

Sensitivity and specificity of a Bayesian single trial analysis for time varying neural signals

Jeff T. Mohl^{1,2,3}, Valeria C. Caruso^{1,2,3,4,5}, Surya T. Tokdar^{1,6}, J. M. Groh^{1,2,3,4*}

¹Duke Institute for Brain Sciences, Duke University, Durham, NC 27708

²Center for Cognitive Neuroscience, Duke University

³Department of Neurobiology, Duke University

⁴Department of Psychology and Neuroscience, Duke University

⁵Center for Human Growth and Development, University of Michigan, Ann Arbor, MI 48109

⁶Department of Statistical Science, Duke University

* Corresponding author, email: jmgroh@duke.edu

Abstract

We recently reported the existence of fluctuations in neural signals that may permit neurons to code multiple simultaneous stimuli sequentially across time¹. This required deploying a novel statistical approach to permit investigation of neural activity at the scale of individual trials. Here we present tests using synthetic data to assess the sensitivity and specificity of this analysis. Data sets were fabricated to match each of several potential response patterns derived from single-stimulus response distributions. In particular, we simulated dual stimulus trial spike counts that reflected fluctuating *mixtures* of the single stimulus spike counts, stable *intermediate* averages, *single* stimulus winner-take-all, or response distributions that were *outside* the range defined by the single stimulus responses (such as summation or suppression). We then assessed how well the analysis recovered the correct response pattern as a function of the number of simulated trials and the difference between the simulated responses to each “stimulus” alone. We found excellent recovery of the mixture, intermediate, and outside categories (>97% correct), and good recovery of the single/winner-take-all category (>90% correct) when the number of trials was >20 and the single-stimulus response rates were 50Hz and 20Hz respectively. Both larger numbers of trials and greater separation between the single stimulus firing rates improved categorization accuracy. The results provide a valid benchmark, and guidelines for data collection, for use of this method to investigate coding of multiple items at the individual-trial time scale.

Introduction

We recently showed that when multiple stimuli are present, some neurons exhibit activity patterns that fluctuate between those evoked by each stimulus alone¹. This dynamic code could allow the representation of all stimuli within the same population of neurons. Such fluctuations may be a widespread phenomenon in the brain, but would be overlooked using conventional analysis methods that investigate mean activity pooled across trials. Of particular interest are cases in which the time-and-trial-pooled responses evoked by multiple stimuli appear to reflect the average of the responses to each stimulus presented in isolation. This phenomenon, known as divisive normalization², has been observed in visual brain areas such as V1 and MT³ as well as other sensory and cognitive domains^{4–9}. However, such responses could either reflect a true averaging of the two stimuli/conditions, producing a consistent stable intermediate level of firing on each trial, or could reflect a dynamic code that flexibly shifts between the individual stimuli across trials.

To evaluate neural responses on a single trial basis, the novel statistical approach introduced in Caruso et al. (2018) characterizes the distribution of spike counts elicited in response to two simultaneous stimuli using Bayesian inference. Here we provide a general assessment of the sensitivity and specificity of that approach by simulating known neural responses as benchmark cases. In particular, we investigate how the analysis performs as we parametrically varied the data sample size (number of trials) and the difference between spike counts across conditions.

We demonstrate that our approach accurately categorizes synthetic neural data into expected categories. The robustness of the results depends heavily on sample size, as well as on firing rate differences between the two single cue conditions. Importantly, the model performs very well under reasonable experimental values (20 trials per condition, 60% firing rate change between conditions). Finally, we show that the model gracefully handles datasets that do not exactly match any of the tested hypotheses. These results demonstrate the viability of the analysis method and provide constraints for interpretation of actual neural data.

Experimental Rationale and Procedures

Neural encoding patterns to be assessed

For simplicity, our approach focused on the case of two simultaneously presented stimuli (dual stimuli) but can be extended to a larger number of stimuli. We consider an experimental setup in which a neuron's response is recorded in 3 interleaved conditions: in the presence of a single stimulus "A", a single stimulus "B", or both stimuli A and B ("AB"). We considered four possible response distributions to dual stimuli, in relation to the distributions observed when only one stimulus is present (Figure 1).

- 1) Neurons might respond to only one of the stimuli, and do so consistently (i.e respond to the same one) across trials. One way this could occur would be if only one stimulus is located in a neuron's receptive field, but it might also apply when both stimuli are in the receptive field (sometimes referred to as a winner-take-all encoding). We label this possibility "single".
- 2) The responses to dual stimuli might be greater than the maximum or less than the minimum of the single-stimulus responses. This category includes enhancement/summation, as well as suppression of the response to one stimulus by another. We refer to this as "outside".
- 3) The responses to dual stimuli are a *consistent* weighted average of the responses evoked by each stimulus alone. Here, the dual stimulus responses are between the bounds set by the two single stimulus responses, and cluster around a stable intermediate value. We refer to this case as "intermediate".
- 4) The responses to dual stimuli may fluctuate such that on each trial the neuron appears to be responding to only one of the two stimuli. We term this possibility "mixture" because it reflects a mixture of two distributions of A and B stimulus responses. This is analogous to a winner-take-all except that the neuron is switching across trials rather than encoding the same stimulus each trial. Like the "intermediate" category, there could be a weighting factor such that a higher proportion of trials favor one stimulus over the other.

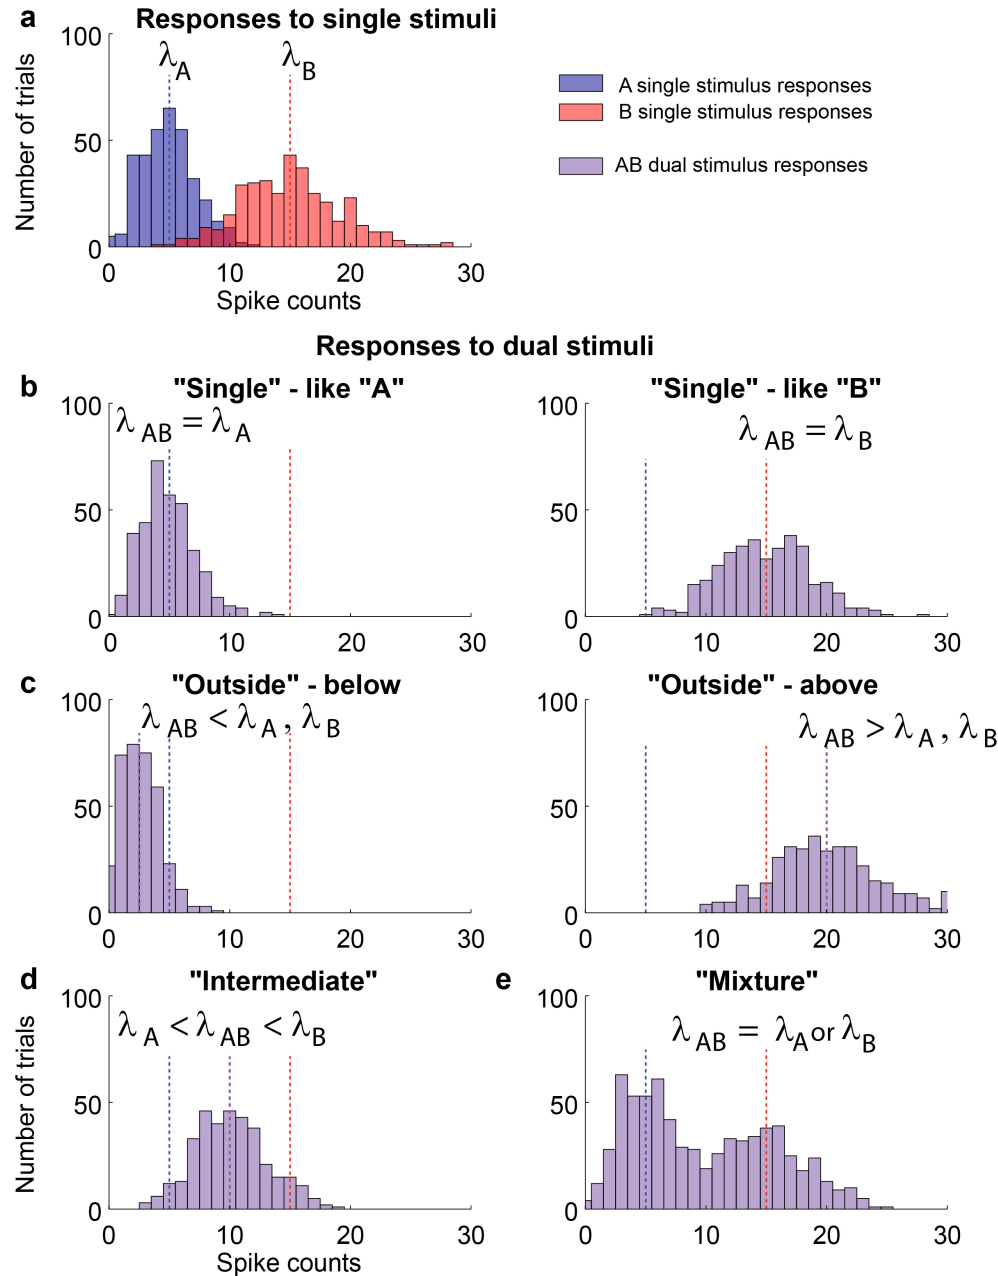


Figure 1

Four possible response patterns to dual stimuli trials, in relation to the responses observed to the component stimuli when presented individually. (a) Single stimulus trials were modeled as evoking spike counts distributed according to a Poisson process with rates λ_A , blue, or λ_B , red. (b) Responses on dual stimulus trials follow a Poisson λ_{AB} matching either λ_A , left, or λ_B , right. (c) The Poisson rate λ_{AB} on dual stimulus trials is less than (left) or greater than (right) those observed in single stimulus trials. In these simulations, λ_{AB} was set to $0.5 * \lambda_A$ (left) or $\lambda_A + \lambda_B$ (right). (d) Spike counts derived from a Poisson process with a rate λ_{AB} between λ_A and λ_B . (e) Spike counts drawn from a mixture of two Poissons with rates λ_A and λ_B .

Model construction, Bayesian model comparison, and synthetic data

These four possibilities can be formalized on the basis of how the spike distributions on combined stimulus trials AB compare to those observed when only A or B are presented alone. If A and B elicit spike counts according to Poisson distributions $Poi(\lambda_A)$ and $Poi(\lambda_B)$, then we can ask which of four competing hypotheses best describe the spike counts observed on combined AB trials:

- a) Single: $F = Poi(\lambda)$ for either $\lambda = \lambda_A$ or $\lambda = \lambda_B$, with λ constant across trials
- b) Outside: $F = Poi(\lambda)$ for some unknown $\lambda \notin (\min(\lambda_A, \lambda_B), \max(\lambda_A, \lambda_B))$
- c) Intermediate: $F = Poi(\lambda)$ for some unknown $\lambda \in (\min(\lambda_A, \lambda_B), \max(\lambda_A, \lambda_B))$
- d) Mixture: $F = \alpha \cdot Poi(\lambda_A) + (1 - \alpha) \cdot Poi(\lambda_B)$ for some unknown $\alpha \in (0,1)$

The plausibility of each of these models was determined by computing the posterior probabilities of each model given the data, with a default Jeffreys' prior¹⁰ on each of the model specific rate (λ) parameters and on the mixing probability parameter (α). Each model was given a uniform prior probability (1/4) and posterior model probabilities were calculated by computation of relevant intrinsic Bayes factors¹¹ (see appendix S1 for a thorough description of the models and model selection strategy).

To evaluate the sensitivity and specificity of this method, we built synthetic neuronal spiking datasets to match each of the four potential encoding strategies tested by the model. Consistent with our previous study¹, we focused on response patterns that could be modeled as deriving from Poisson distributions. In principle, the approach could be extended to other forms of response distributions, but this is beyond the scope of this work.

Data files were generated as spike times drawn using an independent Poisson point process sampled at 1 ms intervals, with constant mean firing rate for 1000 ms (Figure 2a-c). For A and B (single stimulus) trials, Poisson rates were assigned a priori to reflect a range of realistic firing rates for a single neuron presented with different stimuli. AB (combined stimulus) trials for each dataset were generated based on the chosen A and B firing rates and in a manner consistent with one of the four potential hypotheses. For the "single" hypothesis the AB data were generated using a single Poisson with rate λ_{AB} equal to the highest of the component rates, i.e. $\max(\lambda_A, \lambda_B)$. For the "outside" hypothesis, the rate λ_{AB} was set 20% higher than $\max(\lambda_A, \lambda_B)$. For the "intermediate" hypothesis, λ_{AB} was equal to the mean of A and B rates $\lambda_{AB} = 0.5(\lambda_A) + 0.5(\lambda_B)$. For the across trial switching (mixture) model, the data were generated using the same Poisson process, but each trial was randomly chosen to be drawn from either $poi(\lambda_A)$ or $poi(\lambda_B)$ with equal probability. This results in a dataset for which the across trial average firing rate is equal to the average of the λ_A and λ_B rates, but individually each trial is better described as deriving from either the λ_A or λ_B response distributions. Note that it is nearly impossible to tell by visual inspection of a raster plot when a neuron has such a mixed response pattern, even when the trials are sorted as they are in Figure 2c, but the pattern becomes more evident in histograms of the trial-wise spike counts (Figure 2d).

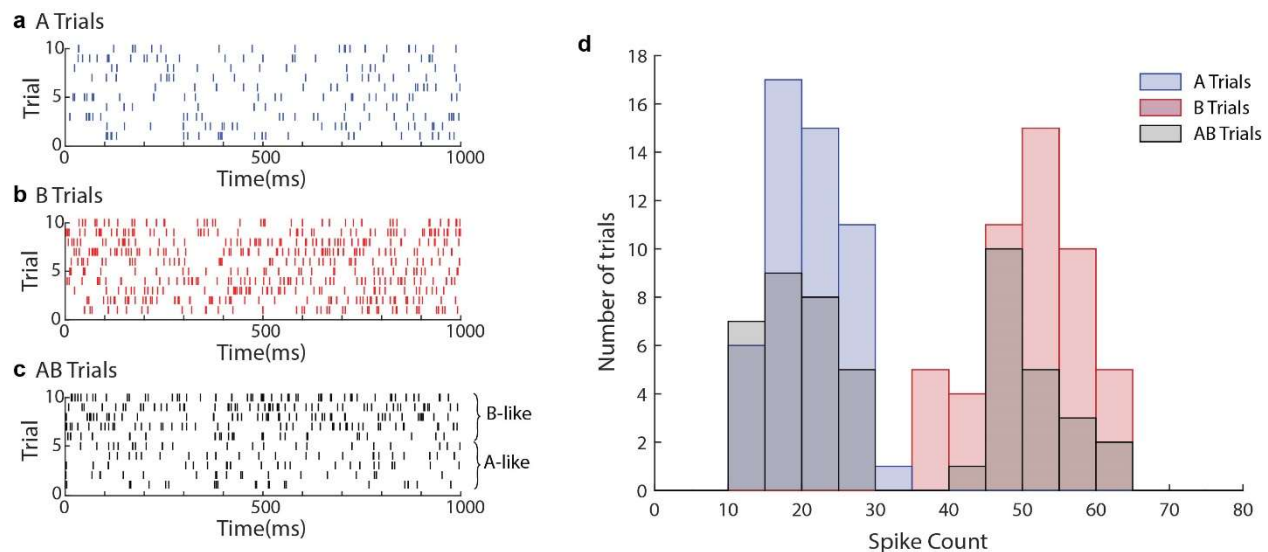


Figure 2

One example synthetic dataset. (a-c) Raster plots for a synthetic dataset built to match the across trial switching hypothesis, where blue rows (a) are single A trials, red rows (b) are single B trials, and black rows (c) are AB trials. AB trials are drawn randomly to match either A or B rates and are sorted so that B-like rates are towards the top of the raster. Even with sorting, this pattern is challenging to see with the naked eye, highlighting the need for analytical methods. (d) Whole trial spike count histogram for 50 repetitions of A, B, and AB trials. From this plot, the bimodality of AB trials for the switching condition is more apparent.

Multiple datasets were generated using this strategy in order to test the power and reliability of the analysis under plausible experimental conditions. These datasets varied both the number of trials per condition (from 5 to 50) and the firing rate separation between A and B trials (with the B rate varying between 30 and 100 Hz, while the A rate was held fixed at 20 Hz). Individual triplet pairs were generated under each of these conditions, analogous to running 100 individual cells through the analysis. This set of parameters was used for all conditions tested, including datasets constructed to not exactly match any of the hypotheses, discussed in the final section of the results.

Code and data availability

Code specific to this paper can be found on GitHub at https://github.com/jmohl/mplx_tests which includes the code used to generate synthetic data for this manuscript. The exact files used to generate plots are available upon request. Source code for the Neural Mixture Model available at <https://github.com/tokdarstat/Neural-Multiplexing>.

Results

Neural Mixture Model accurately characterizes synthetic data built to match hypotheses

The desired analysis outcome is for the output to match the input. That is, data explicitly generated to match the single hypothesis should be correctly labeled as “single”, data generated as

“outside” should be labelled “outside”, etc. Figure 3 illustrates that this is largely the case. The series of simulations shown here involved 20 A trials simulated with $\lambda_A = 20$ Hz, 20 B trials with $\lambda_B = 50$ Hz, and 20 AB trials generated according to the various methods specified above. The “mixture” and “intermediate” categories perform exceptionally well, with 100/100 “mixture” and 99/100 “intermediate” datasets labeled correctly with >95% confidence (dark black bars). This distinction is critical, as these two conditions would produce exactly the same mean rate when averaging across trials, making them indistinguishable using typical neural analysis strategies which average across trials in order to reduce noise. “Single” and “outside” datasets were also correctly labeled in the majority of cases (90/100 and 97/100, respectively), although these hypotheses are not the focus of our analysis as they can be differentiated more easily using simpler statistical methods.

Although the category “single” was correctly identified as the best model for the dataset simulated under the “single” hypothesis 90% of the time, the posterior probability or confidence level did not reach the 95% level observed for the other models. This is due to the narrow definition of this category: response rates on AB trials must be indistinguishable from those occurring on either the A or B trials. All other categories include a range of possibilities which admits this hypothesis as a boundary case (i.e. a weighted average with the weight for A set to 1). Therefore, these models are all fairly competitive in explaining data that is generated to match the “single” case, which explains the low posterior probability of this model. For this reason, it is better to consider the “single” category as reflective of a null hypothesis, where there is no interaction at all between stimuli.

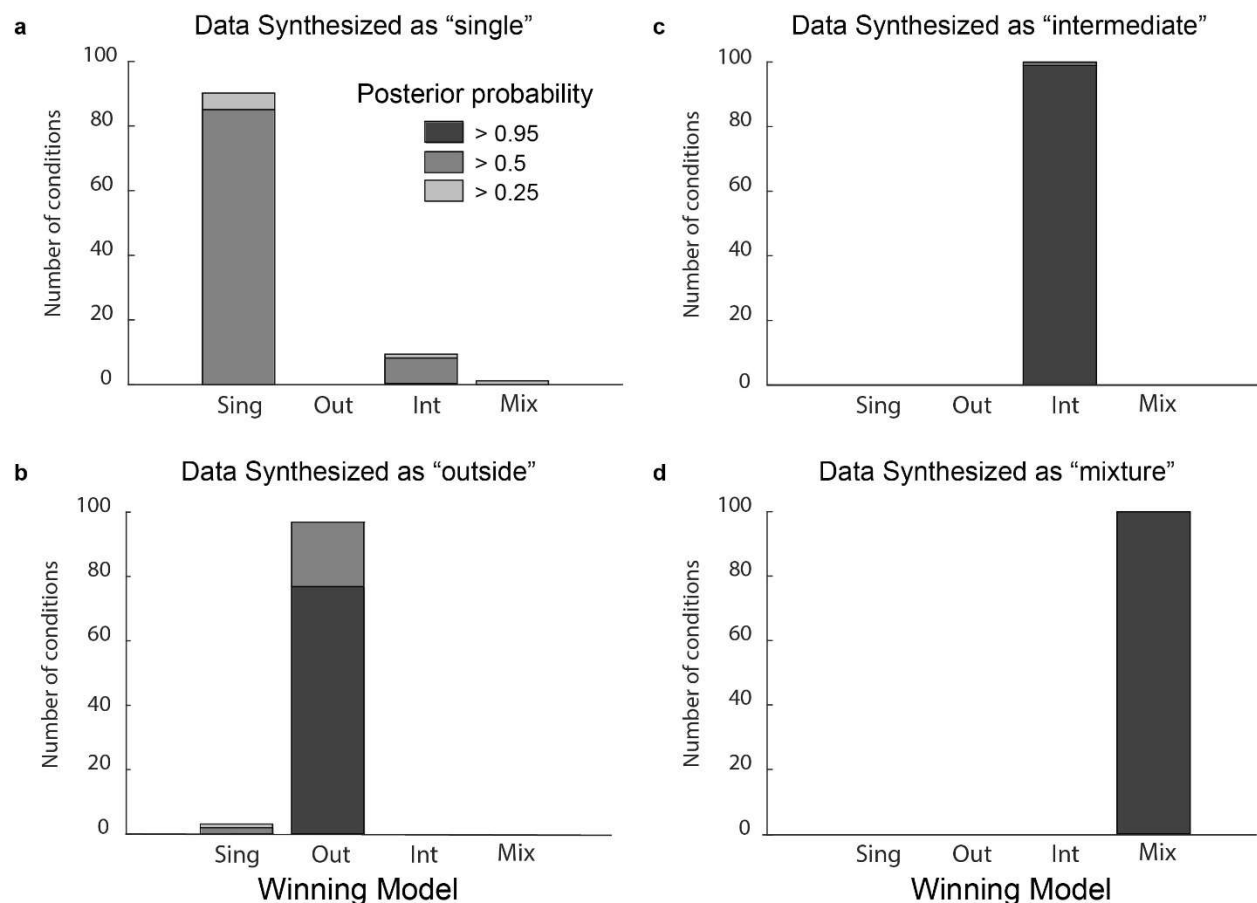


Figure 3

The analysis method correctly categorizes synthetic datasets created to match each model. The shading of the bars indicates the posterior probability with which each individual run of synthetic data ($n=100$) is assigned to a given category. Of particular interest is the very strong separation between intermediate and mixture datasets, as this discrimination is not possible when considering only firing rates averaged across trials. Parameters used for this figure: $\lambda_A = 20\text{Hz}$, $\lambda_B = 50\text{ Hz}$, number of trials in each run = 20 per stimulus condition (60 overall).

Dependence on number of trials and difference between A and B response distributions

The accuracy of this characterization depended on both the amount of data and the difference between the response distributions on A and B trials. The dependence on the number of trials is best appreciated when considering fairly similar A and B response distributions, such as $\lambda_A = 20$ vs $\lambda_B = 30\text{Hz}$ as shown in Figure 4a (right), which depicts the average posterior probability value for the correct model as a function of the number of trials. Even with this modest separation between the A and B response patterns, increasing the number of trials per condition allowed the analysis to better characterize the underlying rates, and therefore better discriminate between competing hypotheses. “Single”, “Intermediate” and “Mixture” had average posterior probability values >0.3 for $N=5$ trials, but performance improves steadily to average posterior probability values of >0.75 for $N=50$ trials. When response distributions are moderately separated, $\lambda_A = 20$ vs $\lambda_B = 50\text{Hz}$ (the same separation used in Figure 3), performance rises more rapidly for all models except “single”. At $N=5$, posterior probability values range from 0.4 for “single” to 0.8 for “mixture”. At $N=30$, posterior probability values equal approximately 1 for “mixture”, “intermediate” and “outside”. Further increasing the firing rate separation to $\lambda_A = 20$ vs $\lambda_B = 100\text{Hz}$ results in very high posterior probabilities of >0.95 even at $N=5$ for “mixture” and “intermediate”; this level is achieved for “outside” at $N=10$.

These figures give a rough sense of the sensitivity of our analysis, demonstrating that the analysis becomes more reliable as more trials per condition are added until reaching asymptote around 30 trials/condition for a 50 Hz vs 20 Hz comparison (Figure 4b). Similarly, increasing the difference in spike count between A and B conditions also improves specificity in the analysis, allowing for accurate characterization with as few as 5 trials (Figure 4c). Although these data were constructed under ideal conditions (the data perfectly matches one of the tested hypotheses), they can be used as a guide for how much data should be collected in order to obtain satisfactory results in a real dataset.

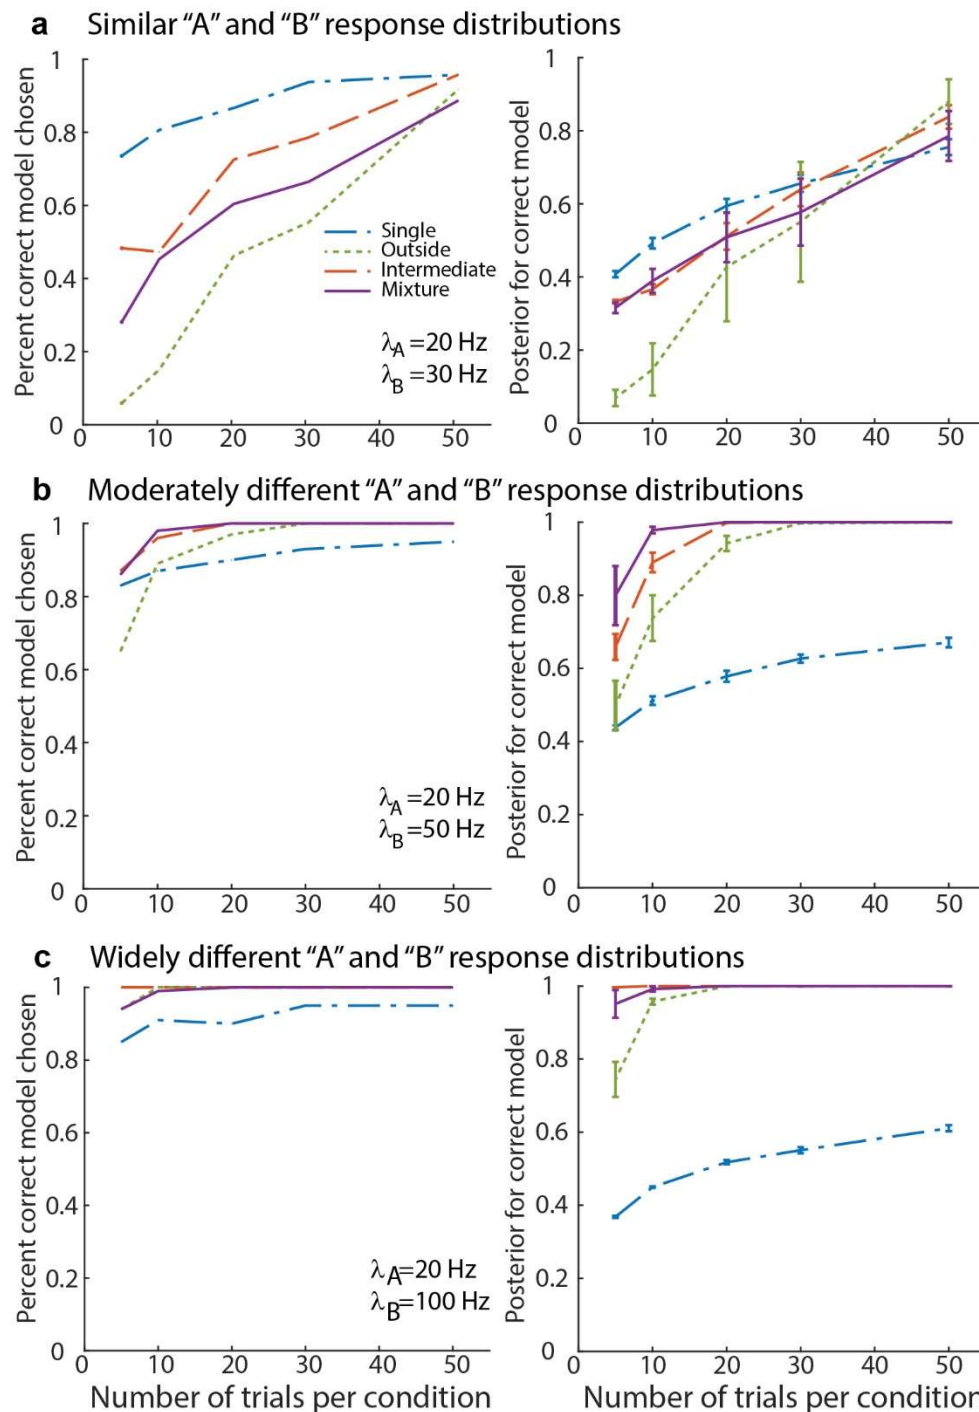


Figure 4

Increasing number of trials or separability of conditions improves accuracy of model comparison. **(a)** Left, percent of triplets which were correctly categorized, split by dataset type, for increasing number of trials per conditions; $\lambda_A = 20$ Hz, $\lambda_B = 30$ Hz. Right, mean and variance of posterior probability for correct model across triplets **(b & c)** same as in **a** but with λ_B set to 50z and 100 Hz respectively. Fewer trials are needed when responses are very different between A and B trials.

Model results are informative even for datasets that do not perfectly match hypotheses

Though this modeling strategy was meant to test between discrete hypotheses, it is unlikely that real neural signals perfectly and uniquely match any one of these scenarios. Therefore, it is important that the analysis accurately reflect deviations from exact hypothesis matches. Here, we consider two potential deviations from the circumstances considered above: weighted averaging of A and B stimuli and incomplete switching between A and B rates.

For weighted averaging datasets the AB trials were generated as in the “intermediate” condition above, except that the AB rate was set to be closer to the A rate than the B rate: $\lambda_{AB} = 0.75 * \lambda_A + 0.25 * \lambda_B$. Because the model returns both a classification and a posterior probability (reflecting the model’s confidence in that classification), we expect that this type of dataset will result in a spread across single and average classifications, but with lower confidence in this assessment. This is indeed the case, as the model returns primarily the intermediate category, with some single winners, but with much lower posterior probabilities than the well matched datasets (Figure 5a, compare with Figure 3c).

We also tested a form of incomplete switching, where stimuli show strong fluctuations but do not quite exactly match the A and B rates. These datasets were generated using the same strategy as the mixture datasets described above, except that A-like or B-like trials were generated with a slightly shifted mean firing rate. Multiple degrees of similarity were tested, but two (80% and 90% similarity) are presented here. From these, the analysis accurately describes a 90% switch as mixture (Figure 5b), but around 80% similarity it begins to interpret many datasets as average (Figure 5c). In some ways this result is expected, as continuing to reduce the similarity would eventually result in a condition that was indistinguishable from true averaging. However, these results demonstrate the high specificity of our analysis for the mixture category, enforcing a strong definition of mixture (literally switching between rates closely matched to the A and B rates, rather than any amount of fluctuation).

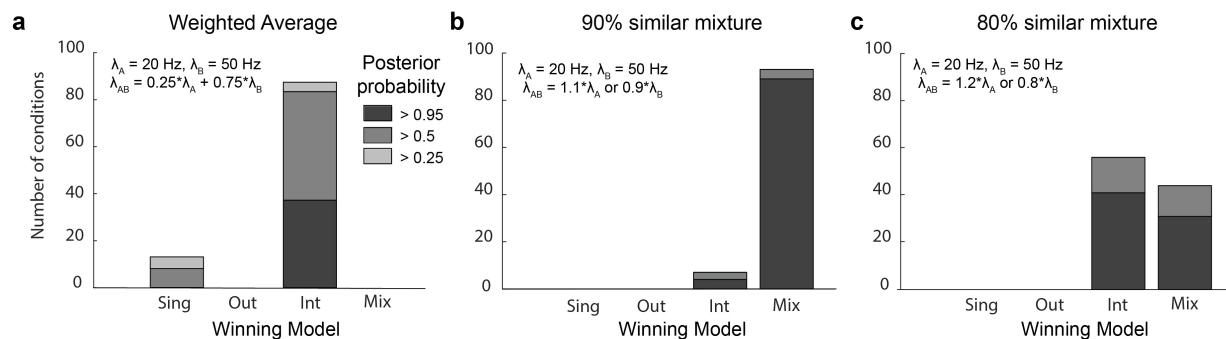


Figure 5

Datasets which do not exactly match the canonical hypotheses are descriptively categorized by model comparison. (a) A dataset generated to reflect weighted average of A and B stimuli is weakly categorized as intermediate with some triplets categorized as singles. (b) Mixture trials generated to alternate between values shifted 10% from the true A and B firing rates are primarily categorized as mixtures. (c) Mixture trials with rates shifted 20% from the true A and B rates are categorized as either mixture or average with equal probability, consistent with the fact that these trials would be much more difficult to discriminate from the true averaging hypothesis. Parameters used for this figure: $\lambda_B = 50\text{Hz}$, $\lambda_A = 20 \text{ Hz}$, number of trials = 20.

Discussion

There is broad interest in understanding the nature and significance of firing patterns in the brain. It is well known that such firing patterns are variable in the face of identical, highly controlled experimental conditions (such as the presentation of the same stimulus in the same context). While many studies have viewed this variability as deleterious “noise” that if unsolved would undermine the ability of the brain to perform its essential tasks^{12–16}, we and others have sought explanations under the possibility that certain forms of such variation may contribute in a positive fashion to brain function^{1,17–24}. In particular, we have successfully modeled variation in whole-trial spike counts to multiple stimuli as being drawn from the observed distributions of spike counts to those same stimuli when presented individually¹.

Here, we benchmarked this analysis on synthetic datasets to provide insight into the sensitivity and specificity of our analysis method as a function of trial counts and the separation between the distributions of spike counts elicited by the component individual stimuli. When response separations are large, e.g. the mean rate for one stimulus condition is 5X the rate for the other, the analysis method can successfully distinguish among the 4 competing hypotheses with as few as $n=5$ trials for each condition ($n=15$ overall). Smaller response separations can be compensated for by collecting more trials to achieve similarly good results. Finally, even when the conditions do not exactly match the assumptions, such as if the component response rates in the “mixture” condition do not precisely match those observed when the single stimuli were presented individually, correct classifications greatly outnumber incorrect ones. Critically, the analysis is conservative against the “mixture” hypothesis in these cases, demonstrating that data which is best fit by this model is truly fluctuating between the responses to single stimuli at the single trial level.

These results suggest that the analysis tested here are suitable for most electrophysiological datasets which match the A, B, AB format. Datasets with at least 10 trials per condition and an average response difference between conditions of approximately 40% (relative to peak rate) can be expected to correctly categorize over 90% of cells, with this accuracy increasing as number of trials increases. Practically, this means that our analysis is particularly well suited for recordings in sensory or motor brain regions with the pronounced tuning and firing rate changes required to differentiate responses at the single trial level.

A situation not tested here is the case in which the response distributions are not derived from Poisson distributions. In our previous work¹ we excluded conditions in which the responses to individual stimuli did not satisfactorily resemble Poisson distributions in order to ensure that our model assumptions were not violated, but this has several downsides. First, it is difficult to have confidence in the success of this exclusion criterion: failing to reject the Poisson assumption is not the same as confirming its validity. Second, a considerable amount of data is excluded in this fashion (as much as 25–50%, depending on dataset, before even considering other exclusion criteria). Finally, there is significant evidence in the literature that spike counts in many brain areas are more variable than would be suggested by a Poisson distribution^{15,25–29}. For all of these reasons, it will be important to both test the model with data sets that violate this assumption and extend the analysis method to include other response distributions such as negative binomials.

Given the broad interest in both noise as a potential limitation on neural representations and in divisive normalization as an elemental computation in sensory processing – with recent suggestions that this process may be impaired in conditions such as autism^{30–32} -- it will be increasingly important to probe patterns of responses at least at the individual trial level. The statistical tools described in the present paper represent an important first step towards uncovering fluctuating patterns in neural activity that may permit greater amounts of information to be encoded in the spike trains of individual and populations of neurons.

1. Caruso VC, Mohl JT, Glynn C, et al. Single neurons may encode simultaneous stimuli by switching between activity patterns. *Nat Commun*. 2018;9(1):2715. doi:10.1038/s41467-018-05121-8
2. Carandini M, Heeger DJ. Summation and division by neurons in primate visual cortex. *Science (80-)*. 1994;264(5163):1333-1336. doi:10.1126/science.8191289
3. Britten KH, Heuer HW. Spatial summation in the receptive fields of MT neurons. *J Neurosci*. 1999;19(12):5074-5084.
4. Kozlov AS, Gentner TQ. Central auditory neurons have composite receptive fields. *Proc Natl Acad Sci*. 2016;113(5):1441-1446. doi:10.1073/pnas.1506903113
5. Ohshiro T, Angelaki DE, DeAngelis GC. A normalization model of multisensory integration. *Nat Neurosci*. 2011;14(6):775-782. doi:10.1038/nn.2815
6. Bao P, Tsao DY. Representation of multiple objects in macaque category-selective areas. *Nat Commun*. 2018;9(1). doi:10.1038/s41467-018-04126-7
7. Louie K, Khaw MW, Glimcher PW. Normalization is a general neural mechanism for context-dependent decision making. *Proc Natl Acad Sci*. 2013;110(15):6139-6144. doi:10.1073/pnas.1217854110
8. Olsen SR, Bhandawat V, Wilson RI. Divisive normalization in olfactory population codes. *Neuron*. 2010. doi:10.1016/j.neuron.2010.04.009
9. Reynolds JH, Heeger DJ. The Normalization Model of Attention. *Neuron*. 2009. doi:10.1016/j.neuron.2009.01.002
10. Berger J. The case for objective Bayesian analysis. *Bayesian Anal*. 2006;1(3):385-402. doi:10.1214/06-BA115
11. Berger J, Pericchi L. The intrinsic Bayes factor for linear models. *Bayesian Stat*. 1996.
12. Shadlen MN, Newsome WT. Noise, neural codes and cortical organization. *Curr Opin Neurobiol*. 1994;4(4):569-579. doi:10.1016/0959-4388(94)90059-0
13. Shadlen MN, Newsome WT. Is there a signal in the noise? *Curr Opin Neurobiol*. 1995;5(2):248-250. doi:10.1016/0959-4388(95)80033-6
14. Zohary E, Shadlen MN, Newsome WT. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*. 1994;370(6485):140-143. doi:10.1038/370140a0
15. Shadlen MN, Newsome WT. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci*. 1998;18(10):3870-3896.
16. Softky WR, Koch C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci*. 1993;13(1):334-350.
17. Akam T, Kullmann DM. Oscillatory multiplexing of population codes for selective communication in the mammalian brain. *Nat Rev Neurosci*. 2014;15(2):111-122. doi:10.1038/nrn3668
18. Hoppensteadt FC, Izhikevich EM. Thalamo-cortical interactions modeled by weakly connected oscillators: Could the brain use FM radio principles? In: *BioSystems*. Vol 48. ; 1998:85-94. doi:10.1016/S0303-2647(98)00053-7
19. Jezek K, Henriksen EJ, Treves A, Moser EI, Moser MB. Theta-paced flickering between place-cell maps in the hippocampus. *Nature*. 2011;478(7368):246-249. doi:10.1038/nature10439
20. Li K, Kozyrev V, Kyllingsbæk S, Treue S, Ditlevsen S, Bundesen C. Neurons in Primate Visual Cortex Alternate between Responses to Multiple Stimuli in Their Receptive Field. *Front Comput Neurosci*. 2016;10:141. doi:10.3389/fncom.2016.00141
21. Lisman JE, Idiart MAP. Storage of 7 ± 2 short-term memories in oscillatory subcycles. *Science (80-)*. 1995;267(5203):1512-1515. doi:10.1126/science.7878473
22. Lisman JE, Jensen O. The Theta-Gamma Neural Code. *Neuron*. 2013;77(6):1002-1016. doi:10.1016/j.neuron.2013.03.007
23. McLelland D, VanRullen R. Theta-Gamma Coding Meets Communication-through-Coherence: Neuronal Oscillatory Multiplexing Theories Reconciled. *PLoS Comput Biol*. 2016;12(10).

- doi:10.1371/journal.pcbi.1005162
24. Meister M. Multineuronal codes in retinal signaling. *Proc Natl Acad Sci.* 2002;93(2):609-614. doi:10.1073/pnas.93.2.609
25. Carandini M. Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS Biol.* 2004;2(9). doi:10.1371/journal.pbio.0020264
26. Gur M, Beylin A, Snodderly DM. Response variability of neurons in primary visual cortex (V1) of alert monkeys. *J Neurosci.* 1997;17(8):2914-2920.
27. Barberini CL, Horwitz GD, Newsome WT. A Comparison of Spiking Statistics in Motion Sensing Neurons of Flies and Monkeys. In: *Motion Vision.* ; 2011:307-320. doi:10.1007/978-3-642-56550-2_16
28. Recanzone GH. Rapidly induced auditory plasticity: the ventriloquism aftereffect. *Proc Natl Acad Sci U S A.* 1998;95(3):869-875.
29. Amarasingham A. Spike Count Reliability and the Poisson Hypothesis. *J Neurosci.* 2006;26(3):801-809. doi:10.1523/jneurosci.2948-05.2006
30. Rosenberg A, Sunkara A. Does attenuated divisive normalization affect gaze processing in autism spectrum disorder? A commentary on Palmer et al. (2018). *Cortex.* 2019;111:316-318. doi:10.1016/j.cortex.2018.06.017
31. Rosenberg A, Patterson JS, Angelaki DE. A computational perspective on autism. *Proc Natl Acad Sci.* 2015;112(30):9158-9165. doi:10.1073/pnas.1510583112
32. Van de Cruys S, Vanmarcke S, Steyaert J, Wagemans J. Intact perceptual bias in autism contradicts the decreased normalization model. *Sci Rep.* 2018;8(1). doi:10.1038/s41598-018-31042-z