



Retinal Metric: A Stimulus Distance Measure Derived from Population Neural Responses

Gašper Tkačik,^{1,*} Einat Granot-Atedgi,² Ronen Segev,³ and Elad Schneidman^{2,†}

¹*Institute of Science and Technology Austria, Am Campus 1, A-3400 Klosterneuburg, Austria*

²*Department of Neurobiology, Weizmann Institute of Science, 76100 Rehovot, Israel*

³*Faculty of Natural Sciences, Department of Life Sciences and Zlotowski Center for Neuroscience, Ben Gurion University of the Negev, 84105 Be'er Sheva, Israel*

(Received 4 August 2012; published 28 January 2013)

The ability of an organism to distinguish between various stimuli is limited by the structure and noise in the population code of its sensory neurons. Here we infer a distance measure on the stimulus space directly from the recorded activity of 100 neurons in the salamander retina. In contrast to previously used measures of stimulus similarity, this “neural metric” tells us how distinguishable a pair of stimulus clips is to the retina, based on the similarity between the induced distributions of population responses. We show that the retinal distance strongly deviates from Euclidean, or any static metric, yet has a simple structure: we identify the stimulus features that the neural population is jointly sensitive to, and show the support-vector-machine-like kernel function relating the stimulus and neural response spaces. We show that the non-Euclidean nature of the retinal distance has important consequences for neural decoding.

DOI: [10.1103/PhysRevLett.110.058104](https://doi.org/10.1103/PhysRevLett.110.058104)

PACS numbers: 87.18.Sn, 87.19.1o, 89.70.Cf

Neural populations convey information about their stimuli by their joint spiking patterns [1]. At the level of single cells, the mapping from stimuli to spikes has often been captured by linear-nonlinear (LN) models [2,3]. Geometrically, a single LN neuron can be viewed as a *perceptron* [4], partitioning the stimulus space into two domains—one of stimuli that evoke spikes and one of stimuli that do not—by a decision boundary, or a hyperplane, determined by the linear feature of the LN model [5–7]. The brain, listening for spikes coming from such a neuron, will thus interpret stimuli as similar insofar as they evoke similar spiking patterns. But how does an interacting population of neurons, as a whole, partition the stimulus space? Conversely, which stimuli are interpreted as different, or similar, by an interacting population?

Answering these questions is fundamental to our understanding of the neural code and depends critically on finding the correct “metric” for sensory stimuli in terms of the information that neural populations carry. Since neurons are noisy, repeated presentations of the same stimulus can result in different neural responses, so the stimulus-to-response mapping of the population needs to be described by the probability distribution, $P(\sigma|s)$ [8]. Two stimuli s_1 and s_2 may be far apart as measured by a chosen distance function, e.g., Euclidean norm $D_2(s_1, s_2) = \|s_1 - s_2\|$, yet they could evoke responses drawn from almost overlapping distributions $P(\sigma|s_1)$ and $P(\sigma|s_2)$, making it nearly impossible for the brain, listening to the spikes σ arriving from the sensory system, to tell those stimuli apart. Conversely, the sensory circuit could be sensitive to specific stimulus changes that have a small Euclidean norm, emphasizing those particular differences as an important feature and encoding it in the neural response. We therefore suggest that the biologically

relevant distance between pairs of stimuli should be derived from the distance between the response distributions evoked by these stimuli [9,10].

To characterize the structure of neural distance in a large population, we recorded extracellularly the activity of $N = 100$ retinal ganglion cells (RGCs) in the tiger salamander using a multielectrode array [11,12]. The retina patch was presented 626 times with a 10 s long segment of spatially uniform flicker with Gaussian distributed luminance drawn independently at 30 Hz [Figs. 1(a) and 1(b)]. Time was discretized into $T = 961$ bins of 10 ms and the joint response of the neurons $\sigma = \{\sigma_a\}$ was represented in each bin by a N -bit codeword, where $\sigma_a = 1$ (0) denoted that the neuron $a = 1, \dots, N$ spiked (did not spike) in that bin. Since direct sampling of the conditional distribution $P(\sigma|s)$ is impractical for such a large population [even with hundreds of repeats, Fig. 1(c)], we inferred a stimulus-dependent maximum entropy (SDME) model for this data that predicts $P(\sigma|s)$ for each time bin, as we report in detail in Ref. [13].

Since only differences in retinal responses can guide the organism’s behavior, the biologically relevant distance between stimuli s_1 and s_2 must be a measure of similarity between their corresponding response distributions. We define the *retinal distance* between the stimuli as the symmetrized Kullback-Leibler distance between the distributions of responses they elicit,

$$D_{\text{ret}}(s_1, s_2) = D_{\text{KL}}^{\text{sym}}(P(\sigma|s_1), P(\sigma|s_2)) \quad (1)$$

where the symmetrized KL divergence is $D_{\text{KL}}^{\text{sym}}(p, q) = 0.5(\sum_x [p(x)\log_2 p(x)/q(x) + q(x)\log_2 q(x)/p(x)])$ [14]. We choose this principled information-theoretic measure because it quantifies the difference between stimuli precisely to the extent that their response distributions are

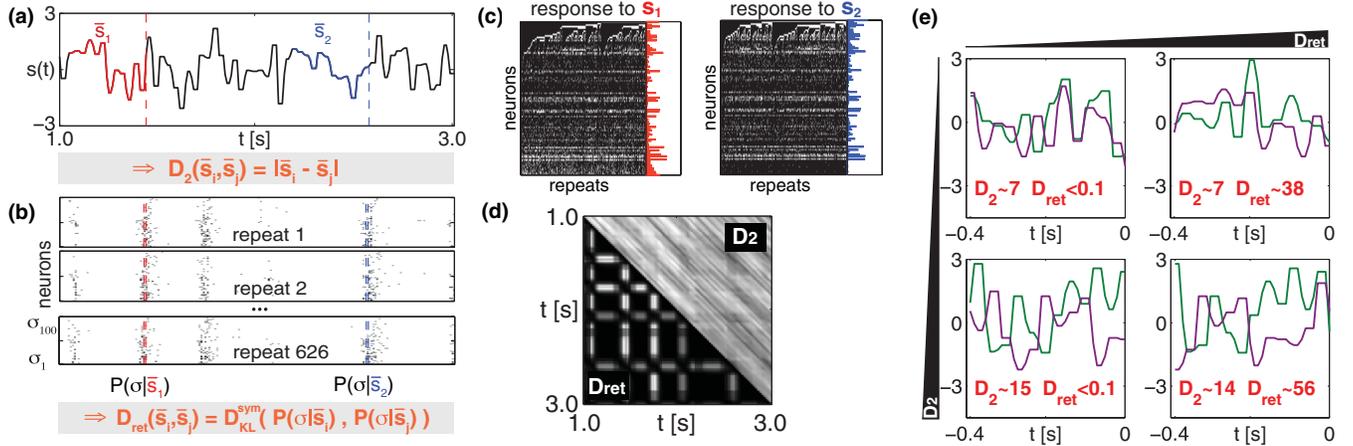


FIG. 1 (color). (a) Stimulus segment with two 400 ms stimulus clips s_1 (red), s_2 (blue). (b) Rasters for 100 RGCs shown for three example repeats; response vectors to s_1 , s_2 shown between dashed lines. (c) Measured response rasters to two highlighted stimulus clips (s_1 , left; s_2 , right; spikes = white dots; colored bars = neuron firing rates). (d) For every pair of time bins in the experiment, the Euclidean distance D_2 between the corresponding stimulus clips is shown in the upper diagonal part of the matrix, and the retinal distance D_{ret} in the lower part. (e) Four pairs of stimulus clips (green and violet) are selected to demonstrate that D_2 (increasing top to bottom) and D_{ret} (increasing left to right) are not monotonically related.

distinguishable [14,15]. Once constructed, the analysis of D_{ret} should help us uncover the fundamental aspects of the stimulus space, in particular, whether the distance between pairs of stimuli is determined by a small number of features or stimulus projections.

Using the SDME model for $P(\sigma|s)$ [16], we computed the retinal distance, D_{ret} , between every pair of stimulus clips in the experiment. Figure 1(d) shows the profound difference between the Euclidean and retinal distance for all the pairs in a two-second interval of the stimulus. In particular, the same value of Euclidean distance can be obtained from very similar neural responses (as measured by D_{ret}), or very different ones, as shown in Fig. 1(e) for selected pairs of stimulus clips.

To further explore the structure of retinal distance, we used multidimensional scaling (MDS) to project the distance matrix $D_{\text{ret}}(s_i, s_j)$, between all pairs ($i, j = 1, \dots, T$) of presented stimulus clips, into a low-dimensional space [17]. This embedding technique does not directly reveal the structure of the stimulus space, but can approximate its effective dimensionality. Every stimulus clip s_i is assigned a K -dimensional point \vec{v}_i in Euclidean space, such that

$$D_{\text{ret}}(s_i, s_j) \approx f(\|\vec{v}_i - \vec{v}_j\|), \quad (2)$$

with $f(\dots)$ being a monotonic function. It is possible to find such accurate mappings for small K : Fig. 2(a) shows the strong correspondence between low-dimensional MDS projections and the retinal distance, for different K values. Figure 2(b) summarizes the MDS performance at low orders in terms of the information it captures about D_{ret} ; higher information values correspond to a tighter relationship with less scatter. Figure 2(c) shows the structure of stimulus space using MDS with $K = 2$, which already captures most of the structure of the stimulus space. The

first coordinate of the MDS projection, $v^{(1)}$, is strongly correlated with the average firing rate in the population: high values correspond to “off-like” stimuli, and small values correspond to flat or “on-like” stimuli that do not drive the neural population well. The increased sensitivity to off-like features is consistent with the known prevalence of off-type cells in the salamander retina and in our dataset. Although on-like stimuli differ substantially in their shapes and thus in their Euclidean distances D_2 , they are largely indistinguishable for the retina. In contrast, groups of

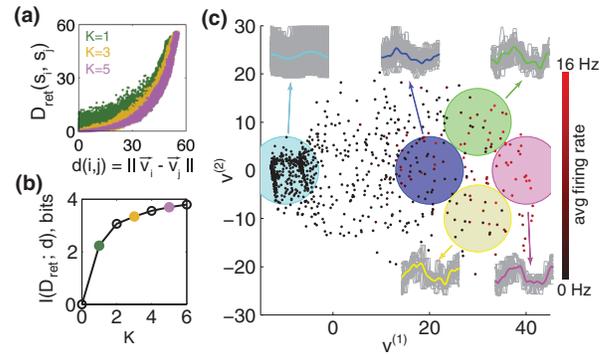


FIG. 2 (color). MDS assigns to each stimulus clip s_i a K -dimensional vector $\vec{v}_i = \{v_i^{(1)}, \dots, v_i^{(K)}\}$. (a) Relationship between D_{ret} (one dot = one pair of stimuli) and the Euclidean norm in MDS space, $d(i, j) = \|\vec{v}_i - \vec{v}_j\|$, gets tighter with K . (b) Information I between D_{ret} and MDS distance d as a function of embedding dimension K . (c) For $K = 2$, all $T = 961$ stimulus clips are shown as points with MDS coordinates $(v^{(1)}, v^{(2)})$; shade of red corresponds to mean population firing rate (scale at right). Five groups of stimuli are denoted by circles (gray lines = individual stimulus clips, color lines = average). Equal distances in the plane correspond to equal D_{ret} .

stimuli sharing the same $v^{(1)}$ coordinate (yellow and green) have a similar shape and evoke a similar population firing rate, yet are distinguishable because they differ along the second coordinate $v^{(2)}$. To the retina, yellow and green groups of stimuli are as distinct from each other as the blue and magenta groups at constant $v^{(2)}$ but different $v^{(1)}$.

Figure 3 presents a model that predicts the mapping of an arbitrary stimulus s into $v^{(1)}$, $v^{(2)}$ and thus allows us to compute D_{ret} . This model, obtained by using simple reverse correlation analysis [3], relies on two coupled linear-nonlinear transformations [Figs. 3(a)–3(d)], and identifies two dominant population-level stimulus features k_1 , k_2 : stimuli are distinguishable only insofar as their projections onto k_1 , k_2 differ. The model accurately predicts D_{ret} [Fig. 3(e)], establishing the relation of Eq. (12) to be $d(i, j) = \|\tilde{v}_i - \tilde{v}_j\| \approx \beta^{-1} \exp\{\alpha^{-1} D_{\text{ret}}(s_i, s_j)\}$ [Fig. 3(d)]. Interestingly, this relation, which we did not assume *a priori*, is exactly the kernel function used in several very successful applications of support vector machine (SVM) classification in machine learning, where one needs to distinguish between “classes” (here, stimulus clips) based on the distributions over “features” (here, neural responses) that they induce [18]. Our findings indicate that the neural population, much like single neurons, performs low-order dimensionality reduction on the incoming stimuli; however, unlike single neurons that can signal only a binary decision in every time bin, the population has access to $\sim 2^N$ states which can encode the variation along the relevant directions with greater precision. This view is consistent with the reported highly redundant code in the retina [19].

Given the failure of the Euclidean metric to predict stimulus similarity, and the success of the low-dimensional model, we asked whether a general quadratic form could explain the retinal distance. We thus looked for an optimal matrix $g_{\mu\nu}$, such that $D_{\text{ret}}(s_i, s_j) \approx \sum_{\mu\nu} (s_i^{(\mu)} - s_j^{(\mu)}) g_{\mu\nu} (s_i^{(\nu)} - s_j^{(\nu)})$, where μ, ν range over all 40 components of stimulus clip vectors s . Using cross-validated least-squares fitting, we found that the optimal $g_{\mu\nu}$ substantially outperformed the Euclidean metric, yet still only captured $\sim 20\%$ of the structure in D_{ret} (Fig. 3(e)). The best-fitting matrix $g_{\mu\nu}$ has a simple structure that is captured by two eigenvectors, matching the pair of population-level stimulus features, k_1 , k_2 , independently inferred in Fig. 3(b). Despite this, the best-fitting static $g_{\mu\nu}$ performs poorly: the eigenvalues corresponding to k_1 , k_2 would have to depend on the stimulus in order to approximate well our D_{ret} model.

Our results carry important implications for stimulus decoding. The accuracy of our model for stimulus similarity enables us to create new stimuli that are similar to each other up to any desired distance. We used Monte Carlo simulation to generate ensembles of full-length stimuli, such that *each* clip from the generated stimulus is less than Θ distant (as measured by D_{ret}) from the corresponding clip in the original stimulus displayed in the experiment [Fig. 4(a)]. The average coherence between the original stimulus and the generated stimuli at different Θ is shown in Fig. 4(b); it broadly peaks between 2–4 Hz, in agreement with reported decoding performance [20]. Our analysis predicts that for small enough Θ , all stimuli from such an ensemble are essentially indistinguishable to the retina.

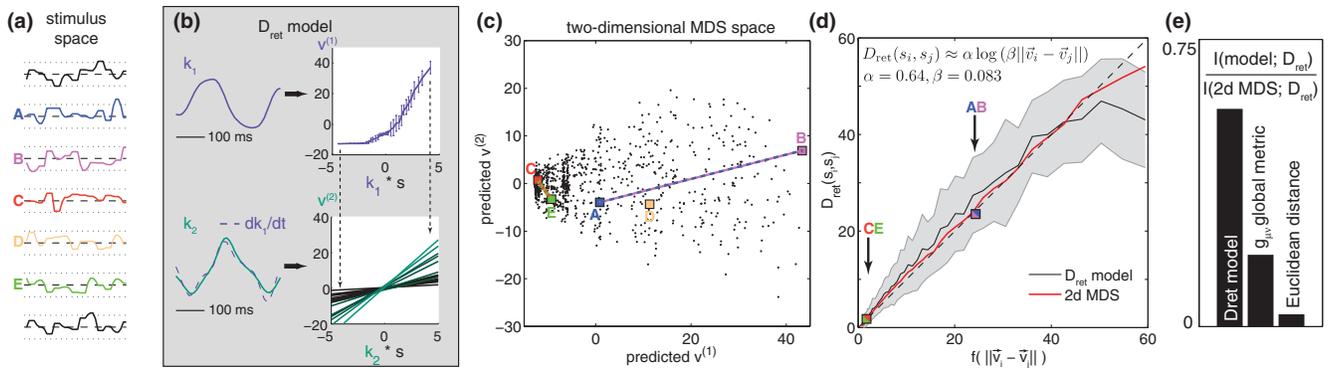


FIG. 3 (color). (a) Sample clips from the stimulus space (five highlighted in color). (b) A model for D_{ret} maps each stimulus clip s to a point $\tilde{v} = (v^{(1)}, v^{(2)})$ in the 2D MDS space. Predictions for \tilde{v} are obtained by filtering the stimulus clip s with a linear filter $k_{1,2}$ and passing the result through a pointwise nonlinearity [for $v^{(2)}$, the transformation is linear with the slope n that depends on the value $k_1 \cdot s$ as shown by dashed arrows; small $k_1 \cdot s$ correspond to low n (dark green), large $k_1 \cdot s$ to high n (bright green)]. The predicted MDS coordinates of s are $[v^{(1)}(k_1 \cdot s), n(k_1 \cdot s)k_2 \cdot s]$. (c) All $T = 961$ stimulus clips represented as points in a plane with coordinates predicted by the model in (b) (colored squares = highlighted stimuli). (d) Given predictions for coordinates $\tilde{v}_i = (v_i^{(1)}, v_i^{(2)})$, the D_{ret} for all ($\sim 4.6 \times 10^5$) pairs of clips (s_i, s_j) can be predicted using a simple fitted relation in the inset (black line = mean binned model predictions vs true D_{ret} , shaded area = std of binned predictions). Red line shows D_{ret} computed from true 2D MDS coordinates of Fig. 2(c). Two highlighted distances from (c) denoted by arrows. (e) The fraction of information about the true D_{ret} captured by the D_{ret} model in (b), by the best-fit global metric $g_{\mu\nu}$, and by the Euclidean distance, normalized by the success of the $K = 2$ MDS.

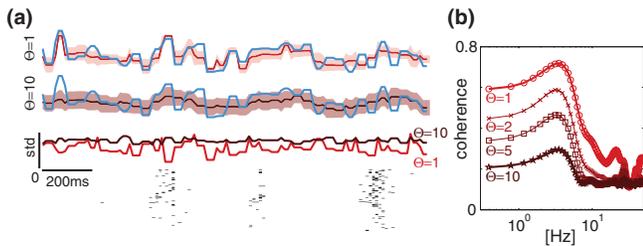


FIG. 4 (color). (a) Two Monte Carlo generated ensembles of long stimuli (ensemble mean = red; std = shade) where each clip $s_i(t)$ of the generated stimulus is less than $\Theta = 1, 10$ distant from the corresponding clip in the original stimulus $s_0(t)$ (blue line). Below: std over the ensembles ($\Theta = 1, 10$), showing how ensembles get more tightly constrained at particular instants corresponding to high neural activity (raster, bottom) and steep stimulus changes. (b) Average coherence, $C(f) = \langle |S_{0i}(f)|^2 / (S_{00}(f)S_{ii}(f)) \rangle_i$, between the original stimulus and the simulated stimuli (with spectra S_{00} and S_{ii} , respectively, and cross-spectrum S_{0i}), increases with decreasing Θ .

In contrast to Euclidean distance, which would allow the generated stimuli to fluctuate around the original waveform equally at every point in time for a given threshold Θ , the retinal distance constrains the possible set of stimuli much more at certain times than at others, reflecting a preference of the retina for encoding specific features in the stimulus. This is in part due to compressive nonlinearities, which squeeze large segments of stimulus space with small feature overlap into a small volume as measured by retinal distance [Fig. 2(c)], while emphasizing stimulus differences with high feature overlap. Consequently, distance models (such as Euclidean distance) with constant metrics are incapable of capturing the characteristics of the retinal distance. We conjecture that decoding approaches based on minimizing standard distance measures (e.g., Ref. [20]) may overly penalize deviations from aspects of the stimulus that are simply not represented in the neural responses, while not emphasizing strongly enough the deviations from population-level stimulus features identified here.

Understanding the neural code crucially depends on our ability to go beyond single-neuron spatiotemporal receptive fields, and identify how many, and which, are the population-level stimulus features encoded by an interacting population. We introduced a novel, biologically relevant distance measure on the space of stimuli based on the activity of large populations of neurons. This approach extended the single-neuron notions of stimulus feature extraction to the neural population, determining how a high-dimensional input space is partitioned and encoded by population responses. Our work thus suggests a principled alternative to arbitrary norms (like the Euclidean norm) for stimulus similarity and decoding, generalizes previous attempts to construct metrics for particular input spaces from neural responses (e.g., Refs. [21,22]), and

complements existing work on the dual problem of constructing relevant spike-train distance measures [23–25].

The approach we presented here will be instrumental in the analysis of upcoming experiments, which allow the recording of large parts of sensory neural circuits, or even of all the cells encoding some parts of the sensory scene [26]. This approach can be immediately applied to other sensory modalities, where it could reveal—much as we have found here—that the neural metric deviates considerably from our intuitive notions of similarity. Moreover, it can be extended to sensory domains where we lack any obvious notion of similarity, e.g., olfaction, for which there exists no natural distance between chemical stimuli [27]. More broadly, as the neural metric is based on the spiking activity itself, this framework can be taken beyond sensory modalities, to study perceptual metrics as well (e.g., Ref. [28]) or used to define neural-based distances for motor behavior that would be critical for neural prosthesis applications [29–31].

This work was supported by an Israel Science Foundation grant (No. 1629/12) and a European Research Council grant (No. 311238 Neuro-Popcode) to E. S.

*gkacik@ist.ac.at

†elad.schneidman@weizmann.ac.il

- [1] F. Rieke, D. Warland, R. R. de Ruyter van Steveninck, and W. Bialek, *Spikes: Exploring the Neural Code* (MIT Press, Cambridge, MA, 1997).
- [2] R. de Boer and P. Kuyper, *IEEE Trans. Biomed. Eng.* **BME-15**, 169 (1968).
- [3] O. Schwartz, J. W. Pillow, N. C. Rust, and E. P. Simoncelli, *J. Vis.* **6**, 484 (2006).
- [4] F. Rosenblatt, *Psychol. Rev.* **65**, 386 (1958).
- [5] R. R. de Ruyter van Steveninck and W. Bialek, *Proc. R. Soc. B* **234**, 379 (1988).
- [6] B. Agüera y Arcas and A. L. Fairhall, *Neural Comput.* **15**, 1789 (2003).
- [7] W. Bialek and R. R. de Ruyter van Steveninck, *arXiv:q-bio/0505003*.
- [8] R. R. de Ruyter van Steveninck, G. D. Lewen, S. P. Strong, R. Koberle, and W. Bialek, *Science* **275**, 1805 (1997).
- [9] R. N. Shepard, *Psychometrika* **22**, 325 (1957).
- [10] M. P. Young and S. Yamane, *Science* **256**, 1327 (1992).
- [11] M. Meister, J. Pine, and D. A. Baylor, *J. Neurosci. Methods* **51**, 95 (1994).
- [12] R. Segev, J. Goodhouse, J. Puchalla, and M. J. Berry II, *Nat. Neurosci.* **7**, 1154 (2004).
- [13] E. Granot-Atedgi, G. Tkačik, R. Segev, and E. Schneidman, *arXiv:1205.6438* [PLoS Comput. Biol. (to be published)].
- [14] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (Wiley, New York, 1991).
- [15] D. A. Butts and M. S. Goldman, *PLoS Biol.* **4**, e92 (2006).

- [16] SDME outperforms conditionally independent models for $P(\sigma|s)$, but the low-dimensional qualitative structure of the stimulus space identified in this Letter can be robustly reproduced with non-SDME models; see Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.110.058104> for [brief description].
- [17] I. Borg and P.J.F. Groenen, *Modern Multidimensional Scaling* (Springer, New York, 2005).
- [18] P.J. Moreno, P.P. Ho, and N. Vasconcelos, *Adv. Neural Inf. Process. Syst.* **16**, 1385 (2003).
- [19] J.L. Puchalla, E. Schneidman, R. A. Harris, and M. J. Berry II, *Neuron* **46**, 493 (2005).
- [20] D.K. Warland, P. Reinagel, and M. Meister, *J. Neurophysiol.* **78**, 2336 (1997).
- [21] C. Curto and V. Itskov, *PLoS Comput. Biol.* **4**, e1000205 (2008).
- [22] J.H. Macke, G. Zeck, and M. Bethge, *Adv. Neural Inf. Process. Syst.* **20**, 969 (2007).
- [23] J.D. Victor and K.P. Purpura, *Network* **8**, 127 (1997).
- [24] Y. Ahmadian *et al.*, Analyzing the Primate Retinal Code with Efficient Model-Based Decoding Methods, SfN 2008 poster, Society for Neuroscience, 2008.
- [25] E. Ganmor, R. Segev, and E. Schneidman, “Semantic organization of a neural population codebook and accurate decoding using a neural thesaurus,” *Cosyne Abstracts 2012*, Salt Lake City, Utah, USA.
- [26] O. Marre, D. Amodei, N. Deshmukh, K. Sadeghi, F. Soo, T. E. Holy, and M. J. Berry II, *J. Neurosci.* **32**, 14 859 (2012).
- [27] R. Haddad, R. Khan, Y.K. Takahashi, K. Mori, D. Harel, and N. Sobel, *Nat. Methods* **5**, 425 (2008).
- [28] J. Freeman and E.P. Simoncelli, *Nat. Neurosci.* **14**, 1195 (2011).
- [29] M. Velliste, S. Perel, M. C. Spalding, A. S. Whitford, and A. B. Schwartz, *Nature (London)* **453**, 1098 (2008).
- [30] C. E. Vargas-Irwin, G. Shakhnarovich, P. Yadollahpour, J. M. K. Mislou, M. J. Black, and J. P. Donoghue, *J. Neurosci.* **30**, 9659 (2010).
- [31] N. G. Hatsopoulos and J. P. Donoghue, *Annu. Rev. Neurosci.* **32**, 249 (2009).