

---

# Easy Batch Normalization

---

Arip Asadulaev<sup>1,2</sup> Alexander Panfilov<sup>3</sup> Andrey Filchenkov<sup>1</sup>

## Abstract

It was shown that adversarial examples improve object recognition. But what about their opposite side, easy examples? Easy examples are samples that the machine learning model classifies correctly with high confidence. In our paper, we are making the first step toward exploring the potential benefits of using easy examples in the training procedure of neural networks. We propose to use an auxiliary batch normalization for easy examples for the standard and robust accuracy improvement.

## 1. Introduction

Samples from the training dataset that the neural network **classifies correctly** with high confidence after the first training epoch are called easy examples (Arpit et al., 2017). Recent results showed that easy examples could be removed from the training dataset without influencing the model generalization (Kishida & Nakayama, 2019). Although, we suppose that we can utilize easy examples to improve neural classifiers’ standard and robust accuracy by enhancing the training procedure.

However, the number of easy examples depends on the chosen classifier and dataset. To avoid this limitation and obtain more easy examples, we propose to use **targeted adversarial attacks** (Yuan et al., 2019) with a ground truth target label to turn all training samples into their easy copies. In Section (3), we show that targeted adversarial attacks actually make samples easy to classify. When we obtain easy examples, **we propose to use an auxiliary batch norm** to address them, similarly as it was proposed for adversarial examples (Xie et al., 2020). Our hypothesis is that network’s **layers trained with easy examples can ”clean up” the adversarial data points** that lie away from the original data

---

<sup>1</sup>ITMO University, Saint-Petersburg, Russia <sup>2</sup>Artificial Intelligence Research Institute, Moscow, Russia <sup>3</sup>Tubingen Universität, Germany. Correspondence to: Arip Asadulaev <aripasadulaev@itmo.ru>.

manifold and help to achieve a practical trade-off between the standard and robust accuracy. We provide experiments and show that easy batch normalization achieves improvement in robust accuracy similar to the original AdvProp (Xie et al., 2020) and even outperforms it in the standard accuracy on FashionMNIST and CIFAR-10 (Krizhevsky, 2009) datasets.

## 2. Background

**Adversarial examples:** Adversarial attacks are intended to “fool” the classifier (Szegedy et al., 2014; Papernot et al., 2017; Yuan et al., 2019; Schott et al., 2019). Using the sample  $x$ , the target label  $y$ , and model with parameters  $\theta$ , we can apply Projected Gradient Descent (PGD) (Madry et al., 2018) iterations to get the adversarial example:

$$x_{i+1} = \text{Proj}_{(x,\varepsilon)} [x_i + \alpha \text{sign}(-\nabla_x L(\theta, x_i, y_i))] \quad (1)$$

Where,  $\text{Proj}_{(x,\varepsilon)}$  is a projection operator onto the  $l_{\text{inf}}$  ball of radius  $\varepsilon$  around the original image  $x_i$ . The underlying data distribution of the adversarial examples is different from the real training data, and training on these examples improves the robust accuracy but decreases the standard one. However, recent results show applications of the adversarial examples for model improvement. The (Xie et al., 2020) proposed to use the adversarial examples in the separate batch normalization (i.e., AdvProp) and showed that it could improve the standard and adversarial accuracy of a model.

**Easy examples:** There is no formal definition for easy examples yet. In our study, we follow the simplest notation proposed in the (Arpit et al., 2017): instances that the model classifies accurately with high confidence are called easy examples.

## 3. Easy batch normalization

The number of easy examples depends on the chosen classifier and dataset. We propose to obtain easy examples by gradient perturbations to avoid this limitation, such examples we called perturbed easy examples. Using gradient descent, we can add “easy features” to the data similarly to how we add adversarial features. The only difference is that

the  $y_i$  in (1) must be equal to the ground truth class label. Such examples are similar to the gradient-based preferable inputs that were used to understand and visualize neural network features (Simonyan et al., 2014; Yosinski et al., 2015; Mahendran & Vedaldi, 2016).

**Do gradients provide easy examples?** With this experiment, we justify that gradient perturbations are indeed adding to the image features that are easy to classify correctly.

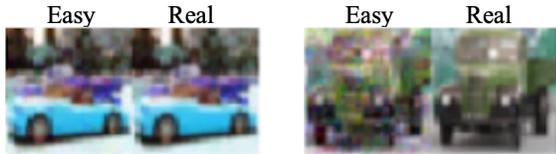


Figure 1: Input class “0”    Figure 2: Input class “1”

As the first step, we trained ResNet-18 one epoch on the CIFAR-10 dataset and collected easy examples for each image by gradient perturbations. Then we took the new randomly initialized ResNet-18 and set up a binary classification task based on the CIFAR-10 dataset. We picked the class “car” as “0” and class “truck” as a class “1”. Then, each sample was connected with its easy version, see Figure 1, 2. After training a new model on these “stacked” samples, we added noise to one of the image parts and evaluated the accuracy. The idea of this experiment is that: if the easy part of the image is actually easy, the model will learn to classify based only on this part and “do not pay attention” to the other part. As a result, the accuracy drops drastically from 99.0  $\rightarrow$  0.0 if we randomize the easy part of the image and do not change 99.0  $\rightarrow$  99.0 otherwise. This result fortifies the proposition that neural networks tend to classify based on easy examples rather than the standard ones.

**Easy batch normalization:** When we obtain easy examples, we propose to use a separate auxiliary batch normalization for them during the training, similarly as it was done for the adversarial examples in AdvProp. We hypothesize that neural network layers that use the easy batch normalization learn to map hidden state statistics back to parts of the data manifold where the classifier network performs well.

## 4. Experiments

**Datasets:** We tested our method on FashionMNIST and CIFAR-10 datasets. These datasets consist of the 50000 training samples and 10000 test samples. We used standard data pre-processing for ResNet models without any data augmentations.

**Settings:** We employed ResNet-18 (He et al., 2016) and ResNet-50 architectures, used in PyTorch (Paszke et al., 2019) AdvProp realization (Han). The ResNets were adapted for the CIFAR-10 similarly to (Zhang et al., 2019)

with wider channels and more parameters than the original. For adversarial and easy examples generation, we used a 1-step  $l_{inf}$  PGD attack similar to the original AdvProp method.

For each experiment, training parameters were identical: 1000 training epochs, SGD optimizer, and cosine annealing learning rate scheduler. The learning rate warm-up schedule was equal to 0.1, 0.0005, and 0.0005 values after the 1, 200, and 450 epochs. To test robust accuracy, a 100-step  $l_2$  PGD attack from robustness library (Engstrom et al., 2019) was used. In addition, robustness was evaluated with a standard version of AutoAttack (AA) (Croce & Hein, 2020) which is a parameter-free and user-independent ensemble of attacks. For AA perturbation size was set to  $\epsilon = 8/255$  for  $l_{inf}$  and  $\epsilon = 0.5$  for  $l_2$  to keep performance comparable with models presented in RobustBench (Croce et al., 2020).

**Results:** See results in Table 1, 2. Easy batch normalization (EBN) meets both goals and improves standard and robust accuracy for most cases, compared to the vanilla AdvProp (AP). The network learns to neutralize adversarial perturbations and maps hidden states back to the data manifold where the network performs well on the training dataset.

Model	Basic	$\epsilon_{l_2} 0.1$	$\epsilon_{l_2} 0.5$	$\epsilon_{l_2} 1.0$	$AA_{l_{inf}}$	$AA_{l_2}$
RN-18 (AP)	94.23	82.82	79.42	76.20	86.87	<b>90.90</b>
RN-18 (EBN)	<b>94.28</b>	<b>86.39</b>	<b>83.27</b>	<b>80.31</b>	<b>86.93</b>	90.89
RN-50 (AP)	94.40	<b>62.03</b>	60.05	57.46	85.28	<b>90.27</b>
RN-50 (EBN)	<b>94.45</b>	63.02	<b>60.65</b>	<b>58.63</b>	<b>85.38</b>	90.20

Table 1: Average results over three random seeds of ResNet (RN) models trained on FashionMNIST

Model	Basic	$\epsilon_{l_2} 0.1$	$\epsilon_{l_2} 0.5$	$\epsilon_{l_2} 1.0$	$AA_{l_{inf}}$	$AA_{l_2}$
RN-18 (AP)	93.62	36.74	29.17	22.82	43.12	<b>77.85</b>
RN-18 (EBN)	<b>93.81</b>	<b>37.49</b>	<b>30.08</b>	<b>24.39</b>	<b>43.75</b>	77.19
RN-50 (AP)	94.55	<b>38.43</b>	29.39	21.96	45.65	<b>78.34</b>
RN-50 (EBN)	<b>94.62</b>	38.01	<b>30.76</b>	<b>23.54</b>	<b>46.47</b>	78.05

Table 2: Average results over three random seeds of ResNet (RN) models trained on CIFAR-10

## 5. Conclusion and future works

We propose a simple remedy to improve the accuracy of classification networks. We experimented with FashionMNIST and CIFAR-10 datasets and showed that our method outperforms basic training settings. Further, we plan to test our method with the recent advancement of the AdvProp method: Fast AdvProp (Mei et al., 2022) and extend our investigation to more complex datasets with more general robustness measurement (Weng et al., 2018). In future work we plan to study the latent space representations of the networks trained with easy batch normalization. Especially we plan to evaluate the idea that **easy examples simplifies decision boundary** and bias neural models to find **lower rank solutions that generalize well** (Huh et al., 2021).

## References

- Arpit, D., Jastrzebski, S., Ballas, N., Krueger, D., Bengio, E., Kanwal, M. S., Maharaj, T., Fischer, A., Courville, A. C., Bengio, Y., and Lacoste-Julien, S. A closer look at memorization in deep networks. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 233–242. PMLR, 2017. URL <http://proceedings.mlr.press/v70/arpit17a.html>.
- Croce, F. and Hein, M. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *International conference on machine learning*, pp. 2206–2216. PMLR, 2020.
- Croce, F., Andriushchenko, M., Sehwag, V., DeBenedetti, E., Flammarion, N., Chiang, M., Mittal, P., and Hein, M. Robustbench: a standardized adversarial robustness benchmark. *arXiv preprint arXiv:2010.09670*, 2020.
- Engstrom, L., Ilyas, A., Salman, H., Santurkar, S., and Tsipras, D. Robustness (python library), 2019. URL <https://github.com/MadryLab/robustness>.
- Han, Y. Pytorch implementation of advprop. URL <https://github.com/tingxueronghua/pytorch-classification-advprop>.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 770–778. IEEE Computer Society, 2016. doi: 10.1109/CVPR.2016.90. URL <https://doi.org/10.1109/CVPR.2016.90>.
- Huh, M., Mobahi, H., Zhang, R., Cheung, B., Agrawal, P., and Isola, P. The low-rank simplicity bias in deep networks. *CoRR*, abs/2103.10427, 2021. URL <https://arxiv.org/abs/2103.10427>.
- Kishida, I. and Nakayama, H. Empirical study of easy and hard examples in CNN training. In Gedeon, T., Wong, K. W., and Lee, M. (eds.), *Neural Information Processing - 26th International Conference, ICONIP 2019, Sydney, NSW, Australia, December 12-15, 2019, Proceedings, Part IV*, volume 1142 of *Communications in Computer and Information Science*, pp. 179–188. Springer, 2019. doi: 10.1007/978-3-030-36808-1\_20. URL [https://doi.org/10.1007/978-3-030-36808-1\\_20](https://doi.org/10.1007/978-3-030-36808-1_20).
- Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, 2009.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018. URL <https://openreview.net/forum?id=rJzIBfZAb>.
- Mahendran, A. and Vedaldi, A. Visualizing deep convolutional neural networks using natural pre-images. *Int. J. Comput. Vis.*, 120(3):233–255, 2016. doi: 10.1007/s11263-016-0911-8. URL <https://doi.org/10.1007/s11263-016-0911-8>.
- Mei, J., Han, Y., Bai, Y., Zhang, Y., Li, Y., Li, X., Yuille, A., and Xie, C. Fast advprop. *arXiv preprint arXiv:2204.09838*, 2022.
- Papernot, N., McDaniel, P. D., Goodfellow, I. J., Jha, S., Celik, Z. B., and Swami, A. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security, AsiaCCS 2017, Abu Dhabi, United Arab Emirates, April 2-6, 2017*, pp. 506–519, 2017. doi: 10.1145/3052973.3053009. URL <https://doi.org/10.1145/3052973.3053009>.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance.pdf>.
- Schott, L., Rauber, J., Bethge, M., and Brendel, W. Towards the first adversarially robust neural network model on MNIST. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019. URL <https://openreview.net/forum?id=S1EHOsC9tX>.
- Simonyan, K., Vedaldi, A., and Zisserman, A. Deep inside convolutional networks: Visualising image classification models and saliency maps. In Bengio, Y. and LeCun, Y. (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Workshop Track Proceedings*, 2014. URL <http://arxiv.org/abs/1312.6034>.

- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. J., and Fergus, R. Intriguing properties of neural networks. In Bengio, Y. and LeCun, Y. (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. URL <http://arxiv.org/abs/1312.6199>.
- Weng, T.-W., Zhang, H., Chen, P.-Y., Yi, J., Su, D., Gao, Y., Hsieh, C.-J., and Daniel, L. Evaluating the robustness of neural networks: An extreme value theory approach. *arXiv preprint arXiv:1801.10578*, 2018.
- Xie, C., Tan, M., Gong, B., Wang, J., Yuille, A. L., and Le, Q. V. Adversarial examples improve image recognition. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pp. 816–825. IEEE, 2020. doi: 10.1109/CVPR42600.2020.00090. URL <https://doi.org/10.1109/CVPR42600.2020.00090>.
- Yosinski, J., Clune, J., Nguyen, A. M., Fuchs, T. J., and Lipson, H. Understanding neural networks through deep visualization. *CoRR*, abs/1506.06579, 2015. URL <http://arxiv.org/abs/1506.06579>.
- Yuan, X., He, P., Zhu, Q., and Li, X. Adversarial examples: Attacks and defenses for deep learning. *IEEE Trans. Neural Netw. Learning Syst.*, 30(9):2805–2824, 2019. doi: 10.1109/TNNLS.2018.2886017. URL <https://doi.org/10.1109/TNNLS.2018.2886017>.
- Zhang, M., Lucas, J., Ba, J., and Hinton, G. E. Lookahead optimizer: k steps forward, 1 step back. *Advances in Neural Information Processing Systems*, 32:9597–9608, 2019.