In D. Reitter & F. E. Ritter (Eds.), Proceedings of the 14th International Conference on Cognitive Modeling (ICCM 2016). University Park, PA: Penn State.

Spiking Neural Model of Supervised Learning in the Auditory Localization Pathway of Barn Owls

Michael O. Vertolli (michaelvertolli@gmail.com)[1]

Terrence C. Stewart (tcstewar@uwaterloo.ca)[2]

[1]Institute of Cognitive Science, Carleton University, 1125 Colonel By Drive, Ottawa, ON K1S 5B6, Canada
[2]Center for Theoretical Neuroscience, University of Waterloo Waterloo, ON N2J 3G1, Canada

Abstract

We propose a large-scale system, with minimal global topological structure, no local internal structure, and a simple online biologically plausible local learning rule that captures supervised learning in the barn owl. We outline how our computational model corresponds to both the underlying neuroscience and the experimental paradigm used in the relevant prism studies of the barn owl. We show that our model is able to capture the basic outcomes of this experimental research despite learning the initial tuning curves, which is not done in other computational models, and a much more restricted time frame relative to the original experimental condition. We outline some variations between our model and the neuroscience outcomes and suggest future extensions in terms of larger models and time frames, more detailed analyses of the learning parameters, and richer model designs.

Keywords: supervised learning, spiking neural network, auditory localization, barn owl

Introduction

Supervised learning in the brain is generally viewed as a learning episode where one neuron ensemble acts as an instructive signal that modulates connectivity and activity in another neuron ensemble (Knudsen, 1994). The auditory localization pathway in the brain has been used as a model system for studying how the brain performs supervised learning both because experience can alter auditory localization and because the pathways are relatively well studied (Knudsen, 2002). The species that has been studied most extensively in this context is the barn owl (*Tyto alba*; Knudsen, Blasdel, & Konishi, 1979).

There are a number of computational models that have been used to compliment the neuroscience research on supervised learning in the barn owl auditory localization pathway (D'Souza, Liu, & Hahnloser, 2010; Fischer, Anderson, & Peña, 2009; Huo & Murray, 2009; Huo, Murray, & Wei, 2012; Witten, Knudsen, & Sompolinsky, 2008). All of these models are small in scale (under 100 neurons), use highly structured systems that are determined by the modelers, and a rather complex learning rule (usually spiking-time dependent plasticity). We propose a larger-scale (and scalable) system, with minimal global topological structure, no local internal structure, and a simple online biologically plausible local learning rule that captures supervised learning in the barn owl.

Neuroscience Background

The optic tectum of the barn owl (analogous to the superior colliculus in humans) hosts the neurophysiological associations between the auditory localization cues and locations in visual space (Knudsen, 2002). Locations in the optic tectum respond maximally to auditory or visual stimuli located at a specific region of space or the receptive field. The associated neurons are tuned to auditory localization cues that correspond to visual field locations (Brainard, Knudsen, & Esterly, 1992).

Audio-visual pairings in the optic tectum are learned through early experience. Studies have shown that exposing juvenile barn owls to prismatic spectacles that displace the visual field horizontally (most commonly 23°) results in a learned displacement in the tuning curves of the corresponding auditory neurons over a number of months (Knudsen, 1985). Learned changes primarily occur in the efferent connections of the external nucleus of the inferior colliculus (ICx): an earlier auditory processing step in the auditory localization pathway that directly connects to the optic tectum (Brainard & Knudsen, 1993). These changes occur on the basis of an error signal projected back from the corrective visual input in the deeper layers of the optic tectum (Peña & DeBello, 2010). It is this re-tuning phenomenon of the auditory ICx neurons on the basis of visual input in the deeper layers of the optic tectum that the proposed model will capture.

Computational Background

We use the Neural Engineering Framework (NEF; Eliasmith & Anderson, 2004). The NEF is used to represent, transform, and add dynamics to vectors of numbers via populations of spiking neurons, their synapses, and recurrence connections in the network.

NEF representations are an n-dimensional extension of the population coding work of Georgopoulos, Schwartz, and Kettner (1986). Neurons in the population are described in terms of an encoder and decoder that translate neural activity (a filtered spike train) to and from the vector space. The activity of a neuron can be expressed as:

$$a = G[\alpha \mathbf{e} \cdot \mathbf{x} + J_{bias}], \quad (1)$$

where G is the activation function, α is a scaling factor, $\mathbf{e}$

is the encoder, $\mathbf{x}$ is the vector to be encoded, and J_{bias} is the background current. The decoded estimate $\hat{\mathbf{x}}$ is described by the following equation:

$$\hat{\mathbf{x}}(t) = \sum_i \mathbf{d}_i a_i(t), \tag{2}$$

where $\mathbf{d}_i$ is the decoder and a_i is the activity of neuron i. Neural activity overall is computed as a spike train:

$$a_i(t) = \sum_s H(e^{-(t-t_s)/\tau_{PSC}}), \tag{3}$$

where H is the Heaviside step function, s is the set of all spikes occurring before the current time, and τ_{PSC} is the post-synaptic time constant for the connection (a neural property). The decoders are computed through a least-squares minimization of the difference between the decoded estimate and the encoded vector:

$$\begin{aligned} \mathbf{d} &= \Upsilon^{-1}\Gamma \\ \Gamma_{ij} &= \int a_i a_j dx \\ \Upsilon &= \int a_j \mathbf{x} dx \end{aligned} \tag{4}$$

We use the Prescribed Error Sensitivity (PES) supervised learning rule described in MacNeil and Eliasmith (2011) that performs the described least-squares minimization online.

$$\Delta\omega_{ij} = \kappa\alpha_j \mathbf{e}_j \cdot \mathbf{E} a_i, \tag{5}$$

where ω_{ij} is the connection weights between the ith and jth neurons, κ is a scalar learning rate, and $\mathbf{E}$ is the error vector that will be minimized. The other symbols are consistent with the previous formulas. This rule is analogous to classic perceptron delta rule with the exception that only a portion of the error signal that each neuron is sensitive to is computed for a given neuron.

The Model

We use a highly simplified model of the auditory input, visual input, ICx, and optic tectum in order to capture supervised learning in the auditory localization pathway of the barn owl. Both the auditory and visual input are described in terms of a 180° arc on the unit circle in the horizontal plane from -1.0 to 1.0.[1] This arc was represented by two values in the network corresponding to the x and y values. Though the input could vary fully over this continuous space, assessment occurred in 4° increments as we describe later.

Three ensembles of neurons describe the basic model: the ICx, the shallow optic tectum, and the deep optic tectum. Each ensemble was comprised of 400 leaky-integrate-and-fire (LIF) neurons. This neuron model is widely used for its flexibility as an approximation of a broad range of neuron models (Koch, 2004). It also functions as a limiting case of more complex models like the Hodgkin-Huxley model (Partridge, 1966). To set the α and J_{bias} parameters (Equation 1) we randomly chose these values such that the resulting ideal tuning curves generated by the LIF neuron model would involve neurons that fired over a range of inputs between ±15° and ±35°.

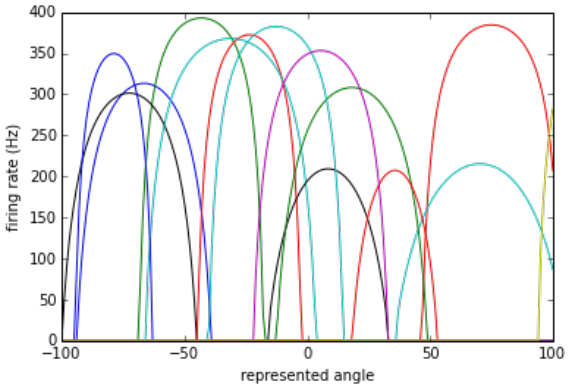

Figure 1: Ideal desired neuron tuning curves (used to generate neuron gain α and bias J_{bias} parameters).

The auditory input enters via the ICx, projects to the shallow optic tectum then to the deep optic tectum. The latter receives the visual input and projects the difference between the shallow optic tectum and the visual input as an error signal back to the ICx connections (see Fig. 2).

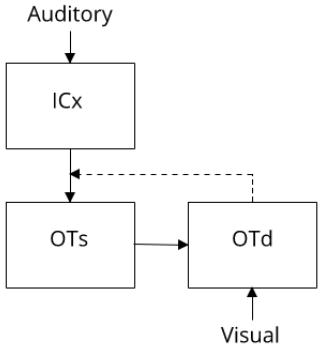

Figure 2: Basic structure of the model, where OTs is the shallow optic tectum, OTd is the deep optic tectum, and the dotted line is an error signal.

Unlike all other computational models of supervised learning in the barn owl, our model did not provide any internal structure. This means that the signal from the ICx to the shallow optic tectum originally had random weights (i.e., computed the function $f(x) = 0$). The system, therefore, had to learn the initial spatial representation prior to learning the modified representation (i.e., the introduction of the prism). It also lends a certain symmetry to the model: initial connections from the ICx to the optic tectum in the early life of the owl (i.e., before experimentation) have an identical learning

[1] This gets rotated to the left by 23° in the second training block.

procedure to subsequent learning. However, this does oversimplify some of the details in the underlying neuroscience (e.g., differences between neurotransmitters used in different learning situations).

In what follows, we describe the training and assessment procedure we used on the model.

Method

The methodological design of the study is based off of the work of Brainard and Knudsen (1998) and Linkenhoker and Knudsen (2002). We outline it, below.

The basic setup is comprised of 4 blocks: two training and two assessment. They occurred in alternating order with one of the training blocks first. Both training blocks lasted 200s (3.33min) of simulation time during which randomly selected stimuli along the 180° arc of the unit circle were presented for 500ms intervals. This resulted in 400 random stimuli being presented during each training block.

The assessment blocks determined the average activity of each neuron in 4° increments along the 180° arc of the unit circle. Each stimuli is presented for 50ms resulting in a total of 2.25s of simulation time across 45 stimuli. No learning occurred during each assessment block. The tuning curves were quantified in terms of the best tuning, which is calculated as the center of the range of values that elicited greater than 50% of the maximum response. The second assessment block used the same best tuning location as its zero point in order to see the difference across assessments.

As many more neurons were examined in the model than in the neuroscience study, we assumed that neurons with the same best tuning location were proximate to one another (consistent with their tonotopic arrangement in the ICx; Knudsen, 2002). Thus, each neuron was normalized relative to the maximum activity elicited by the pool of neurons with the same best tuning location. For example, if 4 neurons had the their best tuning (as described above) 16° left of the zero point, they would all be normalized relative to the maximum activity among them.

The second training block occurred with the prism signal present and a deviation in the visual input of 23° to the left. During the second assessment block the prism was removed, but no learning occurred to stabilize the original output. This was consistent with the original study.

Results

The results of the model were as expected. The model was able to accommodate the prism with appropriate adaptations in the ICx connections. We used the average interaural time difference (ITD; the standard measurement unit in the literature) to arc angle conversion of 2.5*u*s to 1° of the unit circle (Linkenhoker & Knudsen, 2002).

Neurons in the first assessment developed tuning curves approaching a normal curve when normalized and centered to their best tuning (see Fig. 3). The mean tuning curve was an even better approximation of normality (see Fig. 4). The maximum of the mean tuning curve had a score of 0.73.

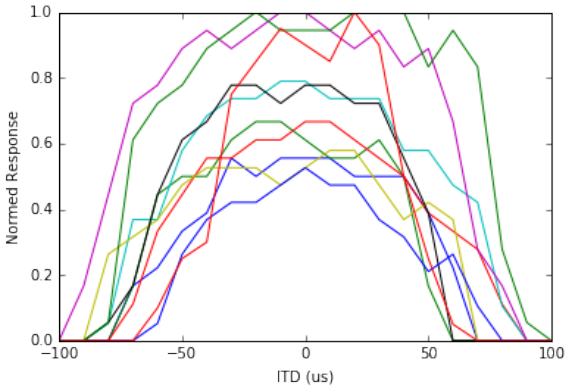

Figure 3: The learned tuning curves of 10 random neurons during the first assessment.

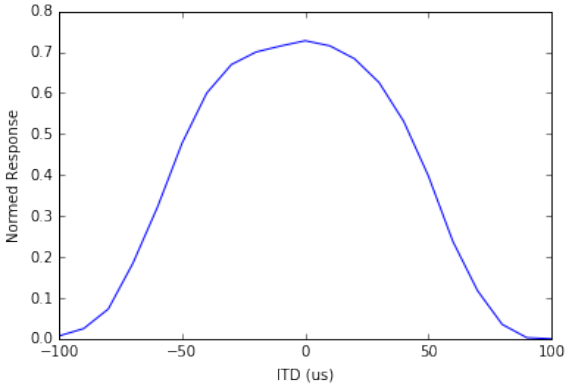

Figure 4: The mean learned tuning curve of 112 random neurons during the first assessment.

Neurons in the second assessment adapted to the prism condition by narrowing their range of activity slightly from the first training condition (see Fig. 6). The maximum activity score was 0.73. Figure 5 shows the neuron tuning curves relative to their own best tuning rather than the best tuning of the initial assessment (Fig. 7). The mean tuning curve of this population achieved a peak angle rotation of 50*u*s and 20° as a consequence of the prism (see Fig. 6).

Discussion

Recall that the goal was to model supervised learning in the localization pathway of barn owls using a simplified spiking neural network and learning rule. Given that the owls in the original empirical study had their initial training (birth until assessment) and re-training (prism adaptation) occur on the order of months instead of minutes, the results are very promising. The average tuning curve was more regular in shape than that expected by the original experiment (see Fig. 4 and 9, respectively). Nevertheless, the model's results are consistent with the corresponding empirical research on the whole.

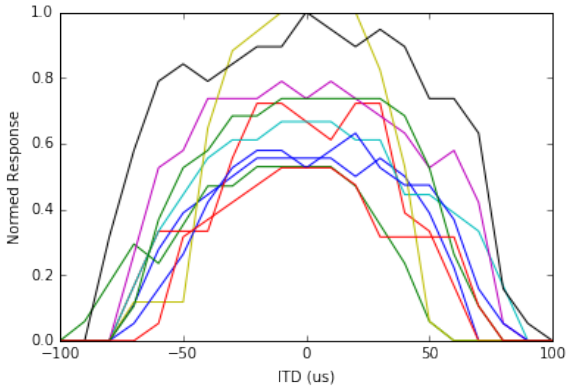

Figure 5: The learned tuning curves of 10 random neurons during the second assessment relative to their own best tuning curves (i.e., within the context of the prism).

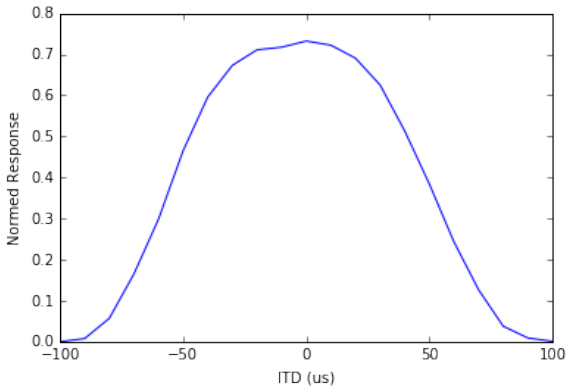

Figure 6: The learned tuning curves of 105 random neurons during the second assessment relative to their own best tuning curves (i.e., within the context of the prism).

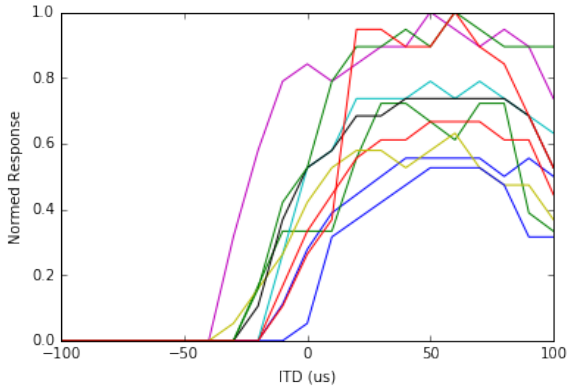

Figure 7: The learned tuning curves of 10 random neurons during the second assessment relative to the best tuning curves from the first assessment.

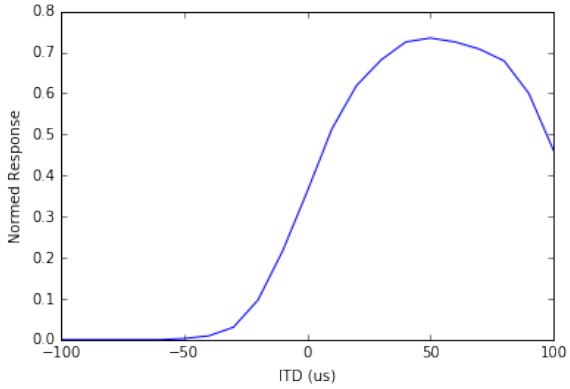

Figure 8: The learned tuning curves of the same 112 neurons from the first assessment during the second assessment relative to the best tuning curves from the first assessment.

In the original study, juvenile owls achieved a median rotation of 43*u*s (17°) with an ideal rotation of 57*u*s for a 23° prism deviation (Brainard & Knudsen, 1998; Linkenhoker & Knudsen, 2002). Our score of 50*u*s is placed in between these two values. We expect that this is a consequence of having our learning rate too high. Lower learning rates require longer time lines, which often require specialized hardware as the model's requirements increase both with the size of the neuron populations and the length of time that is simulated.

The current model is a preliminary prototype to determine whether it would be worthwhile to commit the computational effort to run this model at a much larger scale. Our results suggest that it is worth scaling up the model in future work. This would then allow us to more deeply explore effects of the learning rate as well as the learning trajectory during the prism condition.

The generalizability of this technique is much stronger than all of the other computational models. Many of the models, including Witten et al. (2008), require that the individual neurons are already sensitized to a given location and often only that location. This model learns to represent a distribution of values with a simple learning rule (PES). The system, based on its training input, manages to capture tuning curves that, on average, are very similar to what one would expect a real neuron to have. It is then able to accommodate the prism deviations from within this learned framework.

The current model uses a post-synaptic time constant in the scale of α-amino-3-hydroxy-5-methyl-4-isoxazole propionic acid (AMPA; 5ms), which is known to be a contributing factor for the learned signals (e.g., prior to the prism condition; Knudsen, 2002). Research suggests that there are at least two other neurotransmitters that contribute to supervised learning in the barn owl: *N*-methyl-*D*-aspartate (NMDA) and γ-aminobutyric acid type A ($GABA_A$). AMPA contributes to learning the new signals (during the prism condition) and $GABA_A$ regulates between the two. GABAergic inhibition also seems to have a gating-like effect on the learning process overall. It determines when the error signal is propagated

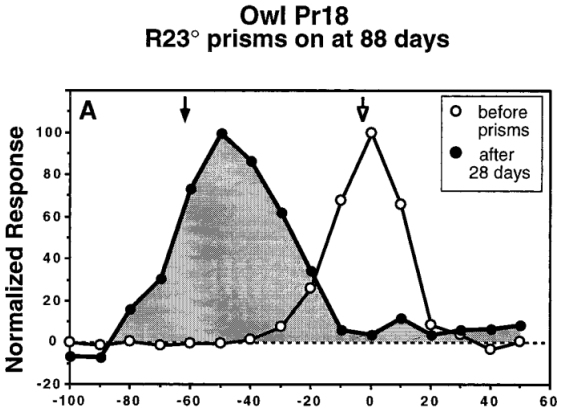

Figure 9: Original graph of the interaural time differences across the prism and no prism conditions, included with permission from Brainard and Knudsen (1998).

from the deeper layers of the optic tectum to the ICx (Peña & DeBello, 2010). We would like to begin incorporating these more detailed aspects into future instantiations of the model.

A number of additional extensions from the neuroscience literature are readily amenable to the model. The learning trajectories associated with larger time frames, in particular, show a number of interesting details. First, juvenile barn owls do not learn to adapt to the prism spectacles until after they are 60 days old, even if they are attached to them at as early as 13 days (Brainard & Knudsen, 1998). Our current research suggests that this property might be a consequence of changes in the initial tuning curves exemplified in Fig. 1. Second, after about 200 days learning tapers off such that effectively no learning occurs after about a year in the standard paradigm (Brainard & Knudsen, 1998; Linkenhoker & Knudsen, 2002). However, as a third detail, learning *is* actually possible in yearling owls and older if a special incremental learning paradigm is used (Linkenhoker & Knudsen, 2002). Adaptations in the learning parameters are particularly amenable to this situation, with a general trend towards higher tolerance learning rates over time. Using dedicated neural simulation hardware, we intend to explore these critical periods and paradigms at full-scale (i.e., day for day of simulation to real time).

Conclusion

Recall that the goal was to model a large-scale system, with minimal global topological structure, no local internal structure, and a simple biologically plausible online local learning rule that captures supervised learning in the barn owl. As a model organism for supervised learning, the barn owl lends itself to computational models of learning. We outlined in broad strokes how the deeper layers of the optic tectum motivates this learning in the neural projections from the external nucleus of the inferior colliculus (ICx) via discrepancies between auditory input and visual input. We outlined the NEF framework, and described how it can be used to map neural activity to vector representations. We also described a simple learning rule for our system (PES).

In subsequent sections, we outlined how our computational model corresponds to both the underlying neuroscience and the experimental paradigm used in the relevant prism studies on the barn owl. We showed that our model is able to capture the basic outcomes of this experimental research despite learning the initial tuning curves (the first training block) and a much more restricted time frame. We outlined some variations between our model and the neuroscience outcomes, mainly in terms of an overly strong learning rate. We then suggested some future extensions of the research in terms of larger models and time frames, more detailed analyses of the learning parameters, and more detailed model designs. Just as the barn owl is a (simple) model organism for the complexity of supervised learning in general, our model also functions as a simple but powerful computational instantiation of the complexity of the model organism.

References

Brainard, M. S., & Knudsen, E. I. (1993). Experience-dependent plasticity in the inferior colliculus: a site for visual calibration of the neural representation of auditory space in the barn owl. *The Journal of Neuroscience*, *13*(11), 4589–4608.

Brainard, M. S., & Knudsen, E. I. (1998). Sensitive periods for visual calibration of the auditory space map in the barn owl optic tectum. *The Journal of Neuroscience*, *18*(10), 3929–3942.

Brainard, M. S., Knudsen, E. I., & Esterly, S. D. (1992). Neural derivation of sound source location: resolution of spatial ambiguities in binaural cues. *The Journal of the Acoustical Society of America*, *91*(2), 1015–1027.

D'Souza, P., Liu, S.-C., & Hahnloser, R. H. (2010). Perceptron learning rule derived from spike-frequency adaptation and spike-time-dependent plasticity. *Proceedings of the National Academy of Sciences*, *107*(10), 4722–4727.

Eliasmith, C., & Anderson, C. H. (2004). *Neural engineering: Computation, representation, and dynamics in neurobiological systems*. MIT press.

Fischer, B. J., Anderson, C. H., & Peña, J. L. (2009). Multiplicative auditory spatial receptive fields created by a hierarchy of population codes. *PloS one*, *4*(11), e8015.

Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, *233*(4771), 1416–1419.

Huo, J., & Murray, A. (2009). The adaptation of visual and auditory integration in the barn owl superior colliculus with spike timing dependent plasticity. *Neural Networks*, *22*(7), 913–921.

Huo, J., Murray, A., & Wei, D. (2012). Adaptive visual and auditory map alignment in barn owl superior colliculus and its neuromorphic implementation. *Neural Networks*

and Learning Systems, IEEE Transactions on, *23*(9), 1486–1497.

Knudsen, E. I. (1985). Experience alters the spatial tuning of auditory units in the optic tectum during a sensitive period in the barn owl. *The Journal of Neuroscience*, *5*(11), 3094–3109.

Knudsen, E. I. (1994). Supervised learning in the brain. *Journal of Neuroscience*, *14*(7), 3985–3997.

Knudsen, E. I. (2002). Instructed learning in the auditory localization pathway of the barn owl. *Nature*, *417*(6886), 322–328.

Knudsen, E. I., Blasdel, G. G., & Konishi, M. (1979). Sound localization by the barn owl (tyto alba) measured with the search coil technique. *Journal of comparative physiology*, *133*(1), 1–11.

Koch, C. (2004). *Biophysics of computation: information processing in single neurons*. Oxford university press.

Linkenhoker, B. A., & Knudsen, E. I. (2002). Incremental training increases the plasticity of the auditory space map in adult barn owls. *Nature*, *419*(6904), 293–296.

MacNeil, D., & Eliasmith, C. (2011). Fine-tuning and the stability of recurrent neural networks. *PloS one*, *6*(9), e22885.

Partridge, L. D. (1966). A possible source of nerve signal distortion arising in pulse rate encoding of signals. *Journal of theoretical biology*, *11*(2), 257–281.

Peña, J. L., & DeBello, W. M. (2010). Auditory processing, plasticity, and learning in the barn owl. *ILAR journal*, *51*(4), 338–352.

Witten, I. B., Knudsen, E. I., & Sompolinsky, H. (2008). A hebbian learning rule mediates asymmetric plasticity in aligning sensory representations. *Journal of neurophysiology*, *100*(2), 1067–1079.