

Auditory Habituation via Spike-Timing Dependent Plasticity in Recurrent Neural Circuits

Richard Veale

Cognitive Science Program, Indiana University
Bloomington, Indiana, USA
Email: riveale@indiana.edu

Matthias Scheutz

Department of Computer Science, Tufts University
Medford, Massachusetts, USA
Email: mscheutz@cs.tufts.edu

Abstract—This paper presents evidence that spiking neuron models of parts of the human auditory system demonstrate habituation to real auditory word stimuli. This is accomplished via the simple addition of a model of spike-timing dependent plasticity to synapses. This result is interesting because the base neural circuit has also been used for pragmatically useful behaviors such as speech recognition. The model increases our understanding of the neural basis for behaviors in the developing human infant by showing that habituation learning can be implemented in the same neural substrate that underlies other types of learning (such as permanent word-learning).

I. INTRODUCTION

Human infants are capable of incredible feats of learning even before they are born. One of these behaviors is habituation to sensory stimuli. Habituation is a type of adaptive learning whereby an organism’s response to a sensory stimulus decreases with successive encounters. Habituation is widely used in behavioral experiments involving human infants to determine whether the infant can discriminate a familiar stimulus versus other stimuli. Despite the widespread use of habituation behavior, convincing biologically accurate neural models of habituation are few in number. Furthermore, they rarely operate on real-world stimuli. However, biologically accurate neural models are beneficial to understanding the developing human infant. Such models are particularly important for understanding how multiple functional behaviors might be implemented in simultaneously in the same neural substrate. Additionally, when such models are implemented in artificial autonomous systems, they can produce natural useful behavior and serve pragmatic purposes.

This paper presents a model as evidence that habituation can be successfully implemented using spike-timing dependent plasticity (STDP) in a recurrent circuit of spiking neurons. The neurons receive real raw sensory input streams and habituate to spoken words. The fact that the model produces habituation behaviors is even more exciting in light of the fact that the exact same circuit (minus STDP) has also been used as a robust speech-recognition system. Thus, it is not an isolated model that *only* reproduces habituation behavior. Rather, this paper is showing how an already-existing model that can perform an independently useful behavior (speech recognition), can *also* demonstrate habituation behaviors. This paper focuses on the details of the habituation behavior.

II. BACKGROUND

Auditory habituation behavior is observable from birth [1], and possibly even prenatally [2] in humans. The “marker” used to identify habituation varies dependent on age and species. For example, in human fetuses heart rate is the measure, whereas in newborns it is head-turning or sucking behavior. However, the pattern remains that it is possible to repeatedly present an auditory stimulus (the “familiar” stimulus), and observe a decrement in the organism’s reaction to that stimulus. Furthermore, when a sufficiently different (“novel”) stimulus is presented, the organisms will respond more than to the familiar stimulus. This “recovery” of response proves that the decrease in response to the repeatedly experienced familiar stimulus was specific to that stimulus. I.e., the decrement was not a general decrement in response to *all* stimuli (e.g. because of boredom from staring at a screen).

This behavioral description raises some questions. What qualifies being “sufficiently different” for a stimulus to be considered a novel type, versus a slightly permuted token of a familiar stimulus? In newborn studies, experimenters use easily quantifiable aspects of auditory stimuli to separate stimulus categories (e.g. square waves of drastically different frequencies, volumes, etc.). However, in older infants it is possible to use even spoken language syllables that differ only in one part of the sound (“dah” versus “gah”)[3]. This raises the interesting question of whether the same substrates are at play in behaviors such as learning words and in auditory habituation. Though information is incomplete, there is evidence that at least the early portions of the sensory processing system are shared between the two behaviors – i.e. there are not two entirely parallel systems which are separate starting from the ears. The evidence for this comes from neuroanatomical, physiological, and ablation studies. From these, we know that the cochlea transforms sound into firing activity in the cochlear nucleus, a collection of neurons in the inner ear. These signals travel up the auditory nerve through several regions of the brainstem and thalamus (superior olivary nucleus (SON), inferior colliculus (IC), medial geniculate nucleus (MGN)) before they reach the primary auditory cortex (A1). While there are collateral signals sent to e.g. the reticular formation directly from some of the low-level areas, these connections do not carry the correct information to uniquely identify auditory

patterns. Thus, they could not mediate habituation to specific auditory word stimuli. It remains possible that these parallel pathways might underlie habituation to more basic aspects of stimuli, e.g. a volume, but this is not addressed further in this paper. Thus, neuroanatomical evidence suggests that there is only one pathway which carries the correct type of information to allow unique identification of auditory patterns.

A spiking neural model based on the auditory processing pathway described above has previously been implemented [4] and used for robust real-time speech recognition by training invariant readout neurons via the liquid state machine (LSM) paradigm [5]. In those experiments, a cochlear model was used to process raw sound into firings in cochlear neurons, which were fed through additional neural areas responding to onsets, offsets, or aggregate activity in each cochlear channel. These neurons then synapsed randomly into a large recurrent circuit with complex temporal dynamics. This circuit served as the “liquid”. Linear readouts were trained in a supervised manner (via linear regression) to respond to invariant properties of the liquid’s response to different auditory word stimulus input categories. This was accomplished by instantiating different projections of the liquid’s state via weight vectors applied to each liquid neuron’s low-pass filtered state. In the model used for these speech recognition experiments, there was no permanent modification of the model’s parameters (weights, etc.), except for the readout weight vectors which were trained only during a supervised “training” phase. In other words, the auditory system was a “static” system. Even though the neural and synaptic models had some short term temporal dynamics, they did not change permanently. In the static model, the response to a word stimulus at given time point would not affect the response to a word stimulus more than a few hundred milliseconds in the future.

The static model is in conflict with what is known about biological systems, especially cortical circuits, which have a large amount of long-term plasticity. Indeed, cortical areas associated with (auditory) memory and habituation¹ (the entorhinal cortex, perirhinal cortex, etc.) have been demonstrated to have long-term changes in synaptic weights via a type of synaptic plasticity known as spike-timing dependent plasticity (STDP) [8]. STDP is a type of plasticity where the change in the efficacy of a synapse closely depends on the relative timings of pre- and post-synaptic spikes.

This paper examines a specific question: what happens when STDP is implemented in the “static” auditory system previously used for word recognition? Specifically, what types of patterns will we see in response to repeated stimulation to familiar versus novel stimuli? If we observe a decrement in response to the repeatedly encountered stimulus greater than the decrement in response to the novel stimuli, then we can conclude that the plastic synapses of the circuit have been selectively modified by experience in a manner

¹Lesion data (albeit in the visual modality) shows that novelty preference is abolished in infant primates who have had the entorhinal cortex, amygdaloid complex, and hippocampus lesioned [6]. Meanwhile, selective hippocampal lesions maintain the ability to establish novelty preferences [7].

specific to the familiar stimulus. This is functionally useful because by observing the response of the circuit to a given auditory stimulus, one can determine above chance whether the stimulus is familiar or is novel. However, it must be kept in mind that there are several unknown factors that may present difficulty in applying this in the real world. For example, words are different lengths, can contain more or less acoustic power and variance of auditory frequencies. Will these variances overcome any regularities in familiarity status of the stimuli? How many stimuli can be simultaneously familiar to the circuit: is it a continuum whereby the response difference between familiar and novel decreases as more stimuli are “remembered”?

In the future, it will also be interesting to test the model in relation to another important phenomenon in language learning: auditory adaptation. Specifically, if an entire “language” is familiarized, would it behave as if that “language” is a familiar stimulus and tokens of a novel language are a novel stimulus? If so, this could serve as an explanation for what causes the auditory adaptation known to occur in human infants. Human infants seemingly gain preference for and expertise in languages present in their ambient environments (even prenatally). Though we do not examine the auditory adaptation via recurrent STDP question in this paper, the model could be used as a strong and simple “foil hypothesis” for a phenomenon that might otherwise be thought to be the result of complex adaptation of more deep-rooted parameters of the neural system. For example, more complex theories might involve modulation of time-constants of neural membranes of A1 to favour temporal properties the language(s) being adapted to, or changes in the early auditory system to favour the specific set of sounds present in the language(s) being adapted to. We plan to examine these possibility in future papers.

III. MODEL

The neural circuit used in the experiments is based on the neural circuit used to perform realtime speech recognition in our lab[4]. A 3-d visualization of the circuit is shown in Figure 1.

The circuit contains a group of cochlear neurons (modelled as exponential-decay PSR current-based leaky integrate-and-fire (LIF) neurons). These cochlear neurons fire at a rate determined by the response of the corresponding channels of the cochlear model². A further set of LIF neurons detect onsets [10] in the cochlear neurons via dynamic synapses using the Markram et al. model of short-term plasticity (STP)[11] (85 neurons, neuron parameters: $V_{reset} = 13.5$ mV, $V_{thresh} =$

²A real-time modification of the Lyon cochlear model [9] was used, parameters: $breakf = 1000$, $qconst = 8.0$, $stepfactor = 0.25$, $sharpness = 5.0$, $notchoffset = 1.5$, $preemphfreq = 300$, $taufactor = 3.0$, producing 85 cochlear channels and thus 85 cochlear neurons, $V_{reset} = 14.5$ mV, $V_{thresh} = 15.0$ mV, $\tau_m = 10.0$ ms, $t_{refract} = 2$ ms, $I_{bg} = 14.5$ nA, PSR parameters: $\tau_E = 0$ ms, $\tau_I = 0$ ms. All LIF neurons in this paper additionally have $R_m = 1$ M Ω , $V_{rest} = 0$ mV. The per-channel spike probability output by the cochlear model is injected as I_{inj} after being multiplied by the constant 5000.0/ms

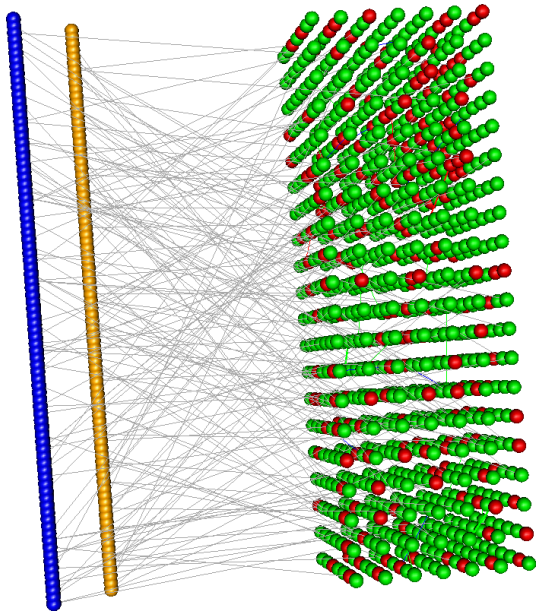


Fig. 1. 3-d visualization of one of the neural circuits used in the experiments. From left: cochlear neurons, onset detectors, liquid neurons (red inhibitory, green excitatory). Only a random 0.5% of synapses are shown.

15.0 mV, $\tau_m = 10.0$ ms, $t_{refract} = 2$, $I_{bg} = 14.5$ nA, PSR parameters: $\tau_E = 3$ ms, $\tau_I = 3$ ms, synapse parameters: $w = 30.0$, $t_{delay} = 1.0$ ms, $U = 0.6$, $D = 0.1$ s, $F = 0.05$ s).

Both the cochlear neurons and the onset detectors project static synapses into 10% of the neurons in a recurrent circuit (the “liquid”, synaptic parameters: $w = 0.045$, $t_{delay} = 1$ ms for those projecting into excitatory liquid neurons and $w = 0.0225$, $t_{delay} = 1$ ms for those projecting into inhibitory liquid neurons). The neurons in the recurrent circuit are Izhikevich neurons with parameters drawn from a Gaussian distribution (neural parameters in format (mean, stddev, min, max), where results outside of min/max were redrawn from a uniform distribution between min and max: $a = (0.02, 0.03, 0.015, 0.080)$, $b = (0.2, 0.01, 0.15, 0.30)$, $c = (-65, 5, -70, -50)$, $d = (8, 2, 0.01, 10.0)$). Liquid neurons implement receptor-type (AMPA, NMDA, GABAa/b) mediated conductances to determine the post-synaptic response (PSR parameters: $\tau_{AMPA} = 5$ ms, $\tau_{NMDA} = 150$ ms, $\tau_{GABAa} = 6$ ms, $\tau_{GABAb} = 150$ ms)[12]. A random 20% of the liquid neurons are selected to be inhibitory. Liquid neurons are arranged at integer points on a 3-dimensional grid. For the experiment, the circuit was made up of 720 neurons, a $7 \times 20 \times 7$ grid. Neurons within the liquid are probabilistically connected by dynamic (STP) synapses with parameters drawn from a Gaussian distribution depending on the type of the pre- and post-synaptic neuron (Excitatory or Inhibitory. synaptic parameters format (w, delay, U, D, F), negative results redrawn from between 0.1 of mean and double mean: (0.075, 2 ms, 0.5, 1.1 s, 0.05 s) for EE, (0.15, 1 ms, 0.05, 0.125 s, 0.120 s) for EI, (-0.0475, 1 ms, 0.25, 0.7 s, 0.02 s) for IE, and (-0.0475, 1 ms, 0.32, 0.144 s, 0.06 s) for II (standard deviation 50%

of the respective means except for weight in which case it is 100% of mean)). The probability of a connection existing between any two liquid neurons (at point a and b) is a function of the distance between the two neurons and their types: $C \cdot e^{(-D(a,b)/\lambda)^2}$, where λ is a global parameter controlling the density of connections ($= 2.0$), $D(\cdot)$ is the Euclidean distance function, and C is a parameter to modulate the probability of a synapse depending on properties of the connected neurons. In our case, $C = 0.3$ if a is an excitatory neuron and b is an excitatory neuron (EE), $C = 0.2$ for excitatory and inhibitory neurons (EI), $C = 0.4$ for inhibitory and excitatory neurons (IE), and $C = 0.1$ for two inhibitory neurons (II).

The STDP model used implements the nearest-neighbor additive STDP equation [13]. Additionally, it implements exponential correction of weight changes [14] to mitigate the tendency of synapses to go to the maximum/minimum allowed weights (STDP parameters: $\mu = 0.019$, $\alpha = 1.05$, $w_{max} = 100$, $\tau = 20$ ms, $\lambda = 0.0001$). STDP is limited to excitatory neuron efferent synapses.

The circuit is implemented and simulated in the NSIM3 simulation engine for spiking neural circuits, and uses the hybrid integration method [15] for integrating the receptor-conductance-mediated Izhikevich neurons ($dt = 1.0$ ms). LIF neurons are updated using the closed-form solution ($dt = 1.0$ ms).

IV. EXPERIMENTS & RESULTS

To test whether the circuit model could support habituation with only the addition of STDP to the synapses, a large set of automated experiments was run. A set of 3 circuits were randomly generated to rule out the possibility that results are specific to any particular circuit configuration. The initial configuration of these circuits (before any STDP-mediated weight changes) was saved so that it could be reinstated for many different experimental trials using the same circuit, with different assignments of stimuli to “familiar” versus “novel” in each trial. This is to rule out the possibility that any observed results are because of particular properties of the familiar stimulus.

The stimulus corpus contained 28 isolated English words spoken 3 times each by an adult male. It is part of a larger corpus used for training speech recognizers for a Human-Robot Interaction task. The corpus is available at http://rveale.com/public/aud_adapt_corpus.tar.gz.

In each trial, a “partition” of the corpus was generated. In a partition, a subset of the word categories (in this paper, always of size 1) was selected to be familiarized, and the rest set aside as “novel” exemplars for testing. The experiment then continued in blocks. In each block, the circuit response to the “familiar” stimulus was measured and compared against the circuit response to a random “novel” stimulus from the ones set aside for testing. Each novel stimulus was only used once within each trial. Additionally, each stimulus in the corpus was used as the “familiar” stimulus an equal number of times. Thus, for a given familiar stimulus A and (different) novel stimuli B, C, D, E , etc., a trial proceeded

as follows (all stimulus presentations inside parentheses are ordered randomly).

(AB)(AC)(AD)(AE)...

The benefit of interleaving familiarization trials and test trials is that it gives a lens into the “trajectory” of habituation. If we had simply done a before-after comparison, with 10 familiar presentations in between, it would not be possible to know modifications to the circuit had all happened within the first encounter with the stimulus, or whether they were spread out over the 10, with gradually changing response, etc.. It is known that habituation generally shows a reverse-exponential trajectory from novel response level to saturated habituation levels [16].

To test STDP as the possible mediator of habituation in the auditory circuit model presented above, all trials were run with the liquids in 4 different “type” conditions: liquid-only STDP (in which only recurrent synapses in the liquid had STDP turned on), liquid-and-input STDP (in which both recurrent liquid synapses and synapses to the liquid from the cochlear model/onset detectors had STDP), input-only STDP (only the inputs to the liquid had STDP), and no STDP (no synapses had STDP). By testing these conditions, it will be possible to see what aspects of the circuit model are important for habituation. It could turn out that only the shape of the input projection matters for habituation, and STDP in the recurrent synapses is not necessary to explain any observed results.

Thus, the experimental design was an $28 \times 4 \times 3 \times 26$ design, of familiar category \times liquid type \times liquid number \times block. The circuit was re-loaded to its initial (weight) state at the beginning of every trial. At the beginning of every word presentation, the circuit’s state (membrane potential, transient synaptic parameters) was randomly reset.

The “response” being measured and compared is the average firing activity (in spikes/second) of the entire liquid over the whole word presentation. This measure is ideal since it is a single value that is simple to calculate. It was chosen over other much more complex measures (such as average state distance between every pairwise liquid response) because of the expense and complexity in calculating them.

Figure 2 shows the results of the experiments in terms of the trajectory of response to familiar versus novel stimuli test stimuli, over 26 familiarization/test blocks for the different liquid types. Figure 3 shows the differences in means for the 4 liquid types between familiar and novel stimuli for their first and last blocks.

A 3-way ANOVA on Familiarity Type \times Block \times Liquid Condition reveals significant interaction between the three conditions (ANOVA, $F(3) = 48.11$, $p < 2^{-16}$). Specifically, the difference between familiar and novel responses after the first block is not significantly different (Tukey’s post-hoc, $p > 0.95$). This is expected, since at that time the circuit has only experienced one instance of the “familiar” stimulus and one instance of a “novel” stimulus.

As expected, in the condition with no STDP there is no significant difference between familiar and novel stimuli after

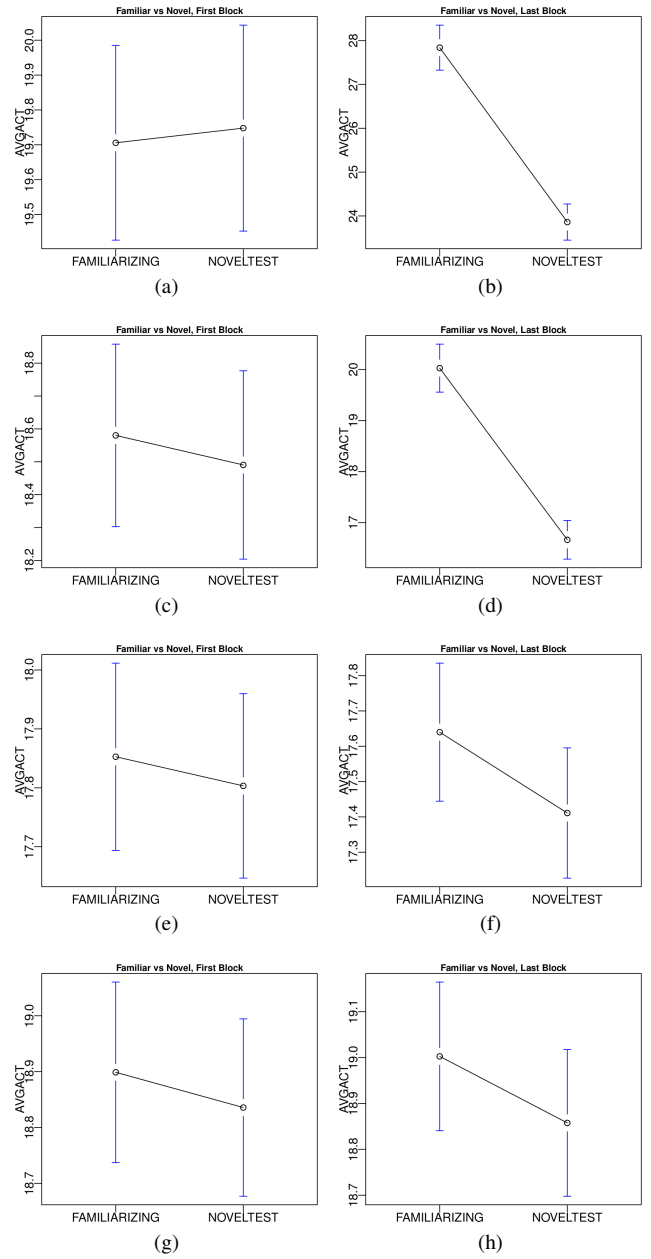


Fig. 3. Comparison of means between the novel and familiar stimuli on the first versus last block, for the 4 different liquid type conditions. Note the difference in scales. (a, b): Input-only STDP condition (c, d): Input and Liquid STDP condition, (e, f): Liquid-only STDP condition, (g, h): no STDP condition

each trial (Tukey’s post-hoc, $p > 0.5$). However, surprisingly the Liquid-only STDP condition likewise shows no significant difference (Tukey’s post-hoc, $p > 0.26$). This result will be discussed further in the context of specific parameter choices.

In the conditions where the input synapses implemented STDP there was a significant difference between the responses to the familiar and novel stimuli. In the Input-only STDP condition, the response to the familiar stimulus was on average 3.97 Hz higher than the response to an arbitrary novel stimulus (Tukey’s post-hoc, $p < 2^{-10}$, bounds: (3.98, 4.69)). Likewise,

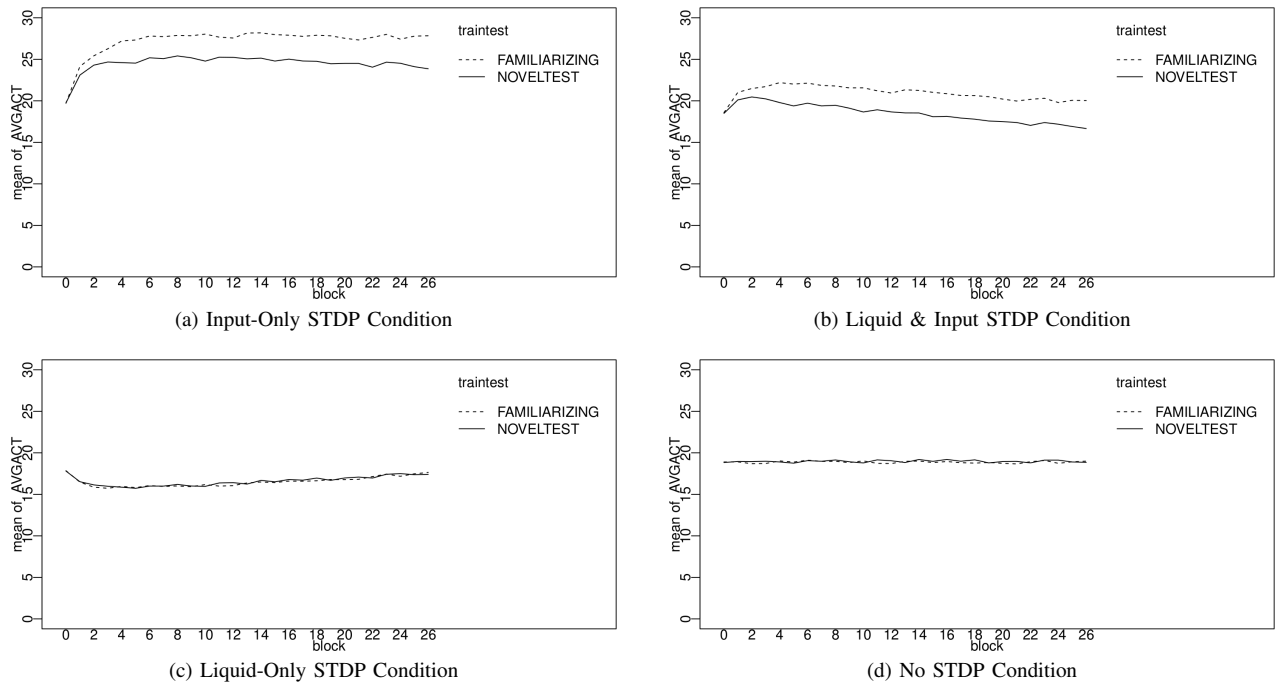


Fig. 2. Trajectory of responses for the four liquid conditions.

in the Liquid-and-Input STDP condition, the response to the familiar stimulus was on average 3.36 Hz higher than a novel response ($p < 2^{-10}$, bounds: (2.69, 4.03)). Note that these differences account for more than 10% of the average firing rate, and thus are quite noticeable.

Overall, the Input STDP conditions separated the responses to the familiar and novel stimuli. As expected, circuits with no learning did not separate the responses at all. A clear trajectory of habituation was observed in the circuits containing STDP in the input synapses, wherein the difference in the responses to the familiar versus novel stimuli increased from being insignificant in the first few familiarization blocks, to being highly significant by the final (26th) block. Since we did not test beyond 26 blocks (i.e., the maximum number of “novel” stimuli in the corpus without repeating them), it is also possible that further habituation would occur had familiarization continued. However, at least in the Input conditions, the trajectory seems to have leveled off by the 26th block.

V. DISCUSSION

The results presented in Section IV demonstrate that it is possible to get robust stimulus-specific decrements in response to real-world auditory stimuli simply by implementing STDP in the input synapses of a cortical-circuit-like recurrent circuit of spiking neurons. A trajectory of habituation was observed wherein the response to the repeatedly experienced familiar stimulus increased significantly in comparison to the response to novel stimuli. Furthermore, it did so despite constant “noise” in the form of novel stimuli interspaced among the familiarization trials (which is what happens in real habituation experiments). Previous experiments in which STDP was

artificially turned off after every familiarization block show the same pattern of results presented in this paper.

Although the response to novel stimuli also changed from baseline along with exclusive familiarization with the familiar stimuli, we hypothesize that this is a “one-time” balancing of the liquid parameters. Future habituation to different familiar stimuli will show a change starting from the response level of the *novel* stimuli in the final block of the experiments.

It is relevant to bring up the fact that in this paper the circuit was not used to control any robotic system or to produce a behavioral response in the sense usually intended by behavioral psychologists. Instead, it was only the “behavior” of the neural circuit that was measured (as if one were recording from neurons in a behaving organism’s brain). However, the transformation is straightforward from the circuit response used as a marker of familiarity in this paper to a more direct (physical) behavioral response. For example, in the case of modulating heart-rate (or sucking behavior), one simply uses the firing rate of the liquid to modulate heart rate (or sucking). In the case of orienting, one uses the firing rate of the liquid to modulate the strength or probability of an orienting response. It is straightforward to include the presented model as part of a larger, embodied system that is acting in the real world, and use the output of the presented model to modulate the overt behavior of that system in a way that matches the behavior of real organisms. Thus, the results should have direct application in producing useful behaviors in autonomous systems. This is especially so since the system already operates on raw sensory data!

The experiments in Section IV only measured the average response of the liquid to each stimulus. In future experiments

we also hope to investigate how the functional computational properties of the circuit are modified by the habituation results we observe, since it is possible that two liquids with the same average firing rate can be performing drastically different computational functions. We are currently running experiments to determine how habituation to a word category affects the performance of the liquid as a speech recognizer both to the familiarized word and to non-familiarized words, though a new (faster) learning method may need to be developed since the linear regression takes an intractable amount of time with the experiment sizes used in this paper. In particular this is important because modifications in the parameters of the circuit caused drastic changes in the trajectories of habituation. The parameter set used to produce the results above is the set that produced the largest difference between familiar and novel stimuli after habituation. We chose to present this data because the goal of this paper was to demonstrate that it is possible to achieve habituation learning in circuits of the type presented. Other parameterizations actually produced different results (with the average firing rate *decreasing* over more stimulus encounters). It is possible that parameterizations exist under which stimulus-specific habituation will occur in the Liquid-only STDP condition. We will explore these more exhaustively in future papers devoted to this topic.

One limitation of the results presented is that we only performed familiarization to one simultaneous word category. The question remains open regarding how the model will respond to simultaneous familiarization to more than one category, and how the response will scale as the number of simultaneous categories increases. It is unclear how to disentangle total experience (to any word) with the ordinal encounter with a given word type. A good method for comparing between results from different numbers of familiar words will need to be developed. Additionally, as discussed in Section I, it will be interesting to explore the possible relationship between simultaneous habituation to a large set of words (e.g. in the same language), and language-specific auditory adaptation. To do this, it will be necessary to compare the response of the liquid between multiple languages. If response properties to novel categories in the familiar language are different than response properties to novel categories in the novel language, then it could be said that a type of auditory adaptation has occurred that is specific to some properties of the familiar language, which is mediated by the exact same STDP mechanisms that beget what we call habituation in this paper.

VI. CONCLUSION

This paper demonstrated that auditory habituation can be observed in a biologically accurate model of the human infant auditory system by simple addition of a type of synaptic plasticity. Since the circuit model is already known to be capable of other useful computations (it has previously been used for speech recognition), the results are relevant since they begin to offer evidence that two different types of learning behavior present in infants (habituation and word-learning) might be mediated by different aspects the same

neural substrates. Though other (neural) models of habituation exist [17], they do not use biologically accurate mechanisms known to exist in the very brain areas where habituation takes place (e.g. spiking neurons, STDP, etc.). The purpose of this paper was to test whether it was possible to achieve habituation in the simplest way possible, by adding a model of plasticity to an already existing model. Though it is possible that this is not how habituation is implemented in the actual neural system of human infants, the model in this paper is a functioning alternative against which future models can be compared.

ACKNOWLEDGMENTS

RV is an NSF Graduate Research Fellow and NSF IGERT.

REFERENCES

- [1] P. R. Zelazo, M. J. S. Weiss, and N. Tarquinio, "Habituation and recovery of neonatal orienting to auditory stimuli," in *Newborn attention: Biological constraints and the influence of experience*, M. J. S. Weiss and P. R. Zelazo, Eds. Ablex Publishing, 1991, pp. 120–141.
- [2] B. S. Kisilevsky and D. W. Muir, "Human fetal and subsequent newborn responses to sound and vibration," *Infant Behavior and Development*, vol. 14, pp. 1–26, 1991.
- [3] W. E. Jeffrey and L. S. Cohen, "Habituation in the human infant," *Advances in Child Development and Behavior*, vol. 6, pp. 63–97, 1971.
- [4] R. Veale and M. Scheutz, "Neural circuits for any-time phrase recognition with applications in cognitive models and human-robot interaction," in *Proceedings of the 34th Annual Conference of the Cognitive Science Society*, N. Miyake, D. Peebles, and R. P. Cooper, Eds. Austin, TX: Cognitive Science Society, 2012, pp. 1072–1077.
- [5] W. Maass, T. Natschläger, and H. Markram, "Real-time computing without stable states: A new framework for neural computation based on perturbations," *Neural Computation*, vol. 14, no. 11, pp. 2531–2560, 2002.
- [6] J. Bachevalier, M. Brickson, and C. Hagger, "Limbic-dependent recognition memory in monkeys develops early in infancy," *NeuroReport*, vol. 4, pp. 77–80, 1993.
- [7] A. Zeamer, E. Heuer, and J. Bachevalier, "Developmental trajectory of object recognition memory in infant rhesus macaques with and without neonatal hippocampal lesions," *The Journal of Neuroscience*, vol. 30, no. 27, pp. 9157–9165, 2010.
- [8] J. Haas, T. Nowotny, and H. Abarbanel, "Spike-timing-dependent plasticity of inhibitory synapses in the entorhinal cortex," *Journal of Neurophysiology*, vol. 96, pp. 3305–3313, 2006.
- [9] M. Slaney, "Lyon's cochlear model," Apple Computer Inc. Cupertino, Ca., Tech. Rep. 13, 1998.
- [10] L. S. Smith and D. S. Fiser, "Robust sound onset detection using leaky integrate and fire neurons with depressing synapses," *IEEE Transactions on Neural Networks*, vol. 15, no. 5, pp. 1125–1134, 2004.
- [11] R. Legenstein, C. Naeger, and W. Maass, "What can a neuron learn with spike-timing-dependent plasticity?" *Neural Comput.*, vol. 17, no. 11, pp. 2337–2382, 2005.
- [12] E. M. Izhikevich, J. a. Gally, and G. M. Edelman, "Spike-timing dynamics of neuronal groups," *Cerebral cortex*, vol. 14, no. 8, pp. 933–944, 2004.
- [13] E. M. Izhikevich and N. S. Desai, "Relating stdp to bcm," *Neural Computation*, vol. 15, pp. 1511–1523, 2003.
- [14] R. Gutig, R. Aharonov, S. Rotter, and H. Sompolinsky, "Learning input correlations through nonlinear temporally asymmetric hebbian plasticity," *The Journal of Neuroscience*, vol. 23, no. 9, pp. 3697–3714, 2003.
- [15] E. M. Izhikevich, "Hybrid spiking models," *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, vol. 368, pp. 5061–5070, 2010.
- [16] G. Shoner and E. Thelen, "Using dynamic field theory to rethink infant habituation," *Psychological Review*, vol. 113, no. 2, pp. 273–299, 2006.
- [17] S. Sirois and D. Mareschal, "An interacting systems model of infant habituation," *Journal of Cognitive Neuroscience*, vol. 16, no. 8, pp. 1352–1362, 2004.