Neural generative models with stochastic synapses capture richer representations

Maruan Al-Shedivat, Emre Neftci, and Gert Cauwenberghs

Abstract

Stochasticity in synaptic vesicle release is one of the major sources of noise in the brain. While the concept of cellular neural noise gave rise to computational models of biological learning such as deep belief networks and algorithms such as spike-sampling, the functional implications of synaptic stochasticity on learning remain unascertained and are often limited to filtering, decorrelation, or regularization. In this work, we approach synaptic stochasticity from the perspective of representations learning showing that it can improve fast concept learning in real situations where labeled data is scarce. We study a two-layer neural network that implements a Boltzmann machine with probabilistic connections. Noisy synaptic strengths lead to a notion of stochastic ensembles of generative models. We demonstrate how such ensembles can be tuned using variational inference methods. Analytically marginalizing synaptic noise for Bernoulli and Gaussian cases, we further use stochastic optimization techniques based on Gibbs sampling for learning the synaptic distributions.

The results have three interesting implications. First, our network does not use noise for mere regularization, as it is used by dropout techniques in artificial networks. Instead, the ensemble that results from synaptic stochasticity is fitted to the data and hence stores the information about its variability. Second, during the learning process only the strongest inhibitory synapses become very reliable while the rest remain unreliable. This relates our model to experimental cortical data. Finally, we demonstrate that knowledge represented in the stochastic ensembles learned in an unsupervised way can leverage the performance of the subsequent classification and enable one-shot learning—the ability to learn categories from a handful of examples. We hypothesize that synaptic stochasticity in the brain may encode large amounts of observed data in such stochastic ensembles and further use them for fast learning from limited discriminative information.

Additional Detail

Animals are able to learn and generalize entire categories of visual patterns from one or a handful of examples. What enables such fast learning? While this capability has been mainly approached from the perspective of high cognitive function [1], we are still to understand it on the neural circuits level. We hypothesize that to enable one-shot learning the brain might use large amounts of observed information encoded in stochastic ensembles stored in the noisy plastic synapses with adjustable release probabilities. Experimental studies confirm that synapses can individually adjust their neurotransmitter release probability dynamically through local field regulation [2]. This suggests that while the synaptic noise can be seen as a stochastic facilitation mechanism, explain Poisson-like spiking [3], or serve for regularization purposes [4], its adaptive nature could offer more intricate learning capabilities.

To understand the implications of the synaptic stochasticity on learning, we study specific computational models on the network level. We consider energy-based models (EBM), and specifically Boltzmann machines, which have been successful in interpreting neural adaptation in a probabilistic generative sense: Such networks can capture the distribution of the input stimuli and generate samples close to the ones observed at the earlier phase [5]. We introduced an extended EBM that results from synaptic unreliability:

$$P(\boldsymbol{v}, \boldsymbol{h}, \theta; \gamma) = \frac{e^{-E(\boldsymbol{v}, \boldsymbol{h}, \theta)}}{Z(\theta)} \pi(\theta; \gamma), \tag{1}$$

where $\boldsymbol{v}$ and $\boldsymbol{h}$ are the sets of visible and hidden variables (neurons), θ is a set of parameters (synapses),

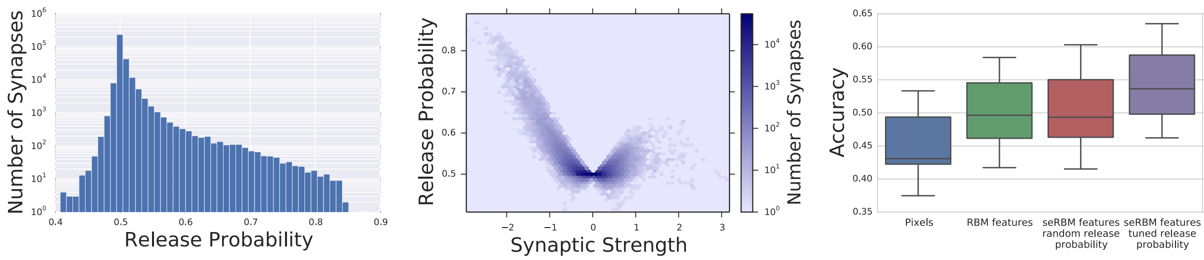

(a) Release probability distribution **(b)** Strength-probability correlation **(c)** One-shot classification accuracy

Figure 1. (a) Most of the tuned synaptic connections have low release probability, and only a small number of them became almost reliable. (b) Correlation between the synaptic strength and the release probability. Strong inhibitory synapses become the most reliable. (c) Accuracy of a logistic regression classifier trained on different data representations. RBM with random noisy synaptic connections leads to a small drop in the subsequent classification performance. When the stochasticity of the synapses is tuned, the representations lead to better generalization and performance of a classifier.

and $E(\boldsymbol{v}, \boldsymbol{h}, \theta)$ is some energy function. We consider synaptic strengths θ as another set of random variables with some parametric marginal distribution $\pi(\theta; \gamma)$. For a fixed θ, EBM represents one single model, and $\pi(\theta; \gamma)$ can be seen as a distribution over the ensemble of such models.

We trained a stochastic ensemble of restricted Boltzmann machines (seRBM) on the MNIST handwritten digits. In case of the Bernoulli synaptic noise—where every synapse was parametrized by a strength and a release probability—the trained model had a distribution of release probabilities with a peak near 0.5 (Figure 1a) which is close to the experimentally measured synaptic release between pyramidal cells in L2/3 area in the cat and rat brains [2]. Moreover, strong inhibitory synapses became more reliable (Figure 1b) which is consistent with the known influence of LTP and LTD on synaptic reliability [6].

We explored the advantage of having a stochastic ensemble of models. For every input stimulus, due to the probabilistic nature of the synapses, the network can generate many different representations. More precisely, for a given input $\boldsymbol{v}$, synapses can sample a model from the conditional distribution $P(\theta \mid \boldsymbol{v})$, and such model can produce a new representation $\boldsymbol{h}$ of the input. When the labeled data are limited while the unlabeled are abundant, a generative stochastic ensemble trained on the unlabeled part can effectively provide an extend number of labeled data representations and eventually leverage the subsequent classification accuracy (Figure 1c).

References

[1] Lake BM, Salakhutdinov R, Gross J, Tenenbaum JB (2011) One shot learning of simple visual concepts. In: Proceedings of the 33rd Annual Conference of the Cognitive Science Society. pp. 2568–2573.

[2] Branco T, Staras K (2009) The probability of neurotransmitter release: variability and feedback control at single synapses. Nature Reviews Neuroscience 10: 373–383.

[3] Moreno-Bote R (2014) Poisson-like spiking in circuits with probabilistic synapses. PLoS CB 10: e1003522.

[4] Wan L, Zeiler M, Zhang S, Cun YL, Fergus R (2013) Regularization of neural networks using dropconnect. In: Proceedings of the 30th International Conference on Machine Learning (ICML-13). pp. 1058–1066.

[5] Reichert DP, Seriès P, Storkey AJ (2013) Charles bonnet syndrome: evidence for a generative model in the cortex? PLoS CB 9: e1003134.

[6] Stevens CF, Wang Y (1994) Changes in reliability of synaptic function as a mechanism for plasticity. Nature 371: 704–707.