research showing that noise correlations can be strongly influenced by more general modulatory processes, such as anesthesia, attention and arousal (Greenberg et al., 2008; Mitchell et al., 2009; Pennartz et al., 2013). It is also supported by the observation that layer 2/3 and layer 5 neurons display relatively high spiking correlations (mean around 0.1), while populations within L4 show correlations around 0.01; almost an order of magnitude smaller (Smith et al., 2012). As has been suggested before, horizontal cortico-cortical connections may be necessary for noise correlations (Ts'o et al., 1986; Smith et al., 2012). Such recurrent circuitry is present in superficial and deep layers, but probably not in layer 4 (Gilbert and Wiesel, 1983; Douglas and Martin, 2004), which may explain the difference in magnitude of noise correlations and provides further evidence that noise correlations might indeed have specific anatomical substrates. Moreover, it has been reported that L2/3 orientation tuned neurons in mouse V1 receive relatively more inputs from neurons that are themselves tuned to that orientation (Ko et al., 2011), which makes it plausible that stimulus tuning depends on specific synaptic connections.

The idea that stimulus tuning and noise correlations are based on different anatomical configurations also explains why we found a dissociation between the spatial dependence of noise correlations on the one hand, and the non-spatial dependence of dPO and signal correlations on the other. A testable prediction that follows from our hypothesis is that synaptic connections between similarly-tuned neurons should be stronger or more reliable than connections between differently-tuned neurons; meaning that the EPSPs they elicit have higher amplitudes or should occur in closer temporal proximity to the EPSPs of other

similarly-tuned synapses connected to the same neuron, which would lead to more efficacious temporal summation. Because those synapses might also be originating from layer 4 neurons, rigorously testing the hypothesis would require simultaneous patching of layer 4 and layer 2/3 neurons *in vivo*.

To conclude, we have provided in-depth descriptions of several analysis methods that can be used to investigate population coding, including several of their drawbacks and advantages, and we have suggested how to avoid problems and biases conveyed by some of these methods, as applied to V1 data. The analyses reveal the importance of input reliability for reading population codes, suggest a role for response heterogeneity and decorrelation in population coding, and show there is a difference in anatomical intersomatic distance dependence between pairwise noise correlations and differences in preferred orientation. We have interpreted these results within a neurophysiological framework and presented testable predictions that follow from the hypotheses that Hebbian plasticity might lead to spiking variability normalization, and that noise correlations and properties underlying directional tuning arise from segregated anatomical substrates.

ACKNOWLEDGMENTS

This work was supported by the Netherlands Organization for Scientific Research-Excellence grant for the Brain & Cognition project 433-09-208 and the EU FP7-ICT grant 270108.

REFERENCES

- Ainsworth, M., Lee, S., Cunningham, M. O., Traub, R. D., Kopell, N. J., and Whittington, M. A. (2012). Rates and rhythms: a synergistic view of frequency and temporal coding in neuronal networks. *Neuron* 75, 572–583. doi: 10.1016/j.neuron.2012.08.004
- Andermann, M. L., Kerlin, A. M., Roumis, D. K., Glickfeld, L. L., and Reid, R. C. (2011). Functional specialization of mouse higher visual cortical areas. *Neuron* 72, 1025–1039. doi: 10.1016/j.neuron.2011.11.013
- Averbeck, B. B., Latham, P. E., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* 7, 358–366. doi: 10.1038/nrn1888
- Baddeley, R., Abbott, L. F., Booth, M. C. A., Sengpiel, F., Freeman, T., Wakeman, E. A., et al. (1997). Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proc. R. Soc. Lond. B Biol. Sci.* 264, 1775–1783. doi: 10.1098/rspb.1997.0246
- Benucci, A., Saleem, A. B., and Carandini, M. (2013). Adaptation maintains population homeostasis in primary visual cortex. *Nat. Neurosci.* 16, 724–729. doi: 10.1038/nn.3382
- Bi, G., and Poo, M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spat. Vis.* 10, 433–436. doi: 10.1163/156856897X00357
- Butts, D. A., Weng, C., Jin, J., Yeh, C.-I., Lesica, N. A., Alonso, J.-M., et al. (2007). Temporal precision in the neural code and the timescales of natural vision. *Nature* 449, 92–95. doi: 10.1038/nature06105
- Ch'ng, Y. H., and Reid, R. C. (2010). Cellular imaging of visual cortex reveals the spatial and functional organization of spontaneous activity. *Front. Integr. Neurosci.* 4:20. doi: 10.3389/fnint.2010.00020
- Cohen, M. R., and Kohn, A. (2011). Measuring and interpreting neuronal correlations. *Nat. Neurosci.* 14, 811–819. doi: 10.1038/nn.2842
- Cohen, M. R., and Maunsell, J. H. R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* 12, 1594–1600. doi: 10.1038/nn.2439
- da Silveira, R. A., and Berry, M. J. II. (2013). High-fidelity coding with correlated neurons. ArXiv13073591 Q-Bio.

- Dayan, P., and Abbott, L. (2001). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. Cambridge, MA: The MIT Press.
- Douglas, R. J., and Martin, K. A. C. (2004). Neuronal circuits of the neocortex. *Annu. Rev. Neurosci.* 27, 419–451. doi: 10.1146/annurev.neuro.27.070203.144152
- Ecker, A. S., Berens, P., Tolias, A. S., and Bethge, M. (2011). The effect of noise correlations in populations of diversely tuned neurons. *J. Neurosci.* 31, 14272–14283. doi: 10.1523/JNEUROSCI.2539-11.2011
- Faisal, A. A., Selen, L. P. J., and Wolpert, D. M. (2008). Noise in the nervous system. *Nat. Rev. Neurosci.* 9, 292–303. doi: 10.1038/nrn2258
- Fano, U. (1947). Ionization yield of radiations. II. The fluctuations of the number of ions. *Phys. Rev.* 72, 26–29. doi: 10.1103/PhysRev.72.26
- Georgopoulos, A. P., Schwartz, A. B., and Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science* 233, 1416–1419. doi: 10.1126/science.3749885
- Gilbert, C. D., and Wiesel, T. N. (1983). Clustered intrinsic connections in cat visual cortex. *J. Neurosci.* 3, 1116–1133.
- Goltstein, P. M., Coffey, E. B. J., Roelfsema, P. R., and Pennartz, C. M. A. (2013). *In vivo* two-photon Ca²⁺ imaging reveals selective reward effects on stimulus-specific assemblies in mouse visual cortex. *J. Neurosci.* 33, 11540–11555. doi: 10.1523/JNEUROSCI.1341-12.2013
- Graf, A. B. A., Kohn, A., Jazayeri, M., and Movshon, J. A. (2011). Decoding the activity of neuronal populations in macaque primary visual cortex. *Nat. Neurosci.* 14, 239–245. doi: 10.1038/nn.2733
- Greenberg, D. S., Houweling, A. R., and Kerr, J. N. D. (2008). Population imaging of ongoing neuronal activity in the visual cortex of awake rats. *Nat. Neurosci.* 11, 749–751. doi: 10.1038/nn.2140
- Guizar-Sicairos, M., Thurman, S. T., and Fienup, J. R. (2008). Efficient sub-pixel image registration algorithms. *Opt. Lett.* 33, 156–158. doi: 10.1364/OL.33.000156
- Hinton, G. E., McClelland, J. L., and Rumelhart, D. E. (1986). “Distributed representations,” in *Parallel Distributed Processing*, eds D. E. Rumelhart, J. L. McClelland, and PDP Research Group (Cambridge, MA: MIT Press), 318–362.
- Hubel, D. H., and Wiesel, T. N. (1974). Sequence regularity and geometry of orientation columns in the monkey striate cortex. *J. Comp. Neurol.* 158, 267–293. doi: 10.1002/cne.901580304
- Hübener, M. (2003). Mouse visual cortex. *Curr. Opin. Neurobiol.* 13, 413–420. doi: 10.1016/S0959-4388(03)00102-8
- Huttenlocher, P. R. (1979). Synaptic density in human frontal cortex—developmental changes and effects of aging. *Brain Res.* 163, 195–205. doi: 10.1016/0006-8993(79)90349-4
- Jacobs, A. L., Fridman, G., Douglas, R. M., Alam, N. M., Latham, P. E., Prusky, G. T., et al. (2009). Ruling out and ruling in neural codes. *Proc. Natl. Acad. Sci. U.S.A.* 106, 5936–5941. doi: 10.1073/pnas.0900573106
- Kang, K., Shapley, R. M., and Sompolinsky, H. (2004). Information tuning of populations of neurons in primary visual cortex. *J. Neurosci.* 24, 3726–3735. doi: 10.1523/JNEUROSCI.4272-03.2004
- Kerr, J. N. D., Greenberg, D., and Helmchen, F. (2005). Imaging input and output of neocortical networks *in vivo*. *Proc. Natl. Acad. Sci. U.S.A.* 102, 14063–14068. doi: 10.1073/pnas.0506029102
- Ko, H., Hofer, S. B., Pichler, B., Buchanan, K. A., Sjöström, P. J., and Mrsic-Flogel, T. D. (2011). Functional specificity of local synaptic connections in neocortical networks. *Nature* 473, 87–91. doi: 10.1038/nature09880
- Kohn, A., and Smith, M. A. (2005). Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.* 25, 3661–3673. doi: 10.1523/JNEUROSCI.5106-04.2005
- Kreile, A. K., Bonhoeffer, T., and Hübener, M. (2011). Altered visual experience induces instructive changes of orientation preference in mouse visual cortex. *J. Neurosci.* 31, 13911–13920. doi: 10.1523/JNEUROSCI.2143-11.2011
- Lamme, V. A. (1995). The neurophysiology of figure-ground segregation in primary visual cortex. *J. Neurosci.* 15, 1605–1615.
- Lee, D., Port, N. L., Kruse, W., and Georgopoulos, A. P. (1998). Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J. Neurosci.* 18, 1161–1170.
- Lee, J., and Maunsell, J. H. R. (2009). A normalization model of attentional modulation of single unit responses. *PLoS ONE* 4:e4651. doi: 10.1371/journal.pone.0004651
- Lehky, S. R., and Sejnowski, T. J. (1990). Neural model of stereoacuity and depth interpolation based on a distributed representation of stereo disparity. *J. Neurosci.* 10, 2281–2299.
- Mitchell, J. F., Sundberg, K. A., and Reynolds, J. H. (2009). Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63, 879–888. doi: 10.1016/j.neuron.2009.09.013
- Montemurro, M. A., and Panzeri, S. (2006). Optimal tuning widths in population coding of periodic variables. *Neural Comput.* 18, 1555–1576. doi: 10.1162/neco.2006.18.7.1555
- Montijn, J. S., Klink, P. C., and van Wezel, R. J. A. (2012). Divisive normalization and neuronal oscillations in a single hierarchical framework of selective visual attention. *Front. Neural Circuits* 6:22. doi: 10.3389/fncir.2012.00022
- Motter, B. C. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *J. Neurophysiol.* 70, 909–919.
- Nicoll, A., and Blakemore, C. (1993). Patterns of local connectivity in the neocortex. *Neural Comput.* 5, 665–680. doi: 10.1162/neco.1993.5.5.665
- Niell, C. M., and Stryker, M. P. (2008). Highly selective receptive fields in mouse visual cortex. *J. Neurosci.* 28, 7520–7536. doi: 10.1523/JNEUROSCI.0623-08.2008
- Nienborg, H., and Cumming, B. (2010). Correlations between the activity of sensory neurons and behavior: how much do they tell us about a neuron's causality? *Curr. Opin. Neurobiol.* 20, 376–381. doi: 10.1016/j.conb.2010.05.002
- Olshausen, B. A., and Field, D. J. (1997). Sparse coding with an overcomplete basis set: a strategy employed by V1? *Vision Res.* 37, 3311–3326. doi: 10.1016/S0042-6989(97)00169-7
- Pagan, M., Urban, L. S., Wohl, M. P., and Rust, N. C. (2013). Signals in inferotemporal and perirhinal cortex suggest an untangling of visual target information. *Nat. Neurosci.* 16, 1132–1139. doi: 10.1038/nn.3433
- Pelli, D. G. (1997). The videotoolbox software for visual psychophysics: transforming numbers into movies. *Spat. Vis.* 10, 437–442. doi: 10.1163/156856897X00366
- Pennartz, C., Montijn, J., and Goltstein, P. (2013). Effects of anesthesia on ensemble patterns of Ca²⁺ activity in mouse visual cortex: a reduction in direction selectivity that is not attributable to increased correlations in ongoing cellular activity. *Soc. Neurosci. Program No.* 160.09/GG17.
- Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E. J., et al. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454, 995–999. doi: 10.1038/nature07140
- Quiroga, R. Q., and Panzeri, S. (2009). Extracting information from neuronal populations: information theory and decoding approaches. *Nat. Rev. Neurosci.* 10, 173–185. doi: 10.1038/nrn2578
- Reynolds, J. H., and Heeger, D. J. (2009). The normalization model of attention. *Neuron* 61, 168–185. doi: 10.1016/j.neuron.2009.01.002
- Ringach, D. L. (2010). Population coding under normalization. *Vision Res.* 50, 2223–2232. doi: 10.1016/j.visres.2009.12.007
- Ringach, D. L., Shapley, R. M., and Hawken, M. J. (2002). Orientation selectivity in macaque V1: diversity and laminar dependence. *J. Neurosci.* 22, 5639–5651.
- Schoups, A., Vogels, R., Qian, N., and Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature* 412, 549–553. doi: 10.1038/35087601
- Seriès, P., Latham, P. E., and Pouget, A. (2004). Tuning curve sharpening for orientation selectivity: coding efficiency and the impact of correlations. *Nat. Neurosci.* 7, 1129–1135. doi: 10.1038/nn1321
- Seung, H. S., and Sompolinsky, H. (1993). Simple models for reading neuronal population codes. *Proc. Natl. Acad. Sci. U.S.A.* 90, 10749–10753. doi: 10.1073/pnas.90.22.10749
- Smith, M. A., Jia, X., Zandvakili, A., and Kohn, A. (2012). Laminar dependence of neuronal correlations in visual cortex. *J. Neurophysiol.* 109, 940–947. doi: 10.1152/jn.00846.2012
- Smith, M. A., and Kohn, A. (2008). Spatial and temporal scales of neuronal correlation in primary visual cortex. *J. Neurosci.* 28, 12591–12603. doi: 10.1523/JNEUROSCI.2929-08.2008
- Stosiek, C., Garaschuk, O., Holthoff, K., and Konnerth, A. (2003). *In vivo* two-photon calcium imaging of neuronal networks. *Proc. Natl. Acad. Sci. U.S.A.* 100, 7319–7324. doi: 10.1073/pnas.1232232100
- Svoboda, K., Helmchen, F., Denk, W., and Tank, D. W. (1999). Spread of dendritic excitation in layer 2/3 pyramidal neurons in rat barrel cortex *in vivo*. *Nat. Neurosci.* 2, 65–73. doi: 10.1038/4569

- Tolhurst, D. D. J., Movshon, J. A., and Thompson, I. D. (1981). The dependence of response amplitude and variance of cat visual cortical neurones on stimulus contrast. *Exp. Brain Res.* 41, 414–419. doi: 10.1007/BF00238900
- Tong, F. (2003). Cognitive neuroscience: primary visual cortex and visual awareness. *Nat. Rev. Neurosci.* 4, 219–229. doi: 10.1038/nrn1055
- Ts'o, D. Y., Gilbert, C. D., and Wiesel, T. N. (1986). Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J. Neurosci.* 6, 1160–1170.
- van Duuren, E., Lankelma, J., and Pennartz, C. M. A. (2008). Population coding of reward magnitude in the orbitofrontal cortex of the rat. *J. Neurosci.* 28, 8590–8603. doi: 10.1523/JNEUROSCI.5549-07.2008
- van Duuren, E., van der Plasse, G., Lankelma, J., Joosten, R. N. J. M. A., Feenstra, M. G. P., and Pennartz, C. M. A. (2009). Single-cell and population coding of expected reward probability in the orbitofrontal cortex of the rat. *J. Neurosci.* 29, 8965–8976. doi: 10.1523/JNEUROSCI.0005-09.2009
- Vinje, W. E., and Gallant, J. L. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287, 1273–1276. doi: 10.1126/science.287.5456.1273
- Vogels, D. R. (1990). Population coding of stimulus orientation by striate cortical cells. *Biol. Cybern.* 64, 25–31. doi: 10.1007/BF00203627
- Wohrer, A., Humphries, M. D., and Machens, C. K. (2013). Population-wide distributions of neural activity during perceptual decision-making. *Prog. Neurobiol.* 103, 156–193. doi: 10.1016/j.pneurobio.2012.09.004
- Yang, T., and Maunsell, J. H. R. (2004). The effect of perceptual learning on neuronal responses in monkey visual area V4. *J. Neurosci.* 24, 1617–1626. doi: 10.1523/JNEUROSCI.4442-03.2004
- Zhang, K., Ginzburg, I., McNaughton, B. L., and Sejnowski, T. J. (1998). Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *J. Neurophysiol.* 79, 1017–1044.

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 03 March 2014; accepted: 14 May 2014; published online: 02 June 2014.

Citation: Montijn JS, Vinck M and Pennartz CMA (2014) Population coding in mouse visual cortex: response reliability and dissociability of stimulus tuning and noise correlation. *Front. Comput. Neurosci.* 8:58. doi: 10.3389/fncom.2014.00058

This article was submitted to the journal *Frontiers in Computational Neuroscience*.

Copyright © 2014 Montijn, Vinck and Pennartz. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.