Biophysically Realistic Models for Learning in Neural Networks

Justin Werfel

Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

http://www.ai.mit.edu

The Problem: How can a distributed system of independent processors, armed with local communication among neighbors and without global knowledge of the structure of the entire network, be trained to perform a desired task? Standard approaches for training such networks suffer from problems of implementation using only local information, and if meant to model systems such as the brain, further have serious difficulties in terms of biological plausibility.

Motivation: Neural networks are widely studied both as a biological model and a potentially useful tool[3]; and there is interest in coordinating networks of independent units, laid out without a predetermined architecture, to perform useful tasks. [1]

The canonical method for training neural networks, backpropagation, does direct gradient descent on an error function [4]. However, it requires an error signal to be sent backwards through the same network that propagates a signal forwards, requiring each synapse to be two-way and multipurpose, and each unit to be able to recognize and transmit two kinds of signals; a biological implementation of these requirements is problematic. Moreover, while the rule is local in space, the adjustment to the weight of a connection between two units depends on the activity of one unit at one time and the error associated with the other unit at a later time; this temporal nonlocality requires each unit to have information about the structure of the entire network, so that it can associate the earlier activity with the error signal after the appropriate propagation delay. This need for global knowledge is a serious obstacle for independent computational units with only local information. What's more, in the case of trajectory learning, backpropagation requires that the network actually run backwards in time. Issues such as these motivate investigations into alternative algorithms.

Previous Work: Williams [5] introduced a class of algorithms he called REINFORCE. The advantages of these algorithms include the following: in adjusting its connections to other units, each unit needs only information that is purely local, both spatially and temporally; a single global reward signal is maximized if each unit acts individually to maximize it; and in the case of trajectory learning, a single accumulator for each unit, keeping a running total of a given quantity, suffices to specify the update when the reward signal is ultimately delivered.

However, one disadvantage is that REINFORCE performs only noisy gradient descent. Williams showed that these algorithms follow the gradient on average, but made no statement about how great the noise in their estimates is, or if and how the noise can be sufficiently minimized that the algorithm is useful in practice.

Approach: REINFORCE algorithms have a learning rule in which the weight adjustment ΔW can be written as

$$\Delta W \propto -(E - b_{ij})e_{ij} \tag{5}$$

where E is the error incurred by the network for a given training example, b_{ij} is a baseline value that may vary locally, and e_{ij} is a quantity called the characteristic eligibility. For at least some exemplars of this class, it can be shown mathematically that an appropriate choice of the baseline guarantees that the error will never increase at any adjustment—a much stronger statement than simply that the average adjustment follows the gradient downhill.

Benchmark tests training feedforward networks to classify handwritten digits further demonstrate that REINFORCE has the capability to perform nearly as well as backpropagation, as shown in Figure 1. Here REINFORCE attains the same levels as does backpropagation on three measures of error—squared error, which defines the landscape on which gradient descent occurs; percent correct classification on the training data set; and percent correct classification on an independent test set—at most a few epochs behind.

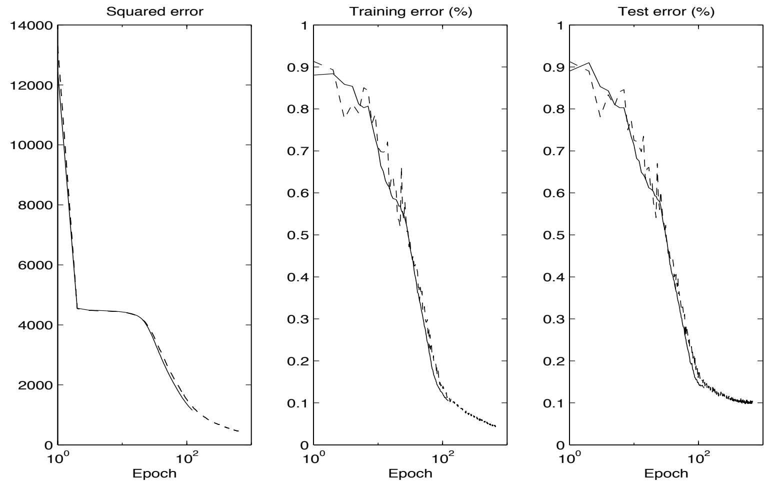

Figure 1: Learning curves for backpropagation (solid) and REINFORCE (dashed), for three error measures, on a digit classification task.

Impact: REINFORCE has not been widely used since its introduction, perhaps in part because of the availablility of established alternatives and the lack of guarantees on its performance; and with a poorly chosen baseline, it can perform quite badly. This work demonstrates that REINFORCE is able to train networks with learning curves directly comparable to those of backpropagation. Its implementation is easier than that of backpropagation, and units in the network need only perform simpler computations. It provides a model by which biological neural systems might plausibly learn, and is potentially useful for training artificial networks of independent processors. In addition, since its operation relies crucially on the presence of noise, it suggests a possible functionality for the stochastic nature of actual neurons.

Future Work: The extent of the success of REINFORCE is surprising; a basic noise analysis suggests that the component of the weight update in the direction of the true gradient ought to be swamped by the orthogonal component, so that REINFORCE might be expected to train a network much more slowly than a deterministic algorithm like backpropagation. Its observed success hence demands a better analysis of how it is able to overcome this apparent noise problem to be so effective.

There are also difficulties extending the theory to the case of trajectory learning, though REINFORCE can be shown to handle trajectories empirically. Additionally, we would like to apply the learning rule to more biophysically realistic models of spiking neurons [2].

Research Support: This work is supported by a National Defense Science and Engineering Graduate Fellowship from the U.S. Department of Defense, with additional support from the David and Lucile Packard Foundation.

References:

[1] H. Abelson, D. Allen, D. Coore, C. Hanson, G. Homsy, T. Knight, R. Nagpal, E. Rauch, G.J. Sussman, and R. Weiss. Amorphous computing. *Communications of the ACM*, 43(5):74–82, May 2000.

[2] P. Bartlett and J. Baxter. Hebbian synaptic modifications in spiking neurons that learn. Technical report, Australian National University Research School of Information Sciences and Engineering, 1999.

[3] J. Hertz, A. Krogh, and R. Palmer. *Introduction to the Theory of Neural Computation*. Addison-Wesley, 1991.

[4] B. Perlmutter. Gradient calculations for dynamic recurrent neural networks: A survey. *IEEE Transactions on Neural Networks*, 6(5):1212–1228, 1995.

[5] R. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8:229–256, 1992.