

---

# Immuno-mimetic Deep Neural Networks (Immuno-Net)

---

Ren Wang<sup>1</sup> Tianqi Chen<sup>1</sup> Stephen Lindsly<sup>1</sup> Cooper Stansbury<sup>1</sup> Indika Rajapakse<sup>1</sup> Alfred Hero<sup>1</sup>

## Abstract

Biomimetics has played a key role in the evolution of artificial neural networks. Thus far, *in silico* metaphors have been dominated by concepts from neuroscience and cognitive psychology. In this paper we introduce a different type of biomimetic model, one that borrows concepts from the immune system, for designing robust deep neural networks. This immuno-mimetic model leads to a new computational biology framework for robustification of deep neural networks against adversarial attacks. Within this Immuno-Net framework we define a robust adaptive immune-inspired learning system (Immuno-Net RAILS) that emulates, *in silico*, the adaptive biological mechanisms of B-cells that are used to defend a mammalian host against pathogenic attacks. When applied to image classification tasks on benchmark datasets, we demonstrate that Immuno-net RAILS results in improvement of as much as 12.5% in adversarial accuracy of a baseline method, the DkNN-robustified CNN, without appreciable loss of accuracy on clean data.

## 1. Introduction

Primarily motivated by cognitive neuroscience, deep neural networks (DNNs) have reached impressive performance on various tasks (Mehdipour Ghazi and Kemal Ekenel, 2016; Zhao, Zheng, Xu, and Wu, 2019; Young, Hazarika, Poria, and Cambria, 2018). The neuro-mimetic point of view has driven both the development of machine learning architectures and machine learning procedures. For example, early inspiration for the convolutional neural network (CNN) came from models of the visual cortex of the brain (Lindsay, 2020) and the popular machine learning frameworks of transfer learning lifelong learning are firmly rooted in cognitive psychology (Thrun & Pratt, 2012; Parisi et al., 2019).

---

<sup>1</sup>University of Michigan, Michigan, United States. Correspondence to: Alfred Hero <hero@eecs.umich.edu>, Indika Rajapakse <indikar@umich.edu>, Ren Wang <renwang@umich.edu>.

However, the extraordinary performance of current DNNs on clean data still have a frustrating lack of robustness to small covariate shifts, e.g., adversarial attacks or drift of test data from their nominal training distributions. In the natural world such robustness is instantiated by another highly evolved biological system: the adaptive immune system. Motivated by analogies between robust machine learning and robustification mechanisms implemented by the immune system, we propose an immuno-mimetic framework for deep neural networks (Immuno-Net) that complements the neuro-cognitive perspective. In this paper, we define the Immuno-Net framework and develop an artificial robust adaptive immune learning system (RAILS), which is inspired by the mechanisms of B-cell affinity maturation in the mammalian immune system.

Similar to the human neural-cognitive system (Elsayed, Shankar, Cheung, Papernot, Kurakin, Goodfellow, and Sohl-Dickstein, 2018), DNNs are vulnerable to small perturbations, e.g., from adversarial attacks (Goodfellow, Shlens, and Szegedy, 2014). Conversely, the mammalian immune system has the inherent capability to defend against a diverse set of attacks (bacterial, viral, fungal, and other infections) due to its built-in discriminator between self and non-self components (Farmer, Packard, and Perelson, 1986). Through mechanisms of selection, hypermutation, and proliferation the adaptive immune system produces B-cells with increasing affinity to the external threat and is able to employ these solutions to further increase robustness by learning from multiple such attacks (Mesin, Ersching, and Victora, 2016; Mesin, Schiepers, Ersching, Barbulescu, Cavazzoni, Angelini, Okada, Kurosaki, and Victora, 2020). Furthermore, the immune system is able to further increase robustness by adaptively learning from multiple attacks (Mesin, Schiepers, Ersching, Barbulescu, Cavazzoni, Angelini, Okada, Kurosaki, and Victora, 2020).

The proposed Immuno-Net RAILS approach emulates the biological immune system *in silico*.

Our framework differs significantly from existing DNN robustification methods such as: outlier detection (Metzen, Genewein, Fischer, and Bischoff, 2017; Feinman, Curtin, Shintre, and Gardner, 2017; Grosse, Manoharan, Papernot, Backes, and McDaniel, 2017), and training-phase DNNs robustification (Madry, Makelov, Schmidt, Tsipras, and Vladu,

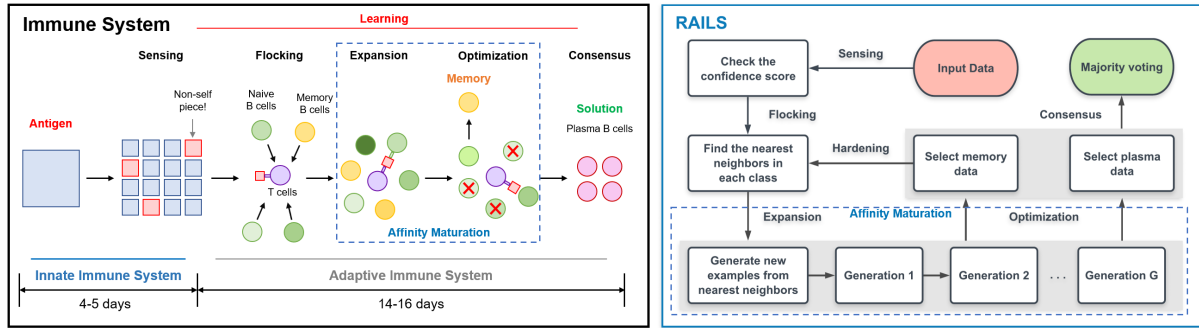


Figure 1. Simplified biological immune system (left) and ImmunoNet RAILS computational workflow (right). Both systems are composed of a five-step process. We consider affinity maturation (De Silva & Klein, 2015) as the combination of expansion and optimization.

2017; Zhang, Yu, Jiao, Xing, Ghaoui, and Jordan, 2019; Cohen, Rosenfeld, and Kolter, 2019) in that the immune system emulation provides a complete end-to-end solution to the adversarial robustness problem.

Specifically, we make the following contributions:

- Immuno-Net RAILS is the first robust adversarial defense framework for DNNs, based on the natural adaptive immune system. In particular: (1) Immuno-Net RAILS emulates the principal mechanisms of the immune response to defend against novel attacks; (2) The learning rates exhibited by RAILS align with those of the immune system; and (3) Immuno-Net RAILS can further improve robustness by emulating adaptive learning (life-long learning) in the immune system.
- We demonstrate that Immuno-Net RAILS improves adversarial robustness of the Deep k-Nearest Neighbor (DkNN) CNN (Papernot and McDaniel, 2018) by 5.62%, 12.5%, 10.32% for the MNIST, SVHN, and CIFAR-10 datasets, respectively.

## 2. Immune System to Computational System

### 2.1. Simplified Immune System

The mammalian adaptive immune system has designed robustness into its biological architecture (Rajapakse and Groudine, 2011). The adaptive immune system is highly complex, so we simplify its learning process into five sub-processes: sensing, flocking, expansion, optimization, and consensus (Cucker and Smale, 2007; Rajapakse and Smale, 2017). As shown in Fig. 1, *sensing* of an attack leads to initial B-cells *flocking* to lymph nodes, and forming temporary structures called germinal centers (De Silva and Klein, 2015; Farmer, Packard, and Perelson, 1986). Germinal centers are populated in the *expansion* and *optimization* phases (together called affinity maturation), where a diverse set of B-cells bearing antigen-specific immunoglobulins divide symmetrically and asymmetrically with mutation and B-

cells with high affinity to the antigen are selected to the next generation (Mesin et al., 2016). Memory B-cells are selected and stored within the evolving process, and can be used to defend against similar future attacks. In the *consensus* phase, plasma B-cells are generated by selecting those B-cells that reach a threshold affinity against the foreign antigen. Plasma B-cells (along with their generated antibodies) represent the optimal solution of the adaptive immune response to a specific attack.

### 2.2. From Immune Response to Computation

Motivated by the simplified biological immune system, we propose a new immuno-mimetic learning strategy - the Immuno-Net Robust Adversarial Immune-inspired Learning System (Immuno-Net RAILS). The comparison in Fig. 1 shows that both the biology in the immune system and the computational biology in Immuno-Net RAILS are composed of the same five-step process. Similar to the B-cell population growth during the clonal expansion, RAILS enlarges the population of candidates then selects high-affinity and moderate-affinity candidates to proliferate to predict the present inputs and defend against future attacks.

Fig. 2 shows the learning curves of Immuno-Net RAILS and the immune system and demonstrates that Immuno-Net RAILS captures key properties of the immune system. The learning curves depict the change of affinity between the population of defenders (B-cells) and the threat (antigen) over time. Note that in the computational system, the antigen is the test data. The initial B-cells (initial candidates) are all from the flocking phase of response to antigen 1 (test data 1), which results in a mismatch to antigen 2 (test data 2), and thus cause the red affinity curves to remain relatively low. One can also see that both green curves meet a small drop at the early stage and monotonically increase thereafter. This two-phase learning process happens due to a diversity vs. selection trade-off; diversity comes from the population size increasing and the randomness in generating new data points, while the selection arises from selection of high-affinity B-cells (newly generated data points).

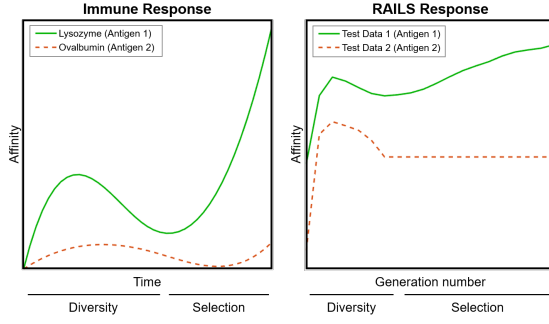


Figure 2. Correspondence between learning rates of Immuno-Net RAILS computational biology model (right) and observed learning rates of *in vitro* B-cells in the natural immune system (left). The flatten region in the red curve (right) corresponds to Immuno-Net RAILS early stopping when there is no solution. The biological experiments were performed in the Rajapakse Lab.

---

#### Algorithm 1 Immuno-Net RAILS

**Require:** Test data point  $\mathbf{x}$ ; Training dataset  $\mathcal{D}_{tr} = \{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_C\}$ ; Number of Classes  $C$ ; Model  $M$  with feature mapping  $f_l(\cdot)$ ,  $l \in \mathcal{L}$ ; Affinity function  $A$ .

##### Sensing

- 1: Check the threat score to detect the threat of  $x$ .

##### Flocking

- 2: **for**  $c = 1, 2, \dots, C$  **do**
- 3: In each layer  $l \in \mathcal{L}$ , find the  $k$ -nearest neighbors  $\mathcal{N}_l^c$  of  $\mathbf{x}$  in  $\mathcal{D}_c$  by ranking the affinity score  $A(f_l; \mathbf{x}_j, \mathbf{x})$ ,  $\mathbf{x}_j \in \mathcal{D}_c$

##### Affinity maturation: Expansion and Optimization

- 4: **end for**
  - 5: **For** each layer  $l \in \mathcal{L}$ , **do**
  - 6: Generate  $\mathcal{P}^{(0)}$  through mutating each of  $\mathbf{x}' \in \text{Union}(\{\mathcal{N}_l^1, \mathcal{N}_l^2, \dots, \mathcal{N}_l^C\}) T/kC$  times.
  - 7: **for**  $g = 1, 2, \dots, G$  **do**
  - 8: **for**  $t = 1, 2, \dots, T$  **do**
  - 9: Select data-label pair  $(\mathbf{x}_p, y_p)$  from  $\mathcal{P}^{(g-1)}$  based on  $\mathbf{P}_{g-1} = \text{Softmax}(A(f_l; \mathcal{P}^{(g-1)}, \mathbf{x})/\tau)$ .
  - 10: Pick all the data  $S_{y_p}$  belonging to class  $y_p$  in  $\mathcal{P}^{(g-1)}$  and calculate  $\mathbf{P}_{y_p} = \frac{\mathbf{P}_{g-1}(S_{y_p})}{\sum \mathbf{P}_{g-1}(S_{y_p})}$
  - 11: Select another pair  $(\mathbf{x}'_p, y_p)$  and apply cross-over and mutation on the two pairs of data. Add the obtained offspring  $\mathbf{x}_{os}$  to  $\mathcal{P}^{(g)}$ .
  - 12: **end for**
  - 13: **end for**
  - 14: Calculate the affinity score  $A(f_l; \mathcal{P}^{(G)}, \mathbf{x})$ ,  $\forall l \in \mathcal{L}$
  - 15: Select the top 5% as plasma data  $S_p^l$  and the top 25% as memory data  $S_m^l$  based on the affinity scores,  $\forall l \in \mathcal{L}$
  - 16: **end For**
  - 17: Obtain the prediction  $y$  of  $\mathbf{x}$  using the majority vote of the  $S_p = \{S_p^1, S_p^2, \dots, S_p^{|\mathcal{L}|}\}$
  - 18: **Output:**  $y$ , the memory data  $S_m = \{S_m^1, S_m^2, \dots, S_m^{|\mathcal{L}|}\}$
- 

### 3. Immuno-Net RAILS Computations

In this section, we describe the implementation of Immuno-Net RAILS as *in silico* analogs to flocking through consensus phases in immune response (Fig. 1). Given the mapping from input to layer  $l$ ,  $f_l: \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$  and  $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^d$ , we first define the affinity  $A(f_l; \mathbf{x}_1, \mathbf{x}_2)$  between  $\mathbf{x}_1$  and  $\mathbf{x}_2$  on layer  $l$ . The affinity can be measured in different ways, e.g, inverse  $\ell_p$  distance, cosine similarity, or inner product. Here we use the negative Euclidean distance measurement, i.e.,  $A(f_l; \mathbf{x}_1, \mathbf{x}_2) = -\|f_l(\mathbf{x}_1) - f_l(\mathbf{x}_2)\|_2$ . Algorithm 1 depicts Immuno-Net RAILS’ five-step workflow. We explain each step in detail in the subsequent section.

#### 3.1. Five-Step Workflow

**Sensing.** This step emulates the immune system’s self/non-self detection and performs the initial discrimination between perturbed inputs and clean inputs. The sensing step can effectively prevent Immuno-Net RAILS from becoming overwhelmed by false positives, i.e., treating clean inputs as adversarial inputs. This is implemented using an outlier detection procedure, for which there are several different methods available (Feinman, Curtin, Shintre, and Gardner, 2017; Xu, Evans, and Qi, 2017). The sensing stage provides a confidence score of the architecture and does not affect Immuno-Net RAILS’s predictions.

**Flocking.** Similar to the B-cells flocking to lymph nodes, Immuno-Net RAILS leverages flocking to provide an initial population for the affinity maturation process. At the implementation level, flocking is used to find the  $k$ -nearest neighbors that have the highest initial affinity score to the input data in each selected layer and from each class. We construct kNN sets  $\mathcal{N}_l^c$  independently for each class  $c$ , thereby ensuring that the initial population for every class is balanced.

$$\begin{aligned} \mathcal{N}_l^c &= \{(\hat{\mathbf{x}}, y_c) | R_c(\hat{\mathbf{x}}) \leq k, (\hat{\mathbf{x}}, y_c) \in \mathcal{D}_c\} \\ \text{Given} \\ A(f_l; \mathbf{x}_i^c, \mathbf{x}) &\leq A(f_l; \mathbf{x}_j^c, \mathbf{x}) \Leftrightarrow R_c(i) > R_c(j) \\ \forall c \in [C], l \in \mathcal{L}, \forall i, j \in [n_c], \end{aligned} \quad (1)$$

where  $\mathcal{L}$  is the set of the selected layers.  $\mathbf{x}$  denotes the input.  $\mathcal{D}_c$  represents the training data from class  $c$  with the size  $|\mathcal{D}_c| = n_c$ .  $R_c: [n_c] \rightarrow [n_c]$  is a ranking function that sorts the indices based on the affinity score. The candidates for flocking might purely come from the original training data or include the memory database which will be clear in the optimization stage introduction.

**Expansion and Optimization.** The expansion step starts from the initial population, and generates new examples (which are offspring) from the existing examples (which are parents) in the population. The “initial B-cells” are nearest neighbors found by (1) in the flocking step, where there are

$kC$  data points in total. We then enlarge the population size to  $T$  by copying each nearest neighbor  $\frac{T}{kC}$  times with random mutations. We call the first  $T$  size population as the 0-th generation. We now introduce the rule of generating new generations. Let  $\mathcal{P}^{(g-1)} = [\mathbf{x}_p(1), \mathbf{x}_p(2), \dots, \mathbf{x}_p(T)] \in \mathbb{R}^{d \times T}$  denote the population in the  $(g-1)$ -st generation. Then the candidates for generating the  $g$ -th generation are selected by

$$\hat{\mathcal{P}}^{(g-1)} = \mathcal{P}^{(g-1)} Z_{g-1}, \quad (2)$$

where  $Z_{g-1} \in \mathbb{R}^{T \times T}$  is a binary selection matrix whose columns are independent and identically distributed draws from  $\text{Mult}(1, \mathbf{P})$ , the multinomial distribution with selection probability vector  $\mathbf{P} \in [0, 1]^T$ . The selection probability for each candidate is obtained through a softmax function.

$$\begin{aligned} \mathbf{P}(\mathbf{x}_i) &= \text{Softmax}(A(f_i; \mathbf{x}_i, \mathbf{x})/\tau) \\ &= \frac{\exp(A(f_i; \mathbf{x}_i, \mathbf{x})/\tau)}{\sum_{\mathbf{x}_j \in S} \exp(A(f_i; \mathbf{x}_j, \mathbf{x})/\tau)}, \end{aligned} \quad (3)$$

where  $S$  is the set of data points and  $x_i \in S$ .  $\tau > 0$  is the sampling temperature that controls sharpness of the softmax operation.

We further leverage *cross-over* and *mutation* operations to increase candidate diversity. In Immuno-Net RAILS, we select two parents from the selected candidates (2) that share the same label, and apply the cross-over operator to combine the two parents together. The cross-over combination is implemented by selecting at random each of its elements (pixels) from the corresponding element of either parent1 or parent2 based on their affinity score. After the combination, mutation is leveraged to generate the offspring  $\mathbf{x}_{os}$  that belongs to the  $(g-1)$ -st generation. This operation randomly and independently mutates the data with probability  $\rho$ , adding uniformly distributed noise in the range  $[-\delta_{\max}, -\delta_{\min}] \cup [\delta_{\min}, \delta_{\max}]$ . Given examples lie in  $[0, 1]^d$ , the resulting perturbation vector is subsequently clipped to satisfy the domain constraint.

After multiple generations, RAILS selects generated examples with high-affinity scores to be plasma data, and examples with moderate-affinity scores are saved as memory data. Similar to plasma B-cells, plasma data represents an optimal solution to the current attack. Memory data can be seen as an analog to memory B-cells, which help warm-start defense responses to similar attacks in future. We set the threshold to be 0.05 and 0.25 for selecting plasma data and memory data, respectively, i.e., top 5% of final generation will become plasma data, and top 25% will become memory data.  $S_p = \{S_p^1, S_p^2, \dots, S_p^{|\mathcal{L}|}\}$  denotes the plasma data we selected from each layer. The memory data is denoted by  $S_m = \{S_m^1, S_m^2, \dots, S_m^{|\mathcal{L}|}\}$ . In this paper, we will only focus on static learning, i.e., only leveraging the plasma data to make predictions.

**Consensus.** All the new examples generated in each generation are associated with labels that are inherited from their parents. The consensus step uses majority voting of the plasma data for prediction of the class label of input  $\mathbf{x}$ .

## 4. Experimental Results

We compare Immuno-Net RAILS to standard Convolutional Neural Network Classification (CNN) and Deep k-Nearest Neighbors Classification (DkNN) (Papernot & McDaniel, 2018) on the MNIST (LeCun et al., 1998), SVHN (Netzer et al., 2011), and CIFAR-10 (Krizhevsky & Hinton, 2009). In addition to the benign test examples, we also generate the same amount of adversarial examples using PGD attack (Madry et al., 2017). The attack strength is  $\epsilon = 60/8/8$  for MNIST, SVHN, and CIFAR-10 by default. The performance will be measured by standard accuracy (SA) evaluated using benign test examples and robust accuracy (RA) evaluated using the adversarial test examples.

Immuno-Net RAILS leverages static learning to make the predictions. The results are shown in Table 1. One can see that RAILS delivers a 5.62%/10.32%/12.5% improvements in RA over DkNN without appreciable loss of SA on the three datasets.

Table 1. SA/RA performance of Immuno-Net RAIL (RAILS) versus CNN and DkNN. Immuno-Net RAILS obtains higher RA for all three datasets.

		SA	RA
MNIST ( $\epsilon = 60$ )	<b>RAILS (ours)</b>	97.95%	<b>76.67%</b>
	CNN	<b>99.16%</b>	1.01%
	DkNN	97.99%	71.05%
SVHN ( $\epsilon = 8$ )	<b>RAILS (ours)</b>	90.62%	<b>48.26%</b>
	CNN	<b>94.55%</b>	1.66%
	DkNN	93.18%	35.7%
CIFAR-10 ( $\epsilon = 8$ )	<b>RAILS (ours)</b>	82%	<b>52.01%</b>
	CNN	<b>87.26%</b>	32.57%
	DkNN	86.63%	41.69%

## 5. Conclusion

Inspired by the immune system, we proposed a new defense framework for deep learning models. Immuno-Net RAILS is able to emulate biological processes with evolutionary programming that model natural selection and affinity maturation of B-cells in the natural immune system. Our experiments show that the Immuno-Net RAILS learning curve mimics the diversity-selection learning phases observed in *in vitro* B-cell affinity maturation experiments in the Rajapakse lab.

In future work, we will explore the mechanisms of the immune system’s adaptive learning (life-long learning) and covariate shift adjustment, which will be consolidated into our computational framework.

## Acknowledgement

We thank Alanawaz Rehemtulla and Sivakumar Jeyarajan for helpful discussions and guidance on the adaptive immune system. This work is supported by the Guaranteeing AI Robustness against Deception (GARD) program from DARPA/I2O.

## References

- Cohen, J., Rosenfeld, E., and Kolter, Z. Certified adversarial robustness via randomized smoothing. In *International Conference on Machine Learning*, pp. 1310–1320, 2019.
- Cucker, F. and Smale, S. Emergent behavior in flocks. *IEEE Transactions on automatic control*, 52(5):852–862, 2007.
- De Silva, N. S. and Klein, U. Dynamics of b cells in germinal centres. *Nature reviews immunology*, 15(3):137–148, 2015.
- Elsayed, G. F., Shankar, S., Cheung, B., Papernot, N., Kurakin, A., Goodfellow, I., and Sohl-Dickstein, J. Adversarial examples that fool both computer vision and time-limited humans. In *Advances in Neural Information Processing Systems*, 2018.
- Farmer, J. D., Packard, N. H., and Perelson, A. S. The immune system, adaptation, and machine learning. *Physica D: Nonlinear Phenomena*, 22(1-3):187–204, 1986.
- Feinman, R., Curtin, R. R., Shintre, S., and Gardner, A. B. Detecting adversarial samples from artifacts. *arXiv preprint arXiv:1703.00410*, 2017.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- Grosse, K., Manoharan, P., Papernot, N., Backes, M., and McDaniel, P. On the (statistical) detection of adversarial examples. *arXiv preprint arXiv:1702.06280*, 2017.
- Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. *Master’s thesis, Department of Computer Science, University of Toronto*, 2009.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, Nov 1998. ISSN 0018-9219. doi: 10.1109/5.726791.
- Lindsay, G. W. Convolutional neural networks as a model of the visual system: past, present, and future. *Journal of cognitive neuroscience*, pp. 1–15, 2020.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- Mehdipour Ghazi, M. and Kemal Ekenel, H. A comprehensive analysis of deep learning based representation for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 34–41, 2016.
- Mesin, L., Ersching, J., and Vitorica, G. D. Germinal center b cell dynamics. *Immunity*, 45(3):471–482, 2016.
- Mesin, L., Schiepers, A., Ersching, J., Barbulescu, A., Cavazzoni, C. B., Angelini, A., Okada, T., Kurosaki, T., and Vitorica, G. D. Restricted clonality and limited germinal center reentry characterize memory b cell reactivation by boosting. *Cell*, 180(1):92–106, 2020.
- Metzen, J. H., Genewein, T., Fischer, V., and Bischoff, B. On detecting adversarial perturbations. In *International Conference on Learning Representations*, 2017.
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. 2011.
- Papernot, N. and McDaniel, P. Deep k-nearest neighbors: Towards confident, interpretable and robust deep learning. *arXiv preprint arXiv:1803.04765*, 2018.
- Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., and Wermter, S. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019.
- Rajapakse, I. and Groudine, M. On emerging nuclear order. *Journal of Cell Biology*, 192(5):711–721, 2011.
- Rajapakse, I. and Smale, S. Emergence of function from coordinated cells in a tissue. *Proceedings of the National Academy of Sciences*, 114(7):1462–1467, 2017.
- Thrun, S. and Pratt, L. *Learning to learn*. Springer Science & Business Media, 2012.
- Xu, W., Evans, D., and Qi, Y. Feature squeezing: Detecting adversarial examples in deep neural networks. *arXiv preprint arXiv:1704.01155*, 2017.
- Young, T., Hazarika, D., Poria, S., and Cambria, E. Recent trends in deep learning based natural language processing. *IEEE Computational Intelligence Magazine*, 13(3):55–75, 2018.
- Zhang, H., Yu, Y., Jiao, J., Xing, E. P., Ghaoui, L. E., and Jordan, M. I. Theoretically principled trade-off between robustness and accuracy. *International Conference on Machine Learning*, 2019.
- Zhao, Z.-Q., Zheng, P., Xu, S.-t., and Wu, X. Object detection with deep learning: A review. *IEEE transactions on neural networks and learning systems*, 30(11):3212–3232, 2019.