

POSTER PRESENTATION

Open Access

Policy gradient rules for populations of spiking neurons

Johannes Friedrich*, Robert Urbanczik, Walter Senn

From Twentieth Annual Computational Neuroscience Meeting: CNS*2011
Stockholm, Sweden. 23-28 July 2011

Population coding is widely regarded as a key mechanism for achieving reliable behavioral decisions given a high neuronal variability. Here, we present two general recipes to derive learning rules from a policy gradient approach for different neural codes and decision making networks, one based on partial integration across feature values, and one based on linear approximation around a target feature. The first technique leads to a tightly code-specific learning rule where details of the code-irrelevant spiking information are integrated away and the code-specificity enters at the synaptic level. The second technique yields modular learning rules which can be weakly code-specific, with a spike-timing dependent base synaptic plasticity rule which is modulated by a code specific population and decision signal. Decisions can be binary, multi-valued, or even continuous-valued. For illustration, we consider a spike count and a spike latency code. We test them on simple model problems and assess the superiority of tight over weak code-specificity with respect to the performance. While code-specific rules increase the performance only marginally when considering a single neuron [1], our tightly code-specific rule designed for population coding can strongly boost performance. Both code-specific learning rules improve in performance with increasing population size as opposed to standard reinforcement learning [2]. For mathematical clarity we presented the rules for an episodic learning scenario. But a biological plausible implementation of a fully online scheme is also possible [2,3].

Published: 18 July 2011

* Correspondence: friedrich@cns.unibe.ch
Department of Physiology, University Bern, Bern, Switzerland
Full list of author information is available at the end of the article

References

1. Sprekeler H, Hennequin G, Gerstner W: **Code-specific policy gradient rules for spiking neurons**. In *Advances in Neural Information Processing Systems* Bengio Y, Schuurmans D, Lafferty J, Williams CKI, and Culotta A 2009, **22**:1741-1749.
2. Urbanczik R, Senn W: **Reinforcement learning in populations of spiking neurons**. *Nature Neuroscience* 2009, **12**(3):250-252.
3. Friedrich J, Urbanczik R, Senn W: **Learning spike-based population codes by reward and population feedback**. *Neural Computation* 2010, **22**(7):1698-1717.

doi:10.1186/1471-2202-12-S1-P111

Cite this article as: Friedrich et al.: Policy gradient rules for populations of spiking neurons. *BMC Neuroscience* 2011 **12**(Suppl 1):P111.

**Submit your next manuscript to BioMed Central
and take full advantage of:**

- Convenient online submission
- Thorough peer review
- No space constraints or color figure charges
- Immediate publication on acceptance
- Inclusion in PubMed, CAS, Scopus and Google Scholar
- Research which is freely available for redistribution

Submit your manuscript at
www.biomedcentral.com/submit

 **BioMed Central**