

Learning Data Augmentation with Online Bilevel Optimization for Image Classification

Saypraseuth Mounsaveng^{*1}, Issam Laradji², Ismail Ben Ayed¹, David Vázquez², Marco Pedersoli¹
¹ÉTS Montreal ²Element AI

1. Implementation details

We use PyTorch to implement our experiments, and Kornia library for the color transformations.

1.1. Model architecture

Tab. 1 shows the BadGAN classifier architecture. Tab. 2 shows the augmenter network for affine and color transformations.

Classifier C	
Input 32x32 Image	
3x3 conv. 96 LReLU(0.2)	
3x3 conv. 96 LReLU(0.2)	
3x3 conv. 96 LReLU(0.2), 0.5 dropout	
3x3 conv. 192 LReLU(0.2)	
3x3 conv. 192 LReLU(0.2)	
3x3 conv. 192 LReLU(0.2), 0.5 dropout	
3x3 conv. 192 LReLU(0.2)	
3x3 conv. 192 LReLU(0.2)	
3x3 conv. 192 LReLU(0.2)	
MLP 10 unit, sigmoid	
10-class Softmax	

Table 1: BadGAN classifier network.

Augmenter A		
Small	Medium	Large
Input n^* dim.	Input 100 dim.	Input 100 dim.
MLP n units	MLP 64 unit	MLP 512 unit
relu, 0.2 dropout	relu, 0.2 dropout	relu, 0.2 dropout
MLP $10 \times n$ units	MLP 32 unit	MLP 1024 unit
relu, 0.2 dropout	relu, 0.2 dropout	relu, 0.2 dropout
-	-	MLP 1024 unit
-	-	relu, 0.2 dropout
-	-	MLP 512 unit
-	-	relu, 0.2 dropout
MLP n units, tanh		

Table 2: Augmenter network for affine and color transformations. For *Small*, n is the number of parameters to learn (6 for affine, 4 for color and 10 when combining both).

1.2. Affine transformation parameters

In our experiments with affine transformations, the augmenter network learns a 2x3 matrix containing the parameters of the affine transformations.

1.3. Color transformation parameters

In our experiments with color transformations, the augmenter network outputs 1 value per transformation. Those values are in a range depending on the transformation: i) Hue in range [-0.5:0.5]; ii) Saturation in range [0:1]; iii) Contrast in range [-1:1] iv) Brightness in range [0:1]. In all cases the zero value corresponds to the identity transformation.

2. Additional experiments

We investigate the influence of augmenter regularization on the model performance. Figure 1 shows the model performance for different weight decay values. This indicates that the model is not sensitive to regularization.




Figure 1: Augmenter regularization. Accuracy for different weight decay values for the augmenter (affine, color and the combination of both) with ResNet18 on CIFAR10.