

# RESTRICTED BOLTZMANN MACHINES PROVIDE AN ACCURATE METRIC FOR RETINAL RESPONSES TO VISUAL STIMULI

**Christophe Gardella<sup>1,2</sup>, Olivier Marre<sup>2</sup> & Thierry Mora<sup>1</sup>**

<sup>1</sup> Laboratoire de physique statistique, CNRS, UPMC and École normale supérieure

<sup>2</sup> Institut de la Vision, INSERM and UPMC

{christophe.gardella, tmora}@lps.ens.fr

olivier.marre@inserm.fr

## ABSTRACT

How to discriminate visual stimuli based on the activity they evoke in sensory neurons is still an open challenge. To measure discriminability power, we search for a neural metric that preserves distances in stimulus space, so that responses to different stimuli are far apart and responses to the same stimulus are close. Here, we show that Restricted Boltzmann Machines (RBMs) provide such a distance-preserving neural metric. Even when learned in an unsupervised way, RBM-based metric can discriminate stimuli with higher resolution than classical metrics.

## 1 INTRODUCTION

The precision with which stimuli are encoded by a population of sensory neurons is still a matter of debate. Sensory neurons process visual stimuli in a highly non-linear manner, and in most cases there is no general model to accurately predict the population response to any given stimulus (Golisch & Meister, 2010). A possible way to estimate the discriminability power of neural populations is to use a neural metric. A metric measure differences in responses triggered by different stimuli, and thus provides a lower bound on how well two stimuli can be discriminated based on responses.

Many definitions of neural metrics have been proposed (Houghton & Victor, 2012). Machens et al. (2003) used a neural metric to measure experimentally the accuracy of grasshoppers auditory receptors, and Narayan et al. (2006) used the same one to study sound-processing cortical neurons in zebra finches cortex. Although these studies focused on single neurons, some information can only be retrieved by considering populations of neurons (Mazor & Laurent, 2005; Fujisawa et al., 2008). Ganmor et al. (2015) proposed a neural metric for population of neurons, learned in a supervised way, so that it preserves distances in the stimulus space. But it requires to learn the distribution of responses to any stimulus, which is not feasible experimentally when the space of stimuli is large. Tkačik et al. (2013) used a similar metric for a high-dimensional stimulus using a model to map stimuli to parametrized response distributions. But such an accurate model is rarely available.

We are interested in deriving a metric directly from the distribution of responses, in an unsupervised way. Recently, using tools from information theory, (Humplik & Tkačik, 2016) showed that Restricted Boltzmann Machines (RBMs) could describe the distribution of responses better than other classical models. Also using information theory, (Zanotto et al., 2017) showed that a variant of RBM could convey information about the stimulus, even when learned in an unsupervised way. But in both cases this was only made possible by the small number of stimuli, and a method to measure the sensitivity of the neural population was still lacking.

We recorded a population of 60 rat retinal ganglion cells (RGCs) stimulated by a bar in motion. We use RBMs to learn the distribution of responses represented by binary strings of spikes and silences. We show that RBMs are capable of reproducing the correlated activity of neurons (see Schwab et al. (2013) for similar results). Even though this learning is unsupervised, the resulting RBM can discriminate visual stimuli with high accuracy: we define a neural metric as a distance between the hidden units of the RBM conditioned by responses, and show that this RBM-based metric can discriminate stimuli based on the neural responses, much better than other classical neural metric.

## 2 METHODS

### 2.1 EXPERIMENTAL SETUP

We analyzed previously published *ex vivo* recordings from rat retinal ganglion cells (RGCs) (Ferrari et al., 2016), i.e. cells that form the output of the retina and send their spikes to the brain. Any visual information accessible to the brain is necessarily represented by RGCs. In brief, we recorded with a multielectrode array a population of 60 RGCs stimulated by a bar in motion (Fig. 1 left). The stimulus is composed of two parts. In the first part, the bar has a Brownian motion with a restoring force, lasting 15331 s. The second part is composed of 391 repetitions of two trajectories of length 0.9 s, called reference trajectories. We also presented perturbations of the reference trajectories, i.e. small changes affecting that trajectory in its middle portion, between 300 and 630 ms. Perturbations varied both in shape and in amplitude: we used 16 different perturbation shapes, each presented at different amplitudes (Fig. 1 right). The amplitude was adapted online: large enough so perturbations could be discriminated from reference trajectories, but small enough so they would not be discriminated perfectly with any metric. The response is considered from the start of the perturbation until 280 ms after its end. We assess the discriminating power of a neural metric by measuring how well it can discriminate responses to perturbations from responses to reference stimuli.

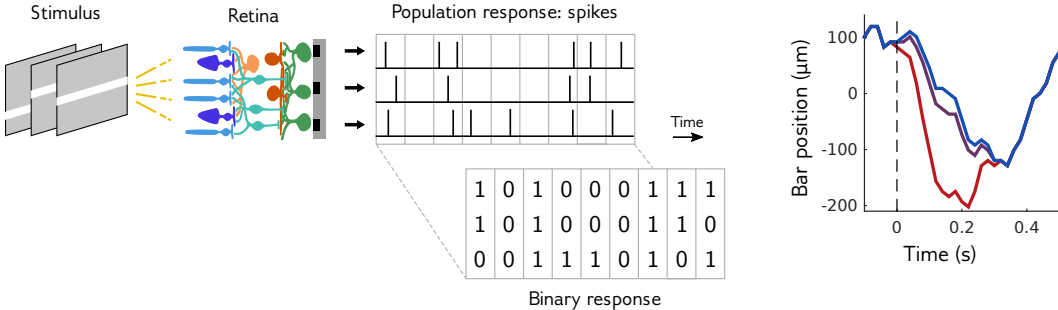


Figure 1: Experimental setup. Left: we stimulated a rat retina with a moving bar. Retinal ganglion cells (in green) were recorded with a multielectrode array. For the computation of the Hamming metric and the RBM, we binarized spike trains in 20 ms time bins. Right: the stimulus consists in a reference stimulus (the bar trajectory in blue), and in perturbations of different shapes and amplitudes. Purple and red perturbations have same shape, with small and large amplitude.

#### 2.1.1 RESTRICTED BOLTZMANN MACHINE-BASED METRIC

The response of  $N$  neurons over time is binarized into  $B$  time bins:  $\sigma_{ib} = 1$  if cell  $i$  spiked during time bin  $b$ , and 0 otherwise (Fig. 1 left).  $\sigma$  is a binary vector of size  $NB$ , labeled by a joint index  $ib$ . We learned a RBM with 40 hidden units, taking as visible units the population response within a single time bin. The inference was done on responses to random bar trajectories, using persistent contrastive divergence (Tieleman, 2008). The RBM-based metric is defined as the  $L_1$  distance between the conditional probabilities of the hidden units  $h$  given responses:

$$d_{\text{RBM}}(\sigma, \sigma') = \sum_b \sum_{j \text{ hidden}} |p(h_j = 1 | \sigma_b) - p(h_j = 1 | \sigma'_b)| \quad (1)$$

### 2.2 DISCRIMINABILITY

To assess the distance-preserving properties of retinal metrics, we measure how well they can discriminate between different stimuli based on their responses. Ideally, responses to different stimuli should be far apart and responses to the same stimulus should be close. As neural responses are noisy, for each stimulus  $s$  we define a probability of responses  $\sigma$ ,  $P(\sigma | s)$ . Given two stimuli  $s$  and  $s'$  and a metric  $d$ , we define the probability of discriminating  $s'$  from  $s$  as:

$$D_d(s, s') = P[d(\sigma, \sigma') > d(\sigma, \sigma'')]_{\sigma, \sigma'' \sim P(\cdot | s), \sigma' \sim P(\cdot | s')} \quad (2)$$

It is equal to 1 if stimuli are perfectly discriminable by neural responses, as measured by metric  $d$ . If stimuli are not discriminable, i.e.  $P(\cdot | s) = P(\cdot | s')$ ,  $D_d$  is equal to 1/2.




Figure 2: RBM reproduces well different statistics of population responses. Left: firing rate of neurons (*i.e.* visible units), in data and predicted by the RBM. Each point corresponds to one neuron. Middle: Pearson correlation between pairs of neurons, in data and predicted by the RBM. Each point corresponds to one pair of neurons. Right: probability distribution of the population rate  $K$  (see Methods), in the data or predicted by independent, Ising and RBM models.

### 3 RESULTS

We trained a RBM to reproduce the distribution of population responses within 20 ms time bins. The RBM could reproduce well the spiking probability of each neuron (called firing rate, Fig. 2 left) and the Pearson correlation between pairs of neurons (Fig. 2 middle). We also tested how well it could reproduce the population rate (Okun et al., 2015; Gardella et al., 2016), defined as the total number of neurons firing within a time bin:  $K = \sum_i \sigma_i$ . We compared with the predictions of a model with independent neurons only reproducing firing rates, and a Boltzmann machine (Ackley et al., 1985) with no hidden units (termed Ising model in the following), which has been extensively used to model neural population statistics (Schneidman et al., 2006; Ferrari, 2016). We found that the RBM outperformed both independent and Ising models at reproducing the distribution of the population rate (Fig. 2 right). RBMs can thus capture well the distribution of neural population responses.

Next we use the RBM to define a neural metric and to discriminate stimuli based on that metric (see Methods). For each of the 2 reference trajectories and each of the 16 perturbation directions, we separate perturbation amplitudes in 3 classes: small, medium and large amplitudes (approximately 100 amplitudes each). We compute the mean discriminability for each class (Eq. 2) for the RBM-based metric and for other metrics often encountered in the literature: the Hamming, van Rossum and Victor Purpura metrics. We found that the RBM-based metric outperformed all three other metrics. Note that during the computation of the RBM-based metric spikes are binned in time, which is likely to be a lossy transformation. Thus it has access to less information than van Rossum and Victor Purpura metrics, which act directly on spike trains. The RBM was trained in an unsupervised way, with a training set corresponding to random bar trajectories, and thus could not learn any stimulus-response relationship. We have thus shown that the RBM, even learned in an unsupervised way, can recover explicitly accurate information about the stimulus from the structure of responses.




Figure 3: RBM-based metric outperforms all three other metrics at discriminating perturbations. Each point represents the mean discriminability for a reference stimulus and a perturbation direction at small, medium or large amplitude, for the RBM-based metric versus another metric.

#### ACKNOWLEDGMENTS

We thank David Schwab for his help with RBMs.

## REFERENCES

- David H Ackley, Geoffrey E Hinton, and Terrence J Sejnowski. A learning algorithm for boltzmann machines. *Cognitive Science*, 9(1):147–169, 1985. ISSN 03640213. doi: 10.1016/S0364-0213(85)80012-4.
- Ulisse Ferrari. Learning maximum entropy models from finite-size data sets: A fast data-driven algorithm allows sampling from the posterior distribution. *Phys. Rev. E*, 94:023301, Aug 2016. doi: 10.1103/PhysRevE.94.023301.
- Ulisse Ferrari, Christophe Gardella, Olivier Marre, and Thierry Mora. Closed-loop estimation of retinal network sensitivity reveals signature of efficient coding, 2016. URL <http://arxiv.org/abs/1612.07712>.
- Shigeyoshi Fujisawa, Asohan Amarasingham, Matthew T Harrison, and György Buzsáki. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nature Neuroscience*, 11(7):823–833, 2008. ISSN 1097-6256. doi: 10.1038/nn.2134.
- Elad Ganmor, Ronen Segev, and Elad Schneidman. A thesaurus for a neural population code. pp. 1–19, 2015. doi: 10.7554/eLife.06134.
- Christophe Gardella, Olivier Marre, and Thierry Mora. A tractable method for describing complex couplings between neurons and population rate. 3(August):1–13, 2016. ISSN 2373-2822. doi: 10.1523/ENEURO.0160-15.2016.
- Tim Gollisch and Markus Meister. Eye Smarter than Scientists Believed: Neural Computations in Circuits of the Retina, 2010. ISSN 08966273.
- Conor Houghton and Jonathan D Victor. Measuring representational distances – the spike-train metrics approach. In G. Kreiman N. Kriegeskorte (ed.), *Visual Population Codes: Toward a Common Multivariate Framework for Cell Recording and Functional Imaging*, pp. chap. 8. MIT Press, Cambridge, MA, 2012.
- Jan Humplik and Gašper Tkačik. Semiparametric energy-based probabilistic models. 2016. URL <http://arxiv.org/abs/1605.07371>.
- Christian K Machens, Hartmut Schütze, Astrid Franz, Olga Kolesnikova, Martin B Stemmler, Bernhard Ronacher, and Andreas V M Herz. Single auditory neurons rapidly discriminate conspecific communication signals. *Nature Neuroscience*, 6(4):341–342, 2003. ISSN 10976256. doi: 10.1038/nn1036.
- Ofer Mazor and Gilles Laurent. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron*, 48(4):661–673, 2005. ISSN 08966273. doi: 10.1016/j.neuron.2005.09.032.
- Rajiv Narayan, Gilberto Graña, and Kamal Sen. Distinct time scales in cortical discrimination of natural sounds in songbirds. *Journal of neurophysiology*, 96(1):252–8, 2006. ISSN 0022-3077. doi: 10.1152/jn.01257.2005.
- Michael Okun, Nicholas a. Steinmetz, Lee Cossell, M. Florencia Iacaruso, Ho Ko, Péter Barthó, Tirin Moore, Sonja B. Hofer, Thomas D. Mrsic-Flogel, Matteo Carandini, and Kenneth D. Harris. Diverse coupling of neurons to populations in sensory cortex. *Nature*, 2015. ISSN 0028-0836. doi: 10.1038/nature14273.
- Elad Schneidman, Michael J Berry, Ronen Segev, and William Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012, 2006. ISSN 0028-0836. doi: 10.1038/nature04701.
- David Jason Schwab, Kristina D. Simmons, Jason S Prentice, and Vijay Balasubramanian. Representing correlated retinal population activity with Restricted Boltzmann Machines. *Cosyne Poster*, 2013.
- Tijmen Tieleman. Training Restricted Boltzmann Machines using Approximations to the Likelihood Gradient. *Proceedings of the 25th International Conference on Machine Learning*, 307:7, 2008. ISSN 21576904. doi: 10.1145/1390156.1390290.

- Gašper Tkačik, Einat Granot-Atedgi, Ronen Segev, and Elad Schneidman. Retinal metric: A stimulus distance measure derived from population neural responses. *Physical Review Letters*, 110 (February):1–5, 2013. ISSN 00319007. doi: 10.1103/PhysRevLett.110.058104.
- Mark C van Rossum. A novel spike distance. *Neural computation*, 13(4):751–763, 2001. ISSN 0899-7667. doi: 10.1162/089976601300014321.
- Jonathan D Victor and Keith P Purpura. Nature and precision of temporal coding in visual cortex: a metric-space analysis. *Journal of Neurophysiology*, 76(2):1310–1326, 1996. ISSN 0022-3077.
- Matteo Zanotto, Riccardo Volpi, Alessandro Maccione, Luca Berdondini, Diego Sona, and Vittorio Murino. Modeling Retinal Ganglion Cell Population Activity with Restricted Boltzmann Machines. 2017. URL <http://arxiv.org/abs/1701.02898>.

## APPENDIX : NEURAL METRICS

We used three neural metrics classically encountered in the literature (Houghton & Victor, 2012): a metric on binned responses, a kernel-based metric, and an edit-length metric. We then present a metric based on a Restricted Boltzmann Machine.

The response of a neuron  $i$  consists in a series of action potentials, or spike train. The spikes are stereotypical, so most of the information in the neural response is contained in the spike times  $(t_{k,i})_k$  (Fig. 1 left).

## HAMMING METRIC

The Hamming metric between responses is the  $L_1$  distance between binned responses:

$$d_{\text{Hamming}}(\sigma, \sigma') = \sum_{b \text{ bin}} \sum_{i \text{ neuron}} |\sigma_{bi} - \sigma'_{bi}| \quad (3)$$

## VAN ROSSUM METRIC

The van Rossum metric is a kernel-based metric (van Rossum, 2001). In order to avoid binning the spike times, which may lose information, we convolve each neuron’s spike train with a kernel  $G$ :  $v_i(t) = \sum_k G(t - t_{k,i})$ . We then take the Euclidean distance between convolved spike trains. Here we set  $G$  as Gaussian with a standard deviation of 10 ms.

$$d_{\text{van Rossum}}(v, v')^2 = \sum_{i \text{ neuron}} \int_t |v_i(t) - v'_i(t)|^2 dt \quad (4)$$

## VICTOR PURPURA METRIC

The Victor Purpura metric (Victor & Purpura, 1996) is an edit-length metric. The distance between two spike trains is the minimal cost necessary to transform a spike train into another using a series of transformations. Deleting or adding a spikes costs +1, whereas moving a spike of  $\Delta t$  has a linear cost  $q\Delta t$ . Here we used  $q = 0.01 \text{ s}^{-1}$ .