

## A Lower Bound on the Detectability of Nonassociative Learning in the Local Bending Reflex of the Medicinal Leech

S. R. LOCKERY AND T. J. SEJNOWSKI<sup>1</sup>

*Computational Neurobiology Laboratory, Salk Institute for Biological Studies and Howard Hughes Medical Institute, P.O. Box 85800, San Diego, California 92186-5800*

Studies of neural mechanisms of learning and memory have focused on large changes at identified synapses. However, memory in distributed processing reflexes could involve widely distributed engrams characterized by small changes at every synapse in the network. To investigate this possibility, we used a neural network optimization algorithm to construct distributed engrams for nonassociative conditioning in a model of the local bending reflex of the medicinal leech (*Hirudo medicinalis*). The model comprised 4 sensory neurons, 10 to 40 interneurons, 8 motor neurons, and up to 480 connections. Synaptic connections in the model were first optimized to reproduce the amplitude and time course of motor neuron synaptic potentials recorded during local bending. This network, which represented the naive state before conditioning, was then reoptimized to the habituated or sensitized state. Following reoptimization, the memory for nonassociative learning was encoded by small changes dispersed across the entire network, and each change made only a small contribution to the learning. Moreover, because the changes were small, resolution of a few tenths of a millivolt, or 3–5% of an average synaptic potential, would be needed to account for half of the nonassociative learning. These results show how difficult distributed engrams can be to detect and provide a likely lower bound on the detectability of nonassociative learning in this and related networks. © 1993 Academic Press, Inc.

A major objective in the cellular analysis of learning and memory is to identify the sites of synaptic plasticity underlying a conditioned change in behavior, a task that has been referred to as a search for the engram (Squire, 1987). Theoretically, the

engram may involve a variety of mechanisms including lasting changes in neuronal activity (Zipser, 1991), the excitability or transfer function of individual neurons (Servan-Schreiber, Printz, & Cohen), the time constants and other parameters governing temporal response properties (Lisberger & Sejnowski, 1992), and the strength of synaptic connections (Hebb, 1949). To date, studies of the synaptic bases of learning and memory have typically focused on large changes at identified synapses (Brown, Chapman, Kairiss, & Keenan, 1988; Byrne, Baxter, Buonomano, Cleary, Eskin, Goldsmith, McClendon, Nazif, Noel, & Scholz, 1991). Comparatively little is known about the different ways in which distributed processing networks can encode the memory for nonassociative learning (Frost, Clark, & Kandel, 1988). At one extreme, the engram could consist of a few large changes in synaptic strength; at the other extreme, it could involve small changes in every synapse. In either case, detection of the engram may be difficult because many synapses must be tested before even a partial account of the learning is achieved.

The local bending reflex of the leech is an advantageous model system in which to study how memories are encoded in distributed processing networks. In response to a moderate dorsal, ventral, or lateral mechanical stimulus, the leech withdraws from the site of contact by contracting longitudinal muscles beneath the stimulus and relaxing those on the opposite side of the body (Kristan, 1982; Lockery & Kristan, 1990a). The nervous system of the leech consists of a chain of segmental ganglia, each with circuitry sufficient for the reflex. Within the ganglion, major input to the reflex is provided by dorsal or ventral pressure sensitive mechanoreceptors, the P cells (Fig. 1A) (Nicholls & Baylor, 1968; Kristan, 1982). The longitudinal muscles are

<sup>1</sup> We thank W. B. Kristan, Jr. and P. Dayan for discussion and comments. Support was provided by The Bank of America–Giannini Foundation for Medical Research, a NIH Postdoctoral Fellowship, The Howard Hughes Medical Institute, and The Mathers Foundation. Address correspondence and reprint requests to Dr. S. R. Lockery, Institute of Neuroscience, Huestis Hall, University of Oregon, Eugene, OR 97403.

controlled by a total of eight types of motor neurons, an excitatory (DE or VE) and inhibitory (DI or VI) type for the dorsal and ventral quadrants on the left and right side of each body segment (Stuart, 1970; Ort, Kristan, & Stent, 1974). The flow of information from sensory neurons to motor neurons is mediated by a layer of interneurons, a subpopulation of which has been identified: those that receive excitatory input from the dorsal P cells and in turn excite dorsal excitatory motor neurons (Lockery & Kristan, 1990a). A detailed physiological analysis of interneuronal connectivity (Lockery & Kristan, 1990b) revealed that dorsal, ventral, or lateral stimuli will activate almost every interneuron, and the activity of each motor neuron is the result of input from many different interneurons. Thus, the reflex appears to be produced by a distributed processing mechanism. Modeling studies support this view by providing working examples of how local bending motor patterns can be produced by interneurons having multiple P cell inputs and motor neurons that receive inputs from many different interneurons (Lockery, Wittenberg, Kristan, & Cottrell, 1989; Lockery, Fang, & Sejnowski, 1990; Lockery & Sejnowski, 1992).

Because the local bending reflex is believed to involve tens of interneurons and hundreds of connections (Lockery & Kristan, 1990b; Lockery & Sejnowski, 1992), it is possible that the changes at individual synapses will be widely distributed and thus too small to detect in practical physiological experiments. Therefore, as a prelude to future physiological experiments, we made a theoretical investigation of this possibility. Specifically, we asked whether widely distributed engrams could exist for the local bending reflex, and if so, what are the possibilities for detecting them? Because we wished to find engrams with the smallest changes possible, the answer provides a likely lower bound on the detectability of learning in the local bending reflex. Furthermore, because of functional and mechanistic parallels between the local bending circuit and other systems, our results could also provide insights into the nature of widely distributed engrams in other systems.

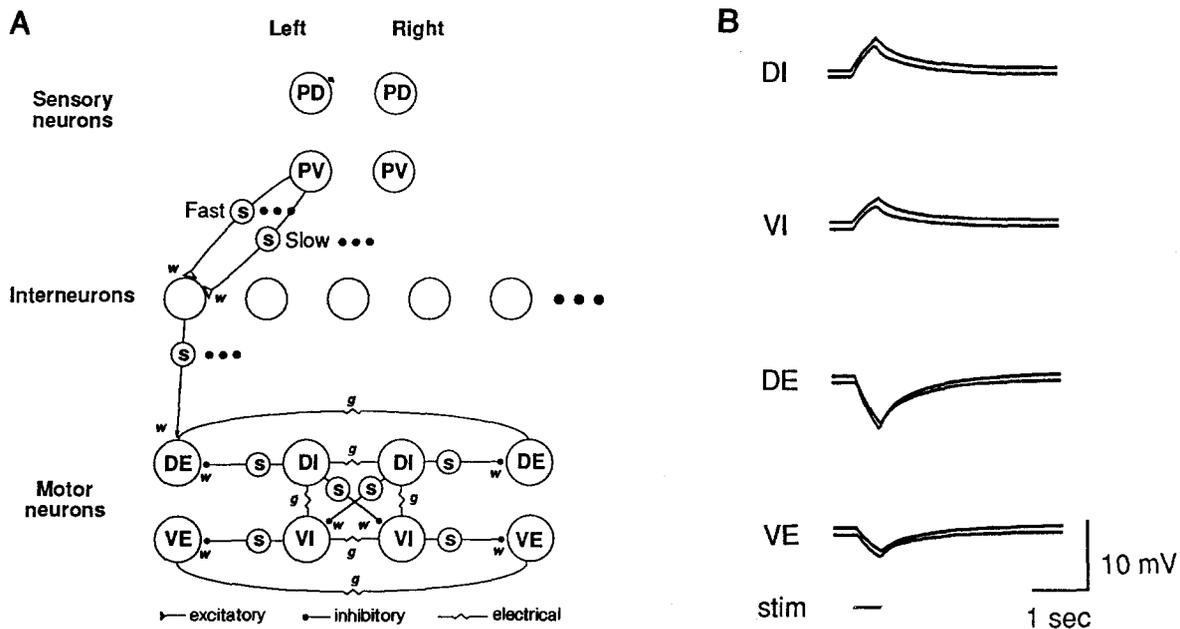
This paper focuses on memories for nonassociative learning—habituation and sensitization in particular. Nonassociative learning differs from associative learning in that statistical correlations between stimuli are not required. In habituation, perhaps the most elementary form of behavioral plasticity, repeated stimulation leads to a decrease in the amplitude or probability of a response (Thompson & Spencer, 1966). In sensitization, the response to a

test stimulus is increased by the delivery of a novel, strong, or noxious stimulus (Peeke & Petrinovich, 1984). Groves and Thompson (1970) identify a second form of sensitization produced by repeated presentation of the test stimulus itself. This form of sensitization has been termed warm-up (Hinde, 1970) to reflect the fact that it commonly reaches its maximum early in the series of repeated test stimuli.

The local bending reflex exhibits habituation and both forms of sensitization. With repeated P cell stimulation, the amplitude of local bending contractions first increases and then decreases, in some cases to below its initial value (Lockery & Kristan, 1991). The increase in response amplitude constitutes warm-up sensitization. The decrease in response amplitude following warm-up could be due to decay of sensitization, habituation, or both. That some preparations lack warm-up, yet exhibit the decline in response amplitude with repeated stimulation (Lockery & Kristan, 1991), supports the view that the decremental process is habituation. Thus, response amplitude in the local bending reflex is governed by simultaneously acting incremental and decremental processes and provides a clear example of the dual-process control of behavior first proposed by Groves and Thompson (1970). Sensitization of the response to P cell stimulation can also be produced by activation of a separate nociceptive pathway that responds to electric shock, mechanical damage, and heat (Lockery & Kristan, 1991).

The model of the local bending reflex employed in this paper was constructed using an optimization algorithm (Pearlmutter, 1989; Rumelhart, Hinton, & Williams, 1986) to adjust synaptic strengths until the model network reproduced the amplitude and time course of motor neuron synaptic potentials recorded in response to dorsal, ventral, and lateral mechanosensory input in naive animals. The model in this state represented the naive network before nonassociative conditioning and is referred to as the naive network (Lockery & Sejnowski, 1992). In the model network produced by this method, the distribution of synaptic potential amplitudes and the number of sensory inputs and motor outputs possessed by individual interneurons closely matched the biological data. Thus, the model provides an accurate representation of the biological network.

The same algorithm also provides a means of finding distributed engrams involving many small changes in synaptic strength. Because the optimization algorithm we chose is based on gradient descent (Press, Flannery, Teukolsky, & Vetterling, 1988), it makes small changes in every connection.



**FIG. 1.** (A) Network model of the local bending reflex. The network contained 4 sensory neurons (P cells), 40 interneurons, and 8 longitudinal muscle motor neurons. Each neuron was represented as a single electrical compartment with an input resistance and membrane time constant. Synapse units (s-units) were inserted between pairs of neurons connected by chemical synapses. The time constant of s-units represented the delays in chemical synaptic transmission. Input from sensory neurons to interneurons was mediated by pairs of s-units with fast and slow time constants. The connections from interneurons to motor neurons (and between motor neurons) were mediated by single s-units with intermediate time constants. Connection strengths ( $w$ ) from sensory neurons to interneurons and from interneurons to motor neurons were adjusted by an optimization procedure to reproduce the amplitude and time course of motor neuron synaptic potentials recorded in response to stimulation of sensory neurons. Motor neurons were also connected by numerous chemical and electrical ( $g$ ) synapses whose strength was determined physiologically. Abbreviations used: PD, P cell with dorsal field; PV, P cell with ventral field; DE, excitator of dorsal muscle; VE, excitator of ventral muscle; DI, inhibitor of dorsal muscle; VI, inhibitor of ventral muscle. (B) Simulated synaptic potentials in four motor neurons in the model network in response to stimulation of the contralateral ventral P cell after synaptic connections had been optimized. Each panel shows the response of a single motor neuron, together with the target response from the data set. For comparison of model and target responses, the latter have been shifted upward by 1 mV. A similar match between model and target synaptic potentials was achieved for the ipsilateral motor neurons and for all other motor neurons in all other patterns. Reproduced, by permission of the publisher, from Lockery and Sejnowski (1992).

This means it can find new sets of connection strengths in which the final connection strengths differ only slightly from the original ones. Taking advantage of this feature of the algorithm, we converted the naive network to the habituated state by reoptimizing its connections to reduce the amplitude of the motor neuron synaptic potentials. The sensitized state, representing the results of either form of sensitization, was produced analogously by reoptimizing the connections of the naive network to increase the motor neuron synaptic potentials. The changes in synaptic strengths following reoptimization to habituated or sensitized levels constituted a theoretical engram for habituation or sensitization. Analysis of the habituated and sensitized states identified salient features of the distributed engram and showed that such engrams can be difficult to detect. This provides a likely lower bound on the detectability of nonassociative learning in this and related networks.

## METHODS

### *Specification of the Model*

Construction of the model has been described in detail previously (Lockery & Sejnowski, 1992); the main points are repeated here.

**Circuit.** The model comprised 4 sensory neurons (P cells), 20 pairs of left-right symmetrical interneurons, and 8 motor neurons (Fig. 1A). The number of interneurons represented, in round numbers, an upper estimate of the total number of local bending interneurons in a single midbody ganglion. A previous study (Lockery & Kristan, 1990b) identified 8 paired and 1 unpaired interneuron contributing to dorsal local bends. In the model networks, the unpaired neuron, which is symmetrical about the midline, can be represented as an additional left-right pair, bringing the total to 9. We assumed that a search for ventral bending interneurons might

yield another 9 pairs, thus 18 pairs in total. Thus, the number was rounded to 20 for convenience. In several simulations, however, the number of interneurons was reduced to test specific hypotheses. Each of the 8 motor neurons in the model represented 1 of 8 types of longitudinal muscle motor neurons, there being 2–4 neurons per type (Ort et al., 1974; Stuart, 1970). These comprise the excitors and inhibitors of dorsal longitudinal muscle, DE and DI, respectively, and the excitors and inhibitors of ventral longitudinal muscle, VE and VI, respectively.

All known chemical and electrical synaptic connections between motor neurons were included and connection strengths (weights) were determined from previous physiological recordings (Granzow, Friesen, & Kristan, 1985). Weights of feedforward connections from sensory neurons to interneurons and from interneurons to motor neurons were adjusted using an iterative optimization procedure (Pearlmutter, 1989) so that the model reproduced the amplitude and time course of motor neuron synaptic potentials recorded in response to single and paired P cell stimulation in eight different patterns (Lockery & Kristan, 1990a). The optimization procedure was not allowed to insert connections between interneurons, since functional connections of this type have not been found; nor was the optimization procedure allowed to insert feedback connections from motor neurons to interneurons, since only one such connection has been identified to date (Friesen, 1989).

*Neurons.* Neurons were modeled as passive, single electrical compartments having in parallel an input resistance and capacitance to ground. Trains of action potentials in sensory neurons were modeled as stepwise increases in presynaptic voltage. Action potentials within trains were not represented individually, since the model is concerned with the rise, fall, and amplitude of the motor neuron synaptic potentials in response to trains of sensory cell impulses, not in the response to individual action potentials.

*Synapses.* Electrical synapses were modeled as ohmic conductances. The synapse between inhibitory and excitatory motor neurons was used as a model for all chemical synapses in the network, since this connection is the most thoroughly studied. Within the ganglion, chemical synaptic transmission between inhibitory and excitatory motor neurons (Granzow et al., 1985), as well as from interneurons to motor neurons and other interneurons in the leech (Angstadt & Calabrese, 1991; Friesen,

1985; Granzow et al., 1985), is a graded function of presynaptic voltage; it does not require action potentials, nor is it substantially affected by them.

In physiological experiments (Granzow et al., 1985) a stepwise increase in presynaptic voltage in DI produced a slowly rising postsynaptic potential in DE (Lockery & Sejnowski, 1992; their Fig. 2A). To account for the long synaptic rise time, synapse units (s-units, Fig. 1A) were inserted between pairs of neurons connected by chemical synapses. The activation level of an s-unit, representing percentage maximum activation of the synapse, varied between zero and one. Steady-state s-unit activation was a sigmoidal (S-shaped) function of presynaptic voltage. The shape of this function was determined using data from a previous study (Granzow et al., 1985). For each s-unit, a single time constant determined the rate at which steady state was approached. The range of this parameter was constrained by previous physiological recordings from pairs of motor neurons (Granzow et al., 1985). Thus, the s-units represented the temporal dynamics of all the steps in transmitter release and production of postsynaptic current.

Single s-units were inserted between interneurons and motor neurons, and between pairs of motor neurons. Preliminary simulations showed that the time course of motor neuron synaptic potentials, which exhibit a fast and slow exponential decay, could not be produced by a local bending model having only one s-unit with a single time constant at each connection. Therefore, two s-units, one with a fast and one with a slow time constant, were inserted between each sensory neuron and interneuron. This was consistent with intracellular recordings of synaptic potentials from interneurons in response to P cell stimulation, which showed fast, slow, and mixed fast and slow synaptic potentials (Lockery & Kristan, 1990b). Only single s-units were inserted between interneurons and motor neurons because multiple components in synaptic transmission at this synapse have not been observed (Lockery & Kristan, unpublished data).

#### *Fixed Parameters*

Neuronal input resistances and time constants were set to values consistent with physiological recordings from neuron somata. s-unit time constants were determined from previous pairwise recordings between pairs of neurons connected by chemical synapses (Granzow et al., 1985) or were fit by hand to reproduce the time course of motor neuron synaptic potentials in response to P cell stimulation

(Lockery & Sejnowski, 1992). Pairwise recordings from motor neurons (Granzow et al., 1985) were used to estimate the strength of the coupling conductance between electrically coupled motor neurons, and the strength of chemical synapses between motor neurons.

#### *Optimization of the Naive Model*

Many techniques are available for optimizing weights and other parameters in neural network models (Hinton, 1987). The simplest is numerical differentiation, in which a parameter is incremented by a small amount and the performance of the network is evaluated. If the performance of the network improves, the change is kept; if performance degrades, the parameter is decreased (Robinson & Arnold, 1990). Backpropagation (Rumelhart et al., 1986) relies on the same principle, but evaluates all parameters simultaneously. Backpropagation is thus simply a means of finding an optimum set of parameters, much like conventional curve fitting techniques. The objective of the optimization procedure was to choose a set of connection strengths to and from the interneurons such that the time course and amplitude of motor neuron synaptic potentials in the model matched those in physiological recordings from naive animals. To construct the set of input-output relations, or data set, for the optimization procedure, the amplitude of a representative motor neuron synaptic potential was digitized for use as a template. This template was scaled to the average peak amplitude for each motor neuron in response to each of eight patterns of single and paired P cell stimulation (Lockery & Kristan, 1990a; their Figs. 4 and 5). The eight patterns used were (1) left dorsal P cell, (2) left ventral P cell, (3) right ventral P cell, (4) right dorsal P cell, (5) both dorsal P cells, (6) both ventral P cells, (7) left dorsal and ventral P cells, and (8) right dorsal and ventral P cells. Patterns 7 and 8 are referred to as lateral stimulation. Initially, connections from sensory neurons to interneurons and interneurons to motor neurons were small and randomly assigned. In accordance with the recurrent backpropagation algorithm (Pearlmutter, 1989), connections in the network were optimized to reduce the total error ( $E$ ) defined as

$$E = \sum_p \sum_k \int_0^T 0.5[V_{pk}(t) - \bar{V}_{pk}(t)]^2 dt,$$

where  $p$  ranges over the input patterns in the data set,  $k$  ranges over the motor neurons in the model, and  $T$  is the duration (1 s) of the voltage trajectories

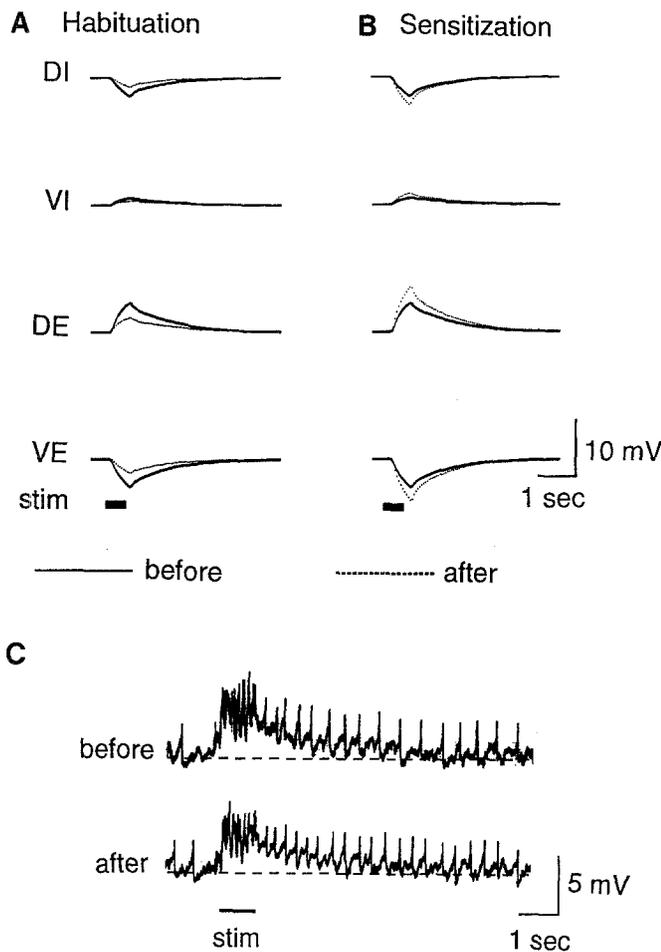
in the simulation. The quantity  $\bar{V}(t)$  is the desired voltage obtained by scaling the template for each motor neuron and each pattern of sensory input.

Connection strengths from sensory neurons to interneurons were constrained to be positive, since no inhibitory connections have been observed between such pairs. In addition, they were constrained to be greater than 1.35 mV (in response to the standard P cell voltage step described below) to ensure that sensory input was widely distributed across all the interneurons. Because most neurons in the leech occur in left-right pairs, each model interneuron on the left was constrained to be the mirror image of a homologue on the right. After optimization, the model interneurons matched their biological counterparts in the number of P cell inputs they received, and the number of motor neurons to which they projected. Thus, the degree to which the model and biological networks process information in a distributed fashion was the same.

#### *Reoptimization of the Naive Model to the Habituated or Sensitized State*

Distributed engrams for habituation and sensitization were found by reoptimizing the model to modified data sets. Extracellular recordings of P-cell-elicited action potentials in motor neurons revealed that firing frequency decreases during habituation training and increases during sensitization training (Lockery & Kristan, 1991). This strongly suggests that the underlying synaptic potentials in the motor neurons change in the same way. However, intracellular recordings of the actual change in the amplitude of each motor neuron synaptic potential in response to each pattern of P cell stimulation are not yet available. Instead, we reoptimized the model to produce synaptic potentials that were decreased in amplitude by 50% for habituation, and increased by 50% for sensitization (Fig. 2).

It is possible that future physiological experiments will show that each synaptic potential is not changed by the same percentage in habituation or sensitization. Perhaps the percentage of change is different for different motor neurons and different patterns of P cell stimulation. In light of this possibility, the value of 50% can be taken to represent the average change in motor neuron synaptic potential. The value of 50% was chosen because it is consistent with the change in amplitude of muscle-tension transients during habituation and sensitization in a semi-intact local bending preparation (Lockery & Kristan, 1991). Because there is a linear



**FIG. 2.** Reoptimization of motor neuron synaptic potentials. (A and B) Simulated synaptic potentials in four motor neurons in the model network in response to stimulation of the two dorsal P cells are shown before (solid line) and after (dotted line) reoptimization to the habituation (A) or sensitization (B) data set. The bars show the duration of the P cell stimulus. Similar results were obtained for all eight motor neurons in all eight patterns of P cell stimulation. (C) Intracellular recordings from a dorsal excitator motor neuron before and after habituation training. The peak response was reduced after training, but the time course was unaffected. The spikes in the records are action potentials conducted electronically from the cell's distal spike initiating zone. Action potentials were omitted from the model because synaptic transmission from interneurons and motor neurons is a graded function of presynaptic voltage and is not affected by the presence or absence of action potentials.

relationship between motor neuron firing rate and peak tension in the same preparation (Mason & Kristan, 1982), it is reasonable to suppose that the underlying synaptic potentials are also changed by a similar percentage.

We further assumed that the time course of synaptic potentials was not affected by training. For habituation, this assumption is based on preliminary recordings from motor neurons before and after

habituation training (Fig. 2C). Equivalent data for sensitization are not yet available. However, preliminary studies show that the time course of instantaneous firing frequency in the motor neurons is similar before and after sensitization (Lockery & Kristan, unpublished data).

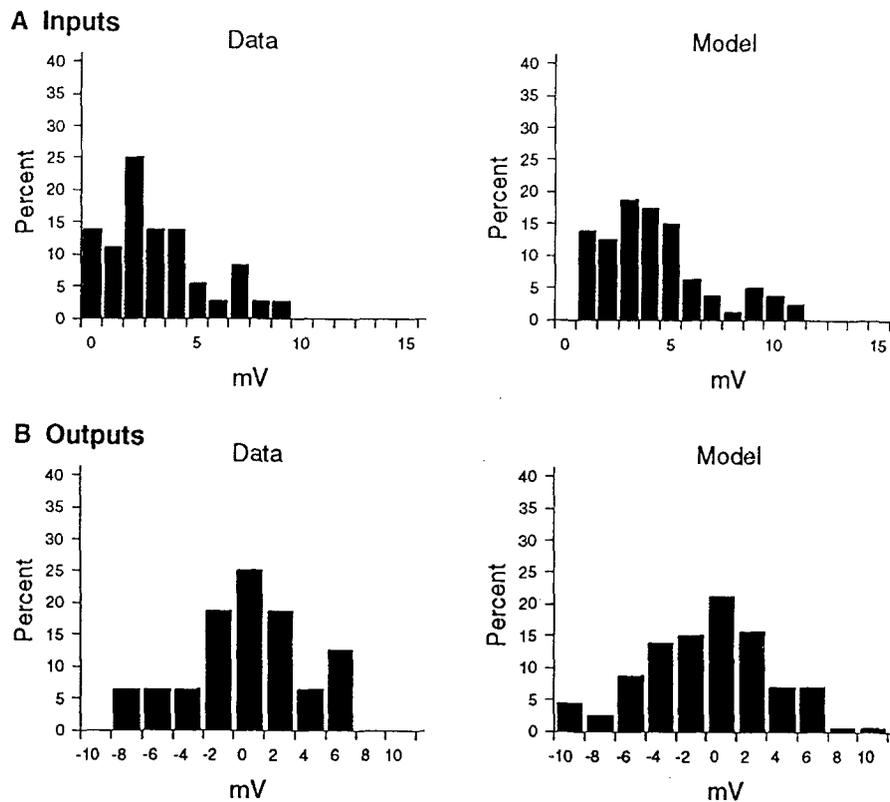
In this study we focused on engrams involving changes distributed widely among sensory neuron to interneuron and interneuron to motor neuron synapses. Because we were interested in a special class of engram in which the changes in synaptic strength are small, we used backpropagation again, this time for the reoptimization process. The choice of backpropagation was appropriate because it is a form of gradient descent which makes small changes in each connection at every iteration and thus favors an engram involving minimal changes in the weights. Although it is a powerful algorithm, backpropagation is not guaranteed to find a solution to every data set. It was therefore of interest to determine whether it could solve the reoptimization task. Reoptimization was allowed to proceed until the average deviation from the habituated or sensitized data set for each motor neuron at each time point was 5% or less. No constraints were placed on the amount each synapse could change, and a synapse could increase or decrease in strength. The absence of such constraints reflects the fact that the range of changes in the actual network is as yet unknown. Indeed, one of the main goals of this study was to evaluate the possibility that there are engrams in which the change might be too small to detect.

#### *Procedural Note*

The optimization procedure was used as a means of rapidly finding new patterns of connectivity within the local bending network, not as a model for the cellular events whereby synaptic strength is altered by learning in the biological system. Thus, this study emphasizes the end point of the reoptimization process, not its intermediate stages.

#### *Changes in Synaptic Strength*

To make possible the comparison of the size of changes in synaptic strength in the model with those in the biological network, care was taken when optimizing the naive network to ensure that the naive model and biological network had similar distributions of synaptic strength. Preliminary simulations showed that the correspondence between model and actual histograms of strengths of connections from P cells depended on the amplitude of



**FIG. 3.** Histograms of synaptic potential amplitudes in interneurons (A) and motor neurons (B) in response to a standard presynaptic stimulus (see Methods). Physiological histograms are shown on the left (data); histograms from a single model network are shown on the right (model). The physiological data are from Lockery and Kristan (1990a). The histograms showed that interneurons in the model functioned in the same operating range as in the biological network. Reproduced, by permission of the publisher, from Lockery and Sejnowski (1992).

the voltage step representing the P cell stimuli which served as input to the network. Step amplitude is a scale factor that was necessary to compensate for the absence of action potentials in the model. The amplitude of the P cell step therefore has no biological significance. Optimizing naive networks with different step amplitudes showed that a step of 10 mV produced the best matching histogram (Fig. 3A). After optimization of each naive network, we examined the distribution of interneuron to motor neuron connection strengths using various presynaptic voltage steps. The best correspondence for histograms of connection strengths to motor neurons (Fig. 3B) was achieved when model interneurons were depolarized with voltage steps of 100 mV. This level of depolarization agrees well with the depolarization produced by the current pulses used in Lockery and Kristan (1990b) to measure the actual connection strengths. Using the same voltage steps, we measured synaptic strengths in the reoptimized networks. Change in synaptic strength was measured by taking the difference between the synaptic strength measured before and

after optimization to the habituated or sensitized data set.

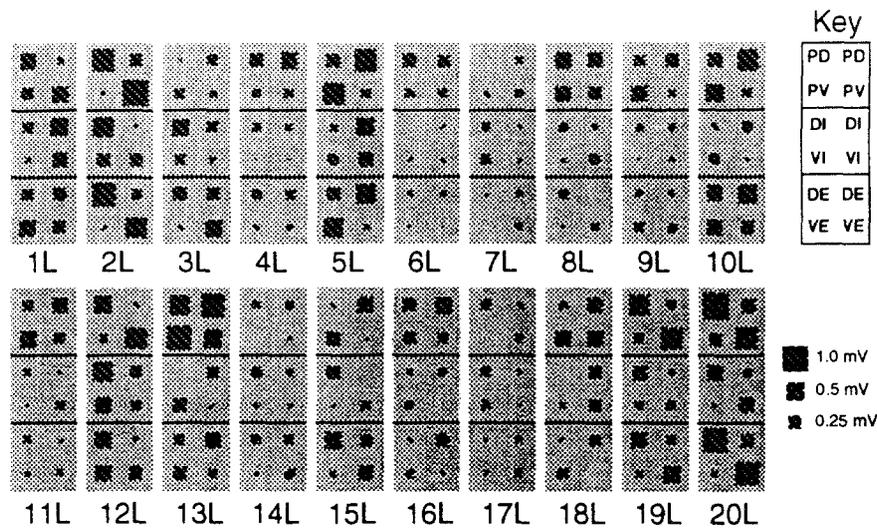
#### Replications

To ensure reliability of the results, six different naive networks were constructed by optimizing from six different sets of random initial connection strengths (Lockery & Sejnowski, 1992). Each of these networks was reoptimized for habituation and sensitization and the results averaged. Averages are presented as means  $\pm$  standard error of the mean.

## RESULTS

#### *Distribution of Changes in Synaptic Strength*

Reoptimizing the local bending network model to the habituated or sensitized data set produced widely distributed engrams in which almost all of the connections in the network were altered (Fig. 4). The figure shows the absolute value of changes in synaptic strength ( $|\text{after reoptimization} - \text{before reoptimization}|$ ) for half the interneurons after reop-



**FIG. 4.** The distributed engram for habituation. Each gray domino shows the absolute value of the changes in connection strength (squares) for the four input connections and eight output connections of a model interneuron. The area of each square is proportional to the absolute value of the change in synaptic strength in millivolts (see scale) at the indicated connection. Only the 20 interneurons on the left side of the network are shown (1L–20L); each of the 20 interneurons on the right side (1R–20R) exhibited mirror image changes to those of its homologue on the left. Similar results were obtained for sensitization. Presynaptic and postsynaptic neurons are indicated in the key, with abbreviations as described in the legend for Fig. 1A.

timizing to the habituated state of one of the six networks; changes in the connections of the other interneurons, which are the left–right homologues of those depicted in the figure, were the mirror image of those shown. Changes in synaptic strength produced by reoptimizing to the sensitized state were opposite in sign to the ones produced by reoptimizing to the habituated state. Histograms of the absolute value of the changes for input and output connections showed that the distributions were skewed toward small changes in synaptic strength (Fig. 5). Across all habituated networks, the average change in synaptic strength was  $0.22 \pm 0.0048$  mV (or  $13.2 \pm 0.979\%$ ); for sensitized networks it was  $0.20 \pm 0.0031$  mV or ( $12.2 \pm 0.710\%$ ). Thus, reoptimization produced engrams involving numerous small changes rather than a few large ones.

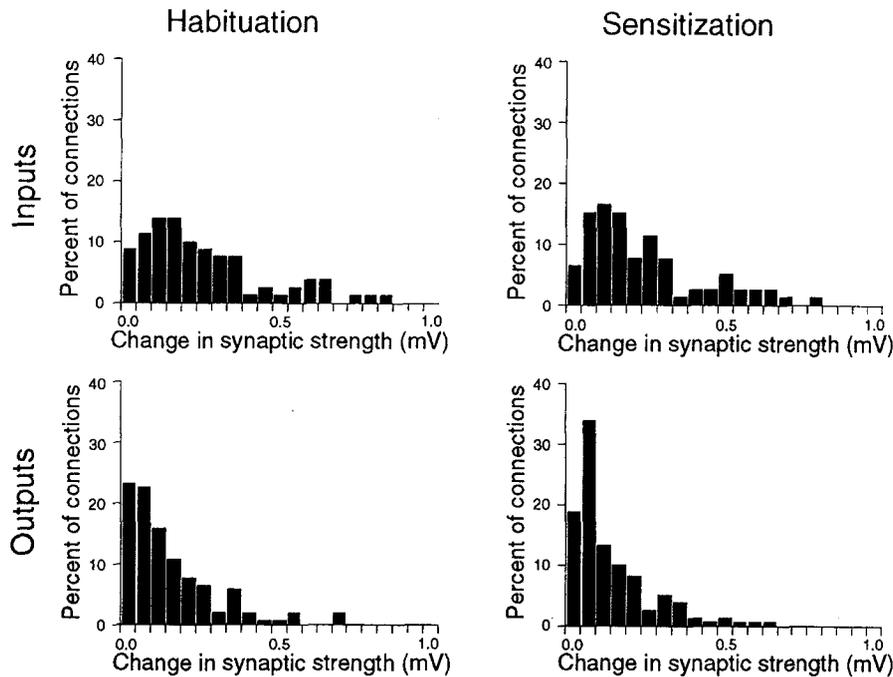
In an intuitively simple hypothesis of nonassociative learning, habituation would be the result of decreases in synaptic strength and sensitization would be the result of increases in synaptic strength. Whether the distributed engrams found by reoptimization followed this rule was determined graphically by making scatter plots of the change in synaptic strength (after reoptimization – before reoptimization) against the synaptic strength before optimization (Fig. 6). As the arrows show, EPSPs that increased in strength fall in the upper right quadrant while EPSPs that decreased in strength fall in the lower right quadrant. For IPSPs the locations are reversed, with increased strengths in the

lower left quadrant and decreased strengths in the upper left quadrant.

Inspection of the scatter plots revealed several trends that were consistent with the simple hypothesis of habituation and sensitization. First, the strong input connections tended to decrease in habituation and increase in sensitization. Thus, most of the strong input connections changed in a way that was consistent with the change in motor output. Second, most of the output connections decreased in synaptic strength in habituation and increased in sensitization. Nevertheless, counterintuitive changes also occurred in both input and output connections; habituation included many increases in the synaptic strength, and sensitization included many decreases in synaptic strength. Thus, in the distributed engrams found by reoptimization, habituation and sensitization were produced by the combined effect of intuitive and counterintuitive changes in synaptic strength.

#### *Effects of Changes in Synaptic Strength at Single Synapses*

Although every connection in the model changed, it was possible that some changes accounted for more of the learning than others. The effectiveness of each connection change was evaluated by determining how close the change at that connection came to producing the reoptimized motor output. A 64-dimensional coordinate system was defined in



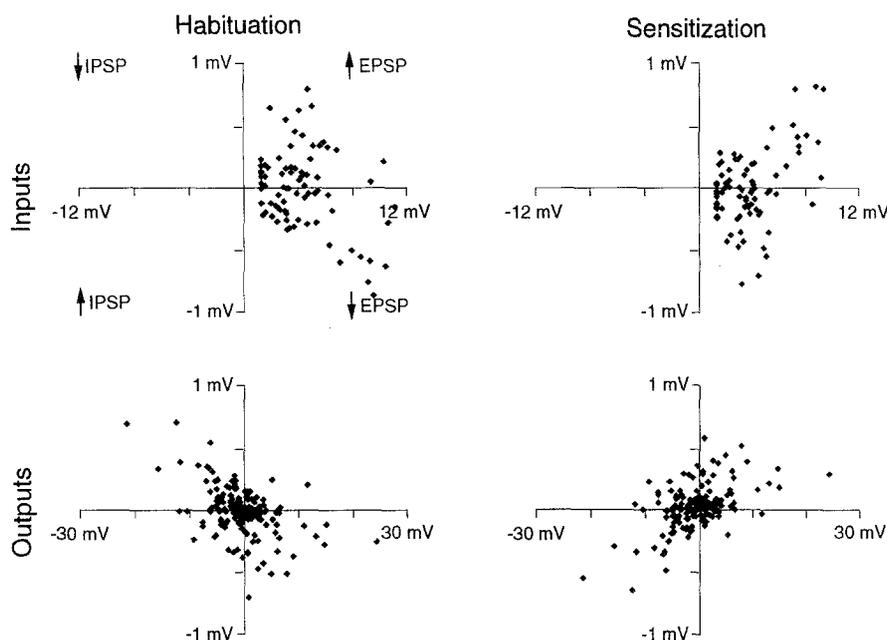
**FIG. 5.** Histograms of the absolute value of the changes in synaptic strength produced by reoptimization (same network as Fig. 4). The top two panels show histograms for connections from sensory neurons to interneurons; the bottom two panels show histograms for the connections from interneurons to motor neurons. Counts of the number of changes in each 0.05 mV bin from 0.0 to 1.0 mV are expressed as a percentage of the total number of connections of that type. The distributions were skewed toward small changes in synaptic strength.

which each axis represented the peak amplitude of the motor neuron synaptic potential for 1 of the 8 motor neurons in 1 of the 8 patterns of P cell stimulation. Naive and reoptimized networks were located at points in this space defined by their synaptic potential amplitudes. In this coordinate system, the change in motor output produced by learning is the vector extending from the point occupied by the naive network to the point occupied by the reoptimized network. The contribution of the change at any one connection to the overall change in motor output is the vector extending from the naive network to an otherwise naive network into which that change has been inserted.

To illustrate, we assume that there is one stimulus pattern and two motor neurons whose output is reduced by 50% in the habituated state (Fig. 7A). In this space, the vector  $w$  represents the change in motor output in the habituated state. The vector  $v$  represents the change in motor output attributable to the change in a single sensory to interneuron, or interneuron to motor neuron connection. To determine  $v$ , the connection of interest was inserted into the otherwise naive network. Each of the eight stimulus patterns was then presented and the peaks of the motor neuron synaptic potentials were de-

termined for each motor neuron and each pattern. The head of  $v$  lies at the point in the coordinate system defined by the peak motor neuron synaptic potentials observed with the change inserted. The effectiveness of the change in a given connection is measured by the relationship between  $w$  and  $v$ . In the example, the connection of interest is responsible for some of the change in motor output; consequently,  $v$  points in the direction of  $w$ , but does not overlie it. A connection whose change was detrimental to the change in output would point away from  $w$ . The effectiveness of each connection was quantified by finding  $p$ , the projection of  $v$  onto  $w$ . Thus, positive values of  $p$  indicate changes that are consistent with the change in motor output produced by learning; negative values of  $p$  indicate changes that are detrimental to the change in motor output. To allow comparison of different networks,  $p$  was normalized to the magnitude of  $w$  ( $\|w\|$ ). The normalized projection gives the fraction of learning accounted for by the associated change in synaptic strength.

In habituation of the six networks, the fraction of learning attributable to individual connections ranged from  $-0.007$  to  $0.081$  with a mean of  $0.0022 \pm 0.00012$ . For sensitization, the range was  $-0.005$



**FIG. 6.** Scatter plots of the changes in synaptic strength of input and output connections of the interneurons shown in Fig. 4. In each panel, the strength of a connection before reoptimization is plotted on the abscissa. The ordinate gives the change in that connection (after reoptimization - before reoptimization). Thus, connections whose strength increased (upward arrows) fall in the upper right and lower left quadrants, while connections whose strength decreased (downward arrows) fall in the lower right and upper left. Input connections were constrained to be positive in the model; hence, there are no points on the left in these two panels. Implementation of this constraint (Lockery & Sejnowski, 1992) caused several connections to have nearly the same value on the abscissa, resulting in small, vertically oriented clusters. These plots show that both habituation and sensitization were produced by the combined effect of increases and decreases in synaptic strength.

to 0.073 with a mean of  $0.0020 \pm 0.00010$ . Thus, no single change accounted for more than 8.1% of the change in motor output and the mean was about 0.2%. That the range of normalized projections includes negative values shows that some changes were detrimental to the change in motor output. However, plotting the distribution of projections for the habituated and sensitized states of individual networks (Fig. 7B) showed that the distributions were strongly biased toward small, positive values. This result indicates that despite the presence of counterintuitive changes in the connections, almost every change contributed positively to the change in motor output.

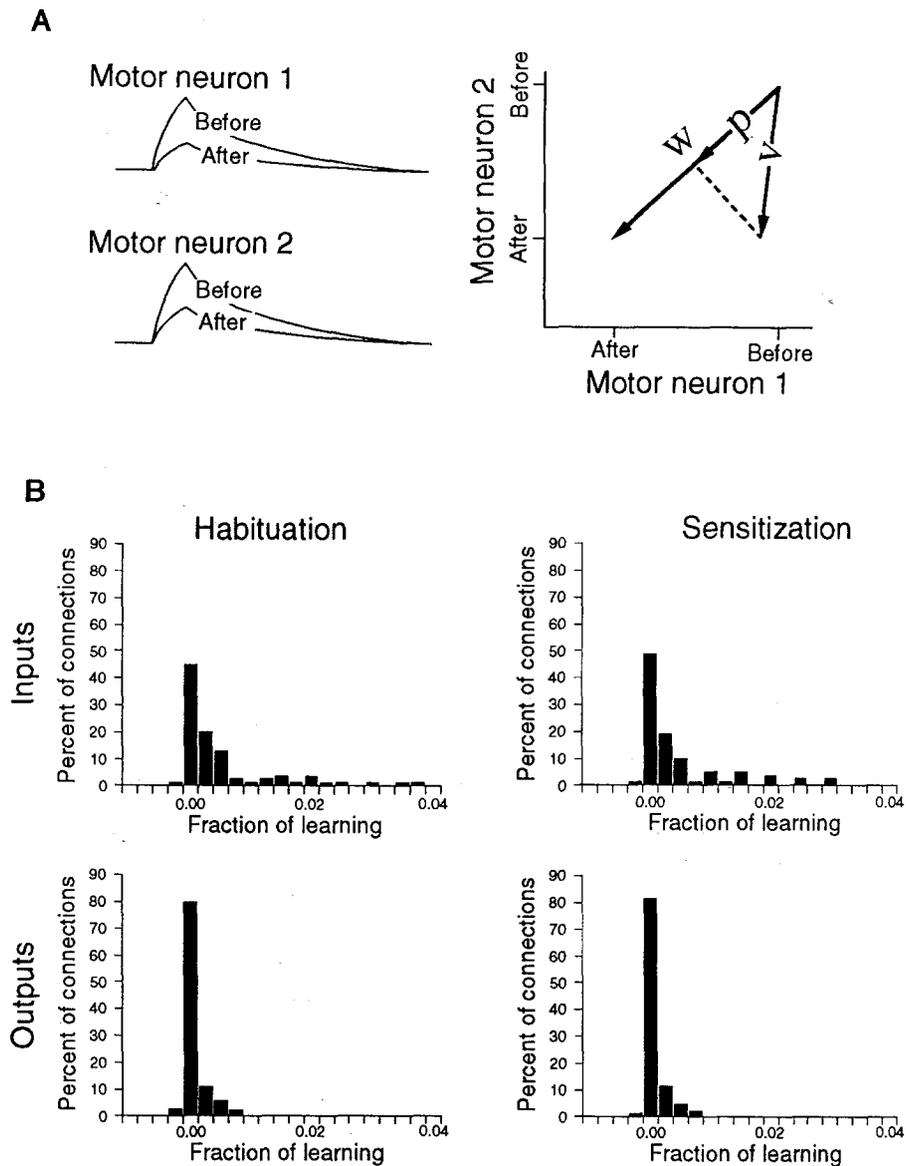
#### *Detecting a Widely Distributed Engram*

Two aspects of the distributed engram complicate the physiological analysis of learning. First, the large number of changes such an engram entails could require testing many or all of the connections to give a complete account of learning; we refer to this as the sampling problem. Second, the individual changes are small and could therefore be hard to

detect; we refer to this as the detection problem. The sampling and detection problems were treated separately in the analysis of the reoptimization engrams.

The sampling problem was investigated by assuming the worst case in which the experimenter has no prior knowledge of which connections are most likely to change, and thus must choose randomly among the 480 connections in the network. Under this assumption, we asked how many connection strength changes would have to be tested to account for half of the habituation or sensitization. The fraction of learning accounted for by a set of changes in connection strength was defined in terms of the method described above for the fraction of learning accounted for by a single change; the vector  $\mathbf{v}$ , representing the effect of all the detectable connection changes, was projected onto the vector  $\mathbf{w}$  representing the total change in output, and the projection  $p$  was calculated and normalized as before. To account for half of the learning, a set of changes in connection strength must have an associated normalized projection of 0.5.

The analysis was carried out by inserting the con-

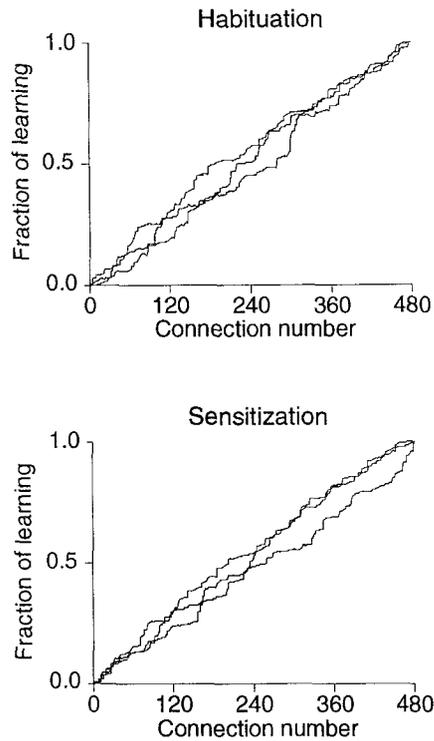


**FIG. 7.** Analysis of the fraction of learning accounted for by individual changes in connection strength. (A) Two-dimensional example of the coordinate system in which each axis represents the peak motor neuron synaptic potential in a network with two motor neurons in response to one pattern of sensory input. In the actual case, there were eight motor neurons and eight patterns of sensory input, thus a 64-dimensional space. The vector  $w$  represents the change in motor output produced by reoptimization. The vector  $v$  represents the change in output attributable to the change in synaptic strength at a single connection. The fraction of learning is found by projecting  $v$  onto  $w$  to obtain the projection  $p$ , which is normalized to the magnitude of  $w$ . (B) Histograms of the fraction of habituation and sensitization accounted for by individual changes in synaptic strength in input and output connections of the interneurons shown in Fig. 4. A single change that accounted for all the learning would have a normalized projection equal to 1. Bin width is 0.0025.

nection strength change at a randomly chosen connection from a habituated or sensitized network into the same network in its naive state, and the fraction of the learning that connection accounted for was calculated. Leaving this change in the network, a second connection strength change was chosen and the fraction of learning accounted for by the two changes together was calculated. This continued un-

til all 480 connections had been chosen. The entire process was repeated 20 times on each of the 6 model networks reoptimized for habituation and sensitization.

Plotting the fraction of learning against the number of changes that had been inserted into the naive network showed that there was, on average, a linear relationship between the number of connections in-



**FIG. 8.** The relation between the number of connections tested and the amount of learning accounted for by the tested connections in the network shown in Fig. 4. Beginning at the origin, a connection strength change was chosen at random from the habituated (above) or sensitized network (below) and was inserted into the same network in its otherwise naive state. The fraction of learning accounted for by this change was determined as in Fig. 7 and another randomly chosen change was inserted. This process was repeated until all changes had been selected, plotting after each insertion the fraction of learning accounted for by the accumulated changes against the number of changed connections. This process was repeated 20 times, each repetition producing a unique curve. For clarity, curves are shown only for the repetitions having the smallest, largest, and the average number of insertions to the point at which half of the learning is accounted for. On average, in this and the other five networks, accounting for half of the learning required sampling half of the connections, but random differences in sampling order produced large variations in the number of connections that had to be sampled to reach this point.

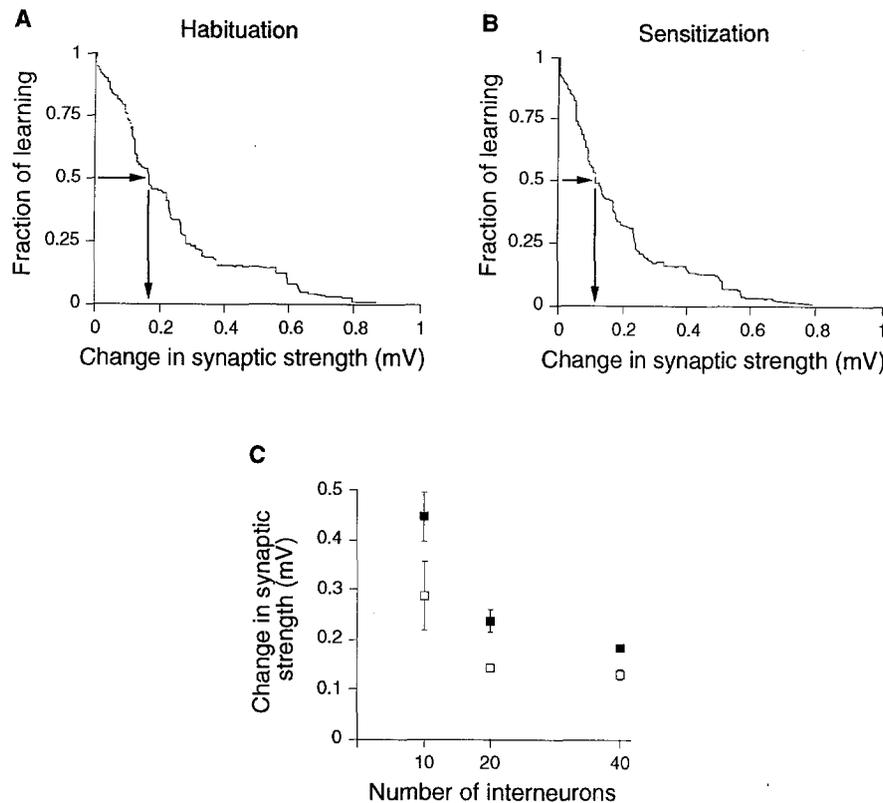
serted into the naive network and the fraction of learning they accounted for (Fig. 8). This would be expected if there is a linear interaction of changes at individual connections, a point that was confirmed by determining whether the sum of all individual projections is equal to  $\|\mathbf{w}\|$ . In the six model networks studied, the average percentage difference between the sum of projections and  $\|\mathbf{w}\|$  was  $7.7 \pm 0.4\%$ ; for sensitization it was  $-4.2 \pm 0.4\%$ . This indicates that the changes in the distributed engram interact approximately linearly. Thus, on the average, to account for half the learning, it would be

necessary to measure the change at approximately half of the connections in the network. However, because of differences in the fraction of learning produced by the change at each connection (Fig. 7B) and the random order in which they were inserted into the naive network, each of the 20 repetitions followed a unique path from 0 to 1 on the ordinate. Thus, the number of changes required to account for half of the learning varied widely across the 20 repetitions and 6 networks. For habituation, this number ranged from 141 to 332; for sensitization, the range was 180 to 319. Thus, in the distributed engrams, the order in which connections are sampled can have a significant effect on the number required to account for half of the learning; in some cases more than two-thirds of the connections in the network must be tested.

The sampling problem might be reduced if the sites of the most effective changes could be anticipated before learning. For example, an experimenter with prior knowledge of the connection strengths in the naive network might suppose that the most effective changes would be those of the strongest connections before conditioning. However, in the distributed engrams produced by reoptimization, there was little or no correlation between the pre-learning amplitude of a connection and the fraction of learning accounted for by that connection. For the six habituated networks, the average coefficient of determination ( $r^2$ ) for input connections was  $0.0058 \pm 0.034$ ; for output connections it was  $0.0050 \pm 0.0028$ . For the six sensitized networks the average  $r^2$  was  $0.0050 \pm 0.0028$ ; for output connections it was  $0.0055 \pm 0.0029$ . Thus, the initial strength of a connection does not necessarily provide a significant clue as to which connections are most profitable to test.

The detection problem was investigated by asking how sensitive the physiological test for changes in synaptic strength would have to be to detect half the learning. Specifically, we supposed the experimenter could find every change greater than or equal to an assumed minimum detectable change, and asked what minimum would yield a set of changes which together account for half of the change in motor output.

The analysis was accomplished by setting the assumed minimum detectable change to successively lower values, beginning with the largest change in connection strength found in the network under consideration. Thus, the largest connection change in the habituated or sensitized network was inserted into the same network in its otherwise naive state, and the fraction of learning accounted for by that



**FIG. 9.** Detectability of habituation and sensitization for the network shown in Fig. 4. (A and B) Graphical determination of the minimum detectable change required to account for half of habituation and sensitization. Beginning on the right, the largest connection strength change was inserted into the otherwise naive network. The fraction of learning accounted for by this change was determined as in Fig. 7 and the second biggest change was inserted. This process was repeated until all changes had been inserted, plotting after each insertion the fraction of learning accounted for by the accumulated changes against the absolute value of the last change inserted. For the examples shown, accounting for half of the learning required detection of changes as small as 0.17 mV for habituation and 0.12 mV for sensitization (arrows). (C) The effect of the number of interneurons in the model network on the minimum detectable change required to account for habituation (solid squares) and sensitization (open squares). There was an inverse relation between the required level of detection and the number of interneurons.

change was determined by calculating the projection  $p$ . Leaving this change in the network, the assumed minimum detectable change was set to the next to largest change in connection strength; this change was then inserted and  $p$  was recalculated. This process was repeated until all changes were inserted.

Plotting the fraction of learning against the minimum detectable change (Figs. 9A and 9B) provided a graphical means of determining how much of the learning could be accounted for at a given level of detectability. Each curve begins at the right with the insertion of the largest change. For completeness, the analysis continued past the point at which half of the learning was accounted for to the point at which all of the learning was accounted for. Figure 9A shows that accounting for half of habituation in the network shown in Fig. 4 required detection of all the changes down to a minimum detectable change of 0.17 mV; accounting for half of sensiti-

zation required detection of all the changes down to 0.12 mV (Fig. 9B). Similar results were found for all six networks. For habituation, the average minimum detectable change required to detect half of the learning was  $0.183 \pm 0.007$  mV; for sensitization it was  $0.129 \pm 0.010$  mV. The values are small with respect to the average synaptic potentials in the model in response to standard presynaptic stimuli (see Methods) which were 3.9 mV for the connections from sensory neurons to interneurons and 4.0 mV for connections from interneurons to motor neurons. Thus, accounting for half the learning in the model requires the ability to detect changes on the order of 3–5% of the amplitude of the average synaptic potential.

We expected that smaller networks, which have fewer connections, would on average require larger changes at each connection to produce the same amount of habituation or sensitization, leading to

a higher minimum detectable change required to account for half of the learning. The assumption of 40 interneurons in the model was based on an upper estimate of the actual number of local bending interneurons. However, it is conceivable that the biological network contains a smaller number of interneurons, with a lower limit of about 10 pairs, based on the prior identification of 8 pairs and 1 unpaired interneuron per ganglion (Lockery & Kristan, 1990b). We therefore wished to determine how quickly the minimum detectable change rises as the number of interneurons is reduced.

The dependence of the minimum detectable change on the number of interneurons in the model was tested by optimizing networks ( $n = 6$ ) with fewer than 40 interneurons to produce normal local bending motor output. Each of these networks was then reoptimized to the habituated and sensitized states, and the average minimum detectable change required to account for half of the learning was evaluated as above. For the 20-interneuron networks, the average minimum detectable change for habituation increased to  $0.238 \pm 0.023$  mV and for sensitization to  $0.144 \pm 0.006$  mV (Fig. 9C). For the 10-interneuron networks, the average minimum detectable change for habituation increased to  $0.446 \pm 0.050$  mV and for sensitization to  $0.287 \pm 0.068$  mV. Thus, the ability to detect small changes in synaptic strength is required to account for half of the learning even when the number of interneurons is reduced well beyond its lower limit.

## DISCUSSION

### *Optimization in Models of Learning*

In a distributed processing network like the local bending reflex, memories might be encoded in distributed engrams involving small changes at every synapse in the network. If so, an experimental search for such an engram might be a difficult undertaking. As an aid to interpretation of future electrophysiological studies of learning and memory in this system, we asked whether such engrams are possible in this network.

We began with the reasonable assumption that nonassociative conditioning will produce, on average, a 50% change in the amplitude of the synaptic potentials of the motor neurons controlling the behavior. We then used a computer algorithm based on gradient descent to search for a set of minimal changes in synaptic strength sufficient to produce the desired change in motor output. This approach

was valid for our purpose, which was to explore the worst case in which all the changes were as small as possible, because we were not trying to find the actual engram. Had we been interested in predicting the actual engram, a better approach would have been to model the actual mechanisms of synaptic plasticity operating at the level of individual synapses. Such an effort is premature in this system, however, because the actual mechanisms of plasticity are not known.

The optimization algorithm is not intended here as a model for the cellular or network level mechanisms of synaptic plasticity. The interest is strictly in possible engrams reached at the endpoint of optimization. Therefore, the question of the biological plausibility of optimization algorithms is as irrelevant as it is in the case of curve fitting techniques such as linear regression, which are widely used in biology and the behavioral sciences. It is worth emphasizing this because the historical tendency to speak of backpropagation in terms of learning algorithms and training has led some to discount the models that happen to use backpropagation (rather than data or trial and error) to select connection strengths or the changes in them. The alternative, to adjust all of the connection strengths by trial and error, is simply not feasible.

### *The Detectability of Nonassociative Learning in the Leech*

Our results set a likely lower bound on the detectability of nonassociative learning in the local bending reflex. Gradient descent favors small changes in connection strength; thus, the engrams provide a lower bound on the detectability of learning. However, we cannot rigorously exclude the possibility of engrams involving even smaller changes. This is because gradient descent algorithms do not necessarily converge to the globally optimal solution. Nevertheless, if the actual engram were similar to the one found by the algorithm, resolution at the level of a few tenths of a millivolt would be required to account for only half of the learning. The limit of resolution that can be achieved in physiological experiments depends on such quantities as the signal-to-noise ratio of the intracellular recordings, the natural variance of synaptic potential amplitudes, and the number of test trials that can be obtained on an individual connection before and after learning. Quantal analysis of the synapse between an inhibitory interneuron and a heart motor neuron in the leech (Nicholls & Wallace, 1978) re-

solved differences as small as 0.25 mV. If differences of this size can be detected in synapses in the local bending reflex, then it would be possible to account for approximately 40% of learning encoded by the minimal engrams presented here (Figs. 9A and 9B), provided that all the connections in the biological network can be tested.

### *Properties of the Minimal Engram*

The model further demonstrates that an engram may involve many counterintuitive changes in synaptic strength. In a simple hypothesis of nonassociative learning, habituation is the result of decreases in synaptic strength and sensitization is the result of increases in synaptic strength. For example, in a linear network, a 50% habituation of the response of all motor neurons to all patterns of sensory input could be produced by decreasing the strength of the input and output connections of all interneurons by half; a 50% sensitization could be produced by a 50% increase in these connections.

The changes in connection strength produced by the reoptimization algorithm, however, differed from this simple model in several respects. First, the changes were most of smaller than 50% (Fig. 6). Second, a substantial number of synaptic strengths increased in habituation and decreased in sensitization. It is conceivable that these counterintuitive changes would be detrimental. In contrast, however, most of the counterintuitive changes in fact contributed positively to the change in output, since a much larger proportion of the changes occupied counterintuitive regions of the scatter plots (Fig. 6) than had negative projection values (Fig. 7). A counterintuitive change could contribute positively to learning if there were compensating effects downstream. For example, an increase in an EPSP from a sensory neuron to an inhibitory interneuron could contribute to habituation of a motor neuron EPSP. However, this cannot be the case in a distributed processing system such as the local bending network because any connection from an interneuron to a motor neuron contributes to EPSPs in response to some patterns of sensory input and IPSPs in response to other patterns of sensory input. This means that a change that ultimately habituates the EPSPs seen in a motor neuron in response to some patterns of sensory input will sensitize IPSPs in same motor neuron seen in response to other patterns of sensory input. Thus, in a distributed processing network, it is the global effect of a change, viewed across all patterns of sensory input, that determines whether it contributes positively to

the learning. This emphasizes the importance of taking the entire input–output function of a reflex into consideration when interpreting the changes in synaptic strength found in the cellular analysis of learning and memory.

The model further demonstrates that an engram can be undetectable in practical terms, even for modest (50%) changes in output in small invertebrate networks. In both habituation and sensitization, the algorithm established engrams consisting of small changes in synaptic strength at almost every connection in the network. Such engrams present several potential problems for the physiological analysis of the learning. First, an account of half of the learning requires measurement of the change at about half (240), and sometimes as many as about two-thirds (319–331) of the connections in the network. Second, even such a partial account of the learning requires detection of changes on the order of 3–5% of the average synaptic potential in the system, whether the model utilizes an upper (20 pairs) or lower (10 pairs) estimate of the number of interneurons in the biological network. Moreover, the lack of a correlation between the strength of a connection in the naive network and the fraction of learning accounted for by the change at that connection shows that focusing on the strongest connections is not necessarily an effective strategy for overcoming the sampling problem in the case of a widely distributed engram.

### *Relevance to Other Systems*

The similarity between the local bending and other networks suggests that it may not be alone in having the capacity to encode memories in minimal engrams. The primary function of interneurons in the local bending reflex is to associate with each sensory stimulus that pattern of motor neuron excitation and inhibition that is required to withdraw from the stimulus (Lockery & Kristan, 1990a). In functional terms, the reflex transforms a four-dimensional input vector encoding stimulus location into an eight-dimensional output vector encoding the associated movement. Moreover, this transformation is achieved by a distributed processing mechanism in a feedforward network (Lockery & Kristan, 1990b). Thus, both functionally and mechanistically, the local bending reflex shares essential features of sensorimotor integration (Anastasio & Robinson, 1990) or spatial transformation (Zipser & Andersen, 1988) in vertebrate systems. Because vertebrate networks often contain many more neurons and connections than the local bending net-

work, it is conceivable that theoretically minimal engrams in these instances are even more widely distributed.

The strategy we have taken in this study might also be useful in gaining intuition about engrams in other distributed systems. For example, the engrams for facts and events in declarative memory tasks may be widely distributed in hippocampal and cortical networks (Rolls, 1989). A worst-case analysis along the lines of the present study could provide an estimate for how difficult it might be to experimentally determine engrams in vertebrate brains.

### REFERENCES

- Anastasio, T. J., & Robinson, D. A. (1990). Distributed parallel processing in the vertical vestibuloocular reflex: Learning networks compared to tensor theory. *Biological Cybernetics*, *63*, 161–167.
- Angstadt, J. D., & Calabrese, R. L. (1991). Calcium currents and graded synaptic transmission between heart interneurons of the leech. *Journal of Neuroscience*, *40*, 185–234.
- Brown, T. H., Chapman, P. F., Kairiss, E. W., & Keenan, C. L. (1988). Long-term synaptic potentiation. *Science*, *242*, 724–728.
- Byrne, J. H., Baxter, D. A., Buonomano, D. B., Cleary, L. J., Eskin, A., Goldsmith, J. R., McClendon, E., Nazif, F. A., Noel, F., & Scholz, K. P. (1991). Neural and molecular bases of nonassociative and associative learning in *Aplysia*. *Annals of the New York Academy of Sciences*, *627*, 124–149.
- Friesen, W. O. (1985). Neuronal control of leech swimming movements: Interactions between cell 60 and previously described oscillatory neurons. *Journal of Comparative Physiology A*, *156*, 231–242.
- Friesen, W. O. (1989). Neuronal control of leech swimming movements II. Motor neuron feedback to oscillator cells 115 and 28. *Journal of Comparative Physiology Series A*, *166*, 205–215.
- Frost, W. N., Clark, G. A., & Kandel, E. R. (1988). Parallel processing of short-term memory for sensitization in *Aplysia*. *Journal of Neurobiology*, *19*, 297–334.
- Granzow, B., Friesen, W. O., & Kristan, W. B., Jr. (1985). Physiological and morphological analysis of synaptic transmission between leech motor neurons. *Journal of Neuroscience*, *5*, 2035–2050.
- Groves, P. M., & Thompson, R. F. (1970). Habituation: A dual-process theory. *Psychological Review*, *77*, 419–450.
- Hebb, D. O. (1949). *Organization of behavior*. New York: Wiley.
- Hinde, R. A. (1970). Behavioral habituation. In G. Horn & R. A. Hinde, (Eds.) *Short-term changes in neural activity and behavior* (pp. 3–40). London: Cambridge University Press.
- Hinton, G. E. (1987). Connectionist learning procedures. *Artificial Intelligence*, *40*, 185–234.
- Jacobs, R. A. (1988). Increased rates of convergence through learning rate adaptation. *Neural Networks*, *1*, 295–307.
- Kristan, W. B., Jr. (1982). Sensory and motor neurons responsible for the local bending response in leeches. *Journal of Experimental Biology*, *96*, 161–180.
- Lisberger, S. G., & Sejnowski, T. J. (1992). Motor learning in a recurrent network model based on the vestibuloocular reflex. *Nature*, *360*, 159–161.
- Lockery, S. R., Fang, Y., & Sejnowski, T. J. (1990). A dynamic neural network model of sensorimotor transformations in the leech. *Neural Computation*, *2*, 274–282.
- Lockery, S. R., & Kristan, W. B., Jr. (1990a). Distributed processing of sensory information in the leech. I. Input–output relations of the local bending reflex. *Journal of Neuroscience*, *10*, 1811–1815.
- Lockery, S. R., & Kristan, W. B., Jr. (1990b). Distributed processing of sensory information in the leech. II. Identification of interneurons contributing to the local bending reflex. *Journal of Neuroscience*, *10*, 1816–1829.
- Lockery, S. R., & Kristan, W. B., Jr. (1991). Non-associative learning in the local bending reflex of the leech. *Journal of Comparative Physiology A*, *168*, 165–177.
- Lockery, S. R., & Sejnowski, T. J. (1992). Distributed processing of sensory information in the leech. III. A dynamical neural network model of the local bending reflex. *Journal of Neuroscience*, *12*, 3877–3895.
- Lockery, S. R., Wittenberg, G., Kristan, W. B., Jr., & Cottrell, G. (1989). Function of identified interneurons in the leech elucidated using neural networks trained by backpropagation. *Nature*, *340*, 648–471.
- Mason, A., & Kristan, W. B., Jr. (1982). Neuronal excitation, inhibition and modulation of leech longitudinal muscle. *Journal of Comparative Physiology A*, *146*, 527–536.
- Nicholls, J. G., & Baylor, D. A. (1968). Specific modalities and receptive fields of sensory neurons in the CNS of the leech. *Journal of Neurophysiology*, *31*, 740–756.
- Nicholls, J. G., & Wallace, B. G. (1978). Modulation of transmission at an inhibitory synapse in the central nervous system of the leech. *Journal of Neurophysiology*, *31*, 740–756.
- Ort, C. A., Kristan, W. B., Jr., Stent, G. S. (1974). Neuronal control of swimming in the medicinal leech. II. Identification and connections of motor neurones. *Journal of Comparative Physiology A*, *94*, 121–154.
- Pearlmutter, B. A. (1989). Learning state space trajectories in recurrent neural networks. *Neural Computation*, *1*, 263–269.
- Peeke, H. V. S., & Petrino, L. (1984). *Habituation, sensitization and behavior*. New York: Academic Press.
- Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. (1988). *Numerical recipes in C*. Cambridge: Cambridge University Press.
- Robinson, D. A., & Arnold, D. B. (1990). Creating pattern generators with learning neural networks. In A. I. Selverston (Ed.), *Short course 3 syllabus* (pp. 45–52). Washington, DC: Society for Neuroscience.
- Rolls, E. T. (1989). Parallel distributed processes in the brain: Implications of the functional architecture of neuronal networks in the hippocampus. In R. G. M. Morris (Ed.) *Parallel distributed processing: Implications for psychology and neuroscience* (pp. 286–308). Oxford: Oxford University Press.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by back-propagating errors. *Nature*, *323*, 533–536.
- Servan-Schreiber, D., Printz, H., & Cohen, J. D. (1990). A network model of catecholamine effects: Gain, signal-to-noise ratio, and behavior. *Science*, *249*, 892–895.

- Squire, L. R. (1987). *Memory and brain*. Oxford: Oxford University Press.
- Stuart, A. E. (1970). Physiological and morphological properties of motoneurons in the central nervous system of the leech. *Journal of Physiology*, **209**, 627–646.
- Thompson, R. F., & Spencer, W. A. (1966). Habituation: a model phenomenon for the study of neuronal substrates of behavior. *Psychological Review*, **73**, 16–43.
- Zipser, D. (1991). Recurrent network model of the neural mechanism of short-term active memory. *Neural Computation*, **3**, 179–193.
- Zipser, D., & Andersen, R. A. (1988). A back-propagation programmed network that stimulates response properties of a subset of posterior parietal neurons. *Nature*, **331**, 679–684.