
Evaluating Adversarial Robustness in Simulated Cerebellum

Liu Yuezhang
HKUST
yliuhl@ust.hk

Bo Li
HKUST
bli@ust.hk

Qifeng Chen
HKUST
cqf@ust.hk

Abstract

It is well known that artificial neural networks are vulnerable to adversarial examples, in which great efforts have been made to improve the robustness. However, such examples are usually imperceptible to humans, thus their effect on biological neural circuits is largely unknown. This paper will investigate the adversarial robustness in a simulated cerebellum, a well-studied supervised learning system in computational neuroscience. Specifically, we propose to study three unique characteristics revealed in the cerebellum: (i) network width; (ii) long-term depression on the parallel fiber-Purkinje cell synapses; (iii) sparse connectivity in the granule layer, and hypothesize that they will be beneficial for improving robustness. To the best of our knowledge, this is the first attempt to examine the adversarial robustness in simulated cerebellum models. We wish to remark that both of the positive and negative results are indeed meaningful—if the answer is in the affirmative, engineering insights are gained from the biological model into designing more robust learning systems; otherwise, neuroscientists are encouraged to fool the biological system in experiments with adversarial attacks—which makes the project especially suitable for a pre-registration study.

1 Introduction

The recent renaissance in deep learning has achieved great success in various challenging supervised learning tasks, which shows the enormous capacity of neural networks. Intriguingly, adversarial examples (i.e., intentional designed small perturbations which could fool machine learning systems) were discovered [1, 2], providing the robustness as an additional criterion. Such adversarial examples are usually undetectable to the human visual system but destructive to artificial neural networks, known as the **robustness gap** between biological and artificial learning systems [3].

Significant efforts have been made in adversarial defense from the engineering perspective [4–6], while on the other side, robustness in the biological system is not fairly understood. Nayebi and Ganguli [7] proposed a defense inspired by saturated neurons, but was later breached with numerically stabilized gradients [8]. Li et al. [9] discovered that regularization fitting neuron responses had a positive effect on robustness, while other studies [10, 11] suggested that adversarial examples could also fool time-limited human based on the behavior experiments. In all, concrete results of adversarial robustness **on the neural circuits level** is lacking investigation.

Under ideal conditions, we should directly study the robustness of a visual system. However, as “an active field of research” [12], it seems to be infeasible to evaluate the robustness of a simulated visual model for the moment—a supervised learning system in nature with detailed understandings ought to be a better candidate for our purpose at the current stage.

The cerebellum is one of the most studied brain areas in computational neuroscience. The standard model, known as the Marr-Albus theory [13, 14], proposed the cerebellum as a perceptron network for movements learning. Although various amendments have been made to the original idea, the

supervised paradigm of cerebellar learning is well established [15–17]. Therefore, as a supervised learning system in nature, it would be intriguing to ask: does the adversarial examples problem also exist in the cerebellum?

We hypothesize that the unique architectures in cerebellar circuits¹ will enhance the robustness, and thus, such adversarial vulnerabilities in artificial neural networks are not expected to exist in the cerebellum model. To test our idea, we propose to simulate the cerebellum and evaluate its robustness against adversarial attacks. Admittedly, addressing such a problem is in some sense betting on both sides: if the results are affirmative, the mechanisms improving robustness will be inspiring for adversarial defense; otherwise, adversarial examples are expected to be discovered in vivo by neuroscientists. Given the significance of both positive and negative results, our work is especially suitable for a pre-register study.

2 Related work

Computational models of cerebellum The long trend for modeling the cerebellum in computational neuroscience was initiated by Marr [13], in which he proposed the cerebellum as a perceptron model under Hebbian learning rule. Albus [14] presented a similar model independently with an emphasis on the depression activity. The proposed plasticity was experimentally verified by Ito et al. [18], and thereby, the Marr-Albus-Ito theory of cerebellum was established. Some of the following studies focused on the connectivity in the cerebellum. Namely, Cayco-Gajic et al. [19] suggested that the sparse connectivity was essential for pattern decorrelation by simulation, and Litwin-Kumar et al. [20] further derived the optimal degrees of connectivity in cerebellum-like structures from a theoretical perspective. Other studies attempted to rebuild the cerebellum and test its capabilities in simulation. Hausknecht et al. [21] evaluated machine learning capabilities of a simulated cerebellum, addressing that such a model can solve simple supervised learning benchmarks such as MNIST. Yamazaki’s group also has a long history in building large scale cerebellar models with either GPU [22] or supercomputer [23], representing the state-of-the-art in computational modeling of the cerebellum. However, as a newly developed topic from the machine learning community, the adversarial robustness of such models has not yet been examined.

Adversarial defense Much effort has been devoted to adversarial defensive techniques. Adversarial training [2] was proposed as an efficient framework to improve the robustness. Madry et al. [4] further equipped the framework with a more advanced attack. Gradient masking effect was discovered [24–26], which showed that a bunch of methods except adversarial training indeed leveraged such effect and provided a false sense of security. Randomization methods are the second type of defensive techniques acquiring comparable success against black-box attacks. Randomness was introduced to input transformation [27], layer activation [28], or feature pruning [5]. Unfortunately, such methods are facing challenges with white-box attacks [29]. Provable defense [30] is another trend aiming at proposing algorithms with guaranteed safety, but usually with the cost of significantly reducing the training accuracy. A recent work [6] attempted to bridge this gap in provable defense, but also benefits from the adversarial training. In all, adversarial training is still the most successful defensive technique developed so far. However, we challenge if such a framework is biologically plausible—no evidence supports that the biological neural circuits process information augmented with adversarial signals repeatedly to enhance its robustness. In pursuit of a machine learning system with human-level robustness, intriguing defensive mechanisms other than adversarial training should be expected in the brain.

3 Methodology

3.1 Anatomy of Cerebellum

We briefly summarize several anatomical evidence to justify our model of cerebellum. Connectivity in the cerebellar cortex is substantially feed-forward: information inputs through the mossy fiber (MF, input neuron) to granule cell (GC, hidden neuron), and subsequently passes through parallel fiber (PF) to Purkinje cell (PC, readout neuron). Another input comes from climbing fiber (CF) to PF-PC synapse. The schematic diagram of cerebellum is shown in Fig 1a. In quantity, corresponding

¹We use the rat cerebellum as an exemplar through our study.

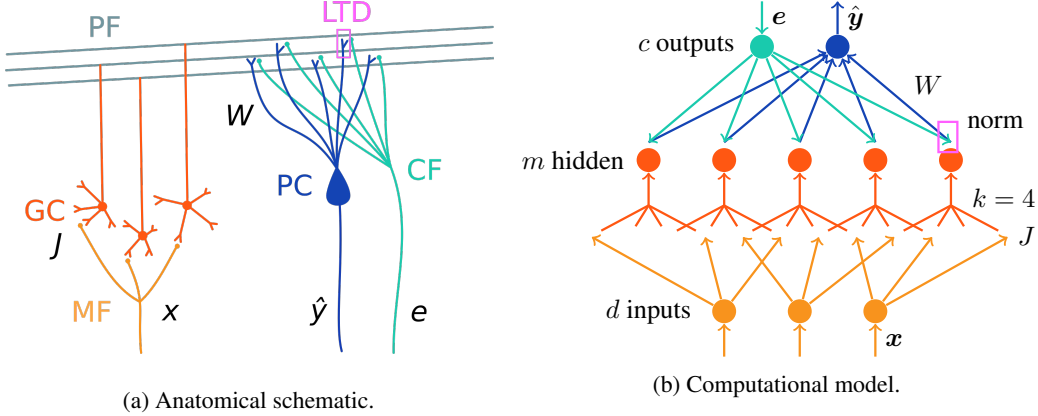


Figure 1: The schematic diagrams of cerebellum. (a) Anatomical schematic. MF: mossy fiber, GC: granule cell, PC: Purkinje cell, CF: climbing fiber, LTD: long-term depression. (b) Computational model. Please note the corresponding relationship between the anatomical and computational model.

to one single PC output, there are approximately $d = 7,000$ MFs, $m = 200,000$ GCs, and one PC in the rat brain [13]. By connectivity, the MF-GC layer is sparsely connected—one GC receives $k = 4$ MF inputs on average [20], while the PF-PC layer is densely connected—up to 200,000 GCs are connected with one single PC output.

The special one-to-one relationship between the CF and PC motivated Marr to predict that the PF-PC synapses are facilitated by the presynaptic and CF (postsynaptic) activities with Hebbian learning rule, and no other synapses are modifiable. Later studies showed that the cerebellar nuclei (CN, the output from PC) and inferior olive (IO, the input to CF) form a close loop to generate instructive signals, and CFs are carrying error signals [15, 17], which essentially established the error propagation mechanism through a single layer in the cerebellar cortex. LTD [31–33] was also discovered on the PF-PC synapse—reduction in the efficacy occurs in excitatory synapses and lasting for hours after receiving a long series of stimulus.

Note that based on the listed facts, backpropagation (BP) through multiple layers is unlikely to be applicable in the cerebellum. The mechanism of error propagation is well established, which relies on the special one-to-one relationship of PC and FC. Since the feedback fibers are not discovered beyond the Purkinje layer, the cerebellum may not propagate the error signals to the granule layer or further, unless different mechanisms [34] have been evolved for the same function of error propagation in the cerebellum.

The absence of multi-layer BP in the cerebellum does not effect applying gradient-based adversarial attacks, as our model could be regard as a combination of fixed random projection and a linear readout trained with (single-layer) BP, i.e., is still end-to-end differentiable (see Sec. 3.2). Our preliminary experiments also confirm that such a setting does not provide additional benefits on robustness comparing to the conventional BP.

3.2 Model Setup

Following Marr’s theory, the cerebellum could be modeled as a simple two-layer rate neural network, as shown in Fig 1b. Consider the inputs $\mathbf{x} \in \mathbb{R}^d$ on MFs, the granule layer responses are given by

$$\mathbf{h} = \sigma(J^T \mathbf{x}), \quad (1)$$

where $J \in \mathbb{R}^{d \times m}$ is the granule synaptic weights matrix which projects inputs to m dimensional space, and $\sigma(\cdot)$ is the ReLU activation function. Note that J is a sparse matrix—each row contains only k non-zero entries. The indices of the non-zero entries are randomly selected (randomly connected, RC). The random connection setting is consider to be acceptable both experimentally [35] and theoretically [20]. The non-zero weights are initialized by Xavier initialization, thus $J_{ij} \sim \frac{1}{\sqrt{k}} \mathcal{N}(0, 1)$, and fixed after initialization.

The Purkinje layer outputs are given by

$$\hat{\mathbf{y}} = W^T \mathbf{z}, \quad \mathbf{z} = \mathbf{h} - \boldsymbol{\mu}, \quad (2)$$

where $W \in \mathbb{R}^{m \times c}$ is the Purkinje synaptic weights matrix which maps the responses to c classes and $\boldsymbol{\mu}$ is the normalization term. Since LTD is discovered in PF-PC synapses, we propose to model it as the normalization technique in machine learning. For the sake of biological plausibility, we consider the simple centralization with moving average, thus $\boldsymbol{\mu}_t = \beta \boldsymbol{\mu}_{t-1} + (1 - \beta) \mathbf{h}_t$.

Here we wish to remark on the reasons we model the LTD as the normalization mechanism. Although the specific functional role of LTD is still unclear, we argue that it shares the same spirit as the “activity-dependent reduction in synapses efficiency through a long time” with normalization techniques in ML, which are usually found critical to reflect biological phenomena [36]. Our preliminary results also suggest such normalization may be useful to enhance robustness.

Now we describe the learning rule in PC. For some training instance (\mathbf{x}, \mathbf{y}) , consider the MSE loss

$$l = \frac{1}{2} \|\hat{\mathbf{y}} - \mathbf{y}\|_2^2 = \frac{1}{2} \sum_{j=1}^c (\hat{y}_j - y_j)^2. \quad (3)$$

By gradient descent, the updates of Purkinje synaptic weights

$$\Delta W \propto \frac{\partial l}{\partial W} = \mathbf{e}^T \mathbf{z}, \quad (4)$$

where $\mathbf{e} = \hat{\mathbf{y}} - \mathbf{y}$ is the error vector transmitted by the CF. Such subtraction error calculation could be supported by the inhibitory connection in the CN-IO loop. Note that the update rule is consistent with Hebbian learning—as synaptic updates are proportional to the product of pre- (\mathbf{z}) and post- (\mathbf{e}) synaptic activities.

4 Experimental Protocol

We now present the experimental protocol on evaluating the adversarial robustness of our cerebellum model. We will focus on three primary characteristics revealed in anatomy as discussed above:

- **Network width.** Why does the cerebellum develop the shallow and wide architecture for supervised learning? Will the shallow and wide structure provide extra benefits compared with deep networks?
- **Long-term depression.** What are the potential benefits of the LTD effect in PF-PC synapses on robustness?
- **Sparse connectivity.** Will the extra sparse connectivity ($k = 4$ vs. $m = 200,000$) in the granule layer provide advantages on robustness compared with fully-connected or other configuration of k ?

All the code, pre-trained models, and scripts necessary to reproduce our results will be made publicly available in the results paper.

Robustness measurement The robustness is defined as the stability against input perturbations. A learning system sensitive to small perturbations will face great challenges in real-world environments with noisy nature. Although we could in general measure robustness with arbitrary disturbances [37], adversarial examples provide a more rigorous evaluation estimating the worst-case robustness [3]. Therefore, we propose to study the adversarial robustness in the proposal, and use the perturbed accuracy under adversarial attacks as the robustness metric.

Primary datasets and training setup To keep the results comparable with machine learning researches, we will conduct our evaluations on the MNIST and CIFAR-10 datasets. Prior studies [21, 38] have already shown that the plain cerebellum-like architectures are capable of solving the MNIST dataset, therefore no additional modifications are needed for our cerebellum model on MNIST. We will use MNIST as the primary benchmark for the experiments in network width, LTD, and sparse connectivity. For the final evaluation of cerebellum-like models and ablative studies, we will include results on both MNIST and CIFAR-10.

However, since the cerebellum does not involve any receptive field (convolution) structures, it could be challenging to apply the model to CIFAR-10 without modification. Here we plan to deploy a visual feature encoder to resolve the problem. Specifically, we will first train a standard VGG-11 model (8 convolution and 3 fully-connected layers) as the baseline. We will apply standard data augmenting, such as random cropping and horizontal flipping during baseline training. After acquiring a desirable visual model of CIFAR-10, we will utilize the 8 convolution layers (discarding FC layers) as the visual feature encoder, mapping the CIFAR-10 dataset to a representation space ($d = 4096$). We will further train our cerebellum model on top of the visual feature encoder, keep the convolution layers fixed while training the cerebellum. Attacking will also be applied to the same model with the visual encoder and the cerebellum backend to keep the perturbation level (ϵ) consistent in the original image space.

Secondary training tasks Despite the MNIST and CIFAR-10, we could also conduct more biologically realistic tasks for modeling cerebellar learning, such as eyelid conditioning [39] or pole balancing through reinforcement learning [21]. However, since such tasks are not widely applicable in adversarial learning research, their significance for evaluating adversarial robustness is relatively limited (for instance, it will be hard to define proper perturbation level and make straightforward comparisons across models). Furthermore, those tasks are found less challenging than MNIST [21]. We here wish to justify our motivation for performing “visual recognition” tasks with the cerebellum model—the consideration is more from adversarial robustness evaluation than the biological plausibility. We will append the experiments on eyelid conditioning and pole balancing as secondary results if time permits.

Threat models We will consider two white-box attacks under l_∞ norms as our threat models: the simple FGSM [2] attack, as well as more advanced PGD [4] attack. For MNIST, we consider L_∞ perturbation $\epsilon = 0.1, 0.3$; for CIFAR-10, we consider L_∞ perturbation $\epsilon = 2/255, 8/255$. The attack settings are standard practices [4, 6] and also recommended by Carlini et al. [3].

Network width The first aspect we propose to study is the effect of network width on adversarial robustness. Following the previous setup, we consider the fully-connected (FC) model in the granule layer with a different number of GCs (m ranges from 1000 up to 50,000). We are expected to observe that the network robustness increases along with its width, as reported by Madry et al. [4]. However, the significance of such an effect is still unclear. Is the wide and shallow structure a decisive factor to improve adversarial robustness in the cerebellum? Or, is it simply a biological drawback as the cerebellum fails to evolve the sophisticated backpropagation algorithm?

Long-term depression To evaluate the effect of LTD, we will repeat the experiments for evaluating the effect of network width. The only difference is that we will apply normalization techniques (moving average) in the PC-FC synapses. We wish to study whether the LTD mechanism will improve the robustness in the cerebellum model.

Sparse connectivity We will next investigate sparse connectivity in the granule layer. We propose to consider different levels of sparsity, namely $k = 1, 2, 4, 10, 50, 200, 784$ (FC). For a fair comparison, we will control the total number of synapses in the models, as in Litwin-Kumar et al. [20]. Since each GC receives k MF inputs, and corresponds to 1 PC and 1 CF synapse for each output (c classes in total), the number of synapses is given by $m(k + 2c)$. For the cerebellum model, $k = 4$, $m = 200,000$, and $c = 10$ is decided by anatomy and datasets, thus the number of synapses is $4.8M$, and the number of GCs is given by $m = 4.8M/(k + 2c)$. In this experiment, we wish to understand if sparse connections in the granule layer will bring advantages to robustness in the cerebellum.

Cerebellum models and ablative studies By with/without the LTD and sparse connectivity, we could eventually evaluate the robustness across four cerebellum-like models: from the base model which is fully connected in the granule layer without LTD ($k = d = 784$, $m = 5970$ on MNIST, $k = d = 4096$, $m = 1166$ on CIFAR-10), to the cerebellum model which is sparsely connected with LTD ($k = 4$, $m = 200,000$). We will also examine whether the parameters decided by the natural cerebellum ($k = 4$) is (near) optimal against adversarial attacks.

Hyperparameters Hyperparameters for training and attacking the cerebellum model are listed in Table 1. The initially proposed values are based on previous studies [38, 4, 6] and our preliminary

experiments. We will tune the hyperparameters during attacking to ensure that our evaluation does not suffer from gradient masking (see discussion below).

Table 1: Initially proposed hyperparameters for evaluating the cerebellum model.

Hyperparameter	Value
Training batch size	10
Epoch	10
Learning rate	$0.5/m$
Optimizer	RMSprop
Optimizer decay	0.99
LTD β	0.99
Seed	{0, 123}
Attacking batch size	10
PGD steps	40
PGD step size	0.01

Gradient masking and robustness evaluation Gradient masking is the major concern in applying gradient-based iterative methods, which will provide a false sense of robustness if not examined properly [3]. The FGSM attack will be used as a sanity test for the PGD attack, as iterative attacks should always perform better than single-step methods. Gradient distribution plots in attacking will be utilized to avoid gradient masking. We will carefully tune the attacking hyperparameters and perform a doubling iterations test to certify that the PGD attack converges (i.e., the fooling rate does not further increase with respect to the number of iterations). We will also conduct an unbounded attack to ensure that our attacking methods have been applied properly.

Random seeds and transferability analysis Since our models involve randomness (especially in the RC granule layer), we propose to report results with different random seeds. We will also perform a transferability attack and analysis across different initialized models [40, 4]. Better robustness against transferable adversarial examples is also expected in the cerebellum model.

References

- [1] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [2] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [3] Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian Goodfellow, Aleksander Madry, and Alexey Kurakin. On evaluating adversarial robustness. *arXiv preprint arXiv:1902.06705*, 2019.
- [4] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- [5] Guneet S Dhillon, Kamyar Azizzadenesheli, Zachary C Lipton, Jeremy Bernstein, Jean Kos-saifi, Aran Khanna, and Anima Anandkumar. Stochastic activation pruning for robust adversarial defense. *arXiv preprint arXiv:1803.01442*, 2018.
- [6] Mislav Balunovic and Martin Vechev. Adversarial training and provable defenses: Bridging the gap. In *International Conference on Learning Representations*, 2020.
- [7] Aran Nayebi and Surya Ganguli. Biologically inspired protection of deep networks from adversarial attacks. *arXiv preprint arXiv:1703.09202*, 2017.

- [8] Wieland Brendel and Matthias Bethge. Comment on "biologically inspired protection of deep networks from adversarial attacks". *arXiv preprint arXiv:1704.01547*, 2017.
- [9] Zhe Li, Wieland Brendel, Edgar Walker, Erick Cobos, Taliah Muhammad, Jacob Reimer, Matthias Bethge, Fabian Sinz, Zachary Pitkow, and Andreas Tolias. Learning from brains how to regularize machines. In *Advances in Neural Information Processing Systems*, pages 9529–9539, 2019.
- [10] Gamaleldin Elsayed, Shreya Shankar, Brian Cheung, Nicolas Papernot, Alexey Kurakin, Ian Goodfellow, and Jascha Sohl-Dickstein. Adversarial examples that fool both computer vision and time-limited humans. In *Advances in Neural Information Processing Systems*, pages 3910–3920, 2018.
- [11] Zhenglong Zhou and Chaz Firestone. Humans can decipher adversarial images. *Nature communications*, 10(1):1–9, 2019.
- [12] Joel Dapello, Tiago Marques, Martin Schrimpf, Franziska Geiger, David D Cox, and James J DiCarlo. Simulating a primary visual cortex at the front of cnns improves robustness to image perturbations. *BioRxiv*, 2020.
- [13] David Marr. A theory of cerebellar cortex. *The Journal of Physiology*, 202:437–470, 1969.
- [14] James S Albus. A theory of cerebellar function. *Mathematical Biosciences*, 10(1-2):25–61, 1971.
- [15] Richard Apps and Richard Hawkes. Cerebellar cortical organization: a one-map hypothesis. *Nature Reviews Neuroscience*, 10(9):670–681, 2009.
- [16] Joseph Chaumont, Nicolas Guyon, Antoine M Valera, Guillaume P Dugué, Daniela Popa, Paikan Marcaggi, Vanessa Gautheron, Sophie Reibel-Foisset, Stéphane Dieudonné, Aline Stephan, et al. Clusters of cerebellar purkinje cells control their afferent climbing fiber discharge. *Proceedings of the National Academy of Sciences*, 110(40):16223–16228, 2013.
- [17] Jennifer L Raymond and Javier F Medina. Computational principles of supervised learning in the cerebellum. *Annual review of neuroscience*, 41:233–253, 2018.
- [18] Masao Ito, Masaki Sakurai, and Pavich Tongroach. Climbing fibre induced depression of both mossy fibre responsiveness and glutamate sensitivity of cerebellar purkinje cells. *The Journal of Physiology*, 324(1):113–134, 1982.
- [19] N Alex Cayco-Gajic, Claudia Clopath, and R Angus Silver. Sparse synaptic connectivity is required for decorrelation and pattern separation in feedforward networks. *Nature communications*, 8(1):1–11, 2017.
- [20] Ashok Litwin-Kumar, Kameron Decker Harris, Richard Axel, Haim Sompolinsky, and LF Abbott. Optimal degrees of synaptic connectivity. *Neuron*, 93(5):1153–1164, 2017.
- [21] Matthew Hausknecht, Wen-Ke Li, Michael Mauk, and Peter Stone. Machine learning capabilities of a simulated cerebellum. *IEEE transactions on neural networks and learning systems*, 28(3):510–522, 2016.
- [22] Tadashi Yamazaki and Jun Igarashi. Realtime cerebellum: A large-scale spiking network model of the cerebellum that runs in realtime using a graphics processing unit. *Neural Networks*, 47:103–111, 2013.
- [23] Hiroshi Yamaura, Jun Igarashi, and Tadashi Yamazaki. Simulation of a human-scale cerebellar network model on the k computer. *Frontiers in Neuroinformatics*, 14:16, 2020.
- [24] Nicolas Papernot, Patrick McDaniel, Arunesh Sinha, and Michael Wellman. Towards the science of security and privacy in machine learning. *arXiv preprint arXiv:1611.03814*, 2016.
- [25] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. Ensemble adversarial training: Attacks and defenses. *arXiv preprint arXiv:1705.07204*, 2017.

- [26] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. *arXiv preprint arXiv:1802.00420*, 2018.
- [27] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, and Alan Yuille. Mitigating adversarial effects through randomization. *arXiv preprint arXiv:1711.01991*, 2017.
- [28] Xuanqing Liu, Minhao Cheng, Huan Zhang, and Cho-Jui Hsieh. Towards robust neural networks via random self-ensemble. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 369–385, 2018.
- [29] Anish Athalye, Logan Engstrom, Andrew Ilyas, and Kevin Kwok. Synthesizing robust adversarial examples. *arXiv preprint arXiv:1707.07397*, 2017.
- [30] Eric Wong, Frank Schmidt, Jan Hendrik Metzen, and J Zico Kolter. Scaling provable adversarial defenses. In *Advances in Neural Information Processing Systems*, pages 8400–8409, 2018.
- [31] Masao Ito and Masanobu Kano. Long-lasting depression of parallel fiber-purkinje cell transmission induced by conjunctive stimulation of parallel fibers and climbing fibers in the cerebellar cortex. *Neuroscience letters*, 33(3):253–258, 1982.
- [32] Masao Ito. Long-term depression. *Annual review of neuroscience*, 12(1):85–102, 1989.
- [33] Tomoo Hirano. Long-term depression and other synaptic plasticity in the cerebellum. *Proceedings of the Japan Academy, Series B*, 89(5):183–195, 2013.
- [34] Jordan Guerguiev, Timothy P Lillicrap, and Blake A Richards. Towards deep learning with segregated dendrites. *Elife*, 6:e22901, 2017.
- [35] Eyal Gruntman and Glenn C Turner. Integration of the olfactory code across dendritic claws of single mushroom body neurons. *Nature neuroscience*, 16(12):1821, 2013.
- [36] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 215–223, 2011.
- [37] Guangyu Robert Yang, Peter Yiliu Wang, Yi Sun, Ashok Litwin-Kumar, Richard Axel, and L.F. Abbott. Evolving the olfactory system. *NeurIPS Workshop on Neuro AI*, 2019.
- [38] Bernd Illing, Wulfram Gerstner, and Johanni Brea. Biologically plausible deep learning but how far can we go with shallow networks? *Neural Networks*, 118:90–101, 2019.
- [39] Javier F Medina and Michael D Mauk. Computer simulation of cerebellar information processing. *nature neuroscience*, 3(11):1205–1211, 2000.
- [40] Nicolas Papernot, Patrick McDaniel, and Ian Goodfellow. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *arXiv preprint arXiv:1605.07277*, 2016.