

UNIVERSITÄT HEIDELBERG



UNIVERSITÄT BERN



Fast inference with spiking networks

Mihai A. Petrovici

Computantional Neuroscience Group Department of Physiology University of Bern Electronic Vision(s) Kirchhoff Institute for Physics University of Heidelberg



Probabilistic computation with spikes

Probabilistic (Bayesian) computing: motivation

get an estimate of the <u>distribution</u> of the sought function (expected value <u>and</u> uncertainty)



Probabilistic (Bayesian) computing: experimental evidence



Probabilistic (Bayesian) computing: experimental evidence



Berkes et al. (2011)



Probabilistic (Bayesian) computing in machine learning





Requirements for probabilistic inference

A system that performs probabilistic inference has to

- \rightarrow represent probability distributions $p(z_1, z_2, ...)$
- \rightarrow calculate posterior (conditional) distributions $p(z_1, z_2, ... | z_k, z_{k+1}, ...)$
- \rightarrow evaluate marginal distributions $p(z_1, z_2) = \sum_{z_3, z_4, \dots} p(z_1, z_2, z_3, \dots)$





Representation of probability distributions

full state space



analytic / parametric

$$p(\boldsymbol{z}) = \frac{1}{Z} \exp\left[\frac{1}{2}\boldsymbol{z}^T \boldsymbol{W} \boldsymbol{z} + \boldsymbol{z}^T \boldsymbol{b}\right]$$



Sampling vs. parametric representation

temporal aspects:

- increasingly correct representation
- anytime computing



computational complexity aspects:

- computation of contitionals is simple
- marginalization is free



Spike-based encoding of an ensemble state

- $z_k = 1 \iff$ neuron has spiked in $[t \tau, t)$
 - ightarrow spike pattern encodes states $\, z^{(t)} \,$



Emulation of Boltzmann machines

$$z_k = 1 \iff$$
 neuron has spiked in $[t - \tau, t)$

 \rightarrow spike pattern encodes states $z^{(t)}$



Neural computability condition $u_k = \log \frac{p(z_k = 1 | \mathbf{z}_{\setminus k})}{p(z_k = 0 | \mathbf{z}_{\setminus k})}$ (which is equivalent to a logistic activation function $p(z_k = 1 | \mathbf{z}_{\setminus k}) = \frac{1}{1 + \exp(-u_k)}$). mediated by synaptic weights: Boltzmann distribution over $z_k \in \{0, 1\}$ $u_k = \sum_{k=1}^{K} W_{ki} z_i + b_k$

$$p(\mathbf{z}) = \frac{1}{\mathbf{Z}} \exp\left[\frac{1}{2}\mathbf{z}^T \mathbf{W}\mathbf{z} + \mathbf{z}^T \mathbf{b}\right]$$

Büsing et al. (2011), Petrovici & Bill et al. (2016)

Emulation of Boltzmann machines





$$u_k = \log \frac{p(z_k = 1 | \mathbf{z}_{\backslash k})}{p(z_k = 0 | \mathbf{z}_{\backslash k})}$$

(which is equivalent to a logistic activation function $p(z_k = 1 | \mathbf{z}_{\setminus k}) = \frac{1}{1 + \exp(-u_k)}$).

Idea: Stochasticity by Poisson background



$$\Rightarrow p_{\text{spike}} \approx \operatorname{erf}[\alpha \cdot (u_{\text{eff}} - \langle u_0 \rangle)]$$
$$\approx \sigma[\alpha \cdot (u_{\text{eff}} - \langle u_0 \rangle)]$$

unfortunately, neurons are a bit more complicated...

The diffusion approximation

noise source: Poisson spike trains high background firing rates relatively low synaptic weights

 $\Rightarrow \qquad \text{membrane as Ornstein-Uhlenbeck process} \\ du(t) = \quad \Theta \cdot [\mu - u(t)]dt + \sigma \, dW(t)$

Ricciardi & Sacerdote (1979)







First-passage-time calculations



Moreno-Bote & Parga (2004)

assumption: $\tau_{\rm syn} \gg \tau_{\rm m}$



p(z = 1)

then, the synaptic input appears quasistatic to the membrane

The membrane autocorrelation propagation

$$p(z_{k} = 1) = \frac{t_{k, refractory}}{t_{total}} = \frac{\sum_{n} P_{n} n \tau_{ref}}{\sum_{n} P_{n} \cdot (n \tau_{ref} + \sum_{k=1}^{n-1} \frac{r_{k}}{r_{k}} + T_{n})}$$
Petrovici & Bill et al. (2016)
$$P_{n} = \left(1 - \sum_{i=1}^{n-1} P_{i}\right) \cdot \int_{V_{thr}}^{\infty} dV_{n-1} \ p(V_{n-1}|V_{n-1} > V_{thr}) \begin{bmatrix} V_{thr}}{\int_{-\infty}^{\infty} dV_{n} \ p(V_{n}|V_{n-1})} \end{bmatrix}$$

$$T_{n} = \int_{V_{thr}}^{\infty} dV_{n-1} \ p(V_{n-1}|V_{n-1} > V_{thr}) \begin{bmatrix} V_{thr}}{\int_{-\infty}^{\infty} dV_{n} \ p(V_{n}|V_{n-1}) \langle FPT(V_{thr}, V_{n}) \rangle \end{bmatrix}$$

$$T_{n} = \int_{0}^{\infty} du_{k} \ \tau_{eff} \ln \left(\frac{\varrho - u_{k}}{\vartheta - u_{k}}\right) p(u_{k}|u_{k} > \vartheta, u_{k-1})$$

$$A \qquad H_{thr}^{H,H} \ h_{thr}^{H,H$$

(Fully visible) LIF-based Boltzmann machines



Petrovici & Bill et al. (2016)

Beyond Boltzmann: Spiking Bayesian networks



Probst & Petrovici et al. (2015)

Deep spiking discriminative architectures



Deep pong





Roth, Zenk (2017)

Deep spiking generative architectures



Short-term plasticity enables superior mixing



... so where does the noise come from?

1st approximation: independent Poisson sources



unrealistic in both biological & artificial systems

better: common pool of presynaptic partners

- \Rightarrow correlated inputs
- \Rightarrow deviation from target distribution

Embedded stochastic inference machines

more realistic: sea of noise



Noiseless stochastic computation



ongoing work with Dominik Dold and Ilja Bytschok



Physical emulation of spiking networks

Simulation: size & time



naturesimulationsynaptic plasticitysecondshourslearningdaysyearsdevelopmentyearsmillenniaevolution> millennia> millions of years

Simulation & emulation: energy scaling





Analog neuromorphic hardware



Schemmel et al. (2010)

Analog neuromorphic hardware





Schemmel et al. (2010)

Adaptive Exponential I&F Model $C_m \dot{u} = g_L (u - E_L) + g_{\text{syn}} (u - E_{\text{syn}}) + g_L \Delta_T \exp\left(\frac{V - V_T}{\Delta_T}\right) - w$ $\tau_w \dot{w} = a(V - E_L) - w$

Waferscale integration \rightarrow BrainScaleS system



Schemmel et al. (2010)

The Hybrid Modeling Facility in Heidelberg

4 million AdEx neurons, 1 billion conductance-based synapses, under construction



Hardware is not software...









LIF sampling on accelerated hardware



Petrovici & Stöckel et al. (2015), Petrovici et al. (2017)

Robustness from structure



Petrovici & Schröder et al. (2017)



Outlook / Work in

progress

Ensemble dynamics



$$p(\sigma) = \frac{1}{Z} \exp\left[\frac{1}{2}\sigma^T J \,\sigma + \sigma^T b\right]$$

spiking networks modeling magnetic systems

$$p(z) = \frac{1}{Z} \exp\left[\frac{1}{2}z^T W z + z^T b\right]$$

ongoing work with Andreas Baumbach

Quantum many-body problems



Carleo & Troyer (2017)

Learning rules

 $\dot{w}_k^a = \eta \, \left(\phi(U) - \phi(\alpha V^a) \right) \, \mathrm{PSP}_k$



backprop

$$\Delta w_{ij} \propto \frac{\partial E}{\partial o_j} \frac{\partial o_j}{\partial \operatorname{net}_j} \frac{\partial \operatorname{net}_j}{\partial w_{ij}}$$



ongoing work with Joao Sacramento and Walter Senn

References

- Bishop, C. M. (2006). Pattern recognition. Machine Learning, 128, 1-58.
- Starkweather, C. K., Babayan, B. M., Uchida, N., & Gershman, S. J. (2017). Dopamine reward prediction errors reflect hidden-state inference across time. Nature Neuroscience, 20(4), 581-589.
- Berkes, P., Orbán, G., Lengyel, M., & Fiser, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. Science, 331(6013), 83-87.
- Salakhutdinov, R., & Hinton, G. (2009, April). Deep boltzmann machines. In Artificial Intelligence and Statistics (pp. 448-455).
- Schemmel, J., Briiderle, D., Griibl, A., Hock, M., Meier, K., & Millner, S. (2010, May). A wafer-scale neuromorphic hardware system for large-scale neural modeling. In *Circuits and systems (ISCAS), proceedings of 2010 IEEE international symposium on* (pp. 1947-1950). IEEE.
- Buesing, L., Bill, J., Nessler, B., & Maass, W. (2011). Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons. *PLoS Comput Biol,* 7(11), e1002211.
- Petrovici, M. A., Bytschok, I., Bill, J., Schemmel, J., & Meier, K. (2015). The high-conductance state enables neural sampling in networks of LIF neurons. BMC Neuroscience, 16(1), O2.
- Petrovici, M. A., Bill, J., Bytschok, I., Schemmel, J., & Meier, K. (2016). Stochastic inference with spiking neurons in the high-conductance state. Physical Review E, 94(4), 042312.
- Ricciardi, L. M., & Sacerdote, L. (1979). The Ornstein-Uhlenbeck process as a model for neuronal activity. Biological cybernetics, 35(1), 1-9.
- Brunel, N., & Sergi, S. (1998). Firing frequency of leaky intergrate-and-fire neurons with synaptic current dynamics. Journal of theoretical Biology, 195(1), 87-95.
- Moreno-Bote, R., & Parga, N. (2006). Auto-and crosscorrelograms for the spike response of leaky integrate-and-fire neurons with slow synapses. *Physical review letters, 96*(2), 028101.
- Probst, D., Petrovici, M. A., Bytschok, I., Bill, J., Pecevski, D., Schemmel, J., & Meier, K. (2015). Probabilistic inference in discrete spaces can be implemented into networks of LIF neurons. *Frontiers in computational neuroscience*, 9.
- Leng, L., Petrovici, M. A., Martel, R., Bytschok, I., Breitwieser, O., Bill, J., ... & Meier, K. (2016). Spiking neural networks as superior generative and discriminative models. Cosyne Abstracts, Salt Lake City USA.
- Jordan, J., Tetzlaff, T., Petrovici, M., Breitwieser, O., Bytschok, I., Bill, J., ... & Diesmann, M. (2015). Deterministic neural networks as sources of uncorrelated noise for probabilistic computations. *BMC Neuroscience, 16*(Suppl 1), P62.
- Diesmann, M. (2013). The road to brain-scale simulations on K. Biosupercomput. Newslett, 8(8).
- Petrovici, M. A.*, Stöckel*, D., Bytschok, I., Bill, J., Pfeil, T., Schemmel, J. & Meier, K. (2015). Fast sampling with neuromorphic hardware. Advances in Neural Information Processing Systems (NIPS).
- Petrovici, M. A., Schroeder, A., Breitwieser, O., Grübl, A., Schemmel, J., & Meier, K. (2017). Robustness from structure: Inference with hierarchical spiking networks on analog neuromorphic hardware. *arXiv preprint arXiv:1703.04145*. (To appear in Proceedings of the ISCAS 2017.)
- Petrovici, M. A., Schmitt, S., Klähn, J., Stöckel, D., Schroeder, A., Bellec, G., ... & Güttler, M. (2017). Pattern representation and recognition with accelerated analog neuromorphic systems. *arXiv preprint arXiv:1703.06043*. (To appear in Proceedings of the IJCNN 2017.)
- Carleo, G., & Troyer, M. (2017). Solving the quantum many-body problem with artificial neural networks. Science, 355(6325), 602-606.