Emergence of sparse coding via dendritic computation in a population of canonical visual binary neurons

Ulises Rodríguez-Domínguez and Hideaki Shimazaki

Graduate School of Informatics, Kyoto University, Kyoto, Japan

Abstract

Spike-count histograms of the activity in a population of neurons, e.g. at the visual cortex or hippocampus, exhibit a sparse profile with most of the neurons silent while highly synchronous states have considerably lower probability. The sparse coding, a scheme framed in the Bayesian brain hypothesis, represents natural stimuli by neurons under such conditions. The generative models of sparse coding of natural stimuli explained various features of early sensory systems, e.g., Gabor-like receptive fields of monkey or cat V1 neurons. However, these generative models typically are not neurally grounded, e.g., potentially assuming continuous-valued and negative neural activity rates. Here we propose a non-negative, spiking generative model grounded on a canonical circuit of pyramidal visual cortex binary neurons. The circuit captures the distribution of either the spontaneous (prior) or evoked (posterior) neural activity, where the spontaneous activity distribution exhibits silent neurons (sparse population), active neurons (hyper-active population) or a bimodal combination of both cases, where the former is characterized by alternating shrinking higher-order interactions between the neurons. This is regulated by the balance between dendritic nonlinearities at the proximal regions of a neuron which receive excitatory inputs, and distal inhibitions from interneurons, known to process predictive information. Compared to the non-sparse population, we demonstrate that, under the sparse population regime, the model captures Gabor-like spatial primitives as basis functions, improving the goodness-of-fit to the observed natural image patches. These results promise the construction of more complex spike-based generative models, which will allow us to test the Bayesian brain hypothesis by measuring directly from sensory cortices.

1 Objective

We aim to elucidate how a sparse distribution profile emerges from interacting binary neurons and specific neural mechanisms while also explaining sparse coding under the Bayesian brain hypothesis, which balances simplicity and interpretability.

2 Methods

We build a generative model of a population of binary visual cortical neurons guided by a canonical circuit found in pyramidal cortical neurons. The model for an image patch $\mathbf{y} \in \mathbb{R}^d$ and a hidden state $\mathbf{x} \in \{0,1\}^N$ is

$$\mathcal{P}\left(\mathbf{x}, \mathbf{y} | \boldsymbol{\theta}\right) = \mathcal{P}\left(\mathbf{y} | \mathbf{x}, \boldsymbol{\phi}\right) \mathcal{P}\left(\mathbf{x} | \boldsymbol{\omega}\right) \tag{1}$$

where $\mathcal{P}\left(\mathbf{x}|\boldsymbol{\omega}\right)$ is the probability mass function for the spontaneous activity (or prior), $\mathcal{P}\left(\mathbf{y}|\mathbf{x},\boldsymbol{\phi}\right)$ is the observation (or likelihood) Gaussian model and $\boldsymbol{\theta} = \{\boldsymbol{\phi}, \boldsymbol{\omega}\}$ the parameters. The Gaussian likelihood mean $\boldsymbol{\Phi}\mathbf{x}$ predicts the stimuli, with $\boldsymbol{\Phi}$ an over-complete vector space. We introduce the binary prior

$$\mathcal{P}\left(\mathbf{x}|\boldsymbol{\omega}, \boldsymbol{\Phi}\right) = \frac{h\left(\sum_{i=1}^{N} x_i\right)}{Z} \exp\left[g^{P_Q}\left(\mathbf{x}; f^{P_Q}\right) + g^{P_{NL}}\left(\mathbf{x}; \boldsymbol{\omega}^{P_{NL}}\right) - g^{I}\left(\mathbf{x}; \boldsymbol{\omega}^{I}, \boldsymbol{\Phi}\right)\right], \tag{2}$$

where Z is its partition function, $g^{P_Q}\left(\mathbf{x}; f^{P_Q}\right)$ at proximal dendrites considers the pairwise interactions with their relevance weighted by f^{P_Q}, while higher order interactions (HOIs) also at proximal dendrites are captured by the sigmoidal term $g^{P_{NL}}\left(\mathbf{x}; \boldsymbol{\omega}^{P_{NL}}\right)$. Such nonlinearity underlies alternating shrinking HOIs, required for non-independent sparse distributions [1]. Distal inhibitory input received at apical branches from either SOM or PV interneurons processing predictive information [2, 3] is captured by the term $g^{I}\left(\mathbf{x}; \boldsymbol{\omega}^{I}, \boldsymbol{\Phi}\right)$. The base measure function $h\left(\sum_{i=1}^{N} x_i\right) = 1 \Big/ \binom{N}{\sum_{i=1}^{N} x_i}$ is required for a widespread distribution to emerge for priors with homogeneous probability of k active neurons (see [1]), and our prior only partially departs from homogeneity. This term remains plausible as an excitatory dendritic log-transformation of laterally projected neural activity (see [4]).

3 Results

We learned the basis functions $\boldsymbol{\Phi}$ employing Gibbs sampling and neurally plausible gradient ascent updates under the Expectation-Maximization framework using 10,000 12x12 image patches from natural stimuli (see Figure 1). The average marginal likelihood, which measures the goodness-of-fit, improves if the strength of the distal inhibition is sufficient (Table 1), which corresponds to the sparse population regime.

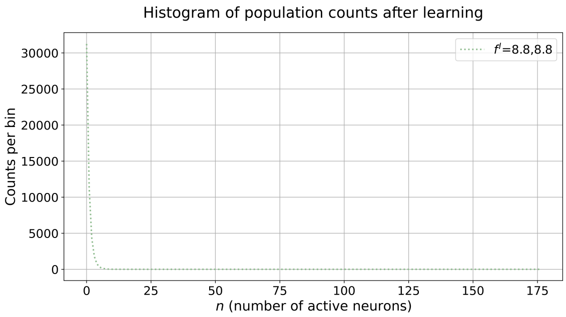

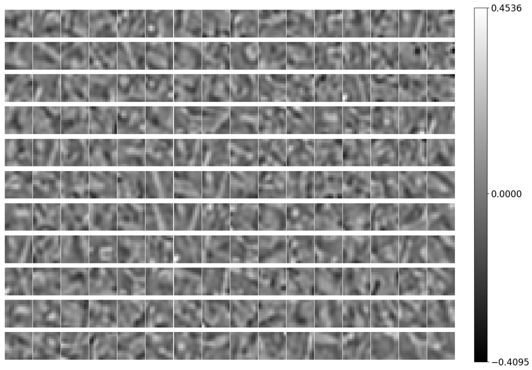

Figure 1: Left: prior distribution of spontaneous population activity (binary spike counts). Right: learned basis functions $\mathbf{\Phi}$ (synaptic weights).

Table 1: Average marginal likelihood (ML) across different conditions (60 per each value for the distal inhibition weights) on test natural image patches after learning. Results should be multiplied by $\left(\sqrt{2\pi\sigma^2}^d\right)^{-1}$ with $d = 144$ the dimension of the vectorized image patches and $\sigma = 0.75$ the standard deviation used for the likelihood.

Avg. strength of distal inhibition	Avg. ML (unnormalized)
1.15	0.0926 (0.0160)
2.25	0.1375 (0.1534)
4.45	0.2378 (0.1253)
8.85	0.3078 (0.1054)
17.65	0.3706 (0.0936)

4 Conclusions

Under the restrictions imposed by the canonical cortical circuit and the biologically plausible dendritic nonlinearities, we found that the sparse distribution profile emerges only if the the balance between the nonlinearities is such that the alternating-shrinking HOIs induced by the sigmoidal proximal nonlinearities dominate. We also found that the evidence (marginal likelihood of data) is improved for the sparse regime, especially when compared to the case when distal inhibitions are insufficient. To the best of our knowledge, we, for the first time, explained emergence of sparse coding using binary neurons by introducing dendritic nonlinear computation to model sparse population activity characterized by their alternating shrinking higher-order interactions.

References

[1] U. Rodríguez-Domínguez, H. Shimazaki, Alternating shrinking higher-order interactions for sparse neural population activity (2023). arXiv:2308.13257.

[2] G. Yang, J. Murray, X. Wang, A dendritic disinhibitory circuit mechanism for pathway-specific gating, Nature Communications 7 (12815) (2016).

[3] L. Hertäg, C. Clopath, Prediction-error neurons in circuits with multiple neuron types: Formation, refinement, and functional implications, Proceedings of the National Academy of Sciences 119 (13) (2022) e2115699119.

[4] P. W. Jones, F. Gabbiani, Logarithmic compression of sensory signals within the dendritic tree of a collision-sensitive neuron, The Journal of neuroscience: the official journal of the Society for Neuroscience 32 (14) (2012) 4923–4934.