Learning a synaptic learning rule

Yoshua Bengio
McGill University, School of Computer Science,
3480 University street, Montréalt Qc, Canada, H3A 2A7

Samy Bengio and Jocelyn Cloutier
Université de Montréal, Département d'Informatique et
de Recherche Opérationelle, Montréal, Qc, Canada, H3C 3J7

Abstract

This paper presents an original approach to neural modeling based on the idea of searching, *with learning methods*, for a synaptic learning rule which is biologically plausible, and yields networks that are able to learn to perform difficult tasks. The proposed method of automatically finding the learning rule relies on the idea of considering the synaptic modification rule as a parametric function. This function has *local* inputs and is the same in many neurons. The parameters that define this function can be estimated with known learning methods. For this optimization, we give particular attention to gradient descent and genetic algorithms. In both cases, estimation of this function consists of a joint global optimization of (a) the synaptic modification function, and (b) the networks that are learning to perform some tasks. The proposed methodology can be used as a tool to explore the missing pieces of the puzzle of neural networks learning. Both network architecture, and the learning function can be designed within constraints derived from biological knowledge.

1 Introduction

One of the major goals of both biological neural networks modeling and artificial neural networks research is to discover better learning rules, to yield networks that can learn more difficult tasks, such as tasks that the brain can handle. Most researchers now accept that animal learning involves changes of synaptic efficacy [Byrn87]. Several researchers have proposed abstract models Connectionist models use fairly simple mechanisms for both the neuron operation and the modification of synapses. They can be used to solve some difficult learning problems, including the problem ("hard learning", [Hint89]) of training hidden layers of the network when reinforcement or supervision is only available to some neurons. However, as has been pointed out by Hinton [Hint89], a mathematically derived algorithm such as backpropagation [Rume86b] does not seem plausible as a biological model for many reasons. Synapses used in backpropagation models permit both forward and backward signals, a functionality not demonstrated to exist in biological synapses. Furthermore, the neurons have to propagate error derivatives through the axon backwards, as well as propagating activity level forwards.

Some models have recently been proposed that seek to narrow the gap between biological models and connectionist models (see for example [Bair90, Bart91]). This is also the goal of the approach presented here. The basic hypothesis of this approach is that it is possible to search with learning algorithms for a synaptic learning rule, and to constrain this rule in order to yield networks that can solve some difficult AI problems, respect aspects of biological knowledge about synaptic mechanisms and behavioral phenomena. The proposed method of automatically finding the learning rule relies on the idea of considering the synaptic modification rule as a parametric function, which has *local* inputs, and is the same in many neurons. The parameters that define this function can be optimized with known learning methods. For this optimization, we give particular consideration to gradient descent and genetic algorithms. In both of these cases, estimation of the synaptic modification rule consists of a joint global optimization of the synaptic modification rule, as well as of networks that learn to perform some tasks with this rule. The architecture of the networks, as well as the learning function can be designed with constraints derived from biological considerations. The proposed methodology can be used as a tool to explore the missing pieces of the puzzle of neural networks learning.

Initial experiments are described in order to assess the feasability of the proposed method in the case of very simple tasks. As a result of these initial experiments, a learning rule is discovered that is similar to backpropagation but is more biologically plausible.

2 Learning the learning rule

Studies on animal learning increasingly agree that learning involves changes in synaptic connections [Hawk89a, Byrn87]. Synaptic transmission is a complex process, in which changes in efficacy can be due to a number of factors, sometimes acting in concert, both at presynaptic and postsynaptic sites. Studies of several forms of learning (*e.g.* classical conditioning and sensitization) suggest that they are related and that complex forms of learning could be obtained by combining mechanisms of elementary forms of learning [Hawk89b]. This helps justifying a model of synaptic change that is a functional combination of several lower level mechanisms (see section 3). In this paper, we do not propose answers to the architectural questions, nor to the reinforcement/supervision problems. Instead, we describe a tool to help search for and refine the synaptic learning rule, given some architectural constraints and a learning criterion.

We consider an automatic method of finding and improving, with learning methods, a synaptic modification function that attempts to satisfy several constraints. To this end, we make the following assumptions: 1) The same rule is used in many neurons [1]. 2) There exists an input/output mapping that corresponds to the learning rule[2]. 3) This mapping can be approximated with a parametric function.

2.1 Gradient Descent

Suppose we want to minimize an error E that depends on parameters θ_i. Then in order to minimize E, the parameters should be updated with

$$\theta_i(t+1) = \theta_i(t) - \epsilon \frac{\partial E(t)}{\partial \theta_i(t)} \tag{1}$$

where ϵ is a very small but positive real number. There exist numerous variants of this method, to make it faster and more resistant to local minima, for example with the use of second-order information and an adaptable $\epsilon(t)$ [Beck89].

Consider a network of neurons and synapses (with strength w_i), and an optimization criteria E, which is a function of the behavior of the network, and that is to be minimized. Let us assume that $\frac{\partial E}{\partial w_i}$ can be computed ([Rume86b] or [Pear89] for a continuous recurrent network). Let the synaptic weight update at time t be defined as follows:

$$w_i(t+1) = w_i(t) + \Delta w_i(t) \tag{2}$$

and let $\Delta w_i(t)$ be a function of local observable quantities, as well as, of a *set of parameters* θ shared by all (or a lot of) synapses:

$$\Delta w_i(t) = \Delta w(local\ variables, \theta) \tag{3}$$

[1] This constraint can be relaxed to several rules for several types of neurons or synapses.

[2] However, a stochastic process may be involved in the computation of the synaptic change, for example as a variable noise term.

For example, a biologically plausible synaptic change function could have as local arguments a measure of presynaptic activity (y_{pre}), the postsynaptic potential (x_{post}), the strength of the synapse (w_i) and a measure of the activity of a facilitatory neuron (or of the concentration of a diffusely acting neuromodulator) (y_{modul}) that modulates the synaptic plasticity:

$$\Delta w_i(t) = \Delta w(x_{post(i)}(t), y_{pre(i)}(t-1), w_i(t), y_{modul(i)}(t), \theta) \tag{4}$$

To perform gradient descent on θ one computes the following derivative:

$$\frac{\partial E}{\partial \theta_j} = \sum_i \frac{\partial E}{\partial w_i(t)} \frac{\partial w_i(t)}{\partial \theta_j} \tag{5}$$

where $\frac{\partial w_i(t)}{\partial \theta_j}$ can be computed recursively ($\frac{\partial w_i(0)}{\partial \theta_j} = 0$):

$$\frac{\partial w_i(t)}{\partial \theta_j} = \frac{\partial w_i(t-1)}{\partial \theta_j} + \frac{\partial \Delta w_i(t)}{\partial \theta_j} \tag{6}$$

To avoid that θ be estimated in function of a particular synapse, neuron or network performing a particular task, it is important that the function $\Delta w(\theta)$ and the parameter set θ that defines it be the same for all (or a large number of) synapses, and that θ be estimated simultaneously with a population of networks learning to perform different tasks.

2.2 Genetic Algorithms

An interesting alternative to gradient based optimization methods is the set of optimization methods called "genetic algorithms" (see [Holl75] and [Gold88]). Genetic algorithms are learning algorithms inspired from several features of biological evolution. They consider a population of solutions to a problem, such as the synaptic modification function, encoded in artificial "chromosomes". Each member of the population is evaluated using an evaluation function. The population undergoes reproduction until a satisfactory performance is attained. During reproduction, "parents" are stochastically chosen to reproduce. This choice favors parents with highest evaluation, *i.e.*, best performance of the evaluation function. Operators such as "cross-over" are applied to the chromosomes of the parents to produce children that are inserted into the population. Domain knowledge can be exploited to create operators which improve the efficiency of the optimization procedure [Whit89].

It should be noted that interesting gradient-descent/genetic-algorithm hybrids have been proposed and could be considered here (see for example [Davi89] or [Whit89]). A coding scheme and a set of genetic operators for $\Delta w(\theta)$ could be designed, based for example on those proposed in [Whit89] to improve neural networks with genetic algorithms.

The advantages of genetic algorithms is that they are quite resistant to the problem of local minima in the optimization criteria and that they don't require a differentiable $\Delta w(\theta)$. However, this is obtained at the cost of improving simultaneously a whole population of solutions instead of just one.

2.3 Methodology

Both these learning methods require the implementation of several neural networks, doing several different tasks. Indeed, it is very important for these tasks to be as varied as possible, because they will constrain the learning power and the generality of the new learning rule.

Figure 1: *Proposed structure of the synaptic learning function $\Delta w(\theta)$: it contains multiple a-priori modules, each representing a different known or hypothesized synaptic plasticity mechanism. A free module may be added to allow the resulting $\Delta w(\theta)$ function to be more powerful.*

Some networks should work on complex learning problems such as pattern or speech recognition problems, to enforce efficiency constraints. There should also be tasks that consist in reproducing behavioral phenomena, such as associative conditioning, to reflect behavioral constraints. And finally, some very small networks may be trained to reproduce neurological phenomena such as habituation, recovery, dishabituation and sensitization.

The synaptic modification function can be implemented as a backpropagation [Rume86a] neural network with constraints, inputs and initial features that are derived from biological considerations, *i.e.*, biological knowledge will be used to bootstrap this function and force it to be plausible (see section 3). This network can have feedback and could be desribed by differential equations (see [Pear89] for the computation of the gradient in that case).

3 Biological Constraints

Biological plausibility constraints can be imposed on the learning rule. For example the rule can be designed in such a way that it uses only information which is local in time and in space: its inputs are locally measurable quantities, sampled in the near past. The rule can have some internal variables which represent a memory of their past values. However, this mechanism should not use an indefinitely large buffer of their past values, as does for example the time-unfolding algorithm for backpropagation recurrent networks [Rume86b].

A biologically plausible implementation of excitatory and inhibitory synapses should allow a different parameter set θ for the $\Delta w(\theta)$ function of excitatory and inhibitory synapses. A generalization of this idea is to allow multiple $\Delta w(\theta)$ functions for the diverse types of synapses that are observed. Various types of synapses, neurotransmitter, pre-, epi- and post-synaptic mechanisms were for example modeled in *Aplysia* [Gard87].

Another way to use biological knowledge is to bootstrap the function $\Delta w(\theta)$ so that it initially has access to a set of *a-priori* subfunctions equivalent to known or hypothesized synaptic modification mechanisms (see Figure 1). These *a-priori* subfunctions may themselves be parametrized functions. The optimization method described in section 3 can be used to search for an optimal combination of these subfunctions, as well as, to tune their parameters. In addition, a *free* subfunction could be included in the system that computes $\Delta w(\theta)$. This free subfunction is initialized with random parameters and it is used to perform computations which the *a-priori* subfunctions and their combination cannot provide, but which can help improving the satisfaction of constraints imposed on $\Delta w(\theta)$. This bootstrapping of $\Delta w(\theta)$ ensures that the networks will at least be able to learn, in the situations in which the *a-priori* subfunctions where shown to work.

An example of a biological model that could be embedded as a module of $\Delta w(\theta)$ is the generalized Hebbian rule proposed in [Done89]:

$$\Delta w_i = \epsilon \; y_{pre(i)} \; k \tag{7}$$

where ϵ is a positive valued learning rate parameter, $y_{pre(i)}$ is the presynaptic activity (at synapse i), and k is a reinforcing signal. That signal can either be the postsynaptic potential $x_{post(i)}$ or the activity of a special training signal, for example a neurofacilitator signal $y_{modul(i)}$, as in the case of the pairing-specific presynaptic facilitation observed in *Aplysia* [Hawk83], or in the case of neuromodulators that may be related to associative learning in *Hermissenda* [Crow89]. In general, several modulating signals should be allowed, corresponding to multiple neuromodulators, as in Figure 1.

Another example of a synaptic modification model that could inspire the design of $\Delta w(\theta)$ is the one proposed in [Klop89]. In this model, synaptic changes depend upon *changes* in pre- and post- synaptic activity, instead of depending upon the absolute levels of these variables. This could be implemented in the design of Figure 1, by allowing multiple delays (and/or recurrence) in the subfunction modules.

4 Conclusion and Initial Experiments

A few simple experiments were performed in order to evaluate the proposed method.

...

We present an original approach to neural modeling, based on the idea of searching *with learning methods* for a parametric synaptic learning rule that is biologically plausible, as well as yielding networks that can learn to perform difficult tasks. The networks architecture, as well as the learning function, can be designed with constraints derived from biological considerations. The method presented may help to bridge the gap between two approaches in neural networks research: mathematically-derived (*e.g.*, connectionist models [Rume86a]) and biologically-faithful models (*e.g.*, [Hawk89a]). This is not only important to improve biological modeling of brain function and learning, but it may also help to discover better artificial neural networks algorithms. By constraining the search for these algorithms to biologically plausible models, we increase the probability that these models will perform and generalize well, when facing tasks on which the brain performs well. It should be noted that the method of learning the learning rule can be viewed as a recursive process. If the resulting learning rule performs better than the original learning method used to learn that rule (*e.g.* gradient descent or genetic algorithms), then that learning rule should be used in a further iteration as a learning method in order to improve itself.

References

[Bair90] Baird B. (1990). Learning with synaptic nonlinearities in a coupled oscillator model of olfactory cortex. To appear in *Analysis and Modeling of Neural Systems*, F.H. Eeckerman ed.

[Bart91] Bartha G.T., Thompson R.F., and Gluck M.A. (1991). Sensorimotor learning and the cerebellum. To appear in *Visual Structures and Integrating Functions*, M. Arbib & J. Ewert eds, Springer-Verlag.

[Beck89] Becker, S. and Le Cun, Y. (1989). Improving the convergence of back-propagation learning with second-order methods. In Touretzky, Hinton and Sejnowski eds., *Proc. of the 1988 Connectionist Summer School*, pp. 29-37, San Mateo. Morgan Kaufmann.

[Beng89] Bengio Y., Cardin R., De Mori R. & Merlo E. (1989). Programmable execution of multi-layered networks for automatic speech recognition. *Communications of the Association for Computing Machinery*, vol. 32, no. 2, pp. 195-199.

[Byrn87] Byrne J.H. (1987). Cellular analysis of associative learning. *Physiological Review*, **67**, pp. 329-439.

[Crow89] Crow T. (1989). Associative learning, memory and neuromodulation in *Hermissenda*. In *Neural Models of plasticity*, J.H. Byrne & W.O. Berry, eds., pp. 1-21.

[Davi89] Davis L. (1989). Mapping neural networks into classifier systems. *Proc. Third International Conference on Genetic Algorithms*, J.D. Shafer ed., Morgan Kaufmann, pp. 375-378.

[Done89] Donegan N.H., Gluck M.A. and Thompson R.F. (1989). Integrating behavioral and biological models of classical conditioning. *Computational Models of Learning in Simple Neural Systems*, Hawkins R.D. & Bower G.H. eds. Academic Press. pp. 109-156.

[Gard87] Gardner D. (1987). Synaptic diversity characterizes Biological Neural Networks. *Proc. IEEE First International Conference on Neural Networks*, San Diego, CA, pp. IV-17 - IV-22.

[Gold88] Goldberg D. (1988). *Genetic Algorithms in Machine Learning, Optimization, and Search.* Addison-Wesley.

[Hawk89a] Hawkins R.D., Bower G.H. (eds.) (1989). *Computational Models of Learning in Simple Neural Systems.* Academic Press.

[Hawk89b] Hawkins R.D. (1989). A biologically based computational model for several simple forms of learning. *Computational Models of Learning in Simple Neural Systems*, Hawkins R.D. & Bower G.H. eds. Academic Press. pp. 65-108.

[Hawk83] Hawkins R.D., Abrams T.W., Carew T.J. and Kandel E.R. (1983). A cellular mechanism of classical conditioning in Aplysia: Activity-dependent amplification of presynaptic facilitation. *Science*, **219**, pp. 400-404.

[Hint89] Hinton G.D. (1987). Connectionist Learning Procedures. *Artificial Intelligence*, vol. 40, Sept. 89, pp. 185-234.

[Holl75] Holland J. (1975). *Adaptation in Natural and Artificial Systems.* University of Michigan Press.

[Klop89] Klopf A.H. (1989). Classical conditioning phenomena predicted by a drive-reinforcement model of neuronal function. In *Neural Models of plasticity*, J.H. Byrne & W.O. Berry, eds., pp. 104-132.

[Pear89] Pearlmutter B.A. (1989). Learning state space trajectories in recurrent neural networks. *Neural Computation* vol. 1, no. 2, pp. 263-269.

[Rume86a] Rumelhart D.E., McClelland J.L. (eds.) (1986). *Parallel Distributed Processing*, volume 1. Bradford Books, MIT Press.

[Rume86b] Rumelhart D.E., Hinton G.E. and Williams R.J. (1986). Learning internal representations by error propagation. Chapter 8 of *Parallel Distributed Processing*, volume 1. Rumelhart D.E. and McClelland J.L. (eds.), Bradford Books, MIT Press.

[Whit89] Whitley D. and Hanson T. (1989). Optimizing neural networks using faster, more accurate genetic search. *Proc. Third International Conference on Genetic Algorithms*, J.D. Shafer ed., Morgan Kaufmann, pp. 391-396.