LEARNING-DEPENDENT ENHANCEMENT OF PERSISTENT ACTIVITY IN THE NEOCORTEX

Eunsol Park, Nicholas J. Audette*& Alison L. Barth
Department of Biological Sciences
Carnegie Mellon University
Pittsburgh, PA 15213, USA
{eunsolp}@andrew.cmu.edu
{albarth}@andrew.cmu.edu

ABSTRACT

Persistent activity in cortical networks has long been of interest to neuroscientists, as it may provide a substrate for working memory. In addition, sustained activity may provide a solution to the credit-assignment problem, whereby delayed reinforcement information can interact with stimulus-driven firing to facilitate long-lasting changes in synaptic weights during learning. Although some cell-autonomous mechanisms can facilitate persistent activity, it is increasingly clear that understanding this phenomenon requires a detailed understanding of cortical circuitry, both in reduced experimental preparation as well as *in vivo*, that can be referenced to and inform each other. We use state-of-the-art methods for cell-type specific identification and activity modulation in acute brain slices of mice to show how persistent activity in neocortical circuits can be initiated by higher-order thalamic inputs and sustained by translaminar activation across the cortical column, and discuss how this may be critical for reinforcement learning. Thus, the mechanisms that initiate and sustain persistent activity may be useful in the design of engineered networks for learning.

1 INTRODUCTION

Reinforcement learning requires solving the notorious "credit assignment" problem, where actions taken by an agent lead to a certain outcome that arrives much later in time. Understanding the mechanisms that allow outcomes to be attributed to their casual actions remains challenging in both machine learning and neuroscience. The brain accomplishes this feat effortlessly, with minimal architectures and low energy costs compared to engineered networks (Schwartz et al., 2019). Understanding how biological networks retain information about recent activity to enable connection updating may provide insights for energy-efficient deep learning models. The precise mechanisms by which a stimulus interacts with delayed reinforcement signals in the brain remain largely unknown. This is non-trivial, since a stimulus and its behavioral consequences are temporally separated by long time intervals (seconds) that dwarf the duration of single action potentials (1-2 ms). Persistent activity initiated by a sensory stimulus is thought to provide a neural substrate for delayed reinforcement signals to update connection weights (Asaad et al., 2017; Lillicrap et al., 2020). How the highly-specified cellular architecture of brain networks facilitate persistent activity remains an area of active investigation. Stimulus-initiated persistent activity has been observed across diverse neocortical areas, including primary sensory cortex as well as higher cortical areas such as prefrontal cortex (Constantinidis & Klingberg, 2016). The conserved laminar organization of the neocortex, with molecularly-defined neuronal cell types and highly-specified connectivity, presents a tractable system to investigate the neuronal circuits that enable persistent activity and reinforcement learning. Recent studies indicate that primary sensory cortex can exhibit sustained activity under certain conditions (Condylis et al., 2020; Cooke et al., 2020; Esmaeili & Diamond, 2019), depending on the input source and the subset of target neurons distributed across the deep and superficial layers of the cortical column. This activity is preserved in acute brain slices (Audette et al., 2019; Chubykin

*New York University, New York, NY 10003, USA

et al., 2013), indicating that the essential circuitry is preserved within local cortical networks. Our data indicate that activation of higher-order thalamic inputs drives asynchronous firing, network disinhibition and sustained translaminar excitation in primary somatosensory cortex that is selectively enhanced in the early stages of learning. Persistent activity cannot be elicited by activation of excitatory neurons within the cortex or by pharmacological suppression of inhibition, suggesting that external thalamic inputs selectively engage cortical circuitry to trigger for persistent activity. Biological circuits may switch between brain states to enable the same neural networks to carry out multiple functions, representing a parsimonious and energy-efficient approach to information processing and learning.

2 Materials and Methods

Animals. Adolescent (P21-29) C57BL6 mice were injected with channelrhodopsin2 (ChR2) in the posterior medial nucleus (POm) of the thalamus. Rbp4-Cre crossed to Ai32 mice on a C57BL6 background were used.

Viral injection. ChR2 tagged with m-cherry or YFP (300-500nL; AAV1-CAG-hChR2(H134R)-mCherry-WPRE/SV4 Catalog No.100054-AAV1, Addgene, Cambridge, MA; AAV2-hSyn-hChR2(H134R)-EYFP, Deisseroth Lab, UNC Vector Core, Chapel Hill, NC) was stereotaxically injected into POm (bregma -1.7, lateral 1.00, depth 3.25mm). Mice were treated with ketofen (5mg/kg, Sigma-Aldrich; St. Louis) after surgery. Mice were recovered in their home cage for 7-13 days before sensory association training (Audette et al., 2019; Bernhard et al., 2020).

Sensory Association Training (SAT). Control animals went through 48 hours of acclimation to the training cage. Animals were trained after 24 hours of acclimation. During the training, trials were initiated by freely moving animals by a nose poke followed at variable delay (0.8-1.2s) by a 500ms gentle air puff (6 psi, 500ms), 500ms was delivered before water delivery (Audette et al., 2019). Animals learned an association between whisker stimulation and water reward.

General electrophysiology. Angled brain slices (45 degree rostro-lateral; $350\,\mu$M thick) that preserve transcolumnnar circuitry were prepared using regular artificial cerebrospinal fluid (ACSF). Modified ACSF (mACSF), identical to regular ACSF but with (in mM) 3.5 KCl, 0.5 $MgSO_4$ and $CaCl_2$ was used for whole-cell patch clamp recording to increase network activity and enable polysynaptic interaction. Using whole-cell current clamp recordings, sub- and suprathreshold responses of L2/3 and L5 pyramidal (Pyr) neurons were recorded while ChR2-expressing POm axons or L5 Rbp4+ Pyr neurons were optogenetically stimulated (5 pulses, 5ms duration, 80ms inter-stimulus interval, 0.05Hz inter-trial interval). Max light intensity at 470nm was measured at 3.25mW. 10 trials of spike data for each cell was binned at 10ms intervals. An average PSTH was generated by averaging all cells of a given population. In a subset of experiments, POm-evoked activity in L2 pyramidal neurons was recorded from acute brain slices transected through L4. Evoked activity was measured as previously described (Audette et al., 2019).

Pharmacology. After recording L2 or L5 Pyr neuron-evoked activity in mACSF, picrotoxin ($100\,\mu$M, Tocris Catalog No.124-87-8) was bath applied for at least 10 minutes prior to data collection.

3 Results

Higher-order thalamus integrates sensory input from the external world with context-specific cues and reinforcement information (Roth et al., 2016; Saalmann et al., 2012; Trageser et al., 2006). In the somatosensory system, stimulation of higher-order thalamus POm can initiate sustained cortical activity *in vivo* (Zhang & Bruno, 2019) as well as in acute brain slices (Audette et al., 2018; 2019). In addition, previous studies have shown that higher-order thalamus inputs are selectively enhanced during association learning (Audette et al., 2019; Pardi et al., 2020). To identify the cortical neurons that are required for POm-initiated sustained activity, we expressed the light-activated channelrhodopsin (ChR2) in specific subpopulations of neurons within the thalamocortical circuit. Optogenetic POm fiber stimulation in acute brain slices from the whisker representation of primary somatosensory cortex (barrel cortex) elicited reliable spiking in L5 Pyr neurons and subthreshold polysynaptic activity in L2 Pyr neurons, consistent with direct POm inputs to these neurons (Au-

dette et al., 2018; Sermet et al., 2019) and reciprocal synaptic connections between L5 and L2 Pyr neurons (Jiang et al., 2015; Lefort et al., 2009). After 24 hrs of SAT, POm fiber stimulation elicited a pronounced increase in evoked firing in Pyr neurons in both L5 and L2 that was present both in the stimulus and post-stimulus period (Figure 1). This persistent activity in L2 was dependent on translaminar interactions, as mechanical transection of the brain slice eliminated this post-stimulus activity (Figure 2). POm stimulation evoked a complex, polysynaptic excitatory postsynaptic potentials (EPSPs) in L2 Pyr neurons that was eliminated after transection, where only the direct POm-mediated EPSP was retained. These data indicate that L2 Pyr neurons receive both direct and indirect excitation from POm fiber stimulation, and that the indirect excitation originates from deep-layer neurons, likely via L5 Pyr neurons. SAT has been shown to rapidly enhance the strength of excitatory POm inputs onto L5 Pyr neurons (Audette et al., 2019). Thus, persistent activity in cortical neurons may be due to more effective POm-evoked spiking of L5 Pyr neurons after SAT.

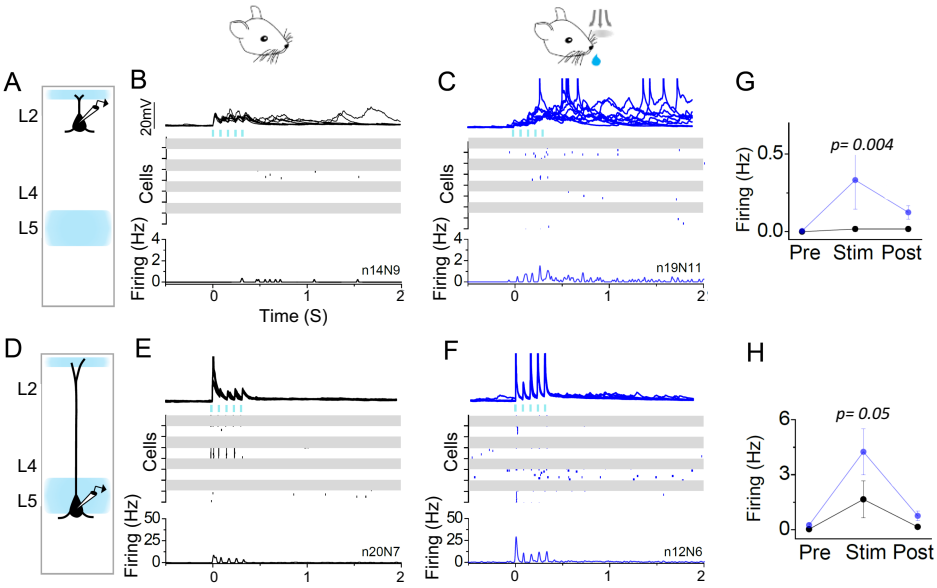

Figure 1: POm-evoked activity circuit activity is rapidly enhanced by SAT. A) Schematic of current-clamp recordings of L2 Pyr neurons in the brain slice. POm afferents express ChR2 (blue). B) Optogenetic stimulation of POm afferents drives firing in L2 Pyr neurons in control animals. (top) 10 overlaid sweeps stimulus-evoked spiking of a single L2 Pyr neuron. (middle) A raster plot from 9 different L2 Pyr neurons. (Bottom) Averaged responses for all recorded L2 Pyr neurons shown as a peri-stimulus time histogram (PSTH; 10ms bin). C) As in (B), but for L2 Pyr neurons after 24 hrs of SAT. D-F) As in (A-C), but for L5 Pyr neurons. G) Average POm-evoked firing frequency across all cells during the 500 ms preceding POm stimulation (pre), during stimulation (stim), and directly following stimulation (post) (black=control; blue=24 hrs of SAT). Averages represented as mean ± SEM. Two-way Repeated Measures ANOVA, Bonferroni post-hoc test p = 0.004 during stimulation. H) As in (G), but for L5 Pyr neurons. p= 0.05 during stimulation. From Audette et al. (2019)
.

To test whether L5 Pyr firing is sufficient to initiate persistent activity, we expressed ChR2 in L5 Pyr neurons in the Rbp-Cre transgenic line and subjected animals to 24 hrs of SAT. Optogenetic stimulation of L5 Pyr neurons did not induce persistent activity in either L2 and L5 Pyr neurons (Figure 3B,E). Our data suggest that although excitatory input from infragranular layers is necessary to maintain POm-evoked persistent activity in L2 (Figure 2), ChR2 activation of L5 Pyr neurons is not sufficient to initiate this activity. L5 Pyr neurons activate strong inhibition in both L2 and L5 (Jiang et al., 2015; Vecchia et al., 2020), and synchronous ChR2-mediated activation of L5 Pyr neurons may activate strong inhibition that blocks persistent activity in both layers. To test whether persistent activity can be induced by L5 Pyr neurons when inhibition is eliminated, we

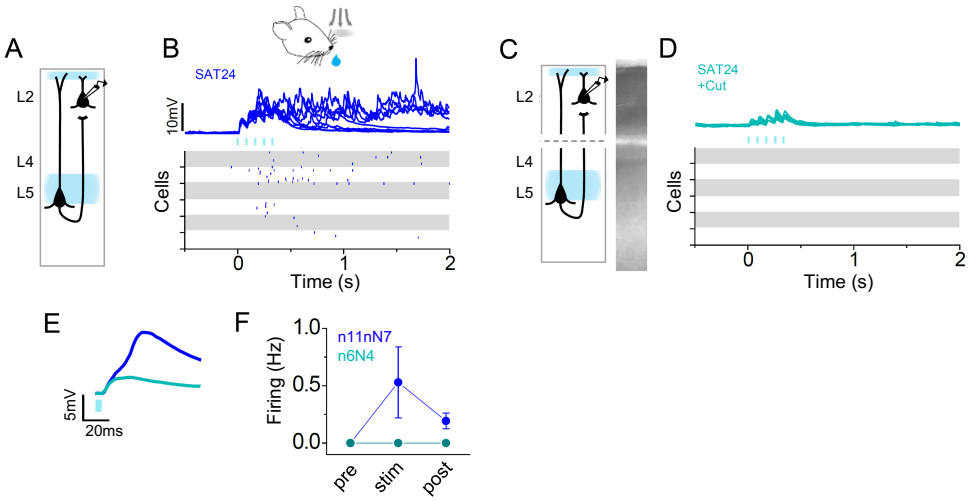

Figure 2: POm-evoked sustained activity in L2 Pyr neurons requires infragranular layers. A) Schematic of current-clamp recordings of L2 Pyr neurons in the brain slice. POm afferents express ChR2 (blue). B) L2 Pyr neuronal response to optogenetic activation of POm fibers shows sustained sub- and suprathreshold activity. (top) 10 overlaid sweeps stimulus-evoked spiking of a single L2 Pyr neuron. (bottom) A raster plot from multiple L2 Pyr neurons. C-D) From the same slice as in (A-B), but where the slice has been transected through L4. Responses are now limited to the stimulus window. E) Averaged responses (10 consecutive sweeps) of pre-transcetion L2 Pyr neurons (dark blue) and post-transection L2 Pyr neurons (light blue) to the first light pulse. Example cells in (B) and (D). F) Average firing frequency across all cells during the 500 ms preceding POm stimulation (pre), during stimulation (stim), and directly following stimulation (post). Averages represented as mean ± SEM. from Audette et al. (2019)

bath-applied the $GABA_A$ receptor antagonist picrotoxin (PTX) and optogenetically stimulated L5 Rbp4+ neurons. PTX did not facilitate prolonged activity in L2 and L5 Pyr neurons (Figure 3C,F). Our data indicate that synchronous, optogenetic activation of L5 Pyr neurons in the absence of cortical inhibition is not sufficient to initiate persistent activity. Taken together, these data indicate that POm afferents uniquely engage the cortical circuit to facilitate persistent activity, a phenomenon that is enhanced during sensory learning.

4 Discussion

Acute brain slices present a tractable model system to evaluate the cellular and synaptic requirements for persistent activity in cortical circuits during learning. Our results indicate that triggering persistent activity is unique to POm fiber activation and cannot be replaced by L5 Pyr firing. How might this be accomplished? POm inputs target a diverse array of cell types distributed across the cortical column, including vasoactive intestinal peptide (VIP) GABAergic neurons that inhibit the spontaneous firing of somatostatin (SST) GABAergic neurons. This VIP to SST disinhibitory motif is ubiquitous in cortical circuits (Lee et al., 2013; Pi et al., 2013) and may be a critical regulator of cortical synaptic plasticity (Williams & Holtmaat, 2019) and learning (Fu et al., 2015). However, L5 Pyr activation during pharmacological suppression of inhibition could not reproduce POm-initiated persistent activity. It is possible that broad and synchronous ChR2-evoked firing in L5 Rbp4+ Pyr neurons yields a markedly different synaptic response in L2 Pyr neurons, possibly due to short-term synaptic depression at L5 to L2 glutamatergic synapses. Alternatively, the use of the Rbp4-Cre driver line to drive ChR2 in our experiments may select a different subset of L5 Pyr neurons than are typically activated by POm afferents.

Recent studies have shown that neurons in prefrontal cortex may maintain some activity that could correspond to an “eligibility trace” for synaptic updating (Lim et al., 2020), which could serve as a solution to the credit assignment problem (Asaad et al., 2017). However, the neural circuit mecha-

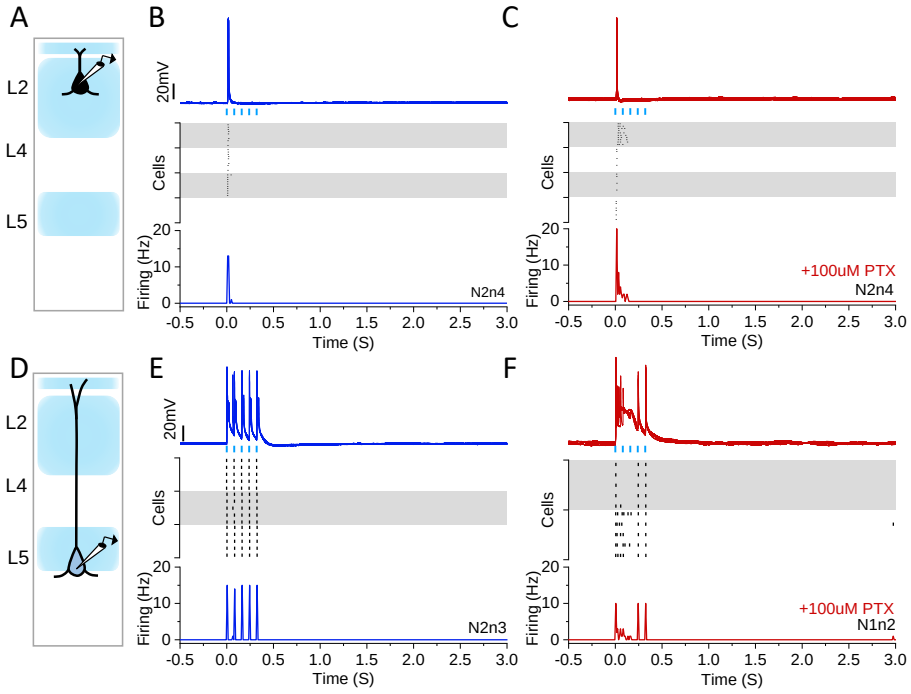

Figure 3: Optogenetic stimulation of L5 Rbp4+ neurons cannot evoke recurrent activity after 24 hrs of SAT. A) Schematic of recordings of L2 Pyr neurons in the brain slice. B) (top) 10 overlaid sweeps stimulus-evoked spiking of a single L2 Pyr neuron. (middle) A raster plot from 5 different L2 Pyr neurons. (Bottom) Averaged responses for all recorded L2 Pyr neurons shown as a peri-stimulus time histogram (PSTH; 10ms bin). C) As in (B), but for L2 Pyr neurons in 100 µM picrotoxin (PTX). D-E) As in (A-B), but for consecutive trials of L5 Rbp4+ Pyr neurons. F) As in (C), but for consecutive trials of L5 Rbp4+ Pyr neurons.

nisms that both initiate and sustain persistent activity are not well-understood. Our data suggest an important role for asynchronous inputs in generating persistent activity in the neocortex. Interestingly, a recent study demonstrated that a more biologically plausible network using asynchronous signal input could outperform conventional artificial neural networks (Uzan et al., 2019). Although artificial neural networks with backpropagation are sufficient to perform tasks such as visual recognition even without a close parallel to brain circuits (Lillicrap et al., 2016), we propose that incorporating temporal dynamics that parallel cortical neural function may improve the performance or energy cost of deep learning models. Biologically-inspired information-transfer networks can use less power and be both more efficient and robust (Navlakha et al., 2015). Future experiments will examine requirements for asynchronous timing and the activation of different neural cell types to evaluate their effect on driving persistent activity across the cortical column. Understanding how neurons take advantage of biophysical and synaptic properties to initiate persistent activity may provide insight into the efficient design of engineered networks for reinforcement learning.

ACKNOWLEDGEMENTS

This work was supported by the AFOSR FA9550-19-1-13022629 and NIH R56 NS119599

REFERENCES

Wael F Asaad, Peter M Lauro, János A Perge, and Emad N Eskandar. Prefrontal neurons encode a solution to the credit-assignment problem. *Journal of Neuroscience*, 37(29):6995–7007, 2017.

Nicholas J Audette, Joanna Urban-Ciecko, Megumi Matsushita, and Alison L Barth. Pom thalamocortical input drives layer-specific microcircuits in somatosensory cortex. *Cerebral cortex*, 28(4):

1312–1328, 2018.

Nicholas J Audette, Sarah M Bernhard, Ajit Ray, Luke T Stewart, and Alison L Barth. Rapid plasticity of higher-order thalamocortical inputs during sensory learning. *Neuron*, 103(2):277–291, 2019.

Sarah M Bernhard, Jiseok Lee, Mo Zhu, Alex Hsu, Andrew Erskine, Samuel A Hires, and Alison L Barth. An automated homecage system for multiwhisker detection and discrimination learning in mice. *Plos one*, 15(12):e0232916, 2020.

Alexander A Chubykin, Emma B Roach, Mark F Bear, and Marshall G Hussain Shuler. A cholinergic mechanism for reward timing within primary visual cortex. *Neuron*, 77(4):723–735, 2013.

Cameron Condylis, Eric Lowet, Jianguang Ni, Karina Bistrong, Timothy Ouellette, Nathaniel Josephs, and Jerry L Chen. Context-dependent sensory processing across primary and secondary somatosensory cortex. *Neuron*, 106(3):515–525, 2020.

Christos Constantinidis and Torkel Klingberg. The neuroscience of working memory capacity and training. *Nature Reviews Neuroscience*, 17(7):438, 2016.

James E Cooke, Julie J Lee, Edward L Bartlett, Xiaoqin Wang, and Daniel Bendor. Post-stimulatory activity in primate auditory cortex evoked by sensory stimulation during passive listening. *Scientific reports*, 10(1):1–14, 2020.

Vahid Esmaeili and Mathew E Diamond. Neuronal correlates of tactile working memory in prefrontal and vibrissal somatosensory cortex. *Cell reports*, 27(11):3167–3181, 2019.

Yu Fu, Megumi Kaneko, Yunshuo Tang, Arturo Alvarez-Buylla, and Michael P Stryker. A cortical disinhibitory circuit for enhancing adult plasticity. *elife*, 4:e05558, 2015.

Xiaolong Jiang, Shan Shen, Cathryn R Cadwell, Philipp Berens, Fabian Sinz, Alexander S Ecker, Saumil Patel, and Andreas S Tolias. Principles of connectivity among morphologically defined cell types in adult neocortex. *Science*, 350(6264), 2015.

Soohyun Lee, Illya Kruglikov, Z Josh Huang, Gord Fishell, and Bernardo Rudy. A disinhibitory circuit mediates motor integration in the somatosensory cortex. *Nature neuroscience*, 16(11): 1662–1670, 2013.

Sandrine Lefort, Christian Tomm, J-C Floyd Sarria, and Carl CH Petersen. The excitatory neuronal network of the c2 barrel column in mouse primary somatosensory cortex. *Neuron*, 61(2):301–316, 2009.

Timothy P Lillicrap, Daniel Cownden, Douglas B Tweed, and Colin J Akerman. Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications*, 7(1): 1–10, 2016.

Timothy P Lillicrap, Adam Santoro, Luke Marris, Colin J Akerman, and Geoffrey Hinton. Backpropagation and the brain. *Nature Reviews Neuroscience*, 21(6):335–346, 2020.

Dong-Hyun Lim, Young Ju Yoon, Eunsil Her, Suehee Huh, and Min Whan Jung. Active maintenance of eligibility trace in rodent prefrontal cortex. *Scientific reports*, 10(1):1–11, 2020.

Saket Navlakha, Alison L Barth, and Ziv Bar-Joseph. Decreasing-rate pruning optimizes the construction of efficient and robust distributed networks. *PLoS Comput Biol*, 11(7):e1004347, 2015.

M Belén Pardi, Johanna Vogenstahl, Tamas Dalmay, Teresa Spanò, De-Lin Pu, Laura B Naumann, Friedrich Kretschmer, Henning Sprekeler, and Johannes J Letzkus. A thalamocortical top-down circuit for associative memory. *Science*, 370(6518):844–848, 2020.

Hyun-Jae Pi, Balázs Hangya, Duda Kvitsiani, Joshua I Sanders, Z Josh Huang, and Adam Kepecs. Cortical interneurons that specialize in disinhibitory control. *Nature*, 503(7477):521–524, 2013.

Morgane M Roth, Johannes C Dahmen, Dylan R Muir, Fabia Imhof, Francisco J Martini, and Sonja B Hofer. Thalamic nuclei convey diverse contextual information to layer 1 of visual cortex. *Nature neuroscience*, 19(2):299–307, 2016.

Yuri B Saalmann, Mark A Pinsk, Liang Wang, Xin Li, and Sabine Kastner. The pulvinar regulates information transmission between cortical areas based on attention demands. *Science*, 337(6095): 753–756, 2012.

Roy Schwartz, Jesse Dodge, Noah A Smith, and Oren Etzioni. Green ai. *arXiv preprint arXiv:1907.10597*, 2019.

B Semihcan Sermet, Pavel Truschow, Michael Feyerabend, Johannes M Mayrhofer, Tess B Oram, Ofer Yizhar, Jochen F Staiger, and Carl CH Petersen. Pathway-, layer-and cell-type-specific thalamic input to mouse barrel cortex. *Elife*, 8:e52665, 2019.

Jason C Trageser, Kathryn A Burke, Radi Masri, Ying Li, Larisa Sellers, and Asaf Keller. State-dependent gating of sensory inputs by zona incerta. *Journal of Neurophysiology*, 96(3):1456–1463, 2006.

Herut Uzan, Shira Sardi, Amir Goldental, Roni Vardi, and Ido Kanter. Biological learning curves outperform existing ones in artificial intelligence algorithms. *Scientific reports*, 9(1):1–11, 2019.

Dania Vecchia, Riccardo Beltramo, Fabio Vallone, Ronan Chéreau, Angelo Forli, Manuel Molano-Mazón, Tanika Bawa, Noemi Binini, Claudio Moretti, Anthony Holtmaat, et al. Temporal sharpening of sensory responses by layer v in the mouse primary somatosensory cortex. *Current Biology*, 30(9):1589–1599, 2020.

Leena E Williams and Anthony Holtmaat. Higher-order thalamocortical inputs gate synaptic long-term potentiation via disinhibition. *Neuron*, 101(1):91–102, 2019.

Wanying Zhang and Randy M Bruno. High-order thalamic inputs to primary somatosensory cortex are stronger and longer lasting than cortical inputs. *Elife*, 8:e44158, 2019.