

Towards Efficient Data-Centric Robust Machine Learning with Noise-based Augmentation

Xiaogeng Liu, Haoyu Wang, Yechao Zhang, Fangzhou Wu, Shengshan Hu

School of Cyber Science and Engineering
National Engineering Research Center for Big Data Technology and System
Services Computing Technology and System Lab
Hubei Engineering Research Center on Big Data Security
Huazhong University of Science and Technology
Wuhan, 430074, China

{liuxiaogeng, wwwwhy, ycz, warkzhou, hushengshan}@hust.edu.cn

Abstract

The data-centric machine learning aims to find effective ways to build appropriate datasets which can improve the performance of AI models. In this paper, we mainly focus on designing an efficient data-centric scheme to improve robustness for models towards unforeseen malicious inputs in the black-box test settings. Specifically, we introduce a noise-based data augmentation method which is composed of Gaussian Noise, Salt-and-Pepper noise, and the PGD adversarial perturbations. The proposed method is built on lightweight algorithms and proved highly effective based on comprehensive evaluations, showing good efficiency on computation cost and robustness enhancement. In addition, we share our insights about the data-centric robust machine learning gained from our experiments.

Introduction

With *deep neural networks* (DNNs) being deployed in more and more fields, how to improve the performance of DNN models has attracted many researchers. A lot of efforts have been made to construct powerful models (Simonyan and Zisserman 2015; He et al. 2016a,b; Zagoruyko and Komodakis 2016) that can run accurately on different data from distinct tasks, and even outperform human’s perception systems. However, the data, which can be regarded as “food” for the *artificial intelligence* (AI) models, has not been sufficiently investigated as well as the model structures or the training functions. Recently, such status brings concerns from some academics and people start to research about the data and call it *the data-centric machine learning* (Miranda 2021).

As illustrated in Fig. 1, the data-centric machine learning is mainly focused on the design of data. The researchers are usually asked to construct elaborately designed datasets for the fixed model structure and loss functions and improve the performance of the models. The data-centric machine learning has many differences with the model-centric machine learning, turning researchers’ eyes on data augmentation (Shorten and Khoshgoftaar 2019) and feature engineering (Dong and Liu 2018).

In this paper, we mainly focus on how to design a dataset that can give robustness for the trained models towards any

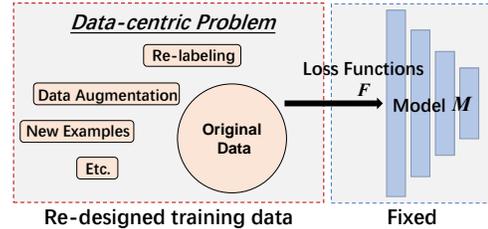


Figure 1: Data-centric machine learning mainly focuses on the design of the dataset rather than model structure or loss function.

potential malicious inputs, i.e., the data-centric robust machine learning in the black-box settings. Besides, we want to keep our solution as simple as possible to achieve efficiency in practice. We first share our analysis about this topic as follows: in general, a data-centric problem can be solved by multiple data augmentation techniques, such as image noise, knowledge distillation with soft label (Hinton, Vinyals, and Dean 2015), augmentation by affine transformations. However, for knowledge distillation with soft label, its reliability has a dependence on the performance of the teacher model while we find it hard to choose (or train) an appropriate teacher for robust learning in the black-box settings. In addition, defensive distillation (Papernot et al. 2016) is proven to be vulnerable to specific adversarial attacks (Carlini and Wagner 2017). For affine transformation, we argue that this method tends to empower the networks with robustness against angle transformation, partial information, etc., which are more likely to occur in the physical world. In this paper, we mainly focus on the digital world where image corruption and adversarial examples are more common and dangerous while affine transformations can hardly provide help on that. Consequently, we believe a more effective and efficient solution is using hard-label data with noise-based data augmentation including conventional image noise and adversarial perturbations (Goodfellow, Shlens, and Szegedy 2015; Madry et al. 2018), forcing the model to learn robust features from corrupted and adversarial images.

Another important and changeable component in data-centric problems is the settings of the (commonly-used) optimizer, especially for the learning rate scheduler. In our method, we choose the cosine annealing schedule with warm

restart (Loshchilov and Hutter 2017). Compared with other schemes and vanilla cosine annealing, this scheduler scans a wider range of learning rate values, provides a better generalization and accuracy, and accelerates the training process of models. In a nutshell, our solution is composed of noise-based data augmentation and cosine annealing learning rate scheduler with warm restart.

We summarize our contributions as follows:

- We introduce a noise-based method to construct a comprehensive training dataset to empower the trained model with strong robustness without dependence on any additional training process or loss function. In addition, this scheme is composed of lightweight algorithms and achieves efficiency in terms of effectiveness and computation cost.
- Specifically, we evaluate the performance of different image noise and adversarial perturbations for data augmentation comprehensively. Based on the concrete quantitative results, we choose Gaussian Noise, Salt-and-Pepper noise, and the PGD perturbations as integrated solutions.
- Our method shows good effectiveness on evaluations and ranks 8-th place out of 3, 691 teams on *AAAI2022 Security AI Challenger¹: Data Centric Robust Learning on ML models*, surpassing the baseline by 20.03% in terms of final score.

Methodology

Noise-based Data Augmentation

Here we highlight four kinds of image noise for data augmentation, including two kinds of conventional image noise (Salt & Pepper, Gaussian) and two kinds of adversarial perturbations (FGSM, PGD).

Salt and Pepper noise. In a corruption image, the Salt&Pepper noise is presented as pure white or black pixel with discrete distribution, which is often caused by sharp and sudden disturbances in the image signal. We argue that it also serves as random masks and makes models learn partial features for robustness.

Gaussian Noise. Adding Gaussian noise is a widely-used method for data augmentation. The probability density function of Gaussian noise follows the normal distribution (i.e, Gaussian distribution).

FGSM. The Fast Gradient Sign Method (FGSM) (Goodfellow, Shlens, and Szegedy 2015) is a single-step gradient-based adversarial attacks scheme, which calculates the gradient sign of output loss from a white-box model and then computes (usually l_∞ -norm) bounded adversarial perturbations as we demonstrate as follows:

$$x_{adv} = x + \eta, \quad \eta = \epsilon \text{sign}(\nabla_x J(\theta, x, y)), \quad (1)$$

where J represents the loss function of the white-box model, ϵ is the perturbation size, and y is the ground-truth label of the clean image x .

PGD. The Projected Gradient Descent (PGD) (Madry et al. 2018) is a strong iterative adversarial attack which generates adversarial examples by:

$$x^{t+1} = \Pi_{x+S} (x^t + \alpha \text{sign}(\nabla_x J(\theta, x, y))), \quad (2)$$

¹<https://tianchi.aliyun.com/competition/entrance/531939/information>

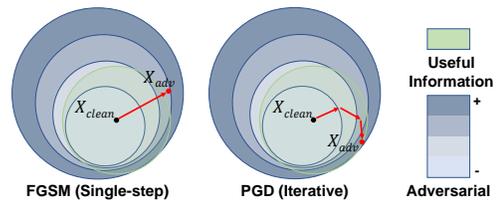


Figure 2: Illustration of “adversarial under-fitting” phenomenon. The darker blue represents the stronger attack ability of an adversarial example. The green region represents a range where the adversarial example can provide comprehensively useful information as training data. Due to the single-step design, the example generated by FGSM is likely to locate on a point with weaker adversarial strength, meanwhile losses more clean features.

where α is the perturbation size per-step, and Π_{x+S} represents the similarity restriction which is usually a clip function in practice.

It can be noted that both FGSM and PGD are calculated on a white-box model, which means a trained model is needed when we construct the new data from the original ones. In our method, we first train the networks on the original clean dataset, then we generate adversarial examples (FGSM, PGD) on them. This leads to a little difference from recent adversarial training (Pang et al. 2021) where adversarial examples are calculated in the training process simultaneously, while our method needs to train a clean model in the first stage due to the specific requirements of data-centric scenario.

Although we cannot modify training process like adversarial training, we argue that the choice of adversarial algorithms in data-centric machine learning is the same as what is widely adopted in adversarial training. In adversarial training, a majority of researchers use PGD (Pang et al. 2021; Zhang et al. 2019; Rice, Wong, and Kolter 2020) to generate adversarial examples, since PGD is proven to have stronger attack strength and unlikely to cause label leaking (Madry et al. 2018). Due to the single-step design, the attack strength of FGSM is usually weaker than PGD in the same perturbation size, and it also tends to create adversarial examples with more conspicuous changes. This means that when it comes to training models straightly on adversarial examples, FGSM may drive the model far from an optimal point where the model can handle different inputs including clean, corrupted, and adversarial images. As Fig. 2 shows, compared with FGSM, PGD is more likely to have stronger adversarial strength and maintain more clean features. We leave a more comprehensive investigation about this in the experiments.

Cosine Annealing and Warm Restart

Cosine annealing combined with warm restart (Loshchilov and Hutter 2017) is a popular and effective technique to solve optimization problems which accelerates the convergence and avoids ill-condition problems. Specifically, as demonstrated in Alg. 1, this scheme sets the learning rate of the optimizer according to the cosine annealing curve and restart after some iterations. In our method, we leverage

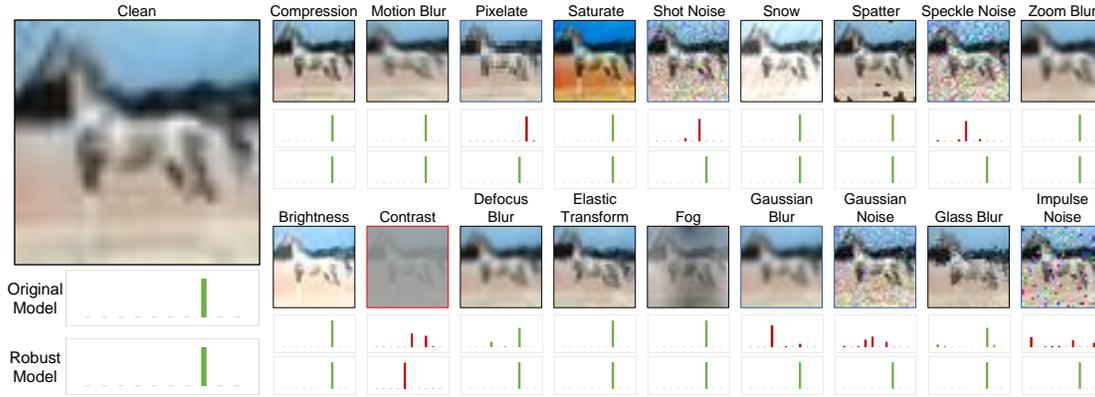


Figure 3: Images from the official test data of AAAI2022 *Security AI Challenger*. The bar charts below each image represent the softmax outputs from the original model (trained on the clean dataset) and the robust model (trained on our augmented dataset). The incorrect results are colored red. The softmax outputs show that our method enhances the robustness of the networks. And the enhanced model can give correct results with high confidence.

Algorithm 1: SGD with CosineAnnealingWarmRestarts

Input: Parameters θ_t ; minimum learning rate γ_{min} ; maximum learning rate γ_{max} ; loss function Φ ; weight decay λ ; momentum μ ; dampening τ ; the number of epochs since the last restart T_{cur} ; the number of epochs between two warm restarts T .

Output: Parameters θ_{t+1} .

```

1: CosineAnnealingWarmRestarts Schedule.
2: if  $T_{cur} == T$  then
3:    $\gamma = \gamma_{min}$ ;
4: else if  $T_{cur} == 0$  then
5:    $\gamma = \gamma_{max}$ ;
6: else
7:    $\gamma = \gamma_{min} + \frac{1}{2}(\gamma_{max} - \gamma_{min})(1 + \cos(\frac{T_{cur}}{T}\pi))$ ;
8: end if
9: Stochastic Gradient Descent.
10:  $g_{t+1} \leftarrow \nabla_{\theta}\Phi(\theta_t)$ 
11:  $g_{t+1} \leftarrow g_{t+1} + \lambda\theta_t$ 
12:  $g_{t+1} \leftarrow \mu g_t + (1 - \tau)g_{t+1}$ 
13:  $\theta_{t+1} \leftarrow \theta_t - \gamma g_{t+1}$ 
14: return  $\theta_{t+1}$ 

```

this technique to get a powerful Stochastic Gradient Descent (SGD) optimizer.

Experiments

Experimental Setting

Implementation details. In our experiments, we construct the training data with four noise-based augmentation methods, i.e., Gaussian noise, Salt&Pepper noise, FGSM (Goodfellow, Shlens, and Szegedy 2015) perturbations, and PGD (Madry et al. 2018) perturbations. For Gaussian noise, we set mean equals 0 and variance equals 0.005. For Salt&Pepper noise, we set the noise ratio to 40%. For FGSM, we set perturbation size to 8/255, 12/255, 16/255 respectively for the comparison study. For PGD, we choose the maximum perturbation equals 8/255, 12/255, 16/255,

epochs equal to 20, and step size equals 0.4/255, 0.6/255, 0.8/255 respectively. All these augmentations will be conducted upon the CIFAR10 (Krizhevsky, Hinton et al. 2009) training data, which contains 50,000 32×32 color images from 10 categories and each category have 5,000 training images.

Models. Following the semi-finals of AAAI2022 *Security AI Challenger*² competition, we use PreActResnet18 (He et al. 2016b) and WideResnet (Zagoruyko and Komodakis 2016) as the test model.

Test dataset. We choose the official test dataset of AAAI2022 *Security AI Challenger* competition, which is a mixed dataset containing clean images, corruption images, and adversarial images. All of these images are extracted or generated from the original CIFAR10 test set. The corruption image set contains 18 types of corruption, and the adversarial image set is generated from 3 different settings by model-ensemble (Dong et al. 2018). We visualize them in Fig. 3.

Optimizer and scheduler. In our experiments, we use a SGD (Robbins 2007) optimizer with the *CosineAnnealingWarmRestarts*³ scheduler to train the test model. Specifically, for SGD, we set the learning rate equals 0.01, momentum equals 0.9, and weight decay equals 5×10^{-4} . For the scheduler, we set $T_0 = 3$, $T_{mult} = 2$, and the minimum learning rate equals 1×10^{-5} . In addition, we set the batch size equals 128 and training epochs equal 185.

Comparison Study

Evaluations of different methods. Tab. 1 shows the comparison study of different noise-based augmentation methods. In this experiment, we evaluate the model trained on the dataset which contains 50,000 images constructed from the same original data with a single augmentation method, e.g., the GS represents the training dataset containing 50,000

²<https://tianchi.aliyun.com/competition/entrance/531939/information>

³https://pytorch.org/docs/master/generated/torch.optim.lr_scheduler.CosineAnnealingWarmRestarts.html

Test Data	Model	CL	GS	S&P	FGSM ₈	FGSM ₁₂	FGSM ₁₆	PGD ₈	PGD ₁₂	PGD ₁₆
Clean	PreActResnet18	0.9922	0.9391	0.9643	0.3888	0.2993	0.1869	0.7302	0.6121	0.4579
Images	WideResnet	0.9923	0.9350	0.9751	0.4518	0.2772	0.2388	0.8776	0.7515	0.5910
Corrupted	PreActResnet18	0.6353	0.6971	0.7162	0.2354	0.1981	0.1419	0.4410	0.4180	0.3818
Images	WideResnet	0.6108	0.6982	0.7265	0.3124	0.2143	0.1638	0.5986	0.4961	0.4315
Adversarial	PreActResnet18	0.3719	0.6433	0.5411	0.2758	0.2149	0.1591	0.7860	0.7500	0.6922
Images	WideResnet	0.3797	0.6506	0.5329	0.3455	0.2387	0.1984	0.8052	0.7754	0.7211
Avg.	Final Score	0.6637	0.7606	0.7427	0.3350	0.2404	0.1815	0.7064	0.6338	0.5459

Table 1: Comparison of noise-based data augmentation methods. Each column represents the training data (50,000 images) augmented by a single method. We highlight the top-3 for each kind of test data and the final score. Notations CL: clean images. GS: images with Gaussian noise. S&P: images with Salt&Pepper noise. The subscript number (8, 12, 16) represents the perturbation size with 8/255, 12/255, 16/255 for FGSM and PGD, respectively. We will use the same notations hereafter.

images only with Gaussian Noise and the FGSM₁₂ represents the training dataset containing 50,000 images only with FGSM perturbations with perturbation size = 12/255. As shown in Tab. 1, Gaussian noise, Salt&Pepper noise, and the PGD perturbations are appropriate methods for robustness enhancement in terms of the final score. Salt&Pepper noise shows superior effectiveness for corrupted images and the PGD perturbations have outstanding performance on the robustness towards adversarial images. Gaussian noise has a relatively balanced accuracy for each kind of image set, leading to its triumph on the final score.

Another non-negligible fact is that FGSM is not competitive to PGD for data-centric robust machine learning in the same perturbation size and this verifies our analysis in Fig 2. We also argue that FGSM in smaller perturbation sizes may be more appropriate for data augmentation.

Evaluations of mixed dataset. According to the performance of different noise-based augmentation methods in Tab. 1, we construct 10 mixed datasets with different proportions of the noise methods as shown in Tab. 2:

Mix.	1	2	3	4	5	6	7	8	9	10
Clean	-	-	-	-	-	-	-	10k	-	10k
GS	10k									
S&P	10k									
FGSM ₈	30k	-	-	-	-	-	10k	10k	-	-
FGSM ₁₂	-	30k	-	-	-	-	10k	10k	-	-
FGSM ₁₆	-	-	30k	-	-	-	10k	-	-	-
PGD ₈	-	-	-	30k	-	-	-	-	10k	10k
PGD ₁₂	-	-	-	-	30k	-	-	-	10k	10k
PGD ₁₆	-	-	-	-	-	30k	-	-	10k	-

Table 2: Contents in 10 different mixed datasets.

As shown in Fig. 4(a), an obvious tendency is that the datasets mixed with PGD perturbations, Gaussian noise, and Salt&Pepper noise will reach a balance between robustness and accuracy as they all have relatively high final scores, among which the dataset with PGD_{8/255} achieves the best performance.

Ablation Study

Fig. 4(b) shows the ablation study of the warm restarts design. In this experiment, we choose three mixed datasets which have the highest final scores in our evaluations and train the test model by cosine annealing scheduler without warm restarts. As illustrated in Fig. 4(b), the warm restarts design can improve the model performance by 2% ~ 3%, which is a significant progress.

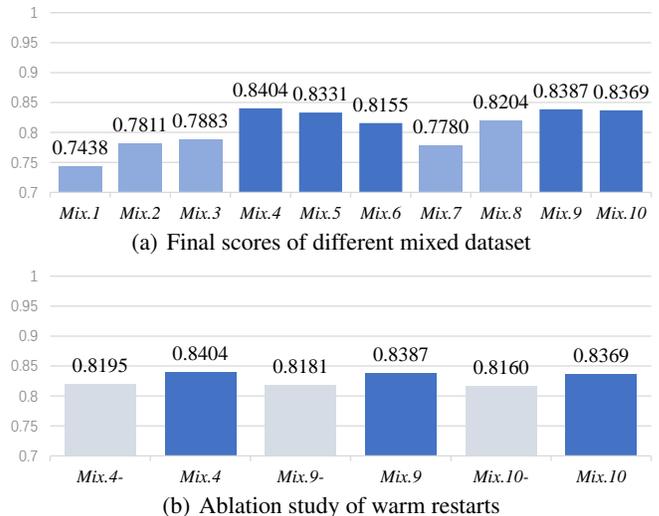


Figure 4: Notations are introduced in Tab. 2. (a) shows the final scores of different mixed datasets, the datasets with PGD are colored darker blue. (b) represents the final scores of the top-3 datasets in Fig. 4(a) with and without (bars with light grey) warm restarts design (but all have cosine annealing), showing that warm restarts leads to about 2% ~ 3% improvement.

Conclusions

- An appropriate scheme to generate adversarial examples is crucial. We argue that the samples generated by iterative algorithms (i.e., PGD) with a relatively small size of perturbations can maintain considerable clean features and avoid insufficiency of adversarial strength.
- Noise-based augmentation can perform better in data-centric settings compared with soft labels or affine transformations. And the appropriate noise type and hyper-parameters can ensure that the model is valid for both clean and perturbed images.
- *CosineAnnealingWarmRestarts* scheduler can avoid overfitting of the neural networks on limited data by restarting the learning rate, and thus enhance the generalization of the model. This method will lead to considerable improvements when it is combined with noise-based data augmentation.

References

- Carlini, N.; and Wagner, D. A. 2017. Towards Evaluating the Robustness of Neural Networks. In *2017 IEEE Symposium on Security and Privacy, S&P 2017*, 39–57.
- Dong, G.; and Liu, H. 2018. *Feature engineering for machine learning and data analytics*. CRC Press.
- Dong, Y.; Liao, F.; Pang, T.; Su, H.; Zhu, J.; Hu, X.; and Li, J. 2018. Boosting Adversarial Attacks With Momentum. In *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'18)*, 9185–9193.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2015. Explaining and Harnessing Adversarial Examples. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR'15)*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016a. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR'16)*, 770–778.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016b. Identity Mappings in Deep Residual Networks. In *Proceedings of the 14th European Conference on Computer Vision (ECCV'16)*, 630–645.
- Hinton, G. E.; Vinyals, O.; and Dean, J. 2015. Distilling the Knowledge in a Neural Network. *CoRR*, abs/1503.02531.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Loshchilov, I.; and Hutter, F. 2017. SGDR: Stochastic Gradient Descent with Warm Restarts. In *Proceedings of the 5th International Conference on Learning Representations (ICLR'17)*.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *Proceedings of the 6th International Conference on Learning Representations (ICLR'18)*.
- Miranda, L. J. 2021. Towards data-centric machine learning: a short review. *lvmiranda921.github.io*.
- Pang, T.; Yang, X.; Dong, Y.; Su, H.; and Zhu, J. 2021. Bag of Tricks for Adversarial Training. In *Proceedings of the 9th International Conference on Learning Representations (ICLR'21)*.
- Papernot, N.; McDaniel, P. D.; Wu, X.; Jha, S.; and Swami, A. 2016. Distillation as a Defense to Adversarial Perturbations Against Deep Neural Networks. In *IEEE Symposium on Security and Privacy, S&P 2016*, 582–597.
- Rice, L.; Wong, E.; and Kolter, J. Z. 2020. Overfitting in adversarially robust deep learning. In *Proceedings of the 37th International Conference on Machine Learning (ICML'20)*, volume 119, 8093–8104.
- Robbins, H. E. 2007. A Stochastic Approximation Method. *Annals of Mathematical Statistics*, 22: 400–407.
- Shorten, C.; and Khoshgoftaar, T. M. 2019. A survey on Image Data Augmentation for Deep Learning. *J. Big Data*, 6: 60.
- Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *3rd International Conference on Learning Representations, (ICLR'15)*.
- Zagoruyko, S.; and Komodakis, N. 2016. Wide Residual Networks. In *Proceedings of the British Machine Vision Conference 2016 (BMVC'16)*.
- Zhang, H.; Yu, Y.; Jiao, J.; Xing, E. P.; Ghaoui, L. E.; and Jordan, M. I. 2019. Theoretically Principled Trade-off between Robustness and Accuracy. In *Proceedings of the 36th International Conference on Machine Learning (ICML'19)*, volume 97, 7472–7482.