IUPUI | DEPARTMENT OF COMPUTER AND INFORMATION SCIENCE
SCHOOL OF SCIENCE
A Purdue University School
Indianapolis

Neurotransmitter Field Theory: A Composite Continuous and Discrete Model

TR-CIS-0315-10

A Technical Report

By

Douglas S. Greer
Mihran Tuceryan

March 15, 2010

Neurotransmitter Field Theory: A Composite Continuous and Discrete Model

Douglas S. Greer and Mihran Tuceryan

Abstract The abstraction of electrical and chemical computations in a unified framework facilitates the analysis of information processing in biological systems. The concept of small molecules functioning as chemical memory is well established for the nucleotides in DNA but not for neural models. We propose a model where neurotransmitters serve not only as a communication mechanism, which transports information in space, but also as a memory mechanism, which stores information in time. Early multicellular organisms may have used chemical memory to encode external physical quantities such as force, energy, or position. Synapses were formed as an evolutionary extension of the endocrine system that transfixed the transmitter between cells and prevented diffusion. Rather than view synapses as discrete points, the concentration of transmitters is modeled as a continuous field, which may represent any function of time, space, or frequency. The dendritic arbor of each neuron, chemically reads an input neurotransmitter field, while the axonal arbor writes and erases the output neurotransmitter field. Discrete neural networks combine with the continuous fields to form a composite paradigm that contains the perceptron model as a special case. The combination of electrical and chemical computations permits an integrated analysis of autoreceptors, dendrodendritic synapses, neuron plasticity, and neuroglia cells. The components of the composite model can interact to create new constructs such as continuous symbols. Computer simulations, which use visual images as the chemical data, demonstrate parametric variations in several neurotransmitter-field algorithms.

Keywords *Neural Coding, Chemical Computation, Autoreceptors, Associative Memory, Neurotransmitter Fields, Continuous Symbols, Plasticity*

1 Introduction

Neurotransmitter fields differ from neural fields in that the underlying state variables are the transmitter concentrations inside synapses rather than neuron action potentials. Neural fields model the behavior of a large number of neurons by taking the continuum limit of discrete neural networks where the continuous state variables represent the mean firing rates. This approach, which attempts to represent physical quantities in the external world by the electrical spikes of neurons, has been studied for some time (Beurle, 1956; Griffith, 1963; Amari, 1977) and extensively reviewed by several authors (Ermentrout, 1998; Vogels et al., 2005). We present a new computational model where information is stored chemically with neurotransmitter molecules rather than electrically with action potentials.

In computer architecture and design, a useful abstraction is to classify components into three basic types: processors that perform arithmetic operations, memories that store information, and switches that move information from one location to another (Siewiorek et al., 1984). This approach focuses on the fundamental properties and removes the

details of the physical implementation. In the case of biological cells, it allows us to combine electrical and chemical computations in a single abstract model.

The massive chemical computations required for development and growth in plants and animals use the nucleotides in DNA and RNA to store information. In neural network models however, neurotransmitters, which are chemically similar to the nucleotide bases, are defined as a communication mechanism, transporting information from the axon to the dendrite. If neurotransmitters are also operating as a memory element, by storing information for some interval of time, then they are acting as *state variables*. This fundamentally alters the computational model that should be used to describe the system. Compared to an action potential, a neurotransmitter molecule is an extremely efficient method for storing information. A single neuron weighs approximately a trillion times that of a small molecule transmitter. Moreover, using action potentials to store information requires a constant expenditure of metabolic energy (ATP), whereas transmitters can store information passively without consuming energy.

A fundament principle of stored-program digital computers, usually attributed to John von Neumann, is that the same memory can be used for both instructions and data. That is, an arbitrary string of bits may be interpreted either as a numerical value or as a machine instruction. Disk drives and memory chips make no distinction between the two types. This same distinction can be applied to biochemical memory. The concept of a connection-strength in the standard neural network model effectively assumes that neurotransmitters are instructions sent from the axon terminal to the dendrite, which causes a change in the membrane polarization. At the synaptic level, there is no way to determine if the change in membrane polarization is actually due to a change in transmitter concentration that is being interpreted as a data value. However, the two different interpretations result in two different computational models. Consequently, both possibilities deserve a thorough investigation.

Many neurotransmitters such as epinephrine, dopamine, and histamine are also hormones in the endocrine system. Hormones do not penetrate the blood-brain barrier, which allows them to serve a distinct role in the central nervous system. There are close similarities between hormones in animals and phytohormones, the chemical messengers in plants. This suggests common computational principles. When cells in the endocrine system release hormones, they intermix, and the resulting concentration will be proportional to the summation of the total amount released minus that absorbed. Other cells interpret hormone levels by their concentration, not the amount produced. Given the chemical similarity, we may assume that in early primitive synapses, the neurotransmitters from multiple cells also combined in the extracellular space, which resulted in a concentration equal to the sum of that contributed by each cell. This is illustrated in Fig. 1a, in a simplified schematic form, with an *archetype synapse* composed of several axon terminals and dendrites. A similar summation still occurs within synapses today. Multiple types of cells, including astrocytes and other neuroglia, locally detect and release transmitters, which intermix inside the synaptic cleft (Haydon and Araque, 2002; Deutch and Roth, 2003; Montana et al., 2004; Silchenko and Tass, 2008). Complex multiple-synapse boutons have been implicated in learning associations (Geinisman et al., 2001).

A path between an axon terminal and a dendritic spine (Fig. 1a) shows that the summation of transmitters in the archetype synapse is separated from the internal

chemical state of both cells by the nonlinear responses of two cell membranes. The two nonlinear responses, which describe the amount of transmitter released, and the sensitivity of the dendrites, vary independently from each other and prevent a linear combination with any summation inside the cells.

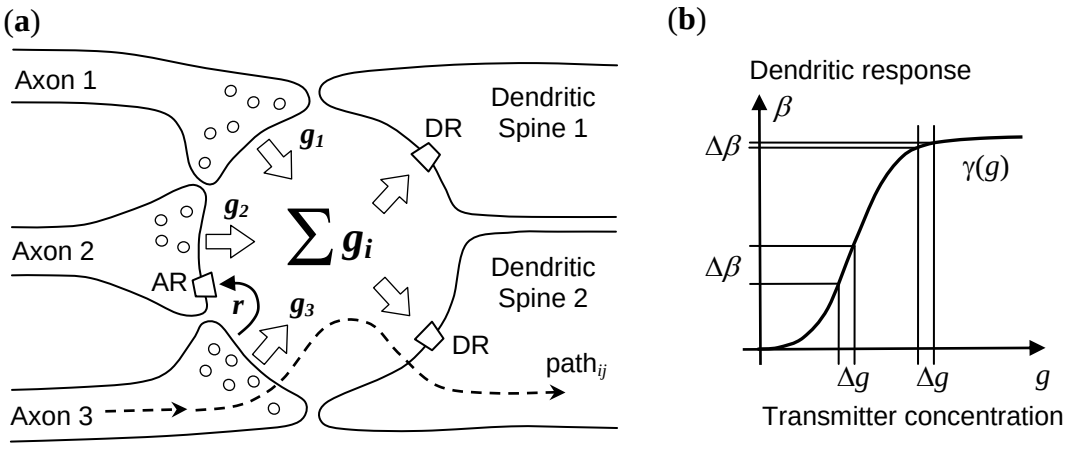

Fig. 1 In an early archetype synapse, (**a**), neurotransmitters intermixed, as do hormones in the endocrine system. The concentration of transmitters within the synapse equals the sum of the amounts contributed by each axon terminal. Dendritic spine receptors (*DR*) as well as axon terminal *autoreceptors* (*AR*) detect the presence of transmitters. The dashed line represents a typical path between an axon terminal and a dendritic spine. A second "regulatory-control" transmitter, r, is shown transmitting information between two axon terminals. Neurotransmitters released by other cells cause the response to an incremental increase to vary (**b**). This prevents the characterization of a single path in isolation

The postsynaptic response, β, as a function of the transmitter concentration, g, is plotted in Fig. 1b. Suppose the presynaptic neuron releases a fixed amount of transmitter shown as Δg. The effect on the postsynaptic neuron will be different when the concentration is near the central inflection point of a sigmoidal transfer function, compared to the effect when it is near the tails. This modulation should be taken into account when describing synaptic plasticity (Abbott and Nelson, 2000; Roberts and Bell, 2002; Shouval et al., 2002). However, it also challenges the notion of a simple connection-strength between neurons. Since a dendrite cannot determine which axon terminal released a transmitter, the nonlinearities prevent the analysis of a single path in isolation. Consequently, using a weight to describe a path between an axon to a dendrite may be an oversimplification. A more accurate model of the interaction occurring between the two cells would include state variables that record the internal synaptic transmitter concentration.

Autoreceptors, such as the one labeled *AR* in Fig. 1a, are receptors on the axon terminal from which the transmitter is released (Carlsson, 1975; Starke et al., 1989; Martin et al., 1991). They are a common regulatory feature on neurons (Deutch and Roth, 2003). In a neural network model, the existence of autoreceptors is paradoxical, since synapses supposedly serve only to transfer a weighted response from the axon to the dendrite. However, in the proposed neurotransmitter field (NTF) model, they are an essential component in the calculations. The importance of molar concentrations in the

analysis of reversible reactions – such as neurotransmitters binding to membrane proteins – justifies the introduction of variables to represent their value within synapses. Including this *synaptic chemical state* allows the effect of the autoreceptor's effect on the axons to be evaluated in the context of a comprehensive model of computation.

While we tend to view an organism as a collection of cells, a large portion of the body content lies outside the cell walls in the extracellular matrix. This includes the collagen in blood vessels and the osseous tissue in bones. The individual cells are fundamentally discrete, whereas the connective tissue forms a continuum. We explore the possibility of an analogous relationship between individual neurons and the neurotransmitters in the extracellular space.

1.1 Neurotransmitter Fields

Primitive nervous systems needed to evaluate the current state of the body and its immediate environment with representations of physical quantities such as mass (external world), position (location of the body surface), energy (visible light and sound vibrations), and force (pressure on the body surface, or the tension on the muscle cross-sections). Time, space, and frequency have the topological structure of n-dimensional Euclidian space ($\mathbf{R}^n$). Since any mathematical differences between mental representations and the actual physical quantities would be persistently removed by natural selection, at a high level of abstraction, the internal representation of these functions can be equated with that of the physical world. With the introduction of neurotransmitter concentration as a state variable, there is an alternative to neuron action potentials as a means of representing fundamental physical properties. Using terminology analogous to electric or magnetic fields, we refer to the chemical concentration of synaptic transmitters viewed as a continuous function in three-dimensional space as a neurotransmitter field or image. This leads to the following conjecture:

Neurotransmitter Field Hypothesis

When multicellular fauna first appeared, organisms began to represent quantities such as mass, force, energy, and position by the chemical concentration of identifiable molecules in the extracellular space. The basic principles of operation developed during this period still govern the central nervous system today.

From this single initial conjecture, we can deduce a series of incremental adaptations, leading to increasing complex structures, combining electrical and chemical computation. Imagine a simple primitive organism with a layer of sensory neurons interfacing to a layer of motor neurons. The interface could be implemented by mapping an external physical quantity to the concentration of a specific hormone in the space between the two layers. Such an interface would suffer from two problems: chemical diffusion due to Brownian motion, and chemical inertia due to the large volume between the layers. These problems would be ameliorated if the cell walls of the two neural layers moved together and formed synapses. This would effectively eliminate diffusion by confining transmitters to the synaptic cleft, and simultaneously reducing the volume of transmitter that must be maintained. The synapses would behave like dots of ink in a photograph, the axons writing and erasing the neurotransmitter "ink," and the dendrites reading it. Even though the synapses are discrete points, they would act as pixels or samples of functions defined in time, space, or frequency. The neural computations would adapt to preserve

the mathematical properties of the continuous neurotransmitter-density representations, since otherwise discretization errors of the continuous physical phenomena would place the organism at a relative disadvantage.

In neural networks, the physical location of a synapse is irrelevant to the computation. However, in neurotransmitter field theory, the location encodes information, the same as the location of a pixel in a digital photograph encodes information. The spatial density of synapses compared to that of neurons also helps to explain the sensory phenomena of hyperacuity.

Using abstract variables for the images, such as $h(x, y, z)$, allows us to evaluate and analyze possible neurochemical computations when the relationship between the operational variables and the actual chemical compounds is unknown, non-specific, or a complex aggregate formula. We present a model where the dendritic arbor encodes the transform of a neurochemical image as an action potential, and the terminal arbor then transforms the action potential back to a neurochemical image. This model is compatible with most experimental data. Population coding has been extensively researched as a method to represent stimuli by using the joint activities of a number of neurons (Wilson and Cowan, 1973; Abbott and Dayan, 1999; Rieke et al., 1999; Gerstner, 2002; Rolls et al., 2004; Averbeck et al., 2006). Sensory inputs, whether visual, auditory, or tactile, transduce a physical quantity to a chemical transmitter concentration, which is then encoded by a dendrite. Thus, the proposed model is consistent with population coding, but also includes the inverse transformation back to a chemical image. Numerous neurosurgery studies, where the stimulation of a single neuron in the hippocampus or amygdala region evokes an entire memory, provide evidence that this latter transformation exists (Penfield, 1958; Vignal et al., 2007; Hamani et al., 2008). Many studies have confirmed early work suggesting that neurons respond to particular input patterns rather than acting as discrete samples or pixels (Hubel and Wiesel, 1962). As an example, a neuron may fire only at a precise time during a particular bird song (Hahnloser R.H., 2002). This is consistent with the proposed NTF model where each action potential encodes a unique chemical image transformation. In the composite model, neurons are discrete entities that do not form a continuum.

The remainder of this paper is organized as follows. Section 2 discusses NTF algorithms as they may pertain to biological cells. Of particular interest is the ability to learn associations. A new computational model that describes presynaptic regulation by autoreceptors is then presented and we discuss how rates of chemical computation may be modulated. We analyze how the secondary messengers inside dendritic trees can also store information, and demonstrate intrinsic plasticity for an association operator. Section 3 briefly covers continuous symbols and their implementation by discrete cells. Section 4, discusses why neural networks are a special case of the more general model, and how the behavior of neurons and neural assemblies can be unified in the mathematical framework of chemical image transformations.

2 Composite Continuous and Discrete Algorithms

In early primitive organisms, the lipids in the cell membrane formed a barrier that maintained differentials in the chemical concentrations between the cell interior and exterior. The *intracellular state* of each neuron was separate and discrete, whereas the surrounding *extracellular state* was continuous and shared by all cells. Therefore, the

natural mathematical model contains both discrete and continuous elements, discrete nodes representing individual cells and continuous images representing the extracellular chemical concentrations. Figure 2 illustrates how neurons interact with chemical images. Viewed in isolation, the discrete elements are similar to a neural network and this framework subsumes the perceptron model as a special case. However, the composite continuous-discrete model allows us to analyze many otherwise confounding neural structures, including autoreceptors, dendrodendritic synapses, and complex anaxonal neurons.

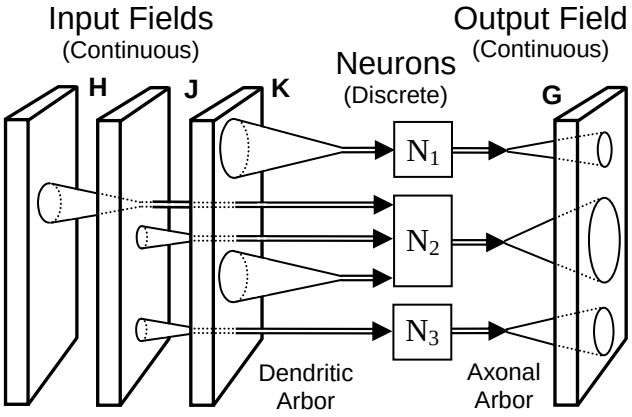

Fig. 2 The composite mathematical model of neural computation shows dendritic and axonal trees connecting the discrete neurons with the continuous neurotransmitter fields. The neuron cell membranes insulate their internal chemistry, creating discrete computational units. The neurotransmitter fields, shown as images, may encode physical quantities like force, mass, energy, or location, defined on continua such as space, time, or frequency. The topological alignment of multiple input images enables multisensory integration and the construction of complex operations

Let $\{N_i\}$ be a set of discrete processing elements (neurons) with their component parts identified by the subscript i. The level of activation, which may be interpreted as the mean firing rate, is denoted as α_i or β_i. We use bold uppercase, e.g. **H** or **G**, for continuous manifolds with the corresponding field values denoted with lowercase $h(x, y, z)$ or $g(x, y, z)$.

The computation performed by a dendritic tree is described by an integral operator that transforms the continuous transmitter concentration to a discrete value. Hilbert spaces, which may be continuous or discrete, are based on inner products (Royden, 1988). A neural network summation is a discrete inner product of two vectors, a vector of input values, and a weight vector. The continuous version of an inner product is the integral of two functions. If $h(x, y, z)$ represents the transmitter concentration in the extracellular space, for each neuron N_i, we define a *measure* $\mu_i(x, y, z)$ that represents the sensitivity of the dendrite to that transmitter. The analogous operations are:

$$\alpha_\iota = \sigma\left(\sum_k \beta_k w_{ik}\right) \;\leftrightarrow\; \alpha_\iota = \sigma\left(\iiint_H h(x,y,z)\, d\mu_i(x,y,z)\right) \qquad (1)$$

where σ is a sigmoidal activation function similar to a hyperbolic tangent. Note that while both equations in Eq. (1) are inner products in a Hilbert space, the left equation is the entire neural network computation, whereas the right equation is only the dendritic

tree operator of the NTF computation. The variable β_i represents the activation values of the input neurons and w_{ik} represents the synaptic connection strengths. In contrast, $h(x, y, z)$ represents the transmitter present in the synaptic cleft and $\mu_i(x, y, z)$ represents the sensitivity of the dendrite to transmitter at a particular location. In the NTF model, the amount of transmitter released by the axon terminals is defined by a separate equation.

In the composite continuous and discrete model, each cell or neuron is discrete, whereas the collective synaptic concentration of transmitters is continuous. If an action potential causes the release of a fixed amount of transmitter into each synapse, the corresponding discrete to continuous transformation is specified by multiplying the neuron activation value, α_i, by a *transmitter function*, $\tau_i(x_2, y_2, z_2)$. In this simple case, the function τ_i represents the amount of neurotransmitter released at each point into the output image – a distinct manifold parameterized with (x_2, y_2, z_2). If **G** is the output manifold, the intermixing of transmitters in synapses (Fig. 1) generates an output field, $g(x_2, y_2, z_2)$, equal to the summation over the neurons:

$$g(x_2, y_2, z_2) = \sum_i \alpha_i \tau_i(x_2, y_2, z_2) \qquad (2)$$

That is, the output image is generated by summing the amount of transmitter released by the active neurons, which for each neuron is the product of its transmitter function τ_i and its level of activation α_i.

2.1 Learning Associations

Neural assemblies that process various sensory inputs or motor control responses must implement unique and specialized computational functions. The ability to remember arbitrary cause-and-effect relationships, and thereby predict likely future events, requires learning associations. Any Boolean or arithmetic operation can be learned by remembering associations, and therefore, the association operator is logically complete. Image associations in the construction of continuous symbols have been proposed as the fundamental operation in the Brodmann areas (Greer, 2009). The association of arbitrary images is facilitated with radial basis functions similar to those employed for pattern recognition in neural networks (Bishop, 2005). Let θ be a bell-shaped curve such as a Gaussian, with the properties: $\theta(v) \geq 0 \ \forall v$, $\theta(v)$ is maximum at the origin, and $\theta(v) \to 0$ as $|v| \to \infty$. So neuron N_i will have a maximal response to a *pattern image* $\xi_i(x, y, z)$, we modify Eq. (1) so that its dendritic response is

$$\alpha_i = \sigma\left(\iiint_H \mathrm{q}\big(h(x, y, z) - \xi_i(x, y, z)\big)\, d\mu_i(x, y, z) \right) \qquad (3)$$

In Eq. (3), the dendrite of each neuron is characterized by two functions: the pattern image to which it is sensitive, ξ_i, and the measure μ_i. For each neuron, we set μ_i to zero outside the cone shown in Fig. 2, and equal to a uniform value inside the cone. The radial basis function, θ, is evaluated pointwise, resulting in the image $\theta[h\text{-}\xi_i](x, y, z)$ that has large values in areas where the input image, $h(x, y, z)$, matches the pattern image, $\xi_i(x, y, z)$. If μ_i is a positive (unsigned) measure, and the sigmoidal activation function σ, is monotonically increasing, then Eq. (3) will have its maximal value when the input h exactly equals ξ_i.

The association transformation for an entire layer of neurons, $\{N_i\}$, connecting the input and output images, is given by combining Eqs. (2) and (3):

$$g(x_2, y_2, z_2) = \sum_i \sigma\left(\iiint_H q(h - \xi_i)\, d\mu_i \right) \tau_i(x_2, y_2, z_2) \qquad (4)$$

where we have omitted the variables of integration for h, ξ_i, and μ_i.

When the dendritic tree of any neuron in Eq. (4) recognizes its particular neurotransmitter pattern, ξ_i, the neuron fires, causing the release of neurotransmitter into the output field. Several neurons with overlapping output regions may all recognize the same input pattern. However, provided the ξ_i, are sufficiently distinct, the corresponding transmitter functions $\tau_i(x_2, y_2, z_2)$ can be easily scaled to sum to an arbitrary output $g(x_2, y_2, z_2)$.

There are effectively two summations in Eq. (4): the explicit sum on i, and the integral over **H.** Because of the second chemical summation in the synaptic cleft, illustrated in Fig. 1, a single layer of neurons in the NTF model can perform the equivalent computation as a two layer neural network (Greer, 2009). Consequently, a few neurons or even a single neuron can encode a transformation or association between two chemical images. The NTF model removes the need for back-propagation algorithms used to determine the weights in two-layer neural networks. New associations are learned by adding new cells, which connect an input image recognized by their dendritic arbor with an output image generated by their terminal arbor.

2.2 Distributed Voting with Autoreceptors

With the inclusion of state variables to express the concentration of neurotransmitters, we can explicitly define autoreceptor calculations. Rather than axons simply releasing a fixed amount of transmitter, as was the case in Eq. (2), we can define computational models where an action potential acts as a signal to try to exert control on the amount of transmitter present in the synapse.

It is possible to construct algorithms where multiple axon terminals attempt to control a single transmitter concentration, and over time, they reach a steady-state consensus. However, if the target level of two axon terminals differs in value, the situation can arise where one terminal is constantly trying to release transmitter and another terminal is constantly trying to remove it. To prevent this, we construct a *distributed-voting system* where any number of axon terminals can participate in the computation of an average value without explicit knowledge of the presence or the activity of other axons. The axon terminals participate in the voting in response to neural spikes. The system outlined is both immediate and efficient, allowing a weighted average to be maintained and updated in a single iteration. This is accomplished using a second *regulatory-control* neurotransmitter, labeled $r(x, y, z)$, that keeps track of the total "votes cast". The votes cast are proportional to the activation levels α_i, which range from zero when the neuron is inactive, to a maximum value during burst firing. The axon terminals cooperate to maintain two *invariants*: the total sum of the votes, $r(x, y, z)$, and the averaged data value, $g(x, y, z)$.

$$r = \sum_i \alpha_i$$
$$g = \frac{\sum_i \alpha_i \tau_i}{\sum_i \alpha_i} \qquad (5)$$

The *data* neurotransmitter value, *g*, equals the weighted average of the target values τ_i, where the weights are the activation levels. To analyze how the invariants are maintained, assume neuron N_0, with a target data value of τ_0, fires, thereby signaling an activation level of α_0. Let r and g be the neurotransmitter concentration following the effect produced by neuron N_0. The new value of the invariants will be:

$$r = \sum_i \alpha_i + \alpha_0 \qquad = r + \alpha_0$$
$$g = \frac{\sum_i \alpha_i \tau_i + \alpha_0 \tau_0}{\sum_i \alpha_i + \alpha_0} = g + \frac{\alpha_0}{r + \alpha_0}(\tau_0 - g). \qquad (6)$$

On the right hand side of Eq. (6), the terms have been rearranged to show how much neurotransmitter must be added to the previous values Eq. (5), to maintain the invariants. It is instructive to analyze the rightmost term for *g*. If $r = 0$, this implies there are no other axons which are attempting to assert the data value, and the amount added will exactly equal the difference between the target value τ_0 and the current value g. That is, the axon terminal of N_0 will completely control the data value. However, an increasingly large value of *r*, implies that other neurons are also trying to assert the data value and Eq. (6) quantifies how the contribution of neuron N_0 to the weighted average will diminish. Thus, the effect of the regulatory-control neurotransmitter, *r*, is to inhibit the release of the data neurotransmitter, *g*, by other axon terminals.

When axon terminals release a fixed amount of transmitter, a mismatch between the correct neuron-firing rate and the correct reuptake rate will cause an error in the transmitter concentration approximately proportional to the imbalance. However, in a distributed-voting scheme, a rate mismatch may have little effect, since an increase in the firing rate simply implies that the target data value, τ_i, will be more strongly asserted. If the transmitter concentration in the synapse is already equal to τ_i, no additional transmitter will be released.

Computational results, illustrated in Fig. 3, show the distinction between axon terminals releasing a fixed amount of transmitter, as defined by Eq. (2), compared to axon terminals participating in a distributed voting system, as defined by Eq. (5). The experimental setup contained three input images with the same output image as shown in the model in Fig. 2. The results were generated using a software simulation written in Java that allows a high-level specification of parameters and records the outputs in XML reports (Greer, 2008). There were effectively three image association processors, each having different input images but sharing a common output image. For illustrative simplicity, the same six image pairs (Fig. 3a) were used in the training of all three associative mappings (H $\rightarrow$ G, J $\rightarrow$ G, and K $\rightarrow$ G). Each association was implemented

with a regular grid of 20×20 neurons, equally spaced and equally sensitive to all three color channels. The images resolution was 128×128 pixels.

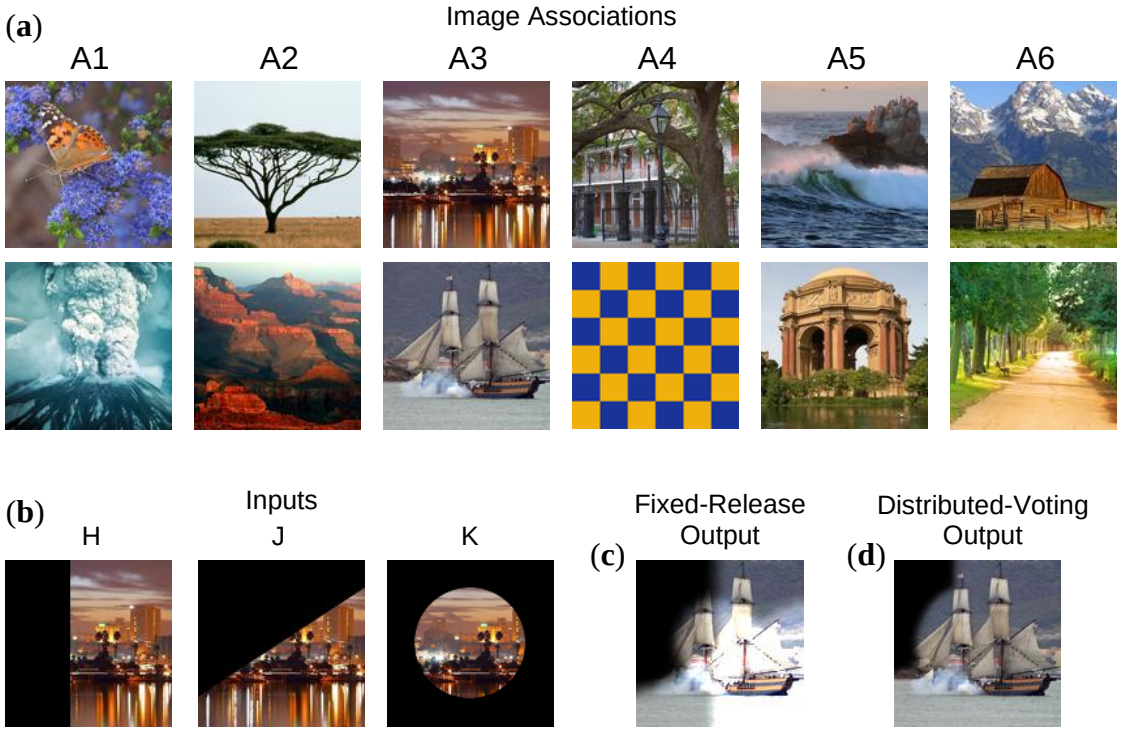

Fig. 3 Results produced with fixed-release and distributed-voting axonal operators in a three-input, single-output image association system. Both systems were trained with the same image associations (**a**) and presented with the same three input images (**b**) (corresponding to images H, J, and K in Fig. 2). When the axon terminals release a fixed amount of transmitter, the result (**c**) shows excessive transmitter as oversaturated image values in regions where multiple input images match the pattern. The output of the distributed-voting operator (**d**) shows how multiple active associations can use autoreceptors to adjust their transmitter release and thereby generate the correct output values

The training phase consisted of creating a grid of neurons whose overlapping receptor and transmitter regions covered the input and output images respectively. The dendritic and axonal arbor of each neuron was 32×32 pixels, clipped for neurons near the image borders. For each image association, a new grid was created and added to the system. To create a smooth transition between regions, the outputs, τ_i, were scaled by a two-dimensional "hill-shaped" function with a standard deviation of approximately six pixels from the neuron center. So that the τ_i functions would sum to unity at each pixel, each grid was "normalized" by dividing by the sum of the scaled values. A Gaussian with a standard deviation of 0.01 was used for θ, where the color components were in the range 0.0 to 1.0. A compressed and shifted hyperbolic tangent function was used for sigmoidal activation function σ, with the position of the inflection point set to 0.7 and the slope at that point set to 5.0.

The image sampling resolution corresponds to the density of synapses and is independent from the resolution of the neuron grid. This demonstrates how a few neurons can generate associations between images with a higher resolution. If one and only one of the three input images contains the recognized pattern, both output methods will correctly

produce the associated image. However, for the inputs shown in Fig. 3b, where the each of the three input images H, J, and K contains some portion of a matching input pattern, the results differ significantly. The output image generated when axon terminals *release fixed amounts of transmitter* (Fig. 3c) produces too much transmitter in regions where multiple input images match their associated pattern (center and bottom right quadrant). However, the weighted voting algorithm (Fig. 3d) produces the correct output even in the regions with multiple matching inputs. This is due to the actions of the autoreceptors, which reduce the amount of transmitter released in accordance with Eq. (6). When g already equals τ_i, no transmitter is released. The recognized output region is also larger when using the distributed-voting operator since the neurons have overlapping axonal arbors, and only one is needed to produce correct pixel values.

There are some similarities and some fundamental differences between common neural-network algorithms and NTF algorithms. For example, during the training phase, time-consuming back propagation methods are not required since learning takes place by adding cells. In the neurotransmitter-field model, neither the training phase nor the testing phase requires a slow (and often error prone) minimization algorithm. In the NTF model, image pixels effectively correspond to synapses rather than neurons. Because the axonal and dendritic arbors are implemented with sub-images, the simulations are memory intensive so testing on large data sets is difficult on a single processor system. However, the algorithms demonstrated are easily parallelizable and scale linearly with image resolution. This allows them to be applied to a large collection of very high-resolution images. For the results shown, the set-up and run time combined was less than fifteen seconds a single processor 2 GHz personal computer.

2.3 Changes in Computation Rates

In digital integrated circuit design, there is an important and fundamental tradeoff between processor speed and energy consumption (Rabaey et al., 1996; Hodges et al., 2003). An action potential consumes a fixed amount of energy and travels at a fixed rate. Consequently, in a neural network model, this tradeoff is static. However, like molecules of ink in a photograph, neurotransmitters can store information passively, without consuming energy. Thus, the NTF model can vary the tradeoff dynamically by simultaneously increasing (or decreasing) both the neuron-firing rate and the transmitter-reuptake rate. This allows an organism to compute rapidly during a fight-or-flight response, or conserve metabolic energy during relatively inactive periods. The lower bound on the rate of computation is limited only by the ability of the cell membranes to trap and hold the transmitters within the synaptic cleft.

Dynamically changing the tradeoff between energy consumption and processing speed is possible when all of the neurons in an entire region simultaneously change their firing rates and reuptake rates. Evidence of this type of widespread modulation, altering the rates of all neurons en masse, occurs during *volume transmission*, the parasynaptic communication that takes place outside synapses (Agnati, 2000; Sykova and Nicholson, 2008). Many neurotransmitters, which spread by non-synaptic chemical diffusion in the cerebrospinal fluid, have been shown to correspond to mental states such as hunger, mood, and alertness (Bach-y-Rita, 2001). These conditions indicate changes in the available energy and psychologically measurable response times.

Variations in Rate Ratios

In the distributed-voting algorithm, the two transmitters, *g* and *r*, are treated differently. The data field, *g*, may be left unchanged unless explicitly modified in response to an action potential. In contrast, the regulatory-control transmitter, *r*, must be periodically removed from a synapse so the weighted average can be recomputed. Consequently, the ratio of the firing rate, to the reuptake rate for the regulatory-control transmitter, is related to mutability. A fast uptake rate allows the data value to change quickly; whereas a slow uptake rate makes the data field relatively stiff, even when neurons impinging on the synapse are firing rapidly.

The effects of modifying the reuptake rate of the regulatory-control transmitter are demonstrated in Fig. 4. This computational experiment starts with the previous results shown in Fig. 3d, and then changes the three input images. Two of the inputs are changed to all black images, but the third input image is changed to the image shown in Fig. 4a, which matches input association A4 (Fig. 3a) except in the rectangular box in the center-left. The system parameters and resolutions were left unchanged. Separate sequences show the effect of a reuptake rate of twenty percent after each iteration (Fig. 4b) and a faster uptake rate of fifty percent (Fig. 4c).

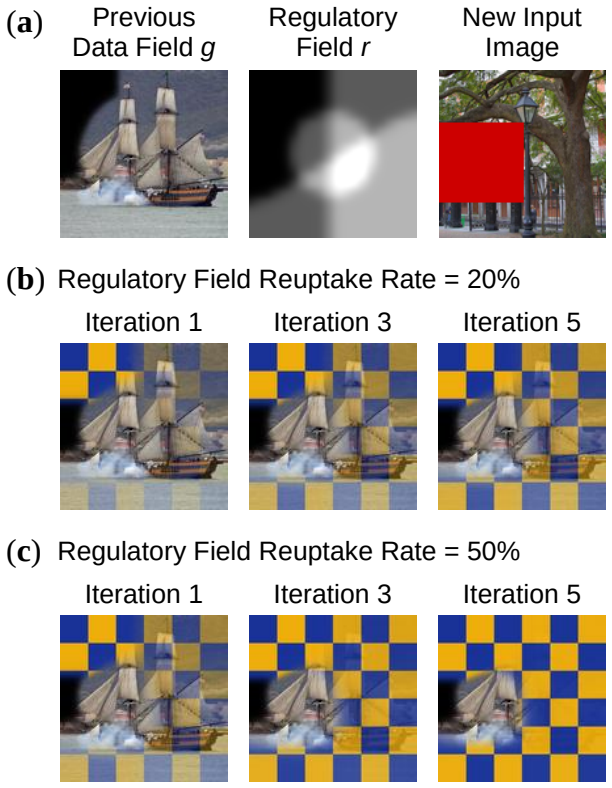

Fig. 4 In a distributed-voting system, the reuptake rate of the regulatory-control neurotransmitter, *r*, influences image mutability. The initial images are the results from the computation shown in Fig. 3. The initial values of the regulatory transmitter, *r*, are zero where none of the input images matched and have maximum values where all three input images matched. A new input image matches the input for a different association (A4 in Fig. 3a). Results of a reuptake rate of twenty percent per iteration (**b**) are slower to change, whereas the results using a fifty percent reuptake rate (**c**) change relatively quickly. In both cases, the upper-left corner where *r* was zero change instantly.

In the prior experiment, none of the input images (Fig. 3b) matched any association pattern in the upper-left portions image. Thus, the regulatory transmitter *r* was zero (black) in this region. When a new input is presented on the inputs, there is an immediate change in this portion of the output image. The slowest rate of change occurs where the most neurons where asserting a value, which resulted in a large value of *r* (lower-right center). Note that the box in the center-left portion of the new input image does not match any of the associations. Therefore, in the corresponding output area, the old data values remain unchanged.

2.4 Intracellular State

The dichotomy between memory and communication that applies to transmitters in synapses also applies to the chemical compounds within a cell. G-protein-coupled receptors and tyrosine kinase receptors activate secondary ion channels internally though pathways within the cytoplasm. *Secondary messengers* (Fig. 5) also influence neuronal growth factors and genetic transcription (Siegelbaum et al., 2000; Schulman and Roberts, 2003). These intracellular transmitters can also be viewed as state variables, storing information in an abstract computational system. As early organisms began to represent physical quantities by the concentration of transmitters in synapses, the intracellular chemistry would also have adapted to improve performance and efficiency. Neurotransmitter field theory allows the effects of local chemical reactions to be analyzed over wide areas. Analysis of secondary neurotransmitters as data, rather than instructions, may also help to establish the connection between working memory and long-term memory.

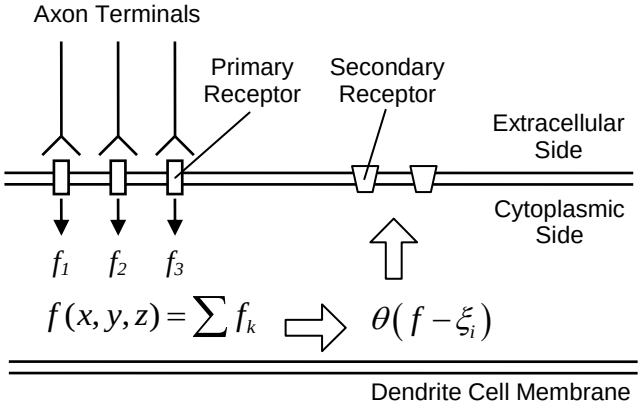

Fig. 5 Intracellular messengers in the cytoplasm can also function as chemical memory. The primary receptors (*rectangular shapes*) bind to transmitters in the extracellular space. The chemical summation is mapped by a nonlinear transfer function $\theta(.)$ to secondary transmitters that depolarize the cell membranes (*trapezoidal shapes*)

The chemical summation illustrated in Fig. 1 can also take place *within* the dendrite rather than the synaptic cleft. A calculation similar to Eq. (4) is shown in Fig. 5, where a secondary messenger, $f(x, y, z)$, responds to an extracellular transmitter through a primary receptor. In an association operation (for neuron N_i), the Gaussian shaped function θ compares the internal chemical sum with a fixed pattern ξ_i and indirectly activates the secondary receptors. Chemical diffusion must be restricted so that the sum is confined to

a relatively localized spatial position (x, y, z). Diffusion inside a neuron is limited by multiple factors including the thin cylindrical shape of the dendrites, and the binding of transmitters to the cell walls. As an example of the latter, diacylglycerol, a secondary messenger, contains hydrophobic chains that remain attached to the cell membrane, as do the hydrophobic residues of the G-proteins. Continuous state variables can be encoded by any identifiable chemical compound with restricted mobility. Moving the summation from the extracellular space to the cell interior removes the requirement that the axon terminals be physically adjacent (Fig. 5) and allows state variables to encode functions specific to a particular neuron, such as $\theta[f\text{-}\xi_i]$.

Forward Validation

In the distributed-voting system of Eq. (5), the data field, g, can be left unchanged until it is explicitly modified. However, the regulatory-control transmitter, r, must be continually removed from a synapse so the data value can be recomputed. For the postsynaptic cells, this allows the implementation of *forward validation* algorithms. When data values become less certain over time this can be detected by values of r near zero. This indicates that there has been no recent activity by the presynaptic neurons and therefore the data value may be less reliable or the margin of error may be larger.

In a neural-network model, when a neuron does not fire, this is interpreted as a zero data value. In contrast, in a distributed-voting model, a zero value of r is interpreted as "data value unknown." Consequently, the portions of an input image where $r > 0$ can be considered more significant or more reliable by the following stage. The measure μ_i in Eqs. (1) and (4) may be used to give more emphasis or significance to those areas of the input image where the data is more likely to be valid. For example, μ_i could be scaled by a monotonically increasing function of r. Regions of the input image where no input neurons are asserting a value would then be effectively ignored. With forward validation, both the data and regulatory-control transmitter fields are used by the postsynaptic neurons in the algorithmic computations.

2.5 Modulating Dendritic Responses

Studies of both vertebrates and invertebrates have shown changes in the intrinsic properties of neurons that are activity-dependent (Turrigiano et al., 1994; Turrigiano et al., 1998; Frick and Johnston, 2005). Intrinsic plasticity, a type of activity dependent modulation, has been linked with learning and memory (Triesch, 2007; Kim and Linden, 2007). Direct interaction between dendritic trees is known to perform complex information-processing functions (Shepherd, 1996). Moreover, dendrites release neurotransmitters back into synapses and form *dendrodendritic synapses* with other dendrites (Ludwig, 2005; Zilberter et al., 2005). Transcellular signaling by membrane-permeable modulators, such as nitrous oxide, occurs not only between neighboring dendrites, but also in a retrograde transmission to the presynaptic cell where they act as a secondary messenger in axon terminals (Siegelbaum et al., 2000). Subthreshold electrical activity between adjacent dendrites has also been documented (Bliss et al., 1995; Shepherd, 2003).

We demonstrate how NTF algorithms that modulate the characteristics of an entire dendritic arbor, or branches of an arbor, can improve performance in a computational

task. A shared NTF acts as an intermediate variable, recording the computational results of one dendritic arbor, which may then be used as an input field by other neurons. As before, a set of neurons was trained to associate image pairs. An overlapping grid of neurons was constructed for each image association and the system comprised the collection of all the grids. Dynamic modulation of the dendritic trees can solve the problem of finding one and only one image association in each local area. In regions where none of the patterns matches the current input, the recognition becomes more lenient, expanding the range of values that are acceptable. In regions where multiple patterns match, the recognition becomes more constrained, reducing the range of acceptable values. Figure 6a illustrates how the function $\theta(.)$ used in Eqs. (3) and (4) can be varied to modulate the dendritic recognition characteristics. The theta function can be constructed by multiplying the two nonlinear functions $\gamma_L(v)$ and $\gamma_R(v)$. Changes to these monotonically increasing and deceasing functions is biologically feasible and allows their product $\theta(v)=\gamma_L(v)\gamma_R(v)$ to be precisely controlled.

The modulation is controlled by a distinct neurotransmitter (labeled *b*) so it will not interfere with the other transmitters, even if it is present in the same synapse. When dendrites sense a pattern, they release transmitter *b*, which causes neighboring neurons to become more restrictive in their recognition. If no pattern can be identified, the value of *b* will be small, and recognition in the neighboring neurons becomes more permissive. Three functions, θ_1, θ_2, and θ_3, shown in Fig. 6b, illustrate the modulation of the receptor characteristics. They have the same mean, ξ_0, but a width that is a function of the recognition field $b(x, y, z)$. Figure 6b shows an input value, h, that is approximately, but not exactly, equal to the pattern value ξ_0. When the modulated theta functions are applied to the difference and mapped by a nonlinear activation function, σ, the changes in the output value may be arbitrarily large.

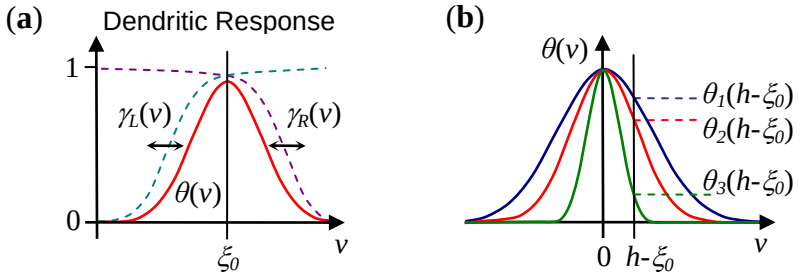

Fig. 6 A theta function (**a**) responsive to a narrow range of values can be constructed from the product of increasing and decreasing nonlinear functions $\gamma_L(v)$ and $\gamma_R(v)$. The width of the theta functions can be modulated (**b**) to change the acceptable inputs in an association operator

A dendrite-modulation experiment was designed and run to test the theta-function modulation shown in Fig. 6b. The results, shown in Fig. 7, used a fixed-transmitter-release axonal operator so the outputs would not depend on the results of previous iterations. The input image used (Fig. 7a) contains portions of two different images that intersect in a saw-tooth pattern. The images are the input patterns for associations A2 and A6 in Fig. 3a with a small amount of additive error. A vertical gradient that varied between 0.004 at the top to 0.025 at the bottom was created and added to the combined image. Initially, the standard deviation of the theta functions was set to 0.01, as in prior

experiments. With this value, approximately the top third of the pixel values fell within the limits of the bell-shaped curve. To cause multiple associations to match near the center of the saw-tooth pattern, the sigmoidal activation function set so the inflection point was at 0.2 with a slope of 20.0. The standard deviation of the theta function was incrementally decreased (or increased) when there were more (or less) than a one recognized association in each region. The change was based on a linear function of the locally summed intracellular activation values, with no recognized association causing an increase of 0.001 in the standard deviation of theta.

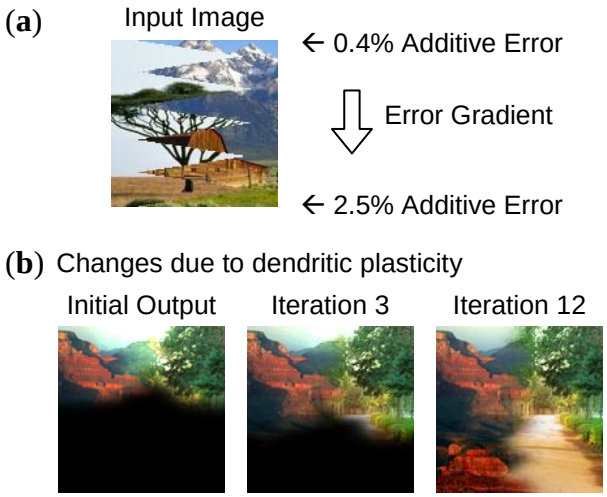

Fig. 7 A simulation of dendritic modulation used to recognize one and only image in each local region. The input (**a**) consists of a combination of two images with additive error increasing vertically towards the bottom. The activation of nearby neurons increases or decreases the range of acceptable values (**b**).

In the initial output image (Fig. 7b), both the left and right associations are recognized in the top-center portion of the image resulting in over-saturated, almost white, pixel values in this region. In progressive iterations, the standard deviation of the neuron theta functions in this area is decreased until locally only a single image is recognized. In the bottom portion of the image, the relative error is large and the width of the theta functions gradually increases until there is at least one matching association. Thus, insufficient recognition causes an increase the permissibility of the receptors, while too much causes a restriction. After twelve iterations, the output consists of the left and right halves of the A2 and A6 images, with each local area containing the closest matching association.

3 Continuous Symbols

The term "symbol" in computer science courses ultimately refers to a string of bits. However, abstract mathematical digits require a concrete physical realization to perform actual computations. Digital computers implement bits as bistable circuits such as the set-reset (SR) flip-flop shown in Fig. 8a. This circuit consists of two recurrently connected NAND gates, which creates a dynamical system with only two stable states, representing "zero" and "one." The state space for the corresponding electrical circuit is shown in Fig. 8b, with the two axes representing the P and Q voltages. The diagram illustrates the system state, $\mathbf{s}(t) = (V_q(t), V_p(t))$, converging along a path to the attractor "A_1". In the

terminology of dynamical systems, the stable states are *fixed-point attractors*. The symbols in a digital computer are not particular voltages, but rather a range of voltages, contained within *attractor basins*. Fixed-point attractors create the stability required for an actual physical realization to operate in the presence of the inevitable noise and transient errors in the inputs.

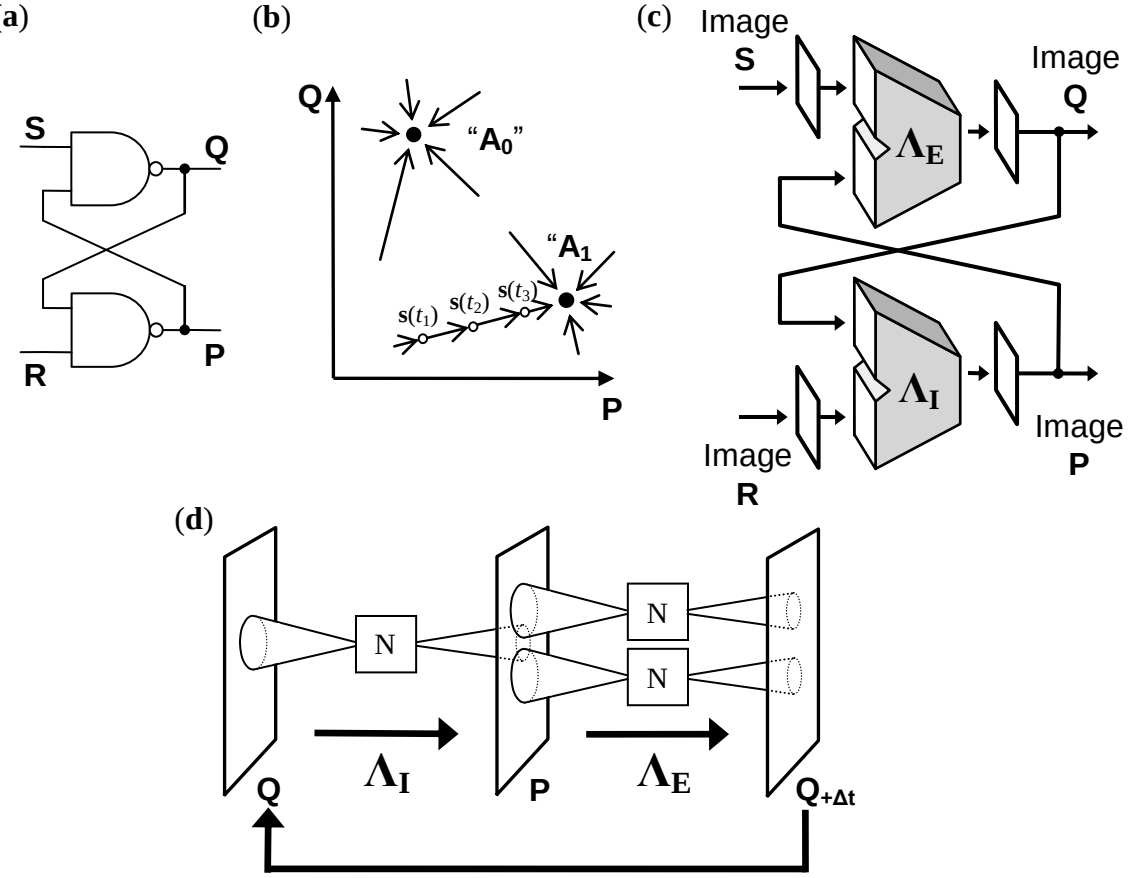

Fig. 8 A single bit, the fundamental discrete symbol, can be created with a set-reset flip-flop (**a**), which can be constructed from two NAND gates connected in a simple feedback loop. The stability of attractors in the state space (**b**) is necessary for the physical realization of a symbol processing system. Analogous principles of operation apply to a psymap (**c**), which can be constructed from two image association processors (Λ_E and Λ_I). A neurotransmitter-field implementation (**d**) shows the internal structure of the two recursively-connected Λ-maps. Even though the dendritic and axonal trees of the neurons (*N*) only cover limited areas, the recurrence allows the association to spread over the entire image

This same principle can be extended to continuous symbols (MacLennan, 1994). In a distributed-voting system, when one input image matches the associated pattern, additional matching images have no effect. Thus, for computational manifolds, the algorithm implements a type of logical "OR" operator. In general, a system capable of associating arbitrary input images with arbitrary output images can be viewed as the continuous version of a digital logic gate. We refer to the structure implementing this association operation as a Λ-*map* (Λ from the Greek word Logikos). Figure 8b shows a system analogous to the SR flip-flop where the bits have been replaced by images and the NAND gates have been replaced by Λ-maps. We label the two Λ-maps Λ_E and Λ_I, and

use the term *psymap* to refer to the combined recurrent image-association system. The topological alignment of multiple images allows inputs from extrinsic sources such as those labeled S and R in Fig. 8c.

Given a collection of image pairs, for each pair (q_j, p_j), it is possible to create a *reciprocal-image attractor* by constructing complementary associations in both Λ-maps so that $q_i = \Lambda_E(p_i)$ and $p_i = \Lambda_I(q_i)$. That is, each Λ-map stores a reciprocal image association. When an attractor is stable, it can serve as a *continuous symbol* in a cognitive system. Analogous to the SR flip-flop, stability implies that all of the images in some neighborhood of an attractor will converge toward it. Convergence is defined using a metric where the images are functions in a topological space (Munkres, 2000). Conceptually, the diagram of the state space shown in Fig. 8b still applies, but the voltage pairs are replaced by image pairs. Since these functions are infinite-dimensional, the behavioral characteristics of a psymap are that of an infinite-dimensional dynamical system (Temam, 1997; Robinson, 2001).

The chemical summation in the synapses allows each Λ-map to be implemented with a single layer of neurons. Consequently, the psymap in Fig. 8c can be implemented with only two layers of neurons. The rightmost image ($Q_{+\Delta t}$,) in Fig. 8d equals the leftmost image (Q) following one iteration of the recurrent computations.

The collection of six image pairs shown in Fig. 3a was used to instantiate and train two Λ-maps with reciprocal associations. The computational results, which used a distributed-voting axonal operator (Fig. 9), demonstrate how the recurrence allows local recognition to spread to an overall global association. In the two examples shown, the initial images are located in one of the attractor basins, and the image sequence converges towards the attractor. The black background does not match a pattern and in these local regions, the output remains unchanged since no neurons are asserting a value. As in the previous experiments the image resolution was 128×128, with the size of the dendritic and axonal arbors set to 32×32 pixels. The parameters were set the same as before, except for the inflection-point of the activation function, where the center-point and slope were set to 0.4 and 20.0 respectively. With this activation function, when 40% of the pixels within the dendritic arbor match, the activation value, α_i, equals 0.5. If no other axonal arbors are asserting a value, even small activation values can cause a locally recognized pattern to spread relatively quickly throughout the reciprocal image.

The NTF model contains both discrete and continuous components. Continuous symbols abstract away the discrete neurons, leaving only the relationships between continuous functions. The discrete portion is still present in the simulations, but is not directly illustrated in the results shown.

Since continuous symbols can remember and recognize representations of physical phenomena without losing information due to a reduction in the topological dimension, image attractors are likely to be present in the CNS. The human cerebral cortex is divided into approximately fifty Brodmann areas based on the thickness or prominence of six cellular layers. In the same area, the layers have the same thickness, but at the boundaries between Brodmann areas, there is a sudden change in the relative thickness. The pyramidal and granule layers both occur in exterior / interior pairs. If we group together the internal and external cortical layers of a single Brodmann area and connect them recursively, the result closely resembles a psymap, where the external layers form the external Λ-map, Λ_E, and the internal layers form the internal Λ-map, Λ_I.

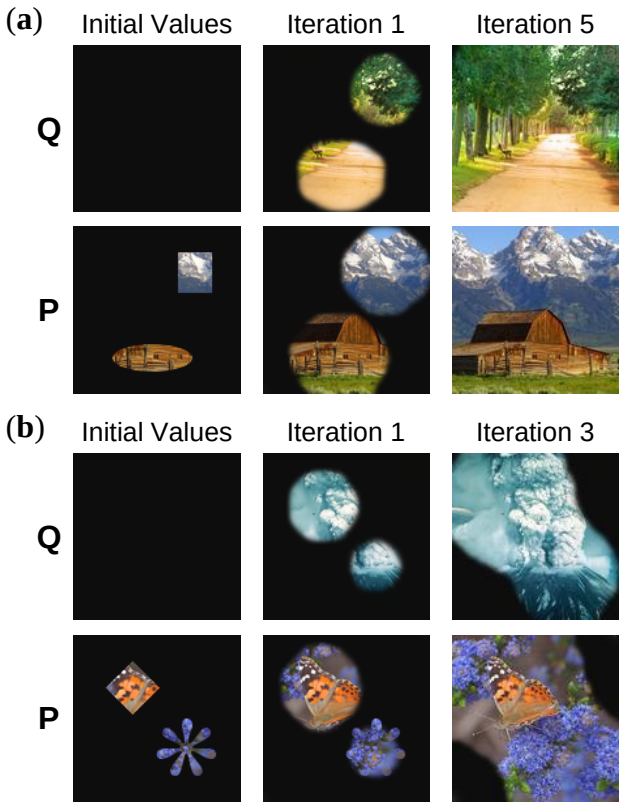

Fig. 9 Results of a simulation showing convergence to one of the stable reciprocal-image-attractors. Starting from different initial values (**a** and **b** *left*), the system converges to different stable states representing unique continuous symbols (A1 and A6 in Fig. 3a). In both cases, the convergence is complete after five iterations. On the third iteration (**b** *right*), the psymap has not yet completely converged to the attractor.

The majority of synaptic connections in the cerebral cortex are between neighboring neurons, and the common interconnection pattern suggests a common principle of operation (Mountcastle, 1978). The cortical columns have a structure that is consistent with one illustrated in Fig. 8d. If we identify each Brodmann area with a single psymap, the entire cerebral cortex can be modeled as a *psymap array* (Greer, 2009). This continuous-symbol processing model predicts that the cellular layers within each Brodmann area will have a constant thickness, since they share the same associations. However, different Brodmann areas have different inputs and store a different number of image associations. Therefore, there should be a distinct transition in cellular thicknesses from one area to the next.

Other mechanisms for implementing continuous symbols may allow the *P* and *Q* images to be collocated in the same physical space. For example, distinct images could be stored with distinct transmitters. For a distributed-voting system, this would require four uniquely identifiable compounds, two data transmitters, and two regulatory-control transmitters.

4 Discussion

Computation can be performed by mechanical, electrical, or chemical means. Many membrane proteins modify the chemical state without directly altering the electrical potentials. Consequently, to interpret neuron behavior, a comprehensive model that goes beyond voltages and currents is required in order to describe, at a more abstract level, how information is transformed, stored, and communicated.

4.1 Composite Model of Computation

The binding of neurotransmitters to receptors, like other reversible reactions, is a function of the concentration levels. Transmitters are removed from the synapse by the axon terminal, not the dendrite. This implies that the synapse is acting as a memory cell controlled by the axons, and not as a simple message passing mechanism. This view is reinforced by the existence of autoreceptors on the axon terminals, which detect the concentration of transmitters in the synapse.

Because they cannot model the biochemical storage of information, artificial neural network models may be an oversimplification. Moreover, anaxonic neurons and many neuroglia cells, such as astrocytes, lack action potentials but detect and release transmitters. In addition, neural networks cannot adequately characterize the bidirectional communication that takes place across cell membranes in both axons and dendrites. With the neurotransmitter field hypothesis, we can define a general model of biological computation that includes neural networks as a restricted special case. By modifying the synaptic transmitter concentration, cells can participate in computations even when they lack action potentials. Figure 10 illustrates an abstract, but realistic, computational model.

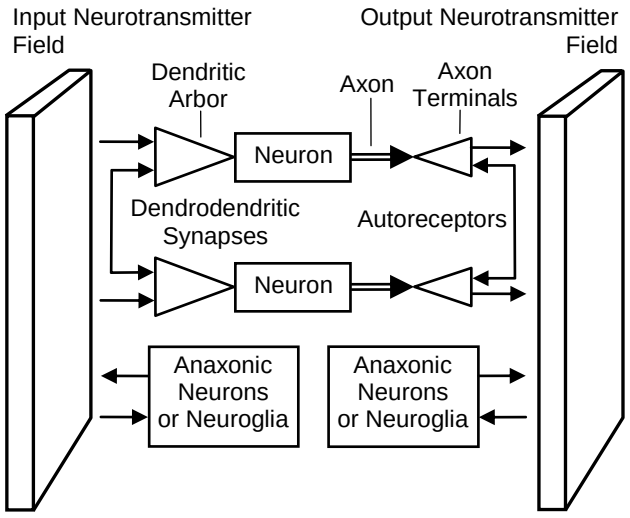

Fig. 10 The neurotransmitter field model includes not only the action potentials, but also the transmitter concentrations within synapses. Except for the axon, most communication in biological cells can take place in either direction across the cell membrane

The thick arrows indicate action potentials, which move information in a single direction down the axons. The thin arrows indicate chemical communication, which takes

place in both directions across the cell membrane. The introduction of chemical state allows the non-electrical computations performed by glial cells and anaxonic neurons to be included in a unified model.

The numerous feedback loops shown in Fig. 10 can be simulated, like other dynamical systems, with a next state transition function. The reactions that occur during an interval of time Δt are a function of the combined electrical and chemical state. Since image pixels are analogous to synapses, many general principles of operation can be tested and analyzed at low or medium resolutions.

4.2 Neural Networks as a Special Case

The neural network model can be generated by restricting the NTF model illustrated in Fig. 10. First, cells that lack action potentials such as anaxonic neurons and neuroglia must be removed from consideration. Second, we must remove dendrodendritic synapses and autoreceptors, which communicate without using action potentials. Only simple synapses with a single axon terminal are permitted, since the mixing of transmitters in complex synapses creates a chemical summation in the synaptic cleft. Similarly, the summation due to mixing of the secondary transmitters inside the dendrites must be excluded from consideration. Finally, and most importantly, neural networks require that synapses are *memoryless*. That is, there is an immediate reuptake by the axon of any transmitter released. Under these circumstances, the NTF model reduces to a neural network model.

A typical axonal tree of a single neuron often connects to another neuron at multiple synapses. This is counterintuitive in the neural network model since a single connection strength should be sufficient. However, in the NTF model, multiple synapses between the same two neurons could be expected, since the axonal tree is asserting an image, and the dendritic tree is recognizing an image.

Dynamical systems are characterized by state information that may be continuous or discrete. When the state of a system is completely specified, future transitions can be predicted. However, when state information is missing, a system may appear to be stochastic, when it is actually deterministic. For example, if we do not know what software a digital computer is executing, its observable sequential behavior is unpredictable. Many descriptive models of nerve impulses include random variables (Dayan et al., 2001; Bell and Sejnowski, 1995). The need to introduce a source of randomness in neural models may in part be due to an oversimplification that excludes the existence of neurochemical memory. We should expect that the concentration of transmitters that surrounds cells must be included in a model in order to predict behavior accurately.

4.3 A Basis for Neural Coding

Since neural networks lack chemical state, neuron spikes must be defined in terms of other spikes. The introduction of chemical memory within synapses expands the available framework for predicting neuron behavior. Image memory is the result of applying the neurotransmitter field hypothesis, thereby extending the chemical state to continuous domains. Biological neurons do not detect the actions of other neurons electrically; they react to the transmitters. Figure 10 captures this relationship by showing how in the NTF

model, action potentials encode chemical image transformations. Neural assemblies also encode chemical image transformations, and this commonality allows the computations of individual neurons to be understood in a larger context.

A mathematical framework for chemical image transformations and computational manifold automata can be defined in the terminology of real analysis and differential geometry (Spivak, 1979; Royden, 1988). For any manifold **H**, the space $L^2(\mathbf{H})$ contains all functions (images) defined on **H** whose square is integrable. These function spaces may serve as a basis for the interpretation of action potentials. For an input manifold **H** and an output manifold **G**, we can represent the effect of a neuron as transforming images in $L^2(\mathbf{H})$ to images in $L^2(\mathbf{G})$. That is, we identify a neuron N_i with a transformation T_i: $L^2(\mathbf{H}) \rightarrow L^2(\mathbf{G})$. The space of all such transformation is extremely large and neurons can only implement a restricted subset of all possible transforms. For example, the association operators described above were parameterized with the images $\xi_i(x, y, z)$ and $\tau_i(x, y, z)$, which reduces the set of possible transforms. The mean firing rate, α_i, is a positive real number ($\alpha_i \in \mathbf{R}_+$). Thus, we can analyze the transformation T_i by decomposing it into two components: a dendritic operator D_i: $L^2(\mathbf{H}) \rightarrow \mathbf{R}_+$ and an axonal operator A_i: $\mathbf{R}_+ \rightarrow L^2(\mathbf{G})$. The dendritic operator, D_i, transforms a function to a real number, and the axonal operator, A_i, transform a real number back to an image. The Cartesian product of these function spaces, $\{T_i\} = \{D_i\} \times \{A_i\}$ also represents a restriction on the set of all possible neuron transforms. In the case of the distributed-voting operator, the autoreceptors in the axon terminals detect the concentration levels in the output image. Thus, the axonal transformation maps the combination of a real number and an output image to new output image. That is, A_i: $\mathbf{R}_+ \times L^2(\mathbf{G}) \rightarrow L^2(\mathbf{G})$. The transform definitions may also include time, or manifolds with time derivatives. These can be used to specify the encoding of action potentials that are functions of temporal changes in the input images.

Physically adjacent neurons typically exhibit very different responses to stimuli. This has confounded efforts to create neural fields from mean-firing rates. However, in the NTF model this discontinuity is expected since spikes represent transforms between chemical image patterns that vary between neurons. Mathematically, different neurons are points sampled from an infinite-dimensional space of possible image transformations; each neuron is sensitive to a particular input image and generates a particular output image. However, since both neural assemblies and individual neurons transform chemical images, image transform analysis binds together the behavior of the whole as the combination of the individual parts.

Neurotransmitters can efficiently retain information over time and thereby form the basis of working memory. The spatial concentration of these small molecules can form internal representations of external physical quantities. This continuous memory can be combined with discrete neural networks to construct an abstract formalism for describing the behavior of neurons. Discrete cellular components may process multiple data and regulatory-control transmitter fields to create new algorithmic constructs such as continuous symbols. The unified composite model of electrical and chemical computation allows the integration of diverse types of cellular interactions and facilitates the analysis of large neural assemblies in a biologically realistic way.

References

Abbott, L. and Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. *Neural Computation*, 11(1):91–101.

Abbott, L. and Nelson, S. (2000). Synaptic plasticity: taming the beast. *Nature Neuroscience*, 3:1178–1183.

Agnati, L. (2000). *Volume transmission revisited*. Elsevier Science.

Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27(2):77–87.

Averbeck, B., Latham, P., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 7(5):358–366.

Bach-y-Rita, P. (2001). Nonsynaptic diffusion neurotransmission in the brain: Functional considerations. *Neurochemical Research*, 26:871–873.

Bell, A. and Sejnowski, T. (1995). An information-maximization approach to blind separation and blind deconvolution. *Neural computation*, 7(6):1129–1159.

Beurle, R. (1956). Properties of a mass of cells capable of regenerating pulses. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 240(669):55–94.

Bishop, C. (2005). *Neural networks for pattern recognition*. Oxford Univ Pr.

Bliss, T., Collingridge, G., Morris, R., Anderson, E., Lynch, G., Baudry, M., Moriyoshi, K., et al. (1995). Subthreshold synaptic Ca2+ signalling in fine dendrites and spines of cerebellar purkinje neurons. *Nature*, 373:12.

Carlsson, A. (1975). Dopaminergic autoreceptors. *Chemical tools in catecholamine research*, 2:219–225.

Dayan, P., Abbott, L., and Abbott, L. (2001). *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. MIT Press.

Deutch, A. and Roth, R. (2003). Neurotransmitters. *In Fundamental Neuroscience (ed. Squire et. al.)*, pages 163–196.

Ermentrout, B. (1998). Neural networks as spatio-temporal pattern-forming systems. *Reports on progress in physics*, 61:353–430.

Frick, A. and Johnston, D. (2005). Plasticity of dendritic excitability. *Journal of neurobiology*, 64(1):100–115.

Geinisman, Y., Berry, R., Disterhoft, J., Power, J., and Van der Zee, E. (2001). Associative learning elicits the formation of multiple-synapse boutons. *Journal of Neuroscience*, 21(15):5568.

Gerstner, W. (2002). *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge Univ Pr.

Greer, D. S. (2008). Sapphire Implementation Notes, http://www.gmanif.com/pubs/TR_CIS_0714-08.html. Technical report, IUPUI.

Greer, D. S. (2009). Images as symbols: An associative neurotransmitter-field model of the Brodmann areas. *Transactions on Computational Science*, V:38–68.

Griffith, J. (1963). A field theory of neural nets: I: Derivation of field equations. *Bulletin of Mathematical Biology*, 25(1):111–120.

Hahnloser R.H., Kozhevnikov A.A., F. M. (2002). An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature*, pages 65–70.

Hamani, C., McAndrews, M., Cohn, M., Oh, M., Zumsteg, D., Shapiro, C., Wennberg, R., and Lozano, A. (2008). Memory enhancement induced by hypothalamic/fornix deep brain stimulation. *Annals of Neurology*, 63(1):119–123.

Haydon, P. and Araque, A. (2002). Astrocytes as modulators of synaptic transmission. *A. Volterra, P. Magistretti, & Haydon (Eds.), Tripartite synapses: Synaptic transmission with glia*, pages 185–198.

Hodges, D., Jackson, H., and Saleh, R. (2003). *Analysis and Design of Digital Integrated Circuits*. McGraw-Hill, Inc. New York, NY, USA.

Hubel, D. and Wiesel, T. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, 160(1):106.

Kim, S. and Linden, D. (2007). Ubiquitous plasticity and memory storage. *Neuron*, 56(4):582–592.

Ludwig, M. (2005). *Dendritic neurotransmitter release*. Springer.

MacLennan, B. (1994). Continuous symbol systems: The logic of connectionism. *Neural networks for knowledge representation and inference*, pages 83–120.

Martin, D., Bustos, G., Bowe, M., Bray, S., and Nadler, J. (1991). Autoreceptor regulation of glutamate and aspartate release from slices of the hippocampal CA1 area. *Journal of Neurochemistry*, 56(5):1647–1655.

Montana, V., Ni, Y., Sunjara, V., Hua, X., and Parpura, V. (2004). Vesicular glutamate transporter-dependent glutamate release from astrocytes. *The Journal of Neuroscience*, 24(11):2633–2642.

Mountcastle, V. (1978). An organizing principle for cerebral function: the unit module and the distributed system. *The mindful brain*, pages 7–50.

Munkres, J. R. (2000). *Topology: A First Course*. Prentice-Hall, 2nd edition.

Penfield, W. (1958). *The excitable cortex in conscious man*. Springfield, Ill., Thomas.

Rabaey, J., Chandrakasan, A., and Nikolic, B. (1996). *Digital integrated circuits: a design perspective*. Prentice Hall.

Rieke, F., Warland, D., and Bialek, W. (1999). *Spikes: exploring the neural code*. MIT Press.

Roberts, P. and Bell, C. (2002). Spike timing dependent synaptic plasticity in biological systems. *Biological Cybernetics*, 87(5):392–403.

Robinson, J. (2001). *Infinite-Dimensional Dynamical Systems*. Cambridge University Press Cambridge, MA.

Rolls, E., Aggelopoulos, N., Franco, L., and Treves, A. (2004). Information encoding in the inferior temporal visual cortex: contributions of the firing rates and the correlations between the firing of neurons. *Biological Cybernetics*, 90(1):19–32.

Royden, H. (1988). *Real Analysis*. Englewood Cliffs, NJ: Prentice-Hall, 3rd edition.

Schulman, H. and Roberts, J. (2003). Intracellular signaling. *In Fundamental Neuroscience. (ed. Squire et. al.)*, pages 259–298.

Shepherd, G. (1996). The dendritic spine: a multifunctional integrative unit. *Journal of neurophysiology*, 75(6):2197–2210.

Shepherd, G. (2003). Information processing in complex dendrites. *In Fundamental Neuroscience. (ed. Squire et. al.)*, pages 319–337.

Shouval, H., Castellani, G., Blais, B., Yeung, L., and Cooper, L. (2002). Converging evidence for a simplified biophysical model of synaptic plasticity. *Biological cybernetics*, 87(5):383–391.

Siegelbaum, S., Schwartz, J., and Kandel, E. (2000). Modulation of synaptic transmission: Secondary messengers. *In Principles of Neural Science. (ed. Kandel et. al.)*, pages 229–252.

Siewiorek, D., Bell, C., and Newell, A. (1984). *Computer Structures: principles and examples*. McGraw-Hill New York.

Silchenko, A. and Tass, P. (2008). Computational modeling of paroxysmal depolarization shifts in neurons induced by the glutamate release from astrocytes. *Biological Cybernetics*, 98(1):61–74.

Spivak, M. (1979). *A Comprehensive Introduction to Differential Geometry*. Publish or Perish.

Starke, K., Gothert, M., and Kilbinger, H. (1989). Modulation of neurotransmitter release by presynaptic autoreceptors. *Physiological reviews*, 69(3):864–989.

Sykova, E. and Nicholson, C. (2008). Diffusion in brain extracellular space. *Physiological reviews*, 88(4):1277.

Temam, R. (1997). *Infinite-dimensional dynamical systems in mechanics and physics*. Springer Verlag.

Triesch, J. (2007). Synergies between intrinsic and synaptic plasticity mechanisms. *Neural computation*, 19(4):885–909.

Turrigiano, G., Abbott, L., and Marder, E. (1994). Activity-dependent changes in the intrinsic properties of cultured neurons. *Science*, 264(5161):974.

Turrigiano, G., Leslie, K., Desai, N., Rutherford, L., and Nelson, S. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, pages 892–895.

Vignal, J., Maillard, L., McGonigal, A., and Chauvel, P. (2007). The dreamy state: hallucinations of autobiographic memory evoked by temporal lobe stimulations and seizures. *Brain*, 130(1):88.

Vogels, T., Rajan, K., and Abbott, L. (2005). Neural network dynamics.

Wilson, H. and Cowan, J. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Biological Cybernetics*, 13(2):55–80.

Zilberter, Y., Harkany, T., and Holmgren, C. (2005). Dendritic release of retrograde messengers controls synaptic transmission in local neocortical networks. *The Neuroscientist*, 11(4):334.