Simple networks and learning rules for spike-timing based computation: learning rules

J. J. Hopfield
Department of Molecular Biology
Princeton University, Princeton NJ 08544-1014

Carlos D. Brody
Cold Spring Harbor Laboratory
PO Box 100, Cold Spring Harbor NY 11724

Abstract

Synaptic plasticity 'rules' enable a neural system to 'learn' or 'adapt'. Plasticity must also be capable of keeping a network in a functional state by repairing the damage that spontaneously takes place due to random synapse-changing processes. This repair should take place during normal operation of the system. We examine a (simulated) network of model neurons that has been designed to carry out olfactory tasks, in which the intensity-invariant recognition of a multiple odors is achieved by appropriate network design. Since the connections which should (and should not) be present are known and understood, it is possible to derive a synaptic plasticity rule, based on the timing of pre- and post-synaptic spikes in the operational designed network, that will maintain a functional network by selecting appropriate synapses to maintain or delete. This derived timing-based plasticity rule is strikingly similar to those experimentally found in LTP. In addition to enabling functional stability of known odor recognition in the appropriately designed system, we find that the same learning rule enables the system to perform single-trial learning of new odors in both unsupervised and supervised circumstances.

Introduction

In the first paper (Brody and Hopfield, paper preceding), we described the principles and behavior of a feed-forward network for which action potential timing and correlations are the essential dynamical computational features. This network implemented a 'many are equal' computing primitive, by synchronizing the action potentials of a group of neurons when their inputs were equal. Using these principles, we constructed a network capable of solving some of the computational problems that arise in the olfactory system. Throughout this second paper, we will refer to that explicitly designed network as the "engineered" solution. The engineered network could invariantly recognize a known odor over an intensity range of 50 and detect a weak known odor in the presence of an unknown background 3 times stronger. It made use of the relative (analog) strengths of the different glomeruli as an important part of its selectivity and background rejection. The operation of this feed-forward network is closely related to the 'computation by synchrony' feedback network introduced two years ago (Hopfield and Brody 2000; 2001), and the present form of network is a surrogate for the more complicated

feedback system. By introducing a common 'broadcast' oscillatory signal to a set of neurons, one of the important roles of the feedback connections can be mimicked without all the complications of feedback. The inclusion of this feature alone is adequate to generate significant computational power based on spike timing.

In the design of this style of network, the stimuli (odors) to be recognized are implicitly described by the pattern of connections between neurons. All the required synapses have the same strength; the analog information about the sensory pattern is itself carried only in the *pattern* **of connections, not in the strengths of the synapses. In our 'engineered' solution, we have selected** *by design* **the appropriate synaptic connection pattern.** *The possible relevance of such computation to biology depends critically on whether a biological plausible learning rule could generate and maintain the 'engineered' pattern of synaptic connections.*

In this paper, we derive a learning rule for this model sensory system based on the requirement that *'learning' must make the system robust and self-repairing in the task that it is performing.* **In its dependence on the timing of pre- and post-synaptic action potentials, the form of the learning rule is strikingly similar to that seen in LTP/LTD experiments. When applied to a novel stimulus in a** *single trial learning* **paradigm, this learning rule produces a set of synaptic connections that are equivalent to those of the 'engineered system'. The same learning rule is successful in** *unsupervised learning* **in an environment where the stimulus environment sequentially contains many novel stimuli in a random time-order.**

Throughout these two papers, we describe a highly selective cell as a 'grandmother cell' (γ-cell). This terminology is only a convenient shorthand for a 'pattern-selective cell'. We do not mean to imply a belief about the representation of high-level concepts by single cells, or the existence of cells devoted to a single concept. We describe an artificial system, but use the terms such as 'odor' and 'glomerulus' in order to focus on an identifiable and real pattern recognition problem.

The selection problem in network repairlearning

In the engineered olfactory system of the previous paper, each of N_g glomeruli was the input for a set of N_r neurons. Each of the N_r repertoire neurons associated with a given glomerulus has a different bias input current, and each receives the same input current from the glomerulus when the glomerulus is activated by an odor. A total of $N = N_g N_r$ neurons potentially make synapses on a single γ-cell (typical values for our simulations were $N_g = 400$, $N_r = 14$, $N = 5600$). When an odor is strongly present, roughly half of the glomeruli are active above threshold level, with different strengths of activity caused by the differing bias currents, and the different strengths of drive which the odor produces for different glomeruli. Typically 0.5 N ~ 2800 repertoire cells are thus driven by the odor, and all N cells driven at least by the bias currents.

The engineering solution chose a single best repertoire cell (from each of the strongly driven glomeruli) to connect to a given γ-cell, and thus made use of ~0.5N_g ~200 synapses. We require a learning rule that will choose similarly; which will describe which new repertoire cells to choose for connections when damage removes some of the original connections. The problem is difficult because all of the repertoire cells are firing at roughly comparable rates (in the phase-locking regime, there is a range of driving currents that produces different spike timings at the same firing rate- see companion paper). *Nevertheless, cConnections from repertoire cells too strongly driven and cells too weakly driven must both be rejected.* **Only input cells for which the combination of odor and bias currents sums to an appropriate intermediate level should receive synaptic reinforcement, and 'have synapses made'. This selection need not be exact; it is only important that the selection be good enough to be able to perform the computational task adequately.**

In searching for a learning rule for robust repair, we will make the assumption of *We will make two simplifying assumptions about the possible form of the learning rule, a stability hypothesis and an assumption of pairwise additivity***.** *robust network capable of self-repairimplemented by experimentally observed synaptic properties***The learning rule will depend on the relative timing of pre- and post-synaptic spike pairs, and be based on the sum of contributions from individual pairs.**

Let the system have connections from all repertoire cells to each γ-cell recognition unit. A sub-set of these, chosen according to the engineering network description, are strong and functionally active. All others are anatomically present but not functionally active, generating no synaptic currents (analogous to experimentally described "silent synapses" (Liao et al., 1995). Suppose that a few of the designed connections become functionally inactive. Expose the system sequentially to a diversity of odors, one of which is the odor known by that γ-cell. The stability hypothesis requires that the learning algorithm will make some new connections functional, selecting them from among the many possible connections in such a way that the functionality of the system, its ability to recognize the given odor, is repaired. The system is large; we do not ask that exactly the same connections that have been lost be restored. However, we do ask that a set of connections functionally equivalent to the original engineered set be relearned: that is, after a long time during which the system has been exposed many times to many odors, and most of the original connections have been replaced, the γ-cell should preserve its ability to recognize and discriminate the same odor that it initially recognized. Similarly, for a pre-synaptic neuron that should *not* **make a connection to a particular γ-cell, we require that when** *any* **single odor is sensed, the repair synaptic plasticity rule should be such as to leave that synapse at zero weight (or return it to zero weight if it had somehow acquired a non-zero weight).**

Another description of this same idea would allow each exposure to an odor remodels all the synapses, making strong synapses out of some of the previous silent ones (or vice versa) according to a 'learning rule'. Again, the requirement is that the learning rule preserves the functionality of the γ-cell in an environment in which

many different odors are (sequentially) present in the environment. This formulation leads to the same learning rule.
notany

Our pairwise additivity hypothesis states that the contribution of a pre-and post-synaptic action potential pairing to the synapse change procedure depends only on the time-difference between the pre- post- pairs of action potentials, and that the effects of all pre-post pairs (both occurring during a time window of relevance to a single trial learning protocol, in our case ~0.5 sec) are additive. Deviations from additivity have been reported, but for the simple task we have chosen, the particular kind of deviations observed (Froemke and Dan, 2002) have little effect.

Constructing a synapse change rule

Figure1 shows the response of the repertoire cells connected to a single glomerulus during a 'sniff' of a known odor, and the action potentials of a γ-cell that recognized that particular odor, in an engineered operational network. Before and after the sniff, the cells are driven only by bias currents, and the γ-cell does not fire. During the sniff, the glomerulus adds a sensory current to these repertoire cells. Only the repertoire cell marked with the asterisk was used in the engineered network, chosen because it had the (bias + sensory current) closest to the average of the chosen (bias + sensory current) across all the other active glomeruli (not shown). It is important that a repertoire cell with the right bias current is chosen. However, there is a range of bias currents within which the resolution of the system does not distinguish. Within this resolution, there are on average about 2 repertoire cells that would have done just as well. Statistically, it does not matter which of these is chosen. The use of 14 repertoire cells per glomerulus corresponds to more resolution (in input current levels) than the system as a whole can utilize. Since the actual spread of peak input currents for the engineered repertoire cells had a width σ_e (previous paper), any repertoire cell within an interval $\pm 2\sigma_e$ would have served about as well as those actually chosen.

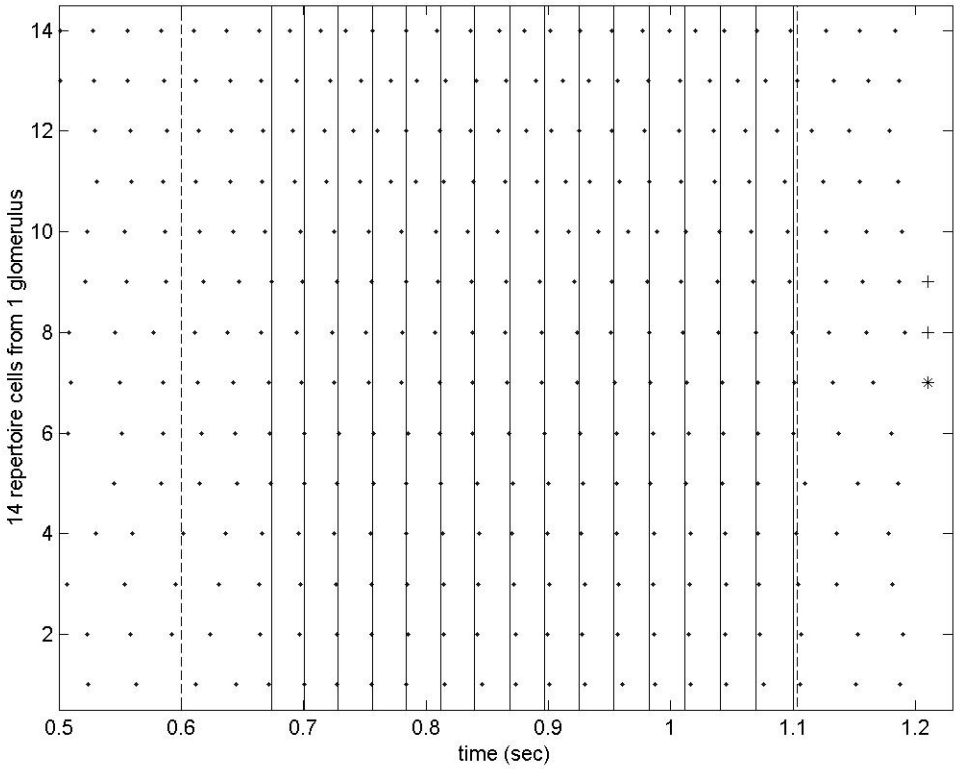

Figure 1. The spike rasters of the set of 14 repertoire cells driven by a single glomerulus. Each row represents a single repertoire cell. The beginning and end of the odor stimulus are indicated by the dashed lines. The rasters are displayed in the order of the bias currents that differentiate the repertoire cells. The γ-cell which recognizes the particular odor fires repetitively while the odor is present; its spikes are indicated by the vertical solid lines. The repertoire cell marked with the asterisk is the only one connected to the γ-cell in the engineered network. Two other cells have bias currents that lie within σ_e of the desired peak current, and would serve equally well in a network. These are identified by the + sign at the right of the rasters.

Our synapse change protocol is based on the timing of the pre- and post- synaptic action potentials of the cells involved (repertoire and γ-cells, respectively). It must be capable of distinguishing between the repertoire cells that are marked (any of which is appropriate for a connection to the γ-cell) and the others, which should not be connected. A single trial with a known odor on the 'engineered' network provides 5600 spike rasters of repertoire cells. Of these, about 540 examples are rasters for neurons whose input currents lie in a ± $2\sigma_e$ range of the optimum current. Any ~ 200 of these will perform almost equally well in driving the γ-cell to selectivelyl recognize that odor. The other subset of 5060 should not be used.

Thus, the problem is one of taking pre- and post- spike timing patterns and correctly classifying them into patterns corresponding to appropriate *vs*. inappropriate connections.

Because there is noise in the system, this classification cannot be completely precise. In a large system it is necessary only that a large majority be correct. From the additivity supposition, when there are several spike pairings j with different values of time difference Δt_j, the quantity M on the basis of which the decision about whether or not a repertoire cell is appropriate is given by

$$M = \sum_j W(\Delta t_j)$$

Here, $W(\Delta t)$ is an unknown mathematical function that describes how the significance of a pre- post- pairing depends on the time difference Δt between the pair. This function will next be determined from the large body of data on 'correct' and 'incorrect' patterns.

To make the conceptual description and computer simulations simple, the possible time intervals of pre-post- synaptic spike pairs are put into timing 'bins' k. The unknown function can then be described in terms of a set of unknown parameters w_k. We have used 29 bins of width 1 ms., with centers in the range -14 to +14 ms. around the time of the relevant γ-cell spike. Pre-post synaptic spike pairs separated by more than 14.5 msec. are presumed to have no effect. (Because all the relevant neurons are firing at about 35 hz, any effects due to pairs of greater than 14.5 ms. separation will be by this procedure actually assigned to a closer pair, a kind of aliasing due to the approximate periodicity of the spike trains) The pairings of a particular pre-synaptic cell and the γ-cell spikes can now be described by a set of integers n_k describing how many pairings occurred within each time bin k.
In these terms, M is given by

$$M = \sum_k w_k n_k$$

The parameters w_k must produce values of M that classify the spike patterns appropriately. If no such set of parameters can be found to accomplish this task, then a linear summation-based learning rule is not adequate for the task.

The spike-pairing classification problem formulated in this fashion is exactly the mathematical problem of pattern classification by a feed-forward 'neural network' having no 'hidden units'. The n_k are the inputs; w_k are the 'weights' and M is the input to the output 'unit'. We have used a procedure that trains an output unit to predict the *probability* **(based on its n_k pattern) that a given repertoire cell belongs to the class of neurons appropriate for connection to the γ-cell (Hopfield, 1987). It makes sense only to attempt to determine the probability; because of noise, it is possible that cells belonging to both classes can sometimes generate the same pattern n_k. We will later, in application, use the value of M to prescribe which connections are made in learning.**

The prediction of the network is taken to be the logistic function.

$$P = 1 / (1 + e^{-M})$$

The weights were determined by iteratively implementing weight change on each example according to the prescription

$$\delta w_j \sim ([1 \text{ or } 0] - P) * n_j \qquad [1 \text{ if a positive example, } 0 \text{ if a negative example}]$$

This procedure minimizes the K-L distance between the network-defined probability and the actual probability distribution without the necessity of defining the actual probability distribution explicitly. In this structure of feed-forward network with the K-L measure of error, there are no local minima when searching for the weights

.

Gradient descent in weight space therefore determines the *unique* **best** w_k

Fig 4 (top) plots the optimal weights w_k **derived from this procedure as a function of the time difference** Δt **between the pre- and post- synaptic spikes. Because the total number of spikes is very similar for all instances, the strength of a connection** w_b **to a bias unit can be traded off against the addition of a constant to each of the weights** w_k**, with no change in the classification performance of the network. The strength of the connection to the bias unit was chosen so that the** $W(\Delta t)$ **goes to zero for large** $|\Delta t|$**, i.e., that pairings too far apart in time will be ineffective. With this choice, it is also the case that** $\Sigma w_k \approx 0$**. Thus the** *emere existence* **of a pre- post pairing at** *some* **point in the interval contains 'no evidence' about the probability of the repertoire cell being among the appropriate set. Only the** *value* **of** Δt **contains evidence of which class a cell belongs to. All decision information is in the timing domain, not in the number of pairs** *per se***.**

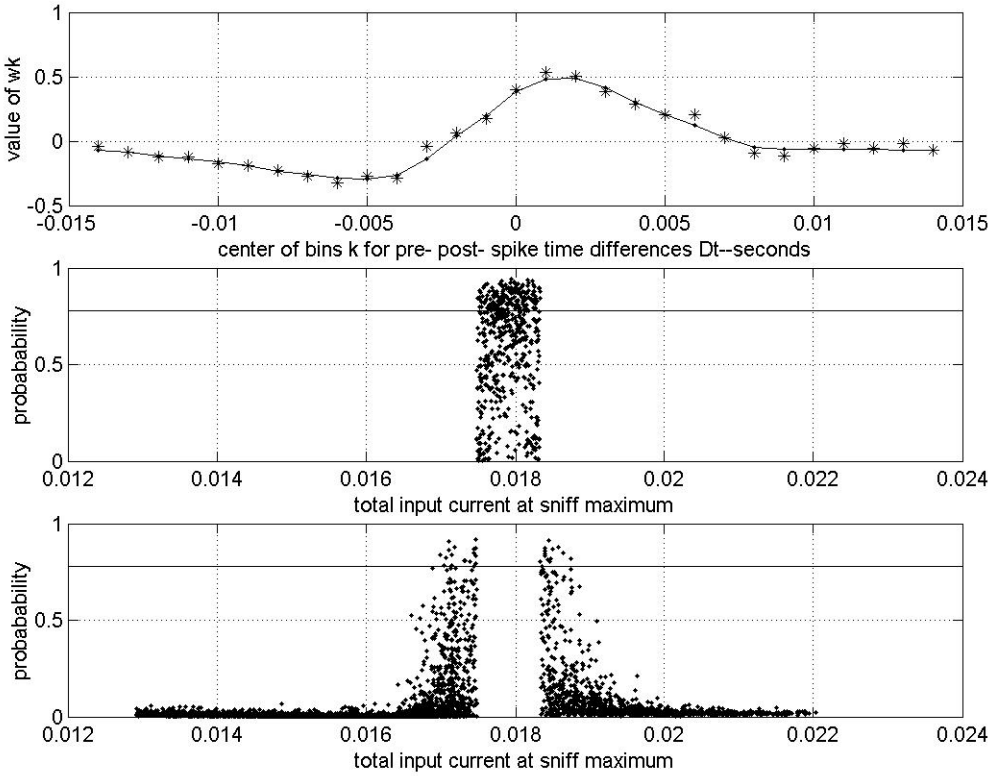

Figure 2. a) The weights w_k determined from 'learning' the classification problem, plotted against the time difference Δt of the pre- post- synaptic spike pairs. Positive values of Δt correspond to the pre-synaptic spike occurring earlier than the post-synaptic spike. The points indicated by * are an evaluation of the weights based on the spike trains during a single sniff. The solid points are the average of 16 such evaluations. The w_k are a discrete representation of an underlying *function*, and the solid line connecting the points indicates the shape of that function. The system had 400 glomeruli, and 14 repertoire cells for each glomerulus. In a single run, there were approximately 540 'good examples' and 5060 'bad examples' in the training.

b,c) The probability prediction that a given repertoire cell is appropriate for a connection to the γ-cell based on the weights shown in a) The ordinate is the driving current of a repertoire cell at the peak of the odor sniff. The engineered solution consists of making synapses of strength "1" to a set of N_g cells having very nearly the same current input (from the bias current and from the olfactory sensory cells). The engineered solution had a mean current of 0.0179 at this peak, and a standard deviation σ_ε = 0.00021

b) Predictions for 540 'positive examples' lying within the $\pm$ $2\sigma_e$ range. The horizontal line is drawn at p = 0.824. c) predictions for 5060 'negative examples' lying outside the +-$2\sigma_e$ range.

The quality of probability estimates available from this procedure can be examined by looking at the probabilities predicted from the using the average weights of Fig 2a. During a sniff, total input currents to repertoire cells grow, peak, and then decay. The value of the peak current for each repertoire cell is plotted as the abscissa in Figs. 2b and 2c. In the designed network, the "correct" repertoire cells, of which there were 200, one from each of 200 glomeruli, had peak currents close to 0.0179. We therefore chose for positive examples spike trains from repertoire cells

with peak input currents lying within the range 0.0179 $\pm 2\sigma_e$. The data shown in Figs 2b,c are from a single trial with a total of 5400 repertoire cells. Combined with the spikes of the γ-cell, the spikes of each repertoire cell generate a spikes-pair pattern vector n_k. Using the weights w_k of Fig. 2a, the abscissa in Figs. 2b and 2c plots the estimated probability that a repertoire cell's spikes-pair pattern vector is that of a positive example. The repertoire cells having input currents within the 'positive example' training range are plotted in Fig 2b (perfect probability estimation would assign to all of these probability 1); the 'negative example' range is plotted in Fig 2c (to which perfect estimation would assign probability 0).

Because of noise and random phasing of the starting situation, some cells do not fall into the correct synchrony pattern within 0.5 seconds, or have badly displaced spikes. But on the whole, 'positive examples' are assigned much higher 'estimated probabilities of being positive examples' than are the negative examples. Virtually all the repertoire cells that receive a high probability score are appropriate for inclusion. In the particular example shown here, the 'engineered solution' had 200 repertoire cells. Of these, 188 lie within the 'good example' region, while 12 lie outside it. (This corresponds well to the fact that the interval for 'good examples was chosen as $\pm 2\sigma_e$ range.) There are 200 repertoire cells that have probabilities above a threshold levelt of 0.781 (drawn in the figure). Of these, 170 are 'good examples' and 30 are 'bad examples'. Thus, given a superfluity of repertoire cells, there is no difficulty in picking a large number of appropriate neurons while including very few poor ones. The standard deviation of the learned solution around its mean is about 1.4 σ_e . Since the engineered system performs well with only 7 glomerular and twice as large a random spread (because 14 glomerular levels represent more resolution than the system can actually use) we anticipate that the learned synaptic connections will perform as well as the engineered one, as will be examined in the following section.

The qualitative nature of the shape of the w_k versus Δt plot in Fig 2a could be anticipated from the data shown in Fig. 1 or in the earlier paper (Brody and Hopfield, previous paper). The neurons in the designed solution on average fire in synchrony when the odor is present, and will induce the post-synaptic cell to fire after the integration of the fast excitatory synaptic current. A positive peak is thus expected near the peak of the integrated synaptic current, which occurs at 3 msec. Repertoire neurons driven with smaller (larger) input currents will fire later (earlier) than this chosen set. To discriminate against these, the positive peak should be relatively narrow, and be flanked by a negative region.

The shape of this synapse choice function has striking qualitative similarity to the shape of the synapse change timing relationship seen in LTP/LTD (Bi and Poo, 1998). Both favor making connections for post-synaptic spikes after pre-, and suppressing connections when post-synaptic spikes occur before pre-; both have a little spillover of positive values into the first few milliseconds of the negative time window. In considering the significance of such comparisons, one must remember that the quantitative aspects should (in order to be optimum) depend on the task

being performed, the synaptic and cell time constants, and the level of noise present. We have found that increasing the noise current in the repertoire cells broadens the positive peak in Fig 2a, and that decreasing the width of the band of positive repertoire cells sharpens this peak. Introducing a delay or α-function type response in the excitatory synaptic current will shift the curve to the right.*shapes that preserve the qualitative features*

Stability of the synapse pattern

To examine for stability and self-correction abilities, we have adopted a drastic protocol for learning based on the descriptions of Fig. 2. Consider a system repeatedly exposed to odors randomly chosen from a set of odors a,b,c,d, ... Begin with an engineered set of synapses, which for N_g = 400 contains about 200 connections. The cell recognizing odor c will not respond to odors a, b, d, ... etc.. We assume that the synapses to γ-cell c will change only when that cell fires, and thus (at least initially) only when odor c is present. When odor c first occurs, γ-cell c and all the repertoire cells produce action potentials. We now *eliminate all the previous synaptic connections* **to γ-cell c, and replace them with connections to the 200 most highly probably repertoire cells, as ranked by the prediction algorithm. With this new set of connections, we iterate the procedure.**

Table 1 illustrates the performance of the system by describing the number of action potentials generated by a γ-cell that is tuned to odor c for a variety of stimuli. There is little difference in the performance between the engineered solution, the first iteration of synapse renewal, and the tenth iteration of synapse renewal for these key tasks. All synapse sets respond well to the target odor over a range of more than 100 in concentration, reject a non-target odor, and detect the target in the presence of a stronger background odor.

odor\synapses	engineered	iteration 1	iteration 10
0.1c	1	4	3
0.3c	13	12	13
1.0c	16	13	15
3.0c	17	12	15
10.0c	12	7	8
30.0c	5	3	4
0.3b	0	0	0
1.0b	0	0	0
10.0b	0	0	0
30.0b	0	0	0
3.0c + 10.0b	6	4	7

Table 1. The number of action potentials generated by a γ-cell during a 0.5 sec. sniff of an odor for three sets of synapses (columns 2,3,4). The first column shows the intensity of the component odors

c and b present during the sniff. This γ-cell is expected to respond strongly only to odors containing c.

At each iteration, there is an overlap of about 60 between the 200 repertoire cells chosen by the learning procedure and the set of 200 original engineered repertoire cells. (The ~60 are a different set each iteration.) The others chosen are almost as good; the standard deviation of the widths of the distribution of the peak currents to the chosen cells at each generation is ~ $1.4\sigma_e$ and does not systematically change with the number of iterations. These facts explain why the performance of the iterated synapse renewal system is very similar to the performance of the engineered system.

One small technical difference will make the odor discrimination capabilities of the learned system slightly less than that of the engineered system. For the engineered system, the 200 repertoire cells were chosen from 200 different glomeruli, while in the learned system, the 200 were chosen as those which appeared 'best' by the learning algorithm *independent of the glomeruli from which they came*. **From random statistics, this will lead to about 126 different glomeruli being used, with the other 74 connections being the result of using more than one repertoire cell from a glomerulus. 200 is a large number for statistical purposes, and it therefore matters rather little that the selection of repertoire cells includes only ~126 different glomeruli.**

From the point of view of cell biology, a similar but slightly simpler alternative procedure could involve picking a threshold in M and using it to determine synapse change, rather than picking the ~$N_g/2$ best connections. For very large N_g, the two procedures should be equivalent. In simulations involving connections to ~200 repertoire cells, iteration with a fixed threshold leads to fluctuations in the number of action potentials generated and thus also to fluctuations in the number of connections made. In our numerical experiments, we found that after wandering in the 200-250 range for several iterations, the number of synapses then sometimes went abruptly to zero. For small networks, stability is best if these fluctuation effects are kept small, either by a mechanism which keeps the average activity of the cell stable to synapse number changes (Turrigiano *et al*.**, 1998), or by a mechanism which keeps the mean synapse number (or total synaptic strength) within a designated range.**

Such a synapse change protocol can easily be described in a cell biology context. Let each pre- post- synaptic action potential pairing generate or consume a quantity of 'alteration factor' according to the corresponding value of w_k. At each synapse, integrate this alteration factor over ~0.5 seconds. If integration over this time period is greater than a threshold value, then make a connection at that synapse (or keep, if it was already there). If the integration is less than the threshold value, then make no connection at that synapse (or eliminate that synapse, if it was previously present).

Single-trial supervised learning of an unfamiliar odor

In the previous section, we have shown that when odor b is presented, a γ-cell selective for odor b can identify a set of connections sufficient to recognize that odor, based on the spike timings of the repertoire cells and the spikes of the γ-cell itself. We now ask whether the same plasticity rules could be used to learn a *new* odor. During a sniff of a new odor, the pre-synaptic cells display the same action potential rasters whether or not they drive a post-synaptic cell. To use the plasticity rule, an appropriate spiking pattern for the post-synaptic γ-cell is required for learning-- in the previous section, this pattern was generated by a set of synaptic connections already present which implicitly 'knew the odor'. We note that the required pattern of γ-cell spikes is not specific to an odor. It is very similar in every γ-cell selective for odor x when a sniff of x is presented. Since this firing is stereotyped and almost periodic at the underlying 35 Hz frequency, it is possible to generate it without any initial connections from repertoire cells to the γ-cell. We will now show that this fact enables the system to achieve single trial learning of an unfamiliar odor, choosing its synapses on the basis of the protocol of the previous section.

Many protocols are capable of driving the γ-cell appropriately for single-trial learning. One of the simplest consists of an input current that comes from the same source as the periodic input current to the repertoire cells. If the periodic current into the repertoire cells is cos(ωt), then the current source for the γ-cell is B(t)*(1+cos(ωt)). Here B(t) describes the overlying intensity envelope that drives all the repertoire cells during a 'sniff', and is used to 'gate' a current that has the same periodic oscillation, at frequency ω, as that which drives all the repertoire cells. The basic 'sniff' had the form of half of a sine-wave lasting 0.5 seconds, and with intensities logarithmically transformed. B(t) was given the same form, but delayed by 70 ms, and represents a signal easily available to a functioning olfactory system. (This delay is not essential, but does improve the quality of single-trial learning).

The learning produced in such a fashion is termed 'supervised' because there is instruction to the system as to *when and where* to learn; some γ-cell(s), those which receive the B(t)*(1+cos(ωt)) current, are designated to learn, and some 'supervisory process' must be designed to do so. Only these selected cells are to receive this gating signal and thus be driven to spike. The choice of *when* to learn is derived from the sniff intensity itself.

Fig 3 compares the spike timings produced by a γ-cell driven by synaptic input from repertoire cells, and a γ-cell driven by the B(t)*(1+cos(ωt)) current. The timing differences between the γ-cell spikes due to synaptic drive and those due to the external gated drive are typically less than 1 ms. Thus spikes produced by the external gated drive can be used as a surrogate spike timing pattern for the learning rule derived in the previous section, and learning of new odors in a single trial may be attempted.

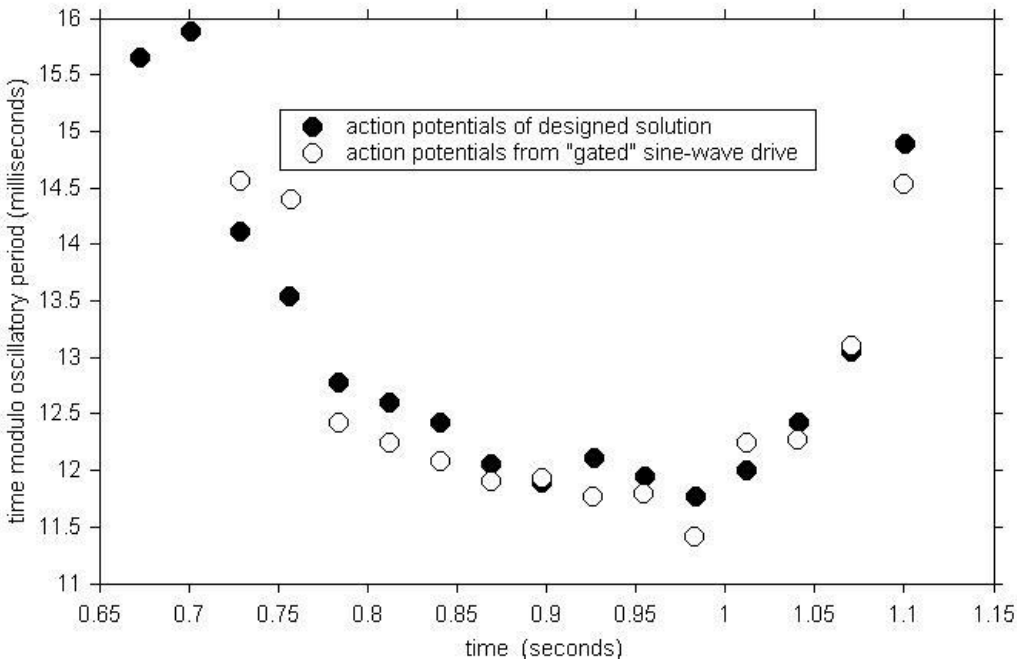

Figure 3. The timing of γ-cell spikes driven by a set of 'engineered' synapses and driven by the single-trial learning protocol. The vertical axis is the phase of spikes with respect to the underlying periodic current, expressed in terms of time {phase = time(modulo period)*2π} in milliseconds. The horizontal axis is real time. The actual sniff occurs between 0.6 and 1.1 sec.

Computer experiments show that this one-shot learning protocol generates a set of connections whose repertoire cell currents are statistically equivalent to the repertoire currents resulting from iterative maintaining the synapses in the manner earlier described. On iteration, it produces data indistinguishable from Table 1.

If the odor to which the system is exposed is not the 'standard intensity' for which the γ-cell spikes were engineered, then all the repertoire cell spikes will be shifted in time by approximately the same amount, earlier in time if the odor is weaker, later if stronger. Such a shift results in a different set of repertoire cells being chosen, but which equally well represents the odor. The shift merely changes the central operating point around which that odor will be recognized over a range of intensities.

Unsupervised learning is also possible without the 70 ms shift in B(t). Without it, the initial synapses are not as good, showing a larger standard deviation around the mean peak current. A second exposure to the same odor, using the self-renewal algorithm of the previous section, produces the usual sharp selectivity.

Unsupervised learning

This same timing rule can be used to implement *unsupervised* learning. In the paradigm of unsupervised learning, no γ-cell is 'instructed' when or what to learn.

In our implementation of this paradigm we start from a situation where there are many γ-cells which initially have random connections to repertoire cells, in such a fashion that the cells are 'broadly tuned' and will respond to many different odors. A synapse modification process is constantly present. Through that process, in an environment of many odors (present one at a time) these initial connections become refined through experience. We will show that eventually each γ-cell has a sharply tuned response, and different γ-cells have tuned themselves to different odors. What a particular γ-cell becomes tuned to in this paradigm depends on the time-history of sensory experience; the presentation of the same set of odors in a different order would produce a different result.

To create broadly tuned cells, we connected each γ-cell with strong synaptic connections to a random set of 5 repertoire cells. The strength of the strong connections was chosen so that the total synaptic strength into the cell was comparable to that which it would have received in the engineered solution; each of these strong synapses was 30 times as strong as the weak synapses in the engineered solution. Such a cell responds to about 1/500 of all random odors (of strength 1.5) by producing more than 14 spikes during the sniff. It produces a much lower rate of noise-induced spiking when no odor is present.

Statistics of the response of a set of 1000 broadly-tuned γ-cells to a sniff of a single odor is shown in Fig. 4. Less than 2% of these cells produced more than 10 spikes during the sniff, but one of them responded with 16 spikes. An essentially identical histogram is produced when a single broadly-tuned cell is exposed to 1000 random odors, because the statistics of the fit of 'random keys in a single random lock' is like the statistics of 'a single random key in many random locks'.

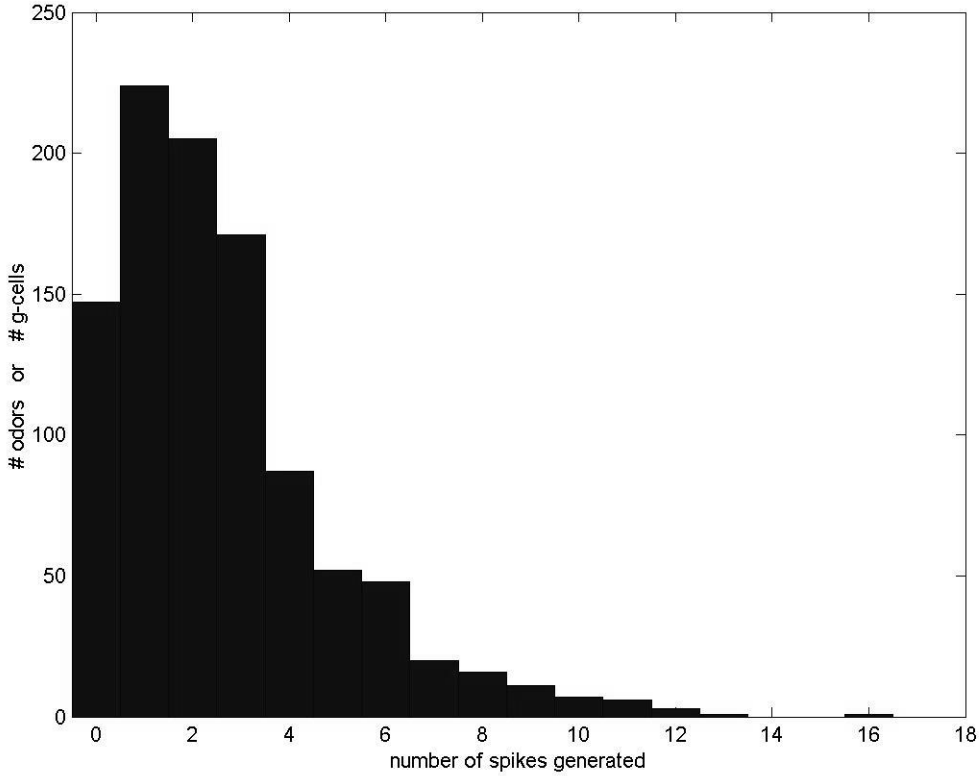

Figure 4. A histogram of the response of 1000 γ-cells, connected to repertoire cells by 5 strong random synapses, when stimulated by a single random odor. The binning on the x-axis is the number of spikes produced. This histogram is the same when many random odors are tested against a single random broadly-tuned γ-cell (y-axis, "# of odors"), or when many random broadly-tuned γ-cells are tested with a single odor (y-axis, "# of γ-cells").

Let the learning procedure used in the previous two sections be implemented when more than 14 spikes are produced by a γ-cell. When such a cell fires strongly to an odor, it does so because its connection pattern, by chance, closely resembles a sub-part of the correct (engineered with 400 glomeruli) connection pattern for that odor. Because it does so, the γ-cell spikes will be close in time to those of the engineered solution. These spikes can be used to implement learning exactly as they were in the case of self-renewal.

The most responsive of the cells produced 16 action potentials. The membrane potential and action potentials for this cell are shown in the lower panel of Fig. 5. The upper panel shows the membrane potential and action potentials produced by an 'engineered' γ-cell connected to 200 repertoire cells, designed to respond to that same odor. Most of the spikes correspond, with occasional missed pairings. The corresponding spikes have a mean shift between the top and bottom panels of 0.0031 sec. and a standard deviation around that shift of 0.0013 sec. The shift means that when the odor is learned, its optimal intensity will be slightly different from the

intensity at which the odor was presented. The small spread around that mean (small in on the scale of the shape of Fig. 2a) indicates that learning on the basis of the spikes of Fig 5 (lower panel) will work well.

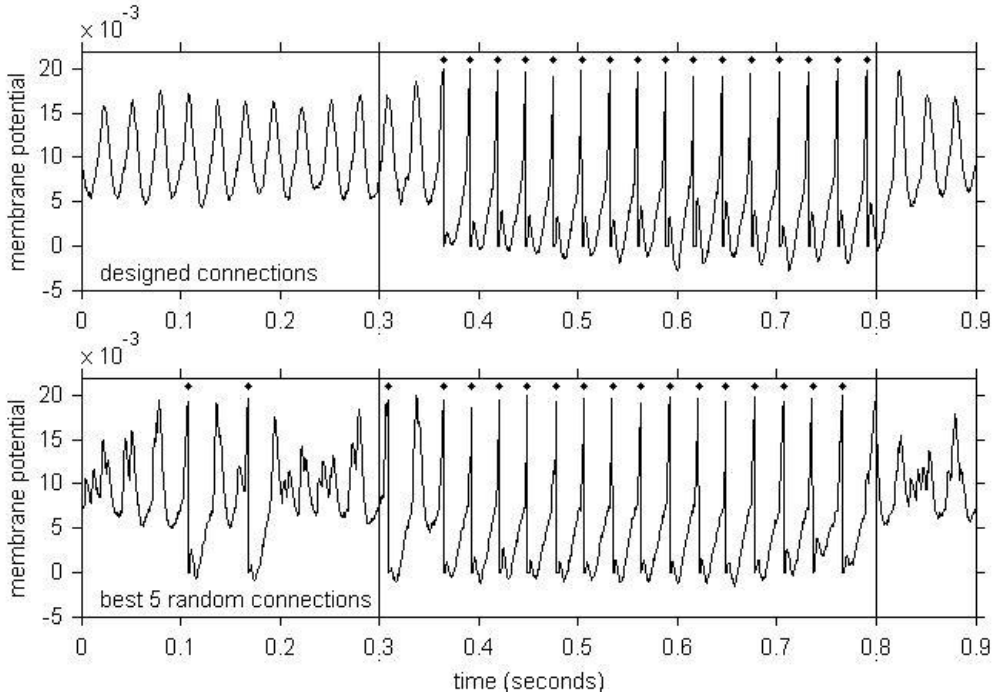

Figure 5. The cell potential and action potentials for an immature cell making 5 strong connections to repertoire cells (lower panel) and the response of an 'engineered' γ-cell (upper panel) making 200 connections for the same odor. The occasional random spiking when the odor is not present in the lower trace is a consequence of the sampling noise due to having 5 strong synaptic connections rather than 200 weaker ones.

The response spikes of the broadly cell are used to select a complete set of connections, converting this broadly tuned cell to a sharply-tuned cell selective for the particular odor which drove it strongly. They produced connections whose spread around ideal values had a standard deviation of 1.6 σ_e. This value is similar to that obtained during iterative synapse renewal, and represents sharp tuning around the selected odor. As expected, with these learned connections to a large number of different glomeruli, this γ-cell is now sharply tuned. It no longer responds to other random odors, *including those many others to which it responded strongly when it was a broadly tuned cell.* **The cell's selection of which odor it will become sharply tuned to is determined by which happens to first drive the cell strongly. The resulting selectivity is the same as that in Table 1. (A system that is to maintain an ability to learn new odors over time would also require a process by which a pool of broadly tuned cells is maintained.)**

The very rapid nature of this unsupervised learning is a result of an algorithm that remodels all synapses at once. Many applications would be better served by a

slower remodeling of the synapses. For example, connections might be made only to those synapses which surpass a fixed threshold (in M) for learning, rather than to all those necessary to generate the a given number of total connections, and eliminate only a fraction of existing synapses rather than all existing synapses in a learning event. A more gradual approach to selectivity will result
.

Conclusion

When synapses have plasticity, computational stability requires that *the normal activity within a network should produce synaptic changes that functionally compensate for spontaneous, noise-induced synaptic change. create new synapses with equivalent functionality to repair damage due to spontaneous loss of synapses***This basic and general principle can be used to derive synapse modification algorithms, based on the observed action potential patterns during normal network function. We have found an optimal synapse modification (learning) rule for a model network that used action potential timing as the basis of its basic 'many are nearly equal' implementation of pattern recognition. The learning rule was shown to lead to stable pattern recognition behavior when iteratively applied to the processing network. The empirically derived has a pre- post- synaptic timing dependence with strong resemblance to the timing rules seen in experimentally observed spike-timing dependent plasticity (e.g., Bi and Poo 1998).**

In addition to enabling functional stability, the timing rule can also be used to rapidly learn *thede novo* **connections suitable for a particular task, defined by the environment, in the context of both unsupervised and supervised learning. This fact, combined with the qualitative similarity between the learning rule derived here and the results of LTP/LTD experiments, demonstrates that a biological system can plausibly implement sophisticated spike-timing based computational algorithms.**

The 'many are almost equal' primitive was used for modeling a computational problem conceptually based on olfaction, using an underlying rhythm in place of feedback (or 'horizontal') connections between the repertoire cells. A time-warp invariant word recognition system can be made on a similar basis. Preliminary results show that in this case also, appropriate connections to a γ-cell can be learned.

Feedback pathways enrich computational dynamics, and enhance network computing ability. Feedback connections, which obviate the necessity of a background rhythms during computation by spontaneously creating an appropriate rhythm, were earlier used (Hopfield and Brody, 2000; 2001) to solve the time-warp speech problem by design. We do not yet have an understanding as to whether it is possible to learn appropriate feedback connections for that computation on the basis of sensory experience.

References

Bi, G.Q. and Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. J. Neurosci. *18*. **10464-10472.**

Brody, C. D. and Hopfield, J. J., previous paper in this journal

Froemke R. C. and Dan, Y. (2002). Spike-timing-dependent synaptic modification induced by natural spike trains. Nature *416*, **433-438.**

Hopfield, J. J. (1987). Learning algorithms and probability distributions in feed-forward and feed-back networks. Proc. Nat. Acad. Sci . *84*, **8429-8433.**

Hopfield, J. J. and Brody, C. D. (2000). What is a moment? "Cortical" sensory integration over a brief interval. Proc. Nat. Acad. Sci. (USA) *97*, **13919-13924.**

Hopfield, J. J. and Brody, C. D. (2001). What is a moment? Transient synchrony as a collective mechanism for spatiotemporal integration. Proc. Nat. Acad. Sci. (USA) *98*, **1282-1287.**

Liao, D. Z., Hessler, N. A. and Malinow, R. (1995). Activation of postsynaptically silent synapses during pairing-induced LPT in CA1 region of hippocampal slice. Nature *375*, **400-404.**

Turrigiano G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C. and Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. Nature *391*, **892-896.**