Hebbian Learning of Temporal Correlations: Sound Localization in the Barn Owl Auditory System

J. Leo van Hemmen and Richard Kempter
Physik-Department der TU München
D-85747 Garching bei München, Germany

Abstract

There exists an unresolved paradox in auditory and electrosensory neural systems in that they encode behaviorally relevant signals in the range of a few microseconds with neurons whose time constants are at least one order of magnitude bigger. Here we focus on the barn owl's auditory system, analyze its azimuthal sound localization, and present results of a modeling study based on computer simulations of a neuron in the laminar nucleus. We argue that three observations resolve the paradox. First, the neuron is driven by *many stochastic* presynaptic signals that arrive more or less coherently. Second, the necessary degree of coherence in the signal arrival times can be attained during ontogenetic development by virtue of an unsupervised Hebbian learning rule. It singles out the 'right' connections with matching delays from a broad distribution of axons with random delays. Third, the learning rule also selects the correct delays from two independent groups of inputs, for example, from the left and right ear.

1 Introduction

Until recently an unresolved paradox existed in auditory and electrosensory neural systems [1, 2]: they encode behaviorally relevant signals in the range of a few microseconds with neurons whose time constants are at least one order of magnitude bigger. In this paper we will explain how the paradox can be resolved. We take the barn owl's auditory system as an example and present results of a modeling study [3] of a neuron in the laminar nucleus that combines inputs from both ears:

(i) A single neuron can be quite a good coincidence detector despite its receiving stochastic input. The reason is that it is driven by *many* presynaptic signals that arrive more or less coherently and that these signals are, to a decent approximation, stochastically independent of each other.

(ii) Coherence presupposes the 'right' underlying hardware which operates with a μs precision. A simple proposal were to generate the hardware by genetic coding but that is hard to believe since thousands of axons require a immense amount of genes that were not able to adapt themselves to fluctuating circumstances such as food. In our opinion, the way out is provided by an unsupervised Hebbian learning rule that selects synapses and, hence, axons with the 'right' delays and suppresses the rest. This is – in a sense – an 'evolutionary' process with a 'survival of the fittest'.

(iii) Combining two groups of input, say from the left and the right ear, one gets the very same tuning. Evaluating the output of many laminar neurons through a population code one arrives at the final μs precision.

In the following we will implement and illustrate these ideas. We do not want to bother the reader by technical details, for which (s)he should consult the literature, in particular [3]. Instead we would like to focus on the essentials and maybe stick to half of the truth, if it leads to an intuitive understanding of the whole truth.

2 Temporal Coding

The importance of temporal coding in neural information processing has been debated vigorously [3]. One of the central issues in this debate is the question of whether neuronal firing can be more precise than the time constants of the neuronal processes involved [6]. We address this problem and study sound localization as it is performed by the barn owl's auditory system.

Barn owls are night hunters, despite the fact that their vision is excellent. Not so good, though, that they can see at night. The barn owl therefore needs a substitute, viz., sound localization. Its typical prey is the field mouse, which is fairly nourishing so that, under normal circumstances, five mice per night suffice. In spring, however, the masculine animal has to feed his offspring, numbering six to seven. That is to say, it has to catch a mouse every ten minutes – and it does so.

When a mouse moves through the grass it produces a broadband spectrum and the barn owl samples the sound between 2 and 8 kHz so as to localize the mouse. Why kiloherz (10^3 Hz)? The speed of sound v is about 330 meters per second (ms^{-1}). Denoting the wave length by λ and the frequency by f, we have $v = f\lambda$ so that for 6 kHz λ amounts to about 5 cm, the size of the mouse. Sitting in a tree, the barn owl has to perform both horizontal and vertical localization so as to fix the direction of its prey. Vertical localization is done by measuring the *intensity difference* between left and right ear. To this end the ears are positioned somewhat asymmetrically and one is 'looking' downwards, the other upwards. For horizontal or azimuthal sound localization, the object of our study, the owl determines the *interaural time difference* or, in short, the ITD.

A barn owl can locate sounds and, hence, prey in a horizontal plane with a precision of 1-2 degrees, a capability which requires a time resolution of less than 5 μs. Figure 1 depicts an interaural time difference and Fig. 2 is an early explanation due to Jeffress [7] of how a neuron could represent a specific direction by being a meeting point of two inputs, one coming from the right, the other one from the left. This idea exploits in an essential way that *single* spikes carry information.

After passing through the cochlea, a kind of 'inverse' piano[1], sound information splits into two pathways, intensity and ITD. The latter continues via the magnocellular nucleus, which produces a phase-locked output (plus a stochastic jitter), and then proceeds to the laminar nucleus, where signals from both ears meet for the first time. The key process in the auditory system is 'phase locking': spikes occur preferentially at a certain phase, termed 'mean phase', of a stimulating tone of frequencies up to 8 kHz. In order to understand the high degree of temporal precision in phase locking, we have to consider both the typical duration τ_s of synaptic input current and the membrane time constant τ_m. For auditory neurons in chicken [10], time constants of synaptic input are already in the range of 200 μs. Even though the passive membrane time constant is about 2 ms, the *effective* membrane time constant of magnocellular and laminar neurons is less than 200 μs because there is an outward rectifying current which is activated above resting potential. This kind of outward rectifying current is commonly found in phase-locking neurons of the auditory system. As a result of the short time constants, the width of single excitatory postsynaptic potentials (EPSPs) in chicken is about 500 - 800 μs at half of the maximum amplitude [10]. We think it is reasonable to assume that laminar neurons in an auditory specialist such as the barn owl are twice as fast as those in chicken and model EPSPs by taking a width of 250 μs.

Let $\mathcal{V}$ be the membrane voltage at the soma of a laminar neuron. When $\mathcal{V}$ reaches the threshold ϑ, a neuron fires. As for the neuronal dynamics, one has to integrate the differential equation $d\mathcal{V}/dt = -\mathcal{V}/\tau_m + I(t)$ with an *effective* membrane time constant $\tau_m =$ 100 μs. An input spike arriving at time t_j^f at a synapse j evokes an input current $I_j^f(t) =$

[1]A piano is an instrument where by pushing a specific key one produces a specific tone. Conversely, by 'inverse piano' we mean a fictitious instrument that 'pushes' the different keys according to the various tones that come in. In this way one obtains a spectral decomposition of the incoming sound.

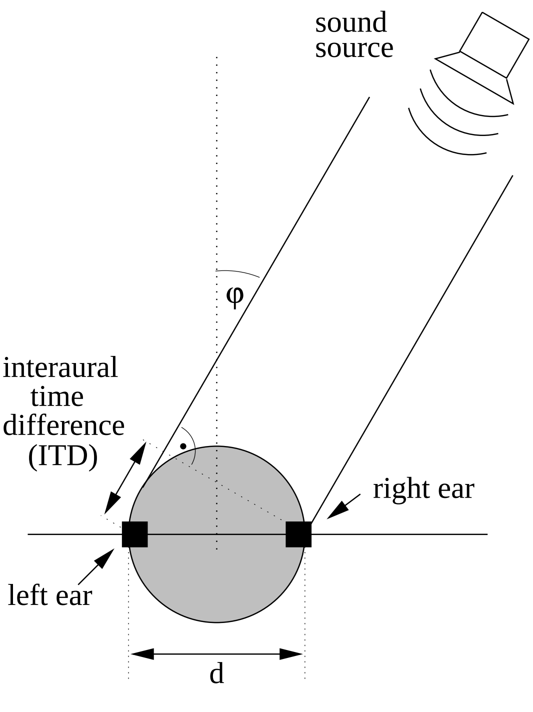

Figure 1: Interaural time difference. Sound coming from a source to the right of the head arrives at the left and right ear at different times. The interaural time difference (ITD) is determined by the azimuthal angle φ, the head's diameter $d \approx 4.5\,$cm and the speed $c \approx$ 331 m/s of sound. For small angles φ we obtain ITD $= \varphi\, d/c$. The barn owl's angular accuracy $\Delta\varphi \approx 1 - 2°$ of sound localization corresponds to a temporal precision ΔITD $= \Delta\varphi\, d/c < 5\,\mu$s.

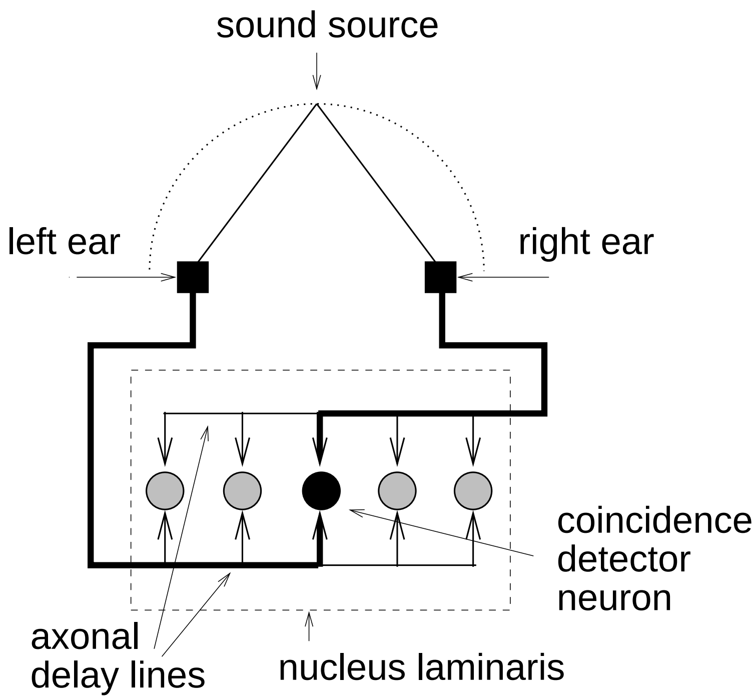

Figure 2:
Jeffress model [7] suggesting a neural circuit in the barn owl's nucleus laminaris (NL) for processing interaural time differences (ITDs). By means of such a circuit an ITD is transformed into a characteristic spatial activity pattern of the neurons (horizontal row of grey/black disks). Starting from the sound source the sound reaches both ears with a azimuth-dependent delay. In the ears the sound is transformed into action potentials (or spikes). These travel along axons that serve as delay lines to the NL. Axons from the left and the right ears converge in the NL at neurons which act as coincidence detectors. If spikes arrive at a coincidence detector neuron "simultaneously", the neuron shows up high activity (dark disk). Temporally displaced inputs cause a weaker response (brighter disks). A comparison of both figures shows that coincident input requires a compensation of a shorter delay from the sound source to one ear by a shorter axonal delay from the other ear. Depending on the azimuthal position of the sound source, there is a characteristic spatial pattern of neuronal activity, which is shown above and can be processed by "higher" brain areas. Though the Jeffress scheme is a nice idea, it has - as most cybernetic schemes – a few drawbacks: (i) there is not a *single* delay line but there are hundreds (altogether many thousands) of delay lines, and (ii) one still has to show that the hardware does what the scheme wants it to do. We will face both problems.

$(1/\tau_s)\,\exp[-(t-t_j^f)/\tau_s]$ for $t > t_j^f$ with $\tau_s = 100\ \mu$s. The total input is $I(t) = \sum_{j,f} J_j I_j^f(t)$ where J_j is the efficacy of synapse j and the sum runs over all synapses and firing times preceding t. For each input channel j, firing times are generated with a periodic probability density $P(t) = (\nu T/\sigma\sqrt{2\pi})\sum_{m=-\infty}^{\infty}\exp[-(t-mT-\Delta_j)^2/2\sigma^2]$ where $\nu = 1$ kHz is a rate, $T = 200\ \mu$s is the period of one of the tones of the stimulus[2], Δ_j is the transmission delay from the ear to a neuron in the laminar nucleus, and $\sigma = 40\ \mu$s is a measure of the temporal jitter. We impose absolute refractoriness through a minimal interspike interval of 0.5 ms.

If the transmission delays Δ_j are drawn from the broad, *mono*modal distribution of Fig. 3a, the one the young barn owl starts with after hatching, then spike arrival is incoherent. If on the other hand the delay distribution is that of Fig. 3c, i.e., a *multi*modal one with the

[2]The cochlea operates like an 'inverse piano' and we now consider one of its 'keys'.

maxima at distances T, then spike arrival is coherent. To fix the effective time constant τ_m of the integrate-and-fire model, we started from a more detailed model [3] with an explicit description of the outward rectifying current [10].

We concentrate on a single frequency channel. The stimulus is a pure tone of, e.g., 5kHz (period $T = 200\ \mu$s). The input from magnocellular neurons to our laminar integrate-and-fire neuron is a more or less periodic sequence of spikes. Presynaptic spikes occur preferentially around a specific mean phase of the tone [9]. The jitter of $\sigma = 40\mu$s takes care of internal sources of variability and noise along the auditory pathway as well as of the bandwidth of frequency tuning of auditory neurons. The mean firing rate is relatively low (667 Hz) so that each presynaptic spike train skips most of the cycles of the 5 kHz tone.

EPSPs of spikes from more than 100 presynaptic neurons that arrive coherently with a common mean phase lead to an oscillating membrane potential $\mathcal{V}$, that builds up, reaches the threshold in the *rising* stage of the cycle and in this way produces a slightly better phase locking than that of the input. The upshot turns out to be (see next section) that a single neuron can reach a precision of about 25 μs, which is not bad at all in view of the time constants involved.

If presynaptic spikes arrive with random phases, however, $\mathcal{V}$ shows aperiodic fluctuations and the output spikes have a *uniform* phase distribution (Fig. 3a). We expect this to be the initial, embryonic, condition before the connections between the magnocellular and the laminar nucleus are tuned during a sensitive period [11]. In adult owls, transmission delays from the ear to the laminar nucleus differ greatly. Their mean value is between 2 and 3 ms, and the standard deviation is about 200–240 μs. Without a tuning of the delays, any phase information would be lost. A Hebbian learning rule [3] proves to be an efficient tuning mechanism.

3 Hebbian Learning

The setup of the previous section implicitly assumed that phase-locked spikes (+ some jitter) impinge upon a laminar neuron. Three weeks after hatching, however, there seems to be a broad continuous distribution of delays and, thus, no phase locking. Genetic coding as a way out is not plausible since there are many thousands of axons that have to be tuned and genetics cannot adapt to a changing environment. How, then, does a barn owl get the right 'hardware'? We suggest by Hebbian learning or, sloppily formulated, "practice makes perfect".

In Hebbian learning, synaptic strength is changed by a small amount ϵ, if presynaptic spike arrival and postsynaptic firing 'coincide'. This simultaneity constraint is implemented by a learning window $W(s)$ where $s = t_j^f - t_n$ is the difference between the arrival time of a presynaptic spike t_j^f at synapse j and the postsynaptic firing t_n. In our model, $W(s)$ has two regimes; see Fig. 3d and [13]. For $s < 0$, $W(s)$ is positive. Thus the efficacy of synapses which are repeatedly active shortly *before* a postsynaptic spike occurs is increased [14]. The weights of synapses that are active shortly *after* the postsynaptic spike originated are *de*creased [15]. (Our theory was earlier, though, than its experimental verification :-) Since depolarization is known to induce potentiation of active synapses [14] and the neuron

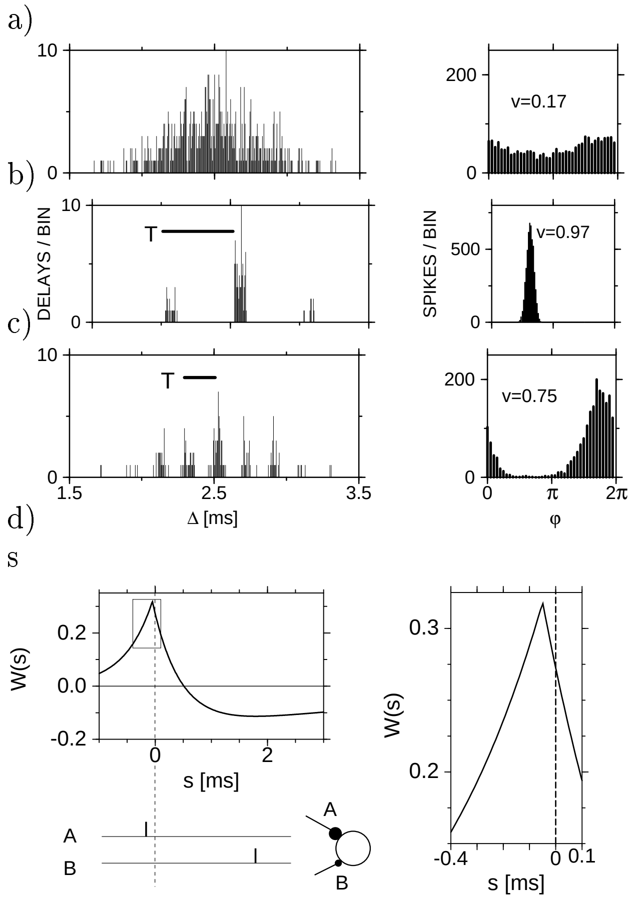

Figure 3: Hebbian learning. In *a*–*c*, the left-hand graphs show synapses binned as a function of the signal transmission delay Δ. On the right-hand side the distribution of output phases is shown in period histograms (bin width 5 μs).
a, Before learning, there are 600 synapses with signal transmission delays (left) drawn from a broad Gaussian distribution (mean 2.5 ms, width 0.6 ms). The output of the laminar model neuron shows *no* phase locking to a 2 kHz (not shown) or a 5 kHz signal (right).
b, After a period of learning while being stimulated by a 2 kHz signal. The 105 synapses which survive learning have delays which differ by *multiples* of the period $T = 500$ μs (scaling bar). The output spikes exhibit phase locking with vector strength [12] $v = 0.97$ corresponding to a temporal precision of 20 μs (right).
c, Same plot as in *b*, after learning a 5 kHz input signal. 154 saturated synapses survive. The output spikes exhibit phase locking ($v = 0.75$) corresponding to a temporal precision of 25 μs. Neither this nor the 20 μs under *b* live up to the required 5 μs; see Section 4.
d, Learning window W, the key notion of the present paper, as a function of the delay $s = t_j^f - t_n$ between a postsynaptic firing time t_n and presynaptic spike arrival t_j^f at synapse j. The graph on the right-hand side shows the boxed region around the maximum on an expanded scale. If $W(s)$ is positive (negative) for some s, the synaptic efficacy is increased (decreased). The postsynaptic firing occurs at $s = 0$ (vertical dashed line). Learning is most efficient, if presynaptic spikes arrive shortly *before* the postsynaptic neuron starts firing as in the 'fat' synapse A. The 'meager' synapse B, which fires *after* the postsynaptic spike originated, is weakened.
Before learning, all synapses have identical weights $J_j = 1$. During learning weights are changed according to $\Delta J_j = \epsilon \sum_f [\gamma + \sum_n W(t_j^f - t^n)]$ with factors $\epsilon = 0.002, \gamma = 0.1$. The sum runs over all spike arrival times $t_j^f < t$ and all postsynaptic firings $t^n < t$. We take $W(s) = 0.3 \exp[(s + 0.05)/0.5]$ for $s < -0.05$ and $W(s) = 0.5 \exp[-(s + 0.05)/0.5] - 0.2 \exp[-(s + 0.05)/5]$ for $s \geq -0.05$ where s is a time in ms. Learning occurs during 3000 seconds in time steps of 5 μs. The synaptic strength saturates at a maximum of 3. A synapse whose weight vanishes is removed. The efficiency of phase locking is quantified by the vector strength. Spike input and electrical time constants are as described in the main text.

is depolarized most of the time between two spikes, we also strengthen each active synapse by a small amount γ, even if no postsynaptic spike occurs. The procedure of continuously strengthening and weakening the synapses automatically (sic!) leads to a normalization of the total input strength feeding into a postsynaptic neuron in a competitive self-organized process.

As a result of learning, a clearly structured distribution of synapses evolves (Figs. 3*b*,*c*).

Delays of the remaining synapses differ approximately by multiples of the period T of the stimulating tone. We emphasize that the period was not previously known to the cell. As is shown by Figs. 3*b* and *c*, the final synaptic pattern of a neuron stimulated by a 2 kHz signal during learning is different from that of a 5 kHz neuron.

The evolution of synaptic weights during learning does *not* depend on the specific shape of the learning window W but only on some generic properties. In particular, the learning window can be nonzero in a domain which extends over several milliseconds (cf. [15]) and therefore is large as compared to the period T of the sound (Fig. 3*d*). Efficient learning depends on the temporal relationship between the process of strengthening the synaptic weights and that of weakening them. The maximum of the window function $W(s)$ should be at $\tilde{s} \approx -\tau_s/2$ where τ_s is the rise time of the postsynaptic potential (Fig. 3*d*). This can be understood through a self-consistency argument. Let us assume that learning has led to a sharply peaked distribution with all synapses having a common delay Δ (mod T). Coherent input arriving at the synapses can trigger a postsynaptic spike with a mean delay of approximately $\tau_s/2$. Because of learning, all synapses which are active slightly *before* the postsynaptic firing will be strengthened. If $\tilde{s} = -\tau_s/2$, the maximal increase of synaptic weights occurs at those synapses which are already strongest.

So far we have considered monaural input, but laminar neurons receive input from both ears and exhibit a sensitivity to the *interaural* time difference (ITD). We divide the synapses into two groups, viz., with input from the left and the right ear. During learning, both ears are stimulated by the same signal and with a fixed ITD. The learning rule selects synaptic connections so that spikes arrive coherently for exactly this ITD. If the same ITD is used *after* learning, the neuron is driven optimally and emits phase-locked spikes. If the ITD does not match, phase locking of the output spikes breaks down and the mean firing rate decreases; see Figs. 4*a* and *b*.

4 Population Coding

Temporal information conveyed by a *single* laminar neuron is limited. The temporal precision of phase locking is about 20-25 μs (Fig. 3) and the ITD tuning curve is only weakly modulated (Fig. 4*a*). Nevertheless, barn owls achieve a behavioral performance corresponding to a temporal precision of 5 μs. How can we understand this remarkable accuracy? The barn owl needs a reaction time of about 100 ms before it starts a head movement [8]. During 100 ms, the firing pattern of a *population* of laminar neurons contains enough information to resolve ITDs with a precision of 5 μs, as we will now show.

Let us consider a *group* of laminar neurons with ITD tuning curves as in Fig. 4*a*, but shifted along the ITD axis so that the optimal responses occur at different ITDs. Neurons are stimulated by a tone with fixed, but unknown, ITD. We estimate the ITD from the neuronal firing pattern by a 'population vector' decoding scheme [16]: Each neuron 'votes' for the ITD corresponding to its optimal response and the votes are weighted according to the number of spikes that occur in a time window of 100 ms. We find that the activity of about 100 independent neurons provides enough information to estimate the ITD with a precision of 5 μs (Fig. 5) – apart from ambiguities that are due to the periodicity of the

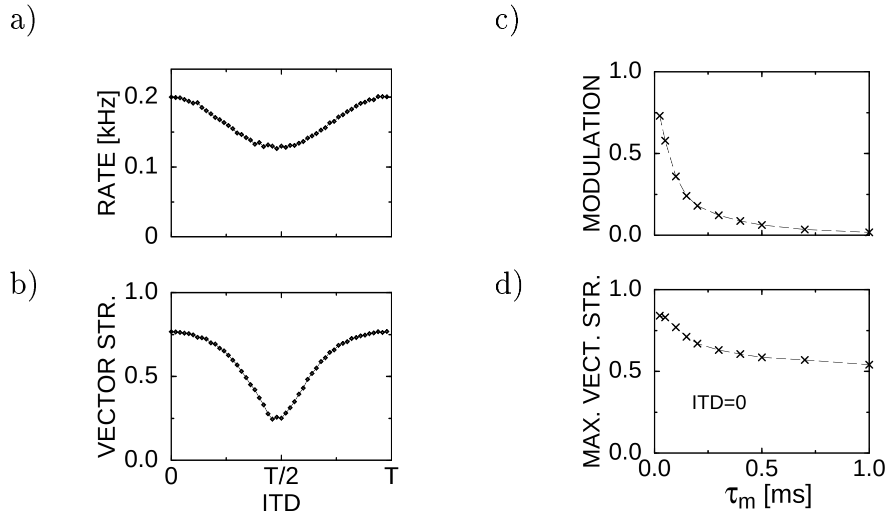

Figure 4: *a* and *b*. Conversion of spike code (input) into rate code (output) through tuning to interaural time difference (ITD). The output rate (*a*) and the vector strength (*b*) of a laminar neuron are shown as a function of the ITD. The neuron has been tuned to a 5 kHz signal ($T = 200\ \mu$s) as described in the text and the legend of Fig. 3. Half of the 154 synapses which survive learning receive signals from the left ear, the others from the right ear. The neuron exhibits best phase locking and maximal output rate $f_{\max}$ for the ITD $= 0$ used during learning. The rate has a minimum $f_{\min}$ for ITD $= T/2$. *c* and *d*. Effect of τ_m. We define the modulation to be $(f_{\max} - f_{\min})/f_{\max}$ and show it as a function of the effective membrane time constant τ_m in *c*. The vector strength at ITD $= 0$ is shown as a function of τ_m in *d*.

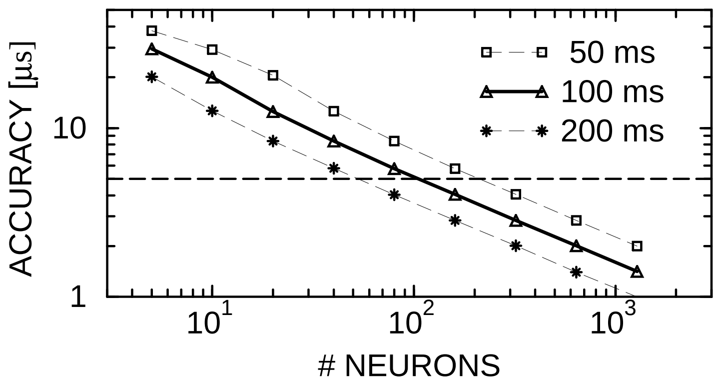

Figure 5: What population coding is good for. The temporal precision that can be achieved in a population of N independent neurons is shown as a function of N under the assumption of population vector coding [16, 19]. The ITD is estimated based on spikes occurring in a time window of $\tau = 50$ ms (squares), 100 ms (triangles), or 200 ms (stars). The accuracy is defined as the average deviation of the estimate from the actual ITD of the stimulus. The horizontal line at 5 μs is the boundary given by the behavioral performance of 1-2 degree of azimuthal angle [8]. For large N, the accuracy scales as $1/\sqrt{\tau N}$ as expected from the central limit theorem.

signal. Weak correlations from common inputs rescale the absolute values but do not alter the conclusions.

In Fig. 5 we took N neurons labeled by $1 \leq k \leq N$ with ITD tuning curves similar to Fig. 4a and approximated by $f_k(x) = \{164 + 36\cos[(x - x_k)2\pi/T]\}$ Hz where x is the

ITD and T the period of the stimulating tone, x_k is the optimal ITD of neuron k, and f_k its mean firing rate in Hz. The values of x_k have a uniform distribution between 0 and T. Spikes of neuron k are generated by a Poisson process with mean rate f_k. In an additional simulation (data not shown) we have confirmed that Poisson statistics give a reasonable assumption and the results of Fig. 5 do not change, if we generate spikes by an integrate-and-fire neuron as described above. The ITD of the stimulus is estimated [16] to be $x^{\text{est}} = (T/2\pi) \arg[\sum_k n_k \, e^{i2\pi x_k/T}]$ where n_k is the number of spikes of neuron k measured in a time window of length τ. We plot the accuracy $\langle |x^{\text{est}} - x|^2 \rangle^{1/2}$ averaged over 20,000 trials. We see that 100μs averaging time needs about 100 neurons as input to get a precision of 5 μs. In reality one finds about 150 neurons, leaving room for additional correlations among the neurons.

The firing precision of 20-25 μs of *single* neurons has been achieved despite the fact that input spikes exhibit a jitter and evoke EPSPs which are ten times broader than the precision just mentioned. In an additional set of simulations, we have systematically varied the effective membrane time constant τ_m and, hence, the width of the EPSP. We find that the temporal precision of neuronal spiking depends only weakly on τ_m (Fig. 4*d*). This is possible because firing always occurs during the *rise* time of a postsynaptic potential. On the other hand, for a 5 kHz signal, ITD tuning breaks down rapidly, if τ_m exceeds 0.1 ms (Fig. 4*c*). Thus modulation of ITD tuning does indeed require short EPSPs. In fact, very short time constants have been measured by Reyes et al. [10]. With the benefit of hindsight we now also see that the stimulating frequency cannot be arbitrarily high. In practice, the barn owl does not go beyond 10 kHz – completely consistent with the above estimates.

5 Discussion

Our results show that the temporal precision of output spikes does not have the membrane time constant itself as an upper bound. Spike timing can achieve a resolution shorter than the rise time of an EPSP, given *coherent* spike arrival – and spike arrival times can be tuned accurately by time-resolved Hebbian learning. The key notion of our theory is the *learning window* describing how a synapse is to change if the presynaptic spike arrives before/after the postsynaptic one and specifying what time scales are involved. It is quite natural that W is positive/negative depending on the presynaptic spike coming too early/too late. If it is slightly earlier, it favors the firing of the postsynaptic neuron so that the synapse is to be strengthened. Coming too late, that's no good and the synapse is to be weakened. This way the mechanism leads to a stable equilibrium. Instead of weeding out synapses by means of W one could think of esoteric mechanisms and allow the synapses to speed up/slow down or the axons to shrink/grow depending on the presynaptic spikes being too late/early. New phenomena do not occur.

The barn owl's neurons reaching a μs precision one may wonder: Why is Nature that good? In the case of a laminar neuron, because of three reasons, all simple biology. First, the neuron receives input through many synapses so that noise is averaged out. Second, it uses a time-resolved Hebbian learning rule singling out the 'right' synapses. In a sense, this is nothing but evolution on a microscale, the 'right' synapses representing the 'survival of the fittest'. Third, combining many neurons performing the same calculation one attains a population code, which is better than that of the individual neurons.

It is tempting to apply the same ideas to information processing in the cerebellum [17], the hippocampus [4], or the cortex [5]. If we multiply time constants by a factor of about 100 and reinterpret our results, then we are led to the conclusion that in areas where effective membrane time constants are in the range order of 10 – 20 ms [18], a temporal code with an accuracy of 1–3 ms should be, and seems to be, possible.

Acknowledgments

R.K. has been supported by the DFG (He 1729/8-1).

References

[1] Carr, C.E. *Ann. Rev. Neurosci.* **16**, 223-243 (1993)

[2] Heiligenberg, W. *Neural Nets in Electric Fish*, Cambridge, Mass. (1991)

[3] Gerstner, W., Kempter, R., van Hemmen, J.L., and Wagner, H. *Nature* **383**, 76–78 (1996) and references quoted therein.

[4] Hopfield, J.J. *Nature* **376**, 33–36 (1995)

[5] Abeles, M., in *Models of Neural Networks II*, Eds. E. Domany, J.L. van Hemmen & K. Schulten (Springer, New York, 1994) 121–140

[6] Softky, W. & Koch, C. *J. Neurosci.* **13**, 334–250 (1993)

[7] Jeffress, L.A. *J. Comp. Physiol. Psychol.* **41**, 35–39 (1948)

[8] Knudsen, E.I., Blasdel G.G. & Konishi, M. *J. Comp. Physiol.* **133**, 1–11 (1979)

[9] Sullivan, W.E. & Konishi, M. *J. Neurosci.* **4**, 1787–1799 (1984)

[10] Reyes, A.D., Rubel, E.W. & Spain, W.J. *J. Neurosci.* **16**, 993–1007 (1996) and **14**, 5352–5364 (1994)

[11] Carr, C.E. in *Advances in Hearing Research*, Eds. G.A. Manley, G.M. Klump, C. Köppl, H. Fastl & H. Oeckinghaus (World Scientific, Singapore, 1995) 24–30

[12] Goldberg, J.M. & Brown, P.B. *J. Neurophysiol.* **32**, 613–636 (1969)

[13] Van Hemmen, J.L., Gerstner, W., Herz, A.V.M., Kühn, R. & Vaas, M. in *Konnektionismus in Artificial Intelligence und Kognitionsforschung*, Ed. G. Dorffner (Springer, Berlin, 1990) 153–162

[14] Bliss, T.V.P. & Collingridge, G.L. *Nature* **361**, 31–39 (1993)

[15] Markram, H., Lübke, J., Frotscher, M. & Sakmann, B. *Science* **275**, 213–215 (1997)

[16] Salinas, E. & Abbott, L.F. *J. Comput. Neurosci.* **1**, 89–107 (1994)

[17] Braitenberg, V. *Network* **4**, 11–17 (1993)

[18] Bernander, Ö., Douglas, R.J., Martin K.A.C. & Koch C. *Proc. Natl. Acad. Sci. USA* **88**, 11569–11571 (1991)

[19] Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. *Science* **233**, 1416-1419 (1986).