Characterizing neural coding performance for populations of sensory neurons: comparing a weighted spike distance metrics to other analytical methods.

G. Marsat

Department of Biology
West Virginia University, Morgantown, WV 26505

Corresponding Author:
G. Marsat
Dept. Biology
West Virginia University
53 Campus Drive
Morgantown, WV 26505
Email: gary.marsat@mail.wvu.edu

Acknowledgments: Funding was provided by West Virginia University and the National Science foundation grant IOS-1557846. Future versions of this article may incorporate the contributions to this research of Jordan Drew, Kevin Daly and Phil Chapman.

Conflict of interest statement: The authors declare no competing financial interests.

ABSTRACT

The identity of sensory stimuli is encoded in the spatio-temporal patterns of responses of the neural population. For stimuli to be discriminated reliably, differences in population responses must be accurately decoded by downstream networks. Several methods to compare the pattern of responses and their differences have been used by neurophysiologist to characterize the accuracy of the sensory responses studied. Among the most widely used analysis, we note methods based on Euclidian distances or on spike metric distance such as the one proposed by van Rossum. Methods based on artificial neural network and machine learning (such as self-organizing maps) have also gain popularity to recognize and/or classify specific input patterns. In this brief report, we first compare these three strategies using dataset from 3 different sensory systems. We show that the input-weighting procedure inherent to artificial neural network allows the extraction of the information most relevant to the discrimination task and thus the method performs particularly well. To combine the ease of use and rapidity of methods such as spike metric distances and the advantage of weighting the inputs, we propose a measure based on geometric distances were each dimension is weighted proportionally to how informative it is. In each dimension, the overlap between the distributions of responses to the two stimuli is quantified using the Kullback-Leibler divergence measure. We show that the result of this Kullback-Leibler-weighted spike train distance (KLW distance) analysis performs as well or better than the artificial neural network we tested and outperforms the more traditional spike distance metrics. We applied information theoretic analysis to Leaky-Integrate-and-Fire model neuron responses and compare their encoding accuracy with the discrimination accuracy quantified through these distance metrics to show the high degree of correlation between the results of the two approaches for quantifying coding performance. We argue that our proposed measure provides the flexibility, ease of use sought by neurophysiologist while providing a more powerful way to extract the relevant information than more traditional methods.

INTRODUCTION

Encoding of sensory signals is typically mediated by the patterned spiking responses of a population of sensory neurons (Stanley, 2013). Depending on the model system, various aspect of this spatio-temporal pattern can represent the identity-information carried by the population response and the reliability in the response reflect the accuracy of encoding (Rieke, 1997). Various methods have been developed to characterize the encoding accuracy and better understand the coding strategy. Three commonly used methods are: 1- Analyses based on information theoretic calculations (Borst and Theunissen, 1999); 2- ROC analysis based on spike metric distances (Victor, 2005) and 3- Pattern classifiers based on artificial neural network (Barrett et al., 2019). The latter is not yet as widely used by traditional neurophysiologist but has proved to be a powerful approach in a broad range of research aimed at reliably recognizing patterns.

In this brief report, we compare these different methods by using three dataset representative of the lower levels of sensory systems. We use recordings from olfactory responses of projection neurons (PNs) in the antennal lobe (AL) of moth (Daly et al., 2016b), pyramidal cell (PC) responses to communication stimuli in the electrosensory lateral line lobe (ELL) of weakly electric fish (Allen and Marsat, 2018) and responses of leaky integrate and fire (LIF) model neurons to frozen white noise in a linear regime. Each system encode the relevant information in different aspect of the population response's spatiotemporal pattern. We ask how accurately we can extract the relevant information from these responses and reliably discriminate between sensory stimuli. We present a new variant of spike distance metric analysis relying on the weighting of inputs and show that it performs most efficiently.

METHODS:

Datasets:

In vivo neural recordings were obtained from previous research and the details of the experiments are given in the corresponding publications (Daly et al., 2016a; Allen and Marsat, 2018). One second long PN recordings from the AL of moth where performed during and after the presentations of 6 different concentrated odors puffs (100 ms long) and one blank stimulus. ELL responses of PCs to electrocommunication stimuli (3 different type 1 chirps – "big chirps) presented on high frequency beats) were recorded in OFF cells and we analyzed a narrow window around the chirp timing.

LIF model responses were created by adjusting the model parameters (threshold, capacitance, input current) to obtain a firing rate modulation linearly correlated with the input. In a first version (figures 2-6) the threshold and capacitance of the neurons in the population were slightly varied by drawing from a normally distributed range of values. Since this heterogeneity could potentially lead to small non-linearities due to threshold and saturation, we used an homogeneous population for the comparisons with information theoretic measures in figures 7-8.

Stimuli consisted of 1 second long, low-pass filtered, frozen white noise and 10 different noise patterns were used as the stimulus set. A population of 100 neurons was created with several repeats of responses to each stimulus were produced by varying the noise included in the LIF model. Noise level and signal to noise ratios were varied in figure 6-8 but were fixed at levels that lead to medium discrimination performance in previous figures.

ROC analyses:

Receiver operating characteristic (ROC) curves were created by quantifying the similarity in spike trains in several different ways. The Euclidean distance (e.g. Staudacher et al., 2009) of the population responses were taken by representing each response in a space were each neural response at each time point is a dimension. Therefore, if we are considering a population of 3 neurons each with response composed of 10 time points, we obtain a space with 30 dimensions. The distances between responses to 2 different stimuli are then compared to the distribution of distances to the same stimulus. These distributions are then used for ROC analysis leading to an error level. The error level can be calculated for populations that vary in size; in the figures we performed the analysis with population sizes that varied from single neurons to one neuron less than our smallest dataset (i.e. 17 neurons).

Another variant of the analysis relied on the spike metric distance described by van Rossum (2001). The application of this method to population responses in past research (e.g. Marsat and Maler, 2010) used population responses averaged across neurons rather than keeping each neurons as separate dimensions. To compare directly the method based on the Euclidian distance described above and this method based on the van Rossum metric we can average the neural responses across neurons before performing the Euclidian distance calculation.

SOM neural net:

Differences in response patterns were also assessed by using an artificial neural network with unsupervised learning that outputs a “self-organized map” (SOM) of the ensemble of inputs. SOM neural network can thus be used to cluster sets of inputs. We used the SOM neural network tool build into Matlab (R2017b) and set it to cluster the array of inputs, from responses to 2 different stimuli, into 2 outputs. In a training phase the network learns to sort the population responses into 2 maximally different sets of outputs without it being specified which response where elicited by what stimulus. A training dataset is then fed to the neural network and the output is evaluated for accuracy: if all responses from one stimulus is sorted into the first output and all responses to the other into the second output, the network performed with 0% error any mismatch leads to a poorer performance level.

KLW method:

The Euclidean distance used with an ROC analysis was enhanced by first weighting each dimension of the Euclidean space. To do so, the distributions of responses (firing rate) for the two stimuli in a given dimension (i.e. a given neuron at a given time point) is compared using the Kullback-Leibler (KL) divergence. The KL values from a set of population responses is then normalized to an average of 1. The firing rate at each time point and for each neuron is then

multiplied with the respective normalized KL value. The KL values can be average across time for a given neuron or even be subject to various time- or firing rate-dependent rules to replicate various biologically plausible decoding mechanisms.

Information measure:

Information content of LIF responses were evaluated using information estimates based on coherence measures (Borst and Theunissen, 1999). To do so, responses from different neurons were averaged and stimulus-response coherence was calculated. Upper-bound coherence was also calculated using the response-response coherence. Information is calculated based on these coherence using the established formula (Borst and Theunissen, 1999).

RESULTS

Using data set from different model systems allow us to make more general statements about the performance of the analysis but also to determine how different coding aspects are revealed by the methods. For example, the PN olfactory responses show a clear diversity in response pattern across neurons (Fig. 1) and it is known that the identity of the neurons active for a given stimulus is key in encoding the stimulus. In contrast, the PC electrosensory responses show a qualitatively more similar response pattern across neurons and it might not be critical to evaluate the neuron-specific response pattern to extract all the information. Another example can be given by comparing the temporal pattern of responses between the olfactory and model responses. For model responses, the detailed temporal structure is key in encoding the stimulus identity and the average firing rate is similar across stimuli. For model neurons it is the contrary, overall firing rate varies largely from neuron to neuron across stimuli but firing rate varies less rapidly across time. The methods discussed in this paper can be used to demonstrate and quantify differences in coding strategies such as these.

Established spike-distances methods.

We first compare two established ways to quantify differences between spike train: the Euclidean distance and the van Rossum spike distance metric. For both with use an ROC analysis to compare the results and display the discrimination error levels as a function of the number of neurons used for the analysis (Fig. 2). The calculation was performed hundreds of time taking different subsets of neurons, repeats and stimuli pairs and averaged; standard deviations displayed shows the variability in the averages across stimuli pairs.

There is an appreciable difference in discrimination performance for the different datasets. The electro-communication stimuli were weak and thus the responses provided a fairly noisy representation of stimulus identity. As a result, the discrimination error levels were high even when the analysis based itself on a population of 16 neurons (Fig. 2A). Olfactory and model responses were less noisy and led to more accurate discrimination (Fig. 2B,C) for analysis based on Euclidean or van Rossum spike distance metrics. For the electrosensory and the olfactory responses, the method based on the Euclidean distance allowed a slightly better discrimination

accuracy whereas for the model responses, the method based on the van Rossum spike distance metric gave significantly better results. The only major difference between the methods as we implemented them is that each neuron was kept as a separate dimension in the Euclidean distance whereas responses where averaged across neurons before performing the van Rossum analysis. If neuron identity is important in encoding stimulus identity, keeping the neurons as separate dimension will allow extracting more accurate information whereas if each neuron is similar in its response pattern, averaging across neurons allows averaging out the noise. To confirm that this is indeed the main difference, we also calculated a Euclidean distance on population responses averaged across neurons and show that discrimination performance is similar to the van Rossum method (Fig. 2).

Comparison with an artificial neural net.

For both the olfactory and electrosensory dataset, we know that the animal is able to discriminate the stimuli as demonstrated by behavioral assays (Daly et al., 2008; Allen and Marsat, 2018). Even though these systems can rely on neural populations larger than tested here, the poor discrimination performance suggest that we might not be extracting as much information as the system actually does. This seem to be particularly obvious for the responses to electrocommunication stimuli that remain at chance levels for all population sizes tested. We therefore tested a methods that has proved to be very efficient in clustering patterns: artificial neural networks. Specifically with used a SOM-type of network that relies on unsupervised learning to change the way it weight input to optimally cluster the data into a fixed number of outputs.

The SOM network performed much better than the other two methods described above (Fig 3). As a result, discrimination was nearly errorless in all three systems for the larger population sizes. This result suggest that the core principle used by SOM network -differential weighting of inputs- allows a much more efficient decoding of the responses. This result is not surprising and can be understood by the fact that neurons or time points that are very noisy and contain very little information about stimulus identity will be weighted down and thus will not influence (i.e. add noise) the discrimination process. More complex neural networks with more layers (i.e. deep neural nets) and a supervised learning component could potentially perform even more reliably in these discrimination tasks. However, the convenience of such analysis methods also decreases as the network complexity increases. Several factors limit the convenience of neural network in analyzing sensory responses. First, the time required for running the analysis becomes unreasonably long when responses with many dimensions (time point X neurons) are used. To illustrate this point, using server-class and overclocked high-performance computers (e.g. 100% usage of 10 cores intel i9 at 5 GHz each), the SOM analysis in Fig. 3C took several days to run and we had to limit the analysis to populations of 12 neurons and less. Also, for more complex network, larger datasets might be required. Moreover, deep neural network perform a transformation on the input that is more opaque to the experimenter and that would take secondary analysis to detail once the network has learned it proper configuration. To benefit from the advantages of weighting the inputs exploited by the neural

nets but keeping the analysis simple we enhance the spike distance metric with a weighting procedure.

KLW measure.

We aim to use the Euclidean distance but weight each dimension according to how informative it is about the difference in the stimulus. In other words the weight has to be proportional to how distinct the distribution of responses are in that dimension. The best measure to characterize how divergent 2 distributions are is the Kullback-Leibler divergence (Kullback and Leibler, 1951). Normalizing the KL divergence to an average of one for each iteration of the analysis allows to keep the overall weight of the inputs unchanged. Also, by performing the calculation with one stimulus of the pair as to focus stimulus then the other one and averaging the two results, our result has the symmetrical quality of a distance. For short we refer to this Kullback-Leibler weighted distance analysis as the KLW method.

The KLW method perform as well or better than the SOM neural net in permitting discrimination of the stimuli (Fig. 4). We note that both analyses rely on weights that can change abruptly from one time point to the next in ways that are not dictated by biologically realistic rules. Indeed, the analysis is aimed at quantifying how well an ideal observer, with sufficient prior knowledge, could discriminated between two stimuli based on the pattern of neural responses. It does not explicitly emulate how well a biologically realistic decoding network could perform. We can however, adjust the way the weights change as a function of time to replicate more biologically plausible mechanism. Various time-dependent or firing-rate dependent mechanisms replicating rules of synaptic plasticity could be implemented. Here we only explore the simplest rule: keeping synaptic weight fixed across time. We average the Kl values across time for a given neurons, thereby each neuron is simply weighted according to its average information about the difference between two stimuli. We can see that the difference between independent and fixed weights across time vary for the 3 model systems tested here (Fig 4). This indicate that in some responses (e.g. model), information about stimuli differences varies across time whereas for other (e.g. olfactory) the informative dimensions are fairly constant in time.

In figure 5, we present two additional alterations to the KLW methods that replicate biologically inspired constraints and that can indicate how an efficient decoding network could be designed and what aspect of the responses carry useful information. We attempted to implement a short-term plasticity rule that allows the weights to change as a function of the preceding firing rate mimicking short-term facilitation and depression. Implementing such a rule could lead to improvement of discrimination performance in the analysis if the biological system is build to encode and decode through such mechanism. In our case the procedure did not lead to improved discrimination suggesting that the rule we implemented would not offer a decoding benefit over having fixed weights over time. We also revisited a concept introduced at the beginning of the result section: having the different neural responses combined or kept as separate dimension. A postsynaptic decoder could receive inputs from all neurons in the population and the postsynaptic potential be combined in the decoding neurons before any

further processing occurs. Averaging responses across neurons before population responses are compared would replicate this scenario. Alternatively, a specific subset of neurons connect to a decoder, and the identity of which neuron is active at what time determines the activation of the decoder. To replicate this decoding scenario we keep neurons as separate dimensions in the analysis. Our results show that some systems would benefit for the former procedure (here the LIF; compare the cyan and blue curves in Fig 5C) and other from the latter (electrosensory and olfactory, Fig 5 A&B).

Noise and information.

We ask how noise affects the discrimination accuracy estimated by the KLW method and how it related to the amount of information carried by the spike trains. To do so, we focus on the LIF model responses to be able to change the amount of noise in the system. We also performed the analysis keeping the noise fixed but changing the signal strength (data not shown), thus similarly affecting the signal-to-noise ratio, and the results were qualitatively similar. We first confirmed that increasing the amount of noise in the model decreased the discrimination performance assessed by our analysis (Fig 6). We use three version of the analysis: a minimalist one with no weighting (i.e. based on a traditional Euclidean distance), a "realistic one" were weights are kept constant across time and the responses from different neurons are averaged before the distances are calculated; and one labeled as "independent" were weight are optimized independently and all dimensions are kept separate.

The use of white noise stimuli conveniently allowed us to calculate the information coding rate about the temporal modulations by simply calculating the coherence between stimulus and response or the response to response coherence. Since we used a linear system in this analysis, the two type of coherences lead to similar estimates of information rate. Information rate was calculated for population responses of varying sizes (similar to the analysis in previous figures) and for neurons with different amount of noise. By plotting the discrimination accuracy as a function of information rate calculated for the same responses (Fig. 7), we show that the 2 are highly correlated. These results demonstrate that the discrimination performance of these analysis is directly related to the information content of the spike train but that the analytic procedure dictates how efficiently the information is used to discriminate between two stimuli.

Noise Correlations and population responses.

Correlations among neurons of a population can be a important aspect of the response that carries information about the stimulus. Stimulus elicited correlations will obviously be taken in account by the KLW method (or other spike distance metrics) since correlated vs uncorrelated responses will lie in different clusters of space. It is less obvious that noise correlations will influence the result of the analysis as it should since we know it can affect the quantity of information carried by the population. To address this issue, we introduced noise correlations in the population of model responses. As expected, these noise correlations decrease the information capacity of the population compared to responses without the noise correlations. More importantly, a proportional decrease in discrimination accuracy was observed (Fig 8).

These results demonstrate that the analysis is adequate to apply to a wide range of neural response types even if they include noise correlations that significantly affect their encoding accuracy.

DISCUSSION & CONCLUSION

We compared different methods of quantifying the accuracy with which a population of neurons encode the identity of the stimulus. We show that the estimate of discrimination performance of a system varies widely based on the method used. Methods that weight inputs provide a clear advantage since they can base the results on information-rich portions of the population response while de-emphasizing noisy portions. To give an every-day life example of this principle, let's imagine we are trying to discriminate between two twins. Differences in their facial features might be so subtle that they are within the noise level of our perceptual judgment. However, if one has a freckle below the left eye and the other one doesn't, it is possible to identify them with 100% reliability. Decoding network in actual biological systems can rely on the same principle, weighting more heavily the input from specific neurons or certain time frames in their responses.

This procedure however requires prior knowledge of the stimuli to discriminate and the responses they elicit. This can be accomplished by a learning process or by having a wide array of decoding units with many possible combinations of synaptic weights. It is important to point out however, that the methods presented here are not aimed explicitly at replicating biologically realistic decoding network. We argue that, by showing that we can discriminate between stimuli with a certain accuracy, the information necessary to accomplish the task is present in the response. Whether or not the biological system can reach this level of performance and how it performs its decoding, are separate question that our analytical tool can only begin to explore. In conclusion we argue that the new KLW method we propose here is a highly efficient way of estimating how accurately sensory neurons encode the stimulus identity to support discrimination and would benefit a wide range of research endeavors.

REFERENCES

Allen, K. M. and Marsat, G. (2018). Task-specific sensory coding strategies are matched to detection and discrimination performance. *J. Exp. Biol.* **221**, jeb170563. DOI: 10.1242/jeb.170563.

Barrett, D. G., Morcos, A. S. and Macke, J. H. (2019). Analyzing biological and artificial neural networks: challenges with opportunities for synergy? *Curr. Opin. Neurobiol.* **55**, 55–64. DOI: 10.1016/j.conb.2019.01.007.

Borst, A. and Theunissen, F. E. (1999). Information theory and neural coding. *Nat. Neurosci.* **2**, 947–957.

Daly, K. C., Carrell, L. A. and Mwilaria, E. (2008). Characterizing psychophysical measures of discrimination thresholds and the effects of concentration on discrimination learning in the moth Manduca sexta. *Chem. Senses* **33**, 95–106. DOI: 10.1093/chemse/bjm068.

Daly, K. C., Bradley, S., Chapman, P. D., Staudacher, E. M., Tiede, R. and Schachtner, J. (2016a). Space Takes Time: Concentration Dependent Output Codes from Primary Olfactory Networks Rapidly Provide Additional Information at Defined Discrimination Thresholds. *Front. Cell. Neurosci.* **9**, 515. DOI: 10.3389/fncel.2015.00515.

Daly, K. C., Bradley, S., Chapman, P. D., Staudacher, E. M., Tiede, R. and Schachtner, J. (2016b). Space Takes Time: Concentration Dependent Output Codes from Primary Olfactory Networks Rapidly Provide Additional Information at Defined Discrimination Thresholds. *Front. Cell. Neurosci.* **9**, 515. DOI: 10.3389/fncel.2015.00515.

Kullback, S. and Leibler, R. A. (1951). On Information and Sufficiency. *Ann. Math. Stat.* **22**, 79–86. DOI: 10.1214/aoms/1177729694.

Marsat, G. and Maler, L. (2010). Neural Heterogeneity and Efficient Population Codes for Communication Signals. *J. Neurophysiol.* **104**, 2543–2555. DOI: 10.1152/jn.00256.2010.

Rieke, F. (1997). *Spikes : exploring the neural code*. MIT Press.

Stanley, G. B. (2013). Reading and writing the neural code. *Nat. Neurosci.* **16**, 259–263. DOI: doi 10.1038/nn.3330.

Staudacher, E. M., Huetteroth, W., Schachtner, J. and Daly, K. C. (2009). A 4-dimensional representation of antennal lobe output based on an ensemble of characterized projection neurons. *J. Neurosci. Methods* **180**, 208–223. DOI: 10.1016/j.jneumeth.2009.03.019.

van Rossum, M. C. (2001). A novel spike distance. *Neural Comput.* **13**, 751–763.

Victor, J. D. (2005). Spike train metrics. *Curr. Opin. Neurobiol.* **15**, 585–592.

FIGURE & LEGENDS

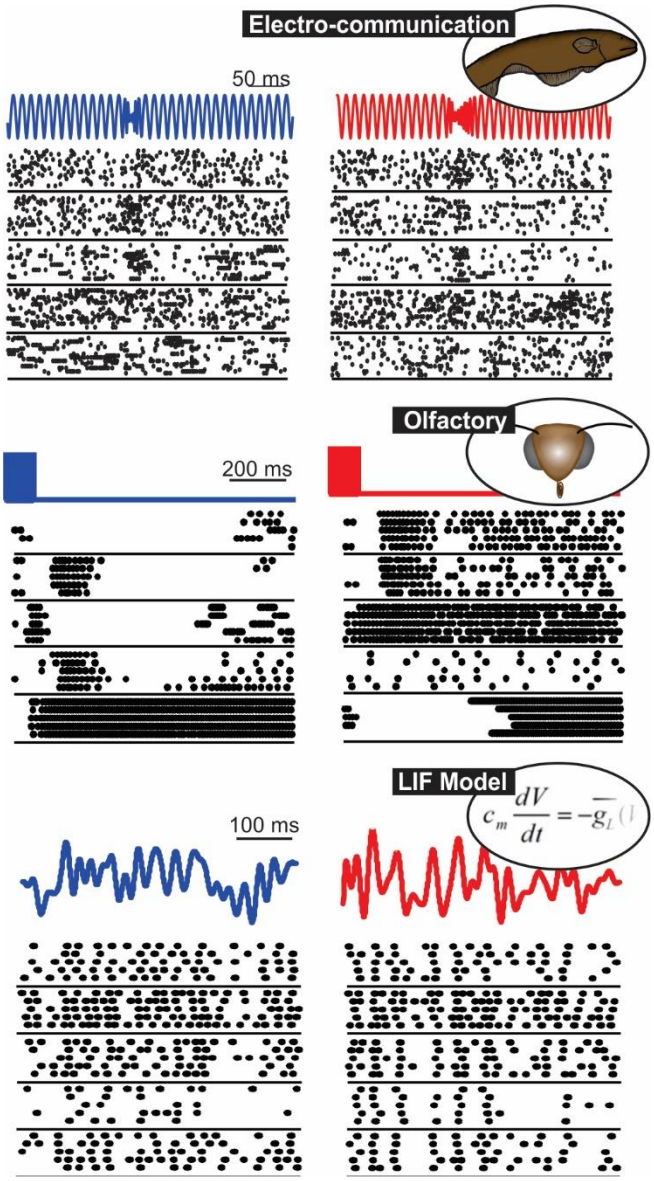

Figure 1: Spatio-temporal patterns of response in 3 systems: the electrosensory system (responses of ELL pyramidal cells to chirps), the moth olfactory system (PN responses to 2 odors) and a population of model neurons (LIF neurons stimulated with frozen white noise). In each panel we show the responses of 5 neurons to repeated presentation of 2 stimuli (blue vs red).

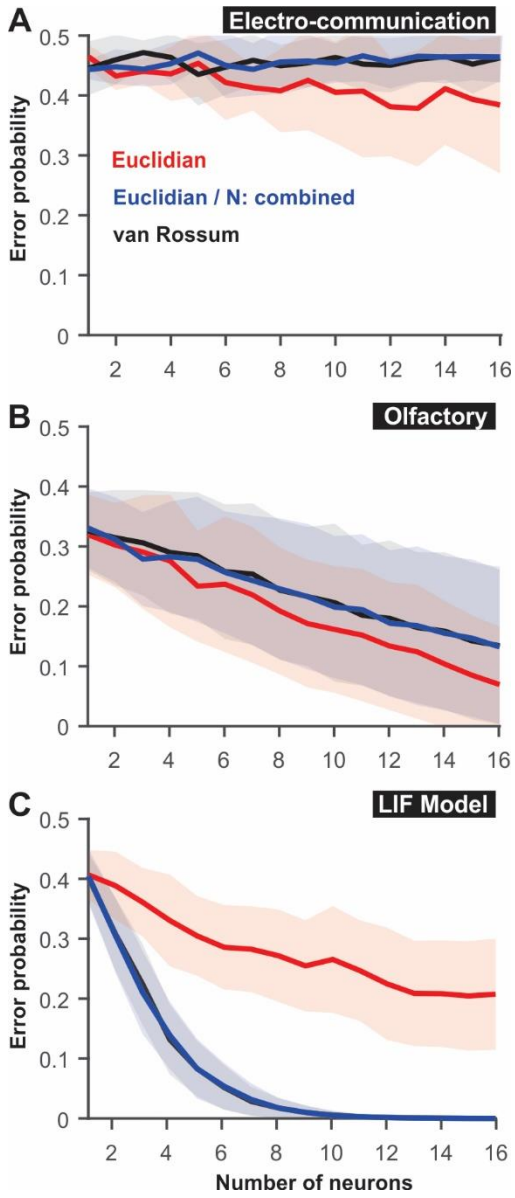

Figure 2: Discrimination accuracy for populations of neurons using established spike metrics. The discriminability (y-axis) of responses to pairs of different stimuli was determined for neural populations of varying sizes (x-axis). We compare methods based on Euclidian distance and the van Rossum spike distance metrics and show that they are essentially identical when the different neurons are not mapped as different dimensions in Euclidian space but combined (e.g. averaged into a population response).

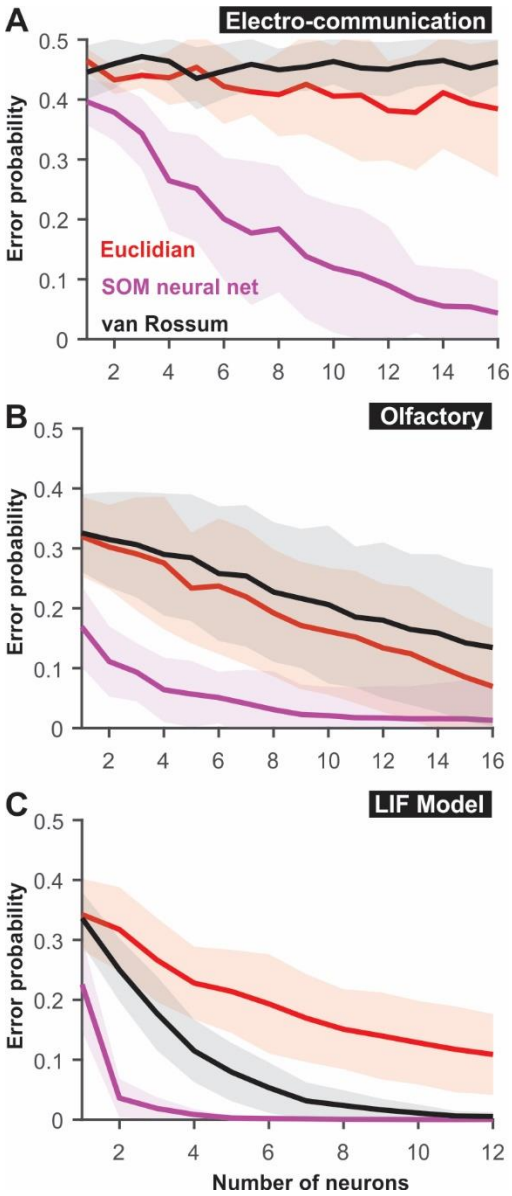

Figure 3: The discriminability of responses can be quantified using a SOM neural net where the spatio-temporally patterned array of inputs is weighted. SOM neural networks uses unsupervised learning to cluster the sets of inputs according to the variability and the patterns present in the dataset. Large intrinsic differences in patterns between responses to 2 stimuli thus lead to reliable clustering.

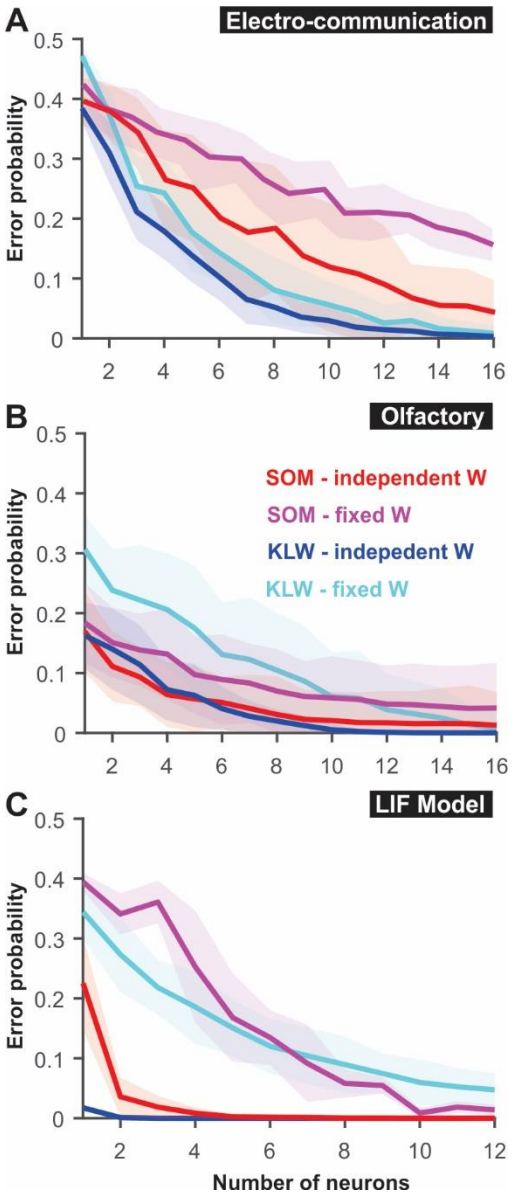

Figure 4: A modified Euclidian distance where each dimension is weighted, performs discrimination analysis as well or better than SOM neural nets. In the "KLW" (Kullback-Leibler Weighted) analysis, each dimension in Euclidian space is weighted by the KL divergence of the response distribution is that dimension. Each dimension (neuron/time bin) can be weighted independently ('independent W'), or a single weight can be set for a given neuron across time bins ('fixed W'). Although using independent weights maximizes the information extracted about the difference in stimuli, using a fixed weight emulates a biologically more realistic decoding network.

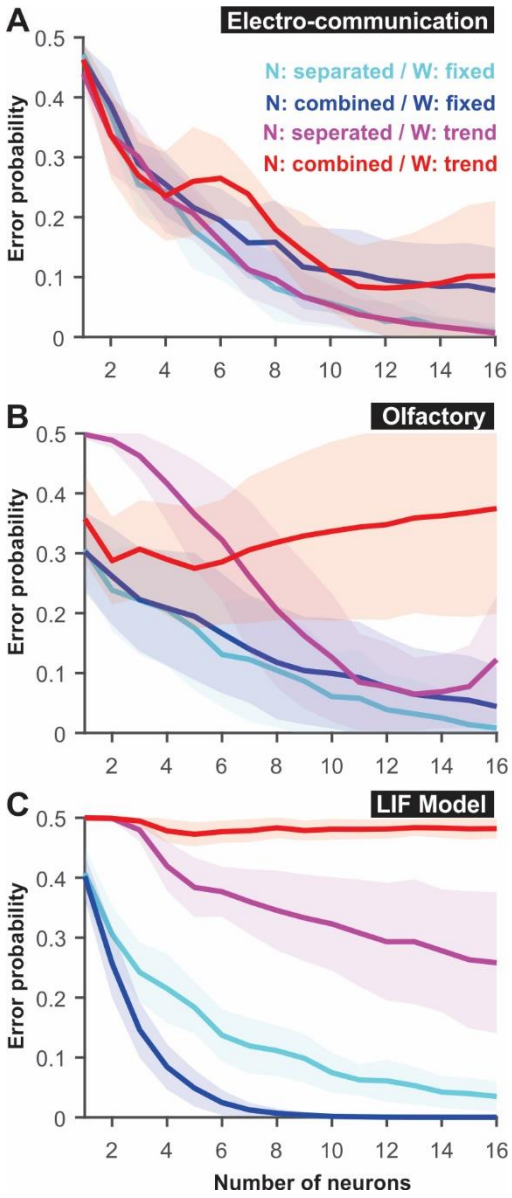

Figure 5: Parameters of a KLW can be altered to reproduce biologically realistic constraints on decoding. The different neurons can be kept as separate dimensions in Euclidian space (‘N: separated’) or combined (e.g. average population response) mimicking the fact that synaptic inputs could combine indiscriminately in the postsynaptic terminals of decoding neurons (‘N: combined’). Weight are calculated independently for each time bin (see Fig 4), but can be set to be fixed across time (‘W: fixed’) for a given neuron or follow a specific function (‘W: trend’). Example of functions that could be implemented include functions that would replicate firing rate -based plasticity such as facilitation and depression.

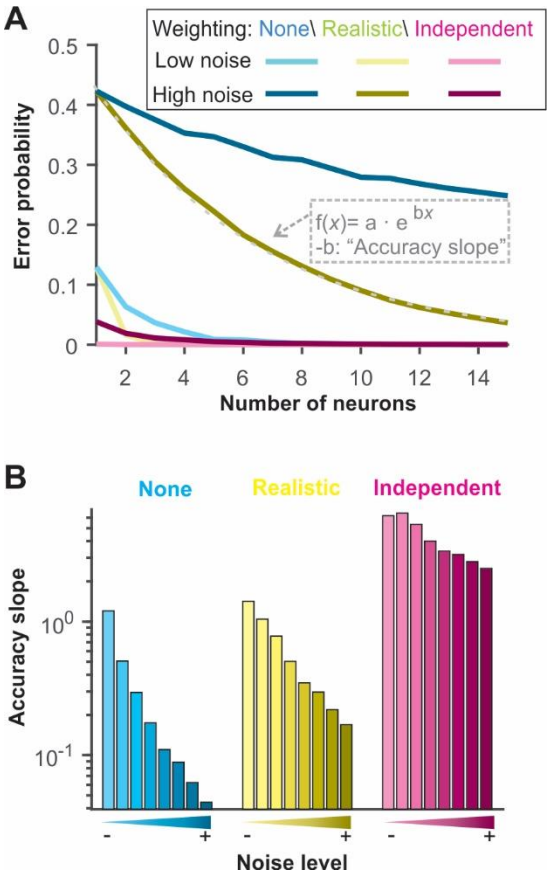

Figure 6: The discrimination accuracy reflects well the signal-to-noise ratio of the neural responses. Varying the amount of noise in our LIF model we quantified the discrimination accuracy using 3 methods. "Basic": Euclidian distance with unweighted dimensions and each neuron kept in separate dimensions; "Realistic": Each neuron is weighted with a fixed weight across time and the different neurons are combined in a single dimension; "Optimized": Both neuron and time bins are weighted independently and kept as separate dimensions. The "Accuracy slope" is the slope of the exponential fit for the curves displayed in A and therefore reflects how quickly the error level decreases as more neurons are included in the analysis. A sharp decrease (high slope) indicate more accurate coding of stimulus identity.

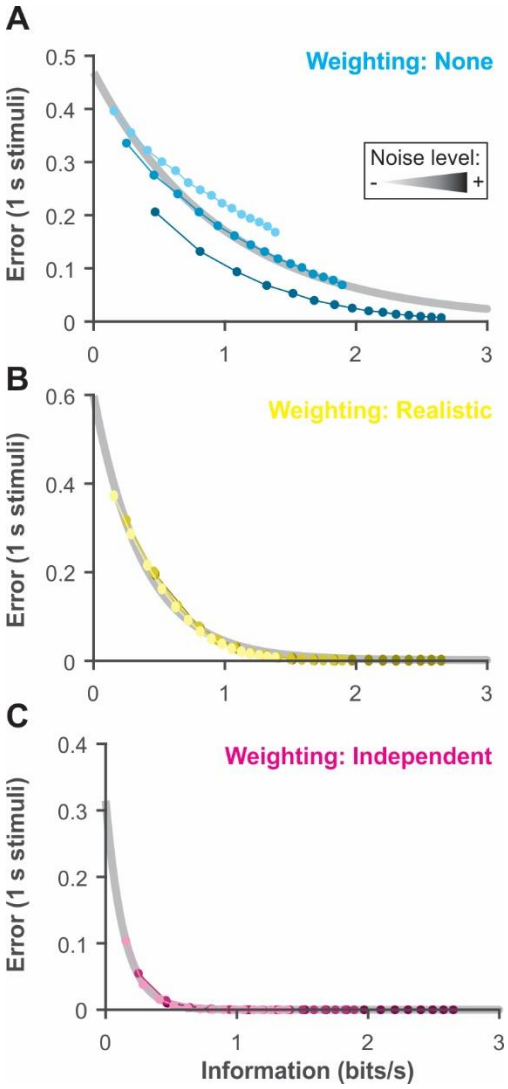

Figure 7: The error rate is proportional to the information content. The coding accuracy was calculated for this population of LIF neurons using information theoretic tools. Using the same neurons and stimuli, we calculated the discrimination accuracy (error probability) using the 3 methods described in Fig 6. The number of neuron included in the analysis was varied as in previous figures (the different points along each curves here) and the amount of noise in the LIF model as varied (varying shape). The exponential relationship between error and information content (grey best fit lines) shows the high correlation between these two measures of coding accuracy (adjusted R-square > 0.98 for B and C).

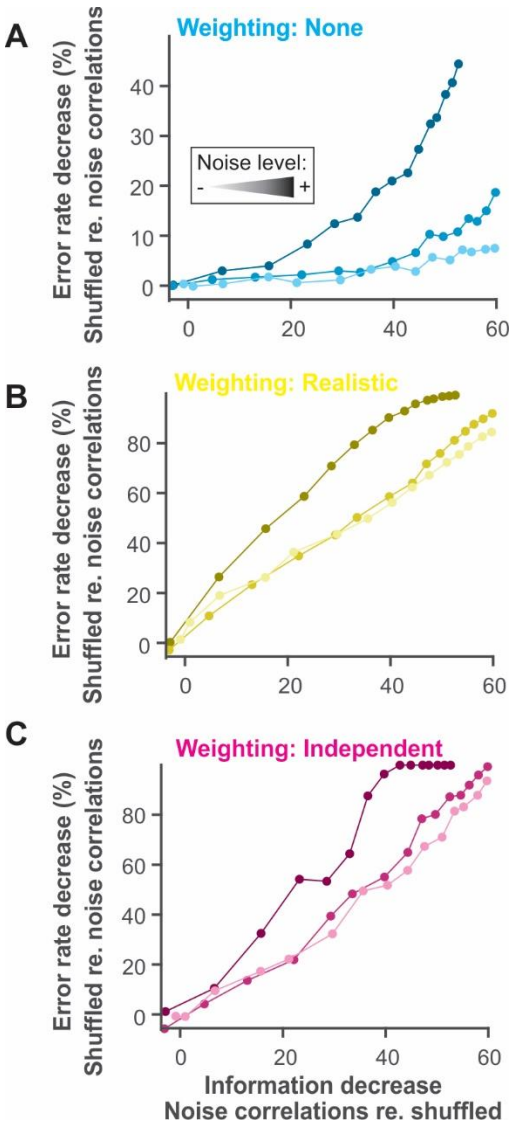

Figure 8: The presence of noise correlation can be taken into account in the discrimination analysis. Noise correlation across responses of different neurons of a population can decrease the total information content of the population compared to responses where the noise is not correlated. An information theoretic analysis of populations of varying sizes (different points along each curves) show this effect here since noise correlation decrease information content compared to identical responses her noise correlations have been shuffled out. This change in information content is mirrored in the discrimination accuracy measure.