Frequency-Spatial Transformation: A Proposal for Parsimonious Intra-cortical Communication

Regev Levi *
Tel-Aviv University

Eytan Ruppin †
Tel-Aviv University

Yossi Matias ‡
AT&T Bell Laboratories

James A. Reggia §
University of Maryland

July 4, 1996

Abstract

This work examines a neural network model of a cortical module, where neurons are organized on a 2-dimensional sheet and are connected with higher probability to their spatial neighbors. Motivated by recent findings that cortical neurons have a resonant peak in their impedance magnitude function, we present a ***frequency-spatial transformation*** scheme that is schematically described as follows: An external input signal, applied to a small input subset of the neurons, spreads along the network. Due to a stochastic component in the neurons' dynamics, the frequency of the spreading signal decreases as it propagates through the network. Depending on the input signal frequency, different neural assemblies will hence fire at their specific resonance frequency. We show analytically that the resulting frequency-spatial transformation is well-formed; an injective, fixed, mapping is obtained. Extensive numerical simulations demonstrate that a homogeneous, well-formed transformation may also be obtained in neural networks with cortical-like "Mexican-hat" connectivity. We hypothesize that a frequency-spatial transformation may serve as a basis for parsimonious cortical communication.

*Address: Dept. of Physiology, Tel-Aviv University, Tel-Aviv, 69978, Israel. E-mail: levi@math.tau.ac.il

†To whom correspondence should be addressed. Address: Depts. of Computer-Science and Physiology, Tel-Aviv University, Tel-Aviv, 69978, Israel. E-mail: ruppin@math.tau.ac.il

‡Address: Room 2D-146, AT&T Bell Laboratories, 600 Mountain Ave., Murray Hill, NJ 07974-0636, USA. E-mail: matias@research.att.com.

§Address: Dept. of Computer Science, University of Maryland, College Park, MD, 20742, USA. E-mail: reggia@cs.umd.edu.

1 Introduction

This paper presents a novel frequency-spatial transformation (FST) scheme, via which an incoming input signal to a given cortical module may activate different neural assemblies, depending on its frequency. We hypothesize that such transformation may take place in the cortex, as a parsimonious way of communication between cortical modules.

Our work is based on viewing the cortex as composed of many relatively small cortical modules [1, 2, 3]. Each cortical module contains thousands of neurons and is composed of numerous ***neural assemblies***, which are sets of neighboring neurons that tend to be activated together. The concept of neural assembly has been at the center of neurobiological thinking for a long time [4, 5]. The idea is that as single neurons themselves do not have sufficient complexity to account for complicated high-level tasks. It is the dynamic aggregation of such neurons into higher-order groups, or neural assemblies, that enables complex cognitive tasks to be carried out. In recent years, the work done in several laboratories using simultaneous, multi-neuron recordings has shown that such neuronal assemblies may be detected, demonstrating context-dependent 'cooperative behavior' between groups of neighboring neurons (see [6] for a comprehensive review).

Figure 1a Figure 1b

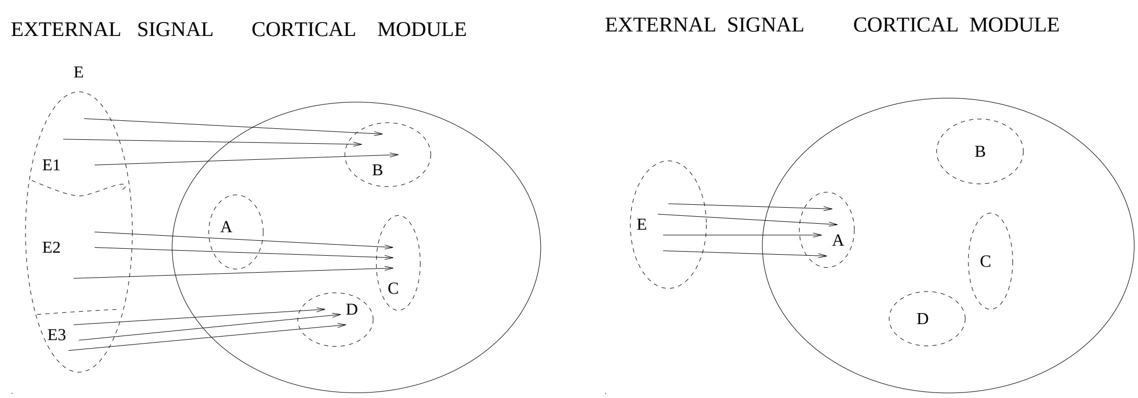

Figure 1: A sketch of a cortical module composed of several neural target assemblies (A,B,C and D) receiving an input signal from external assembly E. Both the segregated wiring (left) and the randomly-connected (right) cases are illustrated.

To describe the proposed frequency-spatial communication scheme, let us concentrate on a single cortical module and consider the scenario sketched in Figure 1. An input signal from an external source E, that may represent a neural assembly from another cortical module or a subcortical source, should activate target assembly B on some occasions, and C or D on others. One way to achieve this, via ***segregated routing*** (Figure 1a), would be to divide E into three segregated groups E_1, E_2, E_3, having specific connections to the different target assemblies. This paper presents another alternative, where the task of mapping an input signal to different assemblies is realized without any precisely tuned segregated routes: Consider a cortical module where the input signal projects (arrives at) only on some specific ***input assembly*** A, as sketched in Figure 1b. Via a frequency-spatial transformation we shall show how an input signal applied only to A can specifically activate several other neural assemblies (such as B,C and D) depending on its frequency.

The existence of a frequency-spatial transformation is naturally expected, as information in the central nervous system seems to be encoded in both spatial and frequency forms [7, 8]. A straightforward possible FST scheme arises if one could assume that each neural assembly has a precisely wired architecture of delay-line connections [9, 10, 11]. However, the existence of such a meticulously wired architecture of delay-lines in the cortex seems implausible [12]. The FST we propose does not require a precisely wired architecture of delay lines. It is based on the finding that cortical neurons have a resonance property: In a recent work, Gutfreund, Segev and Yarom have implemented the impedance magnitude function method, described earlier in [13], to characterize the properties of cortical neurons in the frequency domain [14]. They applied a sinusoidal current to pyramidal neurons in cortical slices at different frequencies, maintaining the overall current, and found that cortical neurons show a clear voltage-dependent ***resonant frequency*** (of 2-6 Hz at resting potential). That is, the voltage response generated by an input current oscillating at the resonant frequency is much stronger relative to the voltage response evoked by input currents in any other frequency.

In addition to the assumption that cortical neurons *in vivo* have a resonant frequency, the proposed FST is based on the following biologically-motivated assumptions:

1. Neurons are randomly connected to each other, with higher probability of connections formed between spatially proximal neurons [15, 16]. In light of their short, local range, we assume for simplicity that all connections have a similar conduction velocity, and that the conductance time over these connections is shorter than the (relative) refractory period.

2. Most cortical neurons have short membrane time constants, of the order of very few inter-spike intervals [8, 17]. Such neurons tend to fire in a synchronous manner [8, 18], and as a first approximation may be regarded as memoryless, i.e., their membrane potential is a function of only the most recently arriving spikes.

3. Neuronal firing has a stochastic component [19]. The rate of spontaneous cortical firing is very low [20].

In Section 2, the FST scheme is first presented and analyzed in a layered network, as a simplified but analytically tractable model. In Section 3 the FST scheme is examined via simulations performed in a more general spatially-connected network, where the neurons are connected primarily to their spatial neighbors via excitatory-inhibitory 'Mexican-hat' connectivity. In the last Section we discuss the relevance of neural resonance to the generation of reverberating activity, and analyze some additional biological and computational issues pertaining to the hypothesized FST scheme.

2 FST in a Layered Network

2.1 Overview

To obtain a general view of the proposed FST scheme, consider a layered network, as illustrated in Figure 2, where each neuron in layer i receives synaptic connections from a randomly selected subset of size s in the preceding layer $i-1$. The network is composed of memoryless, stochastic neurons having a short membrane time constant, that fire with probability p if the number of input spikes they have currently received exceeds a threshold Θ. The input signal is applied to the first layer, denoted the ***input assembly***, and it stimulates each input assembly neuron to fire at frequency f. The resulting neural activity is propagated throughout the layers of the network. As the conductance time along the network connections is shorter than the neurons' refractory period, the input signal propagates through the layers in an essentially feed-forward manner, spreading from the input layer outwards.

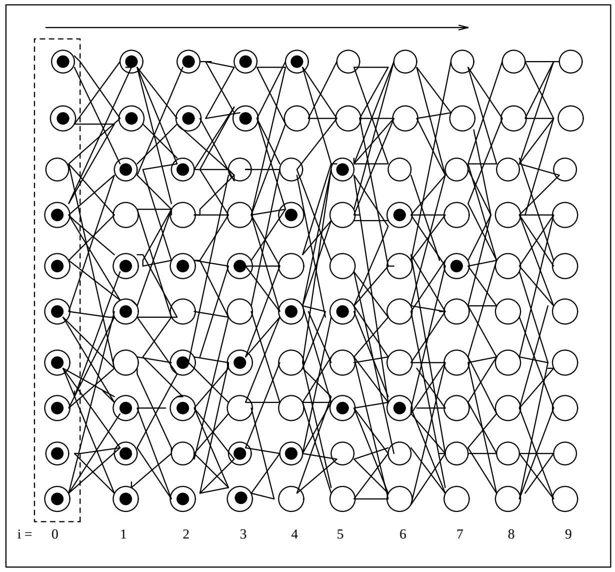

Figure 2: A schematic drawing of an FST in a feed-forward layered network. The input assembly (layer 0) is marked on the left. The spike propagation is illustrated by marking in black the neurons that fire in response to a given spike of the signal train. With every input spike, a similar scenario of neuronal activity, gradually decaying with its propagation along the network is evident, leading to a continuing decrease of the firing frequency measured along the network.

Neurons in the network have a resonant frequency f_r, that in accordance with the biological data [14] is assumed to be lower than the input signal frequency f. As we shall show, the firing frequency decreases as the activation resulting from the input signal propagates through the network. Hence, there exist some ***resonant layer*** where the average firing frequency is near the resonant frequency f_r. The identity of the resonant layer varies with the input signal frequency f; ***the higher f is, f_r occurs in a more distant layer*** - an FST is hence obtained. To demonstrate that the resulting FST is well-formed, we shall now show that:

1. ***The frequency decrease is monotonic***, implying an injective mapping; thus, several distant neural layers (i.e., assemblies) may not be concomitantly activated.

2. ***The mapping is fixed***: input signals with the same frequency, applied to the network at different times, will be mapped to the same layer with high probability, although the neural dynamics have a stochastic component.

2.2 The Frequency Spread

We now elaborate on how the frequency spread is obtained. Let Y_i be the number of spikes received by a certain neuron in layer i. If Y_i exceeds the neuron's threshold Θ, the neuron fires with probability p, and is idle otherwise. Let P_i be the probability of a neuron in layer i to fire in response to an input spike arriving at the input assembly (layer 0) [1]. Each neuron of the input assembly transmits each input spike with some predetermined probability $P_0 > 0$. The number of each neuron's connections with neurons in the preceding layer, s, is taken to be much smaller than the number of neurons in a given layer. There is therefore a very small dependency between neurons in the same layer, and Y_i is approximately a binomial with parameters (s, P_{i-1}). Using the Normal approximation to the Binomial distribution, we have

$$P_i \approx [\mathrm{Prob}\,(Y_i > \Theta)] \cdot p \approx \left[\frac{1}{\sqrt{2\pi}} \int_{t_i}^{\infty} e^{-y^2/2} dy \right] \cdot p \tag{1}$$

where

$$t_i = \frac{\Theta - sP_{i-1}}{\sqrt{sP_{i-1}(1 - P_{i-1})}} \tag{2}$$

is the approximate difference, in standard deviation units, between the expected value of Y_i and the threshold Θ. As evident, the expected firing frequency decreases monotonically along the layers as we move farther from the input assembly, guaranteeing the uniqueness of the frequency-spatial transformation.

2.3 FST Constancy

Let f be the input signal frequency. Let the ***resonant frequency band*** be the interval $[f_r(1-\epsilon), f_r(1+\epsilon)]$ for a predetermined $\epsilon > 0$, and let the ***resonant layer*** be the layer whose neurons' firing frequency is within the resonant frequency band. As we show next, with high probability the FST is accurate, that is, the resonant layer identifies with the ***expected resonant layer*** l, where $f \cdot P_l \sim f_r$ is in the resonant frequency band: Consider some time interval γ that is long enough for the resonant frequency to be evident. The number of spikes X_j that a neuron in layer j fires within the temporal window γ is a binomial with parameters $(f\gamma, P_j)$. Using Chernoff bounds [21] one may estimate the probability that the actual frequency observed at a given layer i will be within an ϵ factor from its expected value $(f \cdot P_i)$. The probability $\bar{P}$ of error, in which a signal is not recognized in the expected resonant layer is bounded by $\bar{P} < e^{-b\epsilon^2 f_r \gamma}$, where $b > 0$ is some constant. Note that $\bar{P}$ exponentially decreases with the resonant frequency f_r, the temporal window γ, and the relative width of the resonant frequency band ϵ.

[1] Note that the probability p is a property of each individual neuron, independent of the network structure, while the probability P_i is a property of the network.

Even when approaching the limit where the resonant frequency band is very narrow, the expected resonant layer still has the highest chance to fire in the resonant frequency. To illustrate this, suppose that k evenly spaced spikes, separated by time intervals of $\Delta_r = 1/f_r$, are required for sustained activity to emerge. Let layer l be the expected resonant layer, $f \cdot P_l \approx f_r$. In response to the propagating input signal, in each time unit a neuron in layer j has f independent events in which it either fires—such event is a ***success***—or it ***fails*** and remains idle. Let $m = \frac{f}{f_r} \approx \frac{1}{P_l}$ and let $Q_m(j)$ be the probability that a sequence of m events consists of one success followed by precisely $m-1$ failures. Let $Q_m^k(j)$ be the probability that such a sequence will repeatedly occur k times, i.e., the probability of generating the spike sequence required for sustained activity. Then, $Q_m(j) = P_j(1-P_j)^{m-1}$, and

$$\frac{dQ_m(j)}{dP_j} = (1-P_j)^{m-2}\,(1-mP_j)\ . \tag{3}$$

By equating to zero, $Q_m(j)$ is maximal at the appropriate layer $j = l$. Moreover, the ratio $\frac{Q_m^k(l)}{Q_m^k(j)}$, $j \neq l$, increases exponentially in k.

2.4 FST With Very Low-frequency Inputs

From a theoretical, computational point of view it is interesting to note that an FST can also be obtained for signals with frequency $f = 1/\Delta$ that is lower than the resonant frequency f_r, i.e., $f_{min} \leq f \leq f_r$, where f_{min} is the neurons' spontaneous firing frequency. To see this, assume that whenever an input assembly neuron receives an input signal spike, it responds by firing a sequence of spikes during a time interval $\Delta_{max} = \frac{1}{f_{min}}$, characterized as a Poisson process with a firing rate of $\frac{c}{\Delta_{max}}$, where c is some constant. The expected number of new spikes in each interval Δ is c, and the probability that such an interval contains more than $c{\cdot}t$ new spikes is exponentially small in t. Thus, the new spikes fall within the subsequent time intervals in a quite homogeneous manner. The input signal is transformed to a sequence of homogeneously dispersed spikes of frequency $f' = f \cdot (c+1) > f_r$, and an FST may now be obtained in the regular manner described previously.

3 FST in a Spatially-connected Network

After analyzing the properties of the proposed FST scheme in the simpler layered network, we now turn to simulations examining the FST scheme in a cortical-like, spatially-connected network, sketched in Figure 3. The network is composed of a 2-dimensional array of 120×120 synchronously updated neurons. The neurons in the network have a cortical-like 'Mexican-hat' excitatory-inhibitory connectivity: Each neuron has 40 proximal excitatory connections of weight $w_p = 1.85$ (to neurons within distance 4) and 20 distal inhibitory connections of weight $w_n = -1$ (to neurons with distance greater than 4 but smaller than 9). The input signal is applied to a spatially continuous input group of neurons consisting of 60 neurons (see Figure 3). If a neuron's input field is above its threshold Θ it fires with probability $p = 0.9$, otherwise it remains quiescent. After it fires each neuron remains refractory for a period of one iteration. As the input signal propagates through the network, consecutive assemblies of synchronously firing neurons are activated, and may be considered as 'layers' in a virtual layered network.

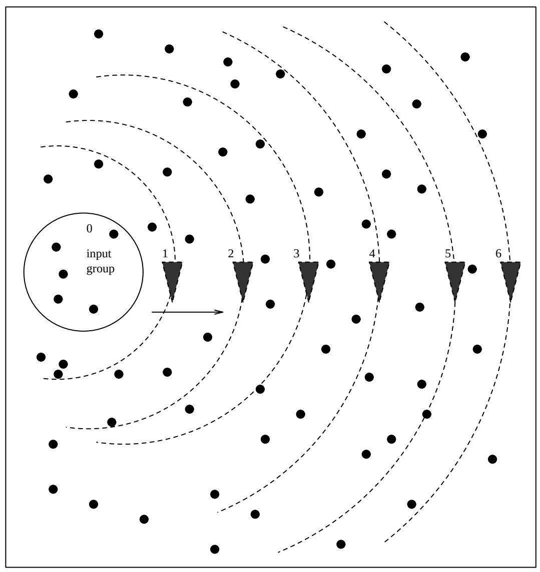

Figure 3: A schematic drawing of a cortical-like random network. For clarity, only the neurons are shown, excluding their connections. The input group is circled on the left. The wedge-shaped icons denote the locations of the simulated 'electrode', measuring the average firing frequency at increasing distances $1d, 2d, 3d\ldots$ The dotted arcs stand for the dynamically generated virtual 'feed-forward layers'.

Figure 4 presents the frequency decrease (the frequency of firing relative to the input signal frequency) along the network, as a function of distance from the input assembly. Using a non-linear sigmoidal threshold function $f(x) = 1/(1 + \exp[-(x - \theta/T)])$, a conventional approximation of the neural input/output firing rate [22, 23], a homogeneous spread of the input frequency along the network can be achieved. The spread and slope of the frequency decrease along the network are primarily a function of the noise level T. With low noise values (in this example, with $T \leq 1.5$) approaching the limit of a deterministic threshold function $Sgn(\cdot)$, the frequency spread achieved does not have a uniform slope, and it decreases only slightly at many locations. As the noise is increased, the rate of frequency decrease along the network is attenuated. If the noise is further increased then, beyond some critical noise value the FST does not remain monotonic any more. Figure 5 presents the mean and standard deviation of the frequency mappings obtained over 20 simulation experiments, held at the same level of noise. As evident, most of the frequency mappings obtained lie fairly close to the mean and the resulting FST is quite deterministic.

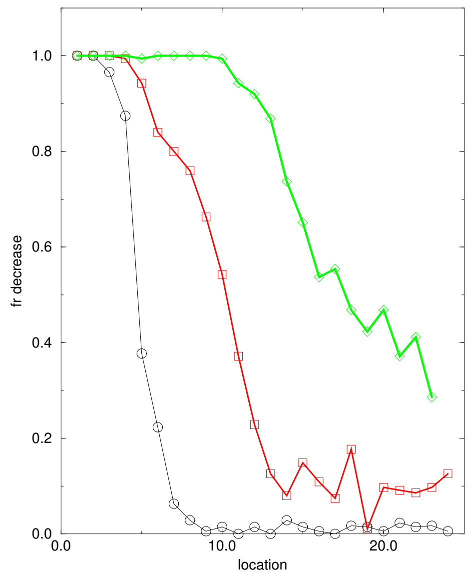

Figure 4: A plot of the firing frequency (f_r) decrease as a function of the distance from the input assembly, in a network with excitatory-inhibitory connectivity. Each neuron has a sigmoidal activation function and a distinct threshold value, determined randomly from the interval $\Theta \in [9, 11]$ and reflecting $\sim 10\%$ perturbations of the mean threshold value. The three curves display (from left to right) the frequency reduction obtained with noise values $T = 1.5, 1.7, 1.9$ correspondingly.

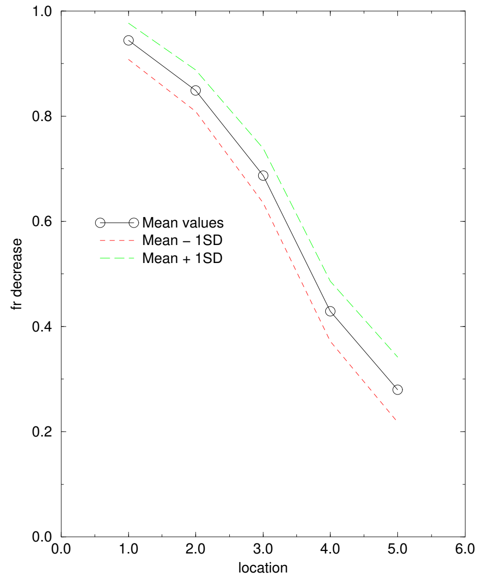

Figure 5: A plot of the firing frequency (f_r) decrease as a function of the distance from the input assembly, in a network with excitatory-inhibitory connectivity. The network parameters are similar to those reported in figure 4, with $T = 1.7$. Both mean and mean $\pm$ standard deviation values are displayed.

Finally, the FST obtained, that is, the location of the assembly firing in the resonant frequency as a function of the input signal firing frequency, is demonstrated in Figure 6. As evident, a fairly linear mapping is obtained over a broad range of input values.

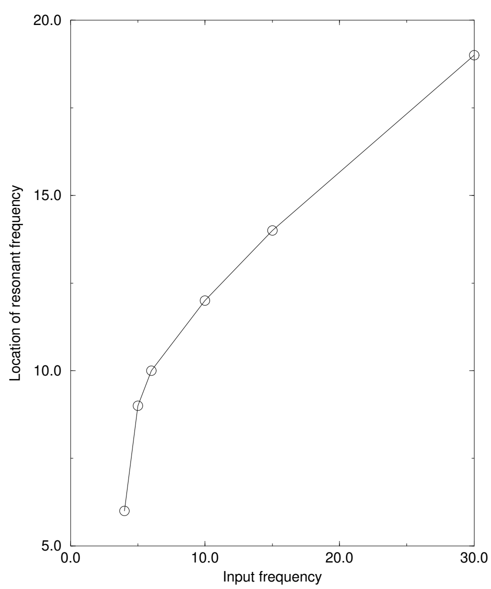

Figure 6: The location of the resonant assembly as a function of the input signal frequency. The neurons' have a sigmoidal activation function with $T = 1.5$, and threshold $\Theta = 9$.

4 Discussion

We have shown that different input frequencies can faithfully evoke the firing of distinct neural assemblies in the resonant frequency, thus constituting an FST scheme. In a simplified layered neural network model we show that an FST is feasible, and that with high probability the transformation is well-formed, i.e., the frequency decrease is monotonic and the resonant frequency is obtained at the expected resonant assembly. The FST scheme is then demonstrated numerically in a model of a cortical network with excitatory-inhibitory connectivity. We propose that such an FST scheme may have a role in intra-cortical communication.

Resonant behavior has been demonstrated experimentally in a wide variety of nerve cells, including peripheral neurons [24] and cortical neurons [14]. The possible important computational role of the resonance property of cortical neurons has already drawn the attention of several researchers. It has been proposed that neural resonance and its underlying subthreshold oscillations in membrane potential may serve as the basis for the conjunctive properties of central nervous system networks, binding between different attributes belonging to the same object [25]. Another computational model using action potential timing for stimulus representation based on membrane potential subthreshold oscillations has been proposed recently [26]. Our paper proposes yet a new computational role for resonance – enabling the cortical realization of an FST.

In the simplified model studied in this paper a fixed resonant frequency is assumed. The experimental data, however, is more complex; the fundamental frequency of resonance in most neurons is low (2-6 Hz), but the resonant frequency is voltage-dependent and may increase (up to levels of 15-20 Hz) if the neuronal membrane is strongly depolarized [14]. In accordance with the FST scheme, this observed increase in the resonant frequency f_r may be advantageous from a computational point of view, since the length of the temporal window (i.e., the membrane time constant) γ decreases as the membrane potential is increased. To see this, recall that the probability that the signal will not be recognized in the expected resonant layer is bounded by $\bar{P} < e^{-b\epsilon^2 f_r \gamma}$. Hence f_r should be increased with the membrane potential to maintain a certain level of correct performance.

The FST scheme demonstrated in this paper maps firing frequencies to spatial locations while the input signal is on. In principle such a mapping may extend beyond the initial period of input application to the network, via the emergence of sustained activity in the neural assembly firing at the resonant frequency. This possibility is in accordance with recent experimental studies that have suggested that the resonance property promotes population activity at the preferred frequency [14]. As manifested in our simulations, the resonant assembly is spatially localized, which further enhances the likelihood of it generating sustained activity. Another simplification of biological reality incorporated in our model is the assumption that the network neurons are memoryless. We believe that the FST scheme will work also when this assumption is relaxed, since as long as the membrane time constants of the neurons are fairly similar over all network (e.g., [17]), the resulting FST will continue to be monotonic. Moreover, it is plausible that adding a memory component to the neurons may enhance the likelihood of the emergence of sustained activity in the resonant assembly. The computational study of these issues requires, however, a much more detailed, compartmental, neural model (such as the one described recently by [27]), and constitutes an important subject for future investigation.

Our work is in line with previous research testifying to the important role noise may have in enhancing neural information processing abilities [28, 29]. It is shown that the extent of memory required to perform an FST (by having an array of precisely connected delay lines) can be significantly reduced by incorporating a stochastic component, both in the neurons' connectivity and in their dynamics. This results in an FST scheme that might play an important role in intracortical communication.

Acknowledgment: We are grateful to Moshe Abeles, Nathan Intrator, Issac Meilijson, and Yehezkel Yeshurun for many stimulating discussions and helpful comments. We thank the referees for their insightful comments.

References

[1] P.S. Goldman and W.J.H. Nauta. Columnar distribution of cortico-cortical fibers in the frontal, association, limbic and motor cortex of the developing rhesus monkey. *Brain Res.*, 122:393–413, 1977.

[2] V.B. Mountcastle. An organizing principle for cerebral function: The unit module and the distributed system. In F.O. Schmitt and F.G. Worden, editors, *The Neurosciences: Fourth Study Program*, pages 21–42. MIT Press, 1979.

[3] J.C. Eccles. The modular operation of the cerebral neocortex considered as the material basis of mental events. *Neuroscience*, 6:1839–1855, 1981.

[4] C. Sherrington. *Man on His Nature. The Gifford Lectures, Edinburgh 1937-1938.* Cambridge University Press, 1941.

[5] D. Hebb. *The Organization of Behavior.* Wiley, 1948.

[6] G.L. Gerstien. Neural assemblies as building blocks of cortical computation. In E.L. Schwartz, editor, *Computational Neuroscience*, pages 179–191. 1989.

[7] H.B. Barlow. Single units and sensation: A neuron doctrine for perceptual psychology. *Perception*, 1:371–394, 1972.

[8] M. Abeles. *Local Cortical Circuits, An Electrophysiological Study*. Springer-Verlag, 1982.

[9] B. Delgutte. Speech coding in the auditory nerve: Processing schemes for vowel like sounds. *J. Acoust. Soc. Am.*, 75:879–886, 1984.

[10] L. Deng, C.D. Geisler, and S. Greenberg. A composite auditory model for processing speech sounds. *J. Acoust. Soc. Am.*, 82:2001–2012, 1987.

[11] D.C. Tam. Temporal-spatial coding transformation: Conversion of frequency-code to place-code via a time-delayed neural network. In *International Joint Conference on Neural Networks*, pages I–130–133, January 1990.

[12] S. Shamma. Spatial and temporal processing in central auditory networks. In C. Koch and I. Segev, editors, *Methods in Neuronal Modeling*. MIT press, 1989.

[13] E. Puil, Y. Gimbarzevsky, and R.M Miura. Quantification of membrane properties of trigeminal root ganglion neurons in guinea pigs. *Journal of Neurophysiology*, 55(5):995–1016, 1986.

[14] Y. Guttfreund, Y. Yarom, and I. Segev. Subthreshold oscillations and resonant frequency in guina-pig cortical neurons: physiology and modelling. *J. of Physiology*, 483.3:621–640, 1995.

[15] V. Braitenberg and A. Schuz. *Anatomy of the Cortex: Statistics and Geometry*. Springer-Verlag, 1991.

[16] M. Abeles. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, 1991.

[17] T.J. Sejnowski. Open questions about computation in cerebral cortex. In J. McLelland and D. Rummelhart, editors, *Parallel Distributed Processing*, volume 2, pages 372–389. MIT Press, 1986.

[18] M. Abeles and G.L.G. Gerstien. Detecting spatio-temporal firing patterns among simultaneously recorded single neurons. *J. Neurophysiol*, 60:909–924, 1988.

[19] H.C. Tuckwell. ***Stochastic Processes in the Neurosciences.*** SIAM, 1989.

[20] M. Abeles, E. Vaadia, and H. Bergman. Firing patterns of single units in the prefrontal cortex and neural network models. ***Network***, 1:13–25, 1990.

[21] H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. ***Annals of Math. Statistics***, 23:493–507, 1952.

[22] H.R. Wilson and J.D. Cowan. Excitatory and inhibitory interactions in localized populations of model neurons. ***Biophysical Journal***, 12:1–24, 1972.

[23] J.C. Pearson, L.H. Finkel, and G.M.Edelman. Plasticity in the organization of adult cerebral cortical maps: A computer simulation based on neuronal group selection. ***Journal of Neuroscience***, 7(12):4209–4223, 1987.

[24] A.C. Crawford and R. Fettiplace. An electrical tuning mechanusm in turtle cochlear hair cells. ***J. of Physiology***, 312:377–412, 1981.

[25] R. R. Llinas, A. A. Grace, and Y. Yarom. In vitro neurons in mammalian layer 4 exhibit intrinsic oscillatory activity in the 10- to 50-hz frequency range. ***Proc. Natl. Acad. Sci.***, 88:897–901, 1991.

[26] J.J. Hopfield. Pattern recognition computation using action potential timing for stimulus representation. ***Nature***, 376:33–36, 1995.

[27] J.L. Hindmarsh and R.M. Rose. Resonance in a model of a mammlian neuron. ***Phil. Trans. R. Soc. Lond. B***, 346:151–163, 1994.

[28] J. Buhmann and K. Schulten. Influence of noise on the function of a 'physiological' neural netrwork. ***Biol. Cybern.***, 56:313–327, 1987.

[29] A.R. Bulsara, R.D. Boss, and E.W. Jacobs. Noise effects in an electronic model of a single neuron. ***Biol. Cybern.***, 61:211–222, 1989.