

This WACV 2021 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

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

Abstract

Data augmentation is a key practice in machine learning for improving generalization performance. However, finding the best data augmentation hyperparameters requires domain knowledge or a computationally demanding search. We address this issue by proposing an efficient approach to automatically train a network that learns an effective distribution of transformations to improve its generalization. Using bilevel optimization, we directly optimize the data augmentation parameters using a validation set. This framework can be used as a general solution to learn the optimal data augmentation jointly with an end task model like a classifier. Results show that our joint training method produces an image classification accuracy that is comparable to or better than carefully hand-crafted data augmentation. Yet, it does not need an expensive external validation loop on the data augmentation hyperparameters.

1. Introduction

Deep learning methods are based on large models in which the number of parameters is much higher than the dimensionality of the input data as well as the number of available samples [20, 46]. In this setting, overfitting is a major problem [49]. Standard regularization techniques applied directly to the model parameters only add very general knowledge about the parameter values, which leads to modest improvement in the final model accuracy [37]. Adding training samples artificially generated by applying predefined transformations to the initial samples, which is referred to as data augmentation, has shown to be a promising regularization technique to increase a model performance [21]. However, the selection of the best data augmentation is challenging and requires specific domain knowledge. Indeed, data transformations can be valid only for specific domains and heuristically chosen transformations, for example by transferring transformations useful in a domain into another can be counterproductive. For in-



Figure 1: **Model training.** In an epoch, the classifier parameters ω are trained in the standard supervised way in the inner loop. Jointly, the data augmentation parameters are trained on the validation set in the outer loop using an online differentiable method.

stance, a data augmentation transformation like horizontal flip is valid for natural images because, in nature, the horizontal mirror of an object is visually still a valid object. However, applied to a dataset containing numbers or letters, it can generate a non-existing symbol or even a different symbol of the alphabet, which would confuse the model training.

A simple way to define the best data augmentation is to use expert knowledge to define the best transformations and their parameters for a given dataset. However, this is not practical, as for each single dataset, an expert should be consulted to obtain a possibly useful set of transformations and their parameters, which is not always possible due to cost constraints or limited expert knowledge availability.

To mitigate this challenge, it is possible to select those transformations heuristically, and then find their optimal parameters using validation data. Conventionally, a hyperparameter search is performed across different sets of transformations, and the one that leads to the best validation accuracy is selected as the optimal set of transformations. This approach is appealing as it allows the model to learn the best transformations directly from the data, but it is not scalable. Given a large range of transformations to test, retraining the

^{*}Corresponding author: saypraseuth.mounsaveng.1@etsmtl.net

[†]Code is available at https://github.com/ElementAI/ bilevel_augment

algorithm every time with a different transformation set is computationally very demanding [9].

In this work, we address the problem of how to learn the data augmentation that maximizes the validation accuracy efficiently by proposing a method based on bilevel optimization. With this framework, we aim at identifying image transformations that minimize the validation loss while training the end task model. However, as described in section 3 and shown in Fig. 1, instead of solving the complete bilevel optimization problem, we approximate it with an online version where in every iteration a new set of transformations is learned and adapted to the learning phase.

We summarize our contributions as follows: i) we propose an online, differentiable approach for learning the optimal data augmentation regime using a validation set. As this method is differentiable, we can efficiently optimize a large transformation network that learns to perform data augmentation automatically; ii) we show that our proposed model using different sets of transformations achieve comparable or better results than conventional methods on five different datasets. Further improvements were shown with our method on a medical imaging dataset where effective transformations are difficult to define.

The remainder of the paper is structured as follows. We conduct a literature review in Section 2. In Section 3 we explain how to approximate our bilevel optimization problem such that we can jointly learn the optimal data augmentation and the classifier. In Section 4 we define the experimental setup. Finally, we present our experimental results and draw conclusions about the presented work in Section 5, and 6, respectively.

2. Related Work

Data augmentation consists in creating new data points from existing ones in order to get a larger training set. It was found to be essential for achieving state-of-the-art image classification results [22].

Data augmentation transformations are usually chosen heuristically based on expert domain knowledge. For natural images, usual transformations are image flip, rotation and color changes [41]. More complex transformations such as occluding parts of an image [12] or blending images [28, 53] seem also to be useful. These transformations can significantly improve the task performance. However, there is no guarantee that they are optimal nor that they are even useful at all. To avoid the manual selection of transformations, recent studies have investigated automatic data augmentation learning. We distinguish those methods between GAN-based and AutoAugment-based approaches.

GAN-based Generative Adversarial Networks (GANs) [18] can generate realistic new samples of a certain dataset or class, thus they can be adapted for data augmentation. Mirza and Osindero [35] and Odena et al.

[38] proposed to generate images conditioned on their class that could be directly used to augment a dataset. CatGAN [48] on the other hand, performs unsupervised and semi-supervised learning as a regularized information maximization problem [26] with a regularization based on the generated samples. Also based on GAN, but directly used for data augmentation, DAGAN [1] conditions the augmented image on the input image. TripleGAN [7] and Bayesian data augmentation [50] train a classifier jointly with the generator. These approaches generate general image transformations, but in practice, it is not as performant as using predefined transformations. TANDA [42] is the only GAN-based approach that uses predefined transformations. It defines a large set of transformations and learns how to combine them to generate new samples that follow the same distribution as the original data. This approach is better, but it is still based on the assumption that the augmented data should follow the same distribution as the original data. Instead, we argue that data augmentation should improve the performance of the classifier, independently from the visual similarity of the generated data.

AutoAugment AutoAugment [9] is a data augmentation method that learns sequences of transformations that maximize the classifier accuracy on a validation set. This objective is better than simply reproducing the same data distribution as in GAN-based models, as it favors transformations that generalize well on unseen data. However, it is computationally expensive as it performs the complete bilevel optimization by training the classifier in the inner loop until convergence for each set of evaluated transformations. Some solutions to reduce the computational cost are proposed in follow-up works. Fast AutoAugment [29] optimizes the search space by matching the density between the training set and the augmented data. Alternatively, Population Based Augmentation (PBA) [23] focuses on learning the optimal augmentation schedule rather than only the transformations. However, even if these approaches reduce the computational cost of AutoAugment, they do not leverage gradient information. Faster AutoAugment [19] does this by combining AutoAugment with a GAN discriminator and considering transformations as differentiable functions. OHL-Auto-Aug [30] uses an online bilevel optimization approach and the REINFORCE algorithm on an ensemble of classifiers to estimate the gradient of the validation loss and learn an augmentation probability distribution. RandAugment [10] goes further by showing that a same performance level as AutoAugment can be obtained by randomly selecting transformations from the predefined pool and just tune the number of transformations to use and a global (same for all transformations) magnitude factor. However, this approach also requires prior knowledge of useful transformations.

Hyperarameter Learning Our work has similarities with techniques used in the hyperparameter optimization field. Hyperparameters tuning is important to obtain the best performances when training neural networks on a given dataset. Classic approaches assume that the learning model is a black-box and use methods like grid search, random search [5, 6]. Bayesian optimisation [47], or a tree-search approach [24]. These approaches are simple but expensive because they repeat the optimization from scratch for each sampled value of the hyperparameters and so are only applicable to low dimensional hyper-parameter spaces. A different line of research is to leverage the gradient of these (continuous) hyperparameters (or hyper-gradients) to perform the hyper-optimization. The first work proposing this idea [3], shows that the implicit function theorem can be used to this aim. [13] was the first work to propose a gradient-based method using a bilevel optimization approach [8] to learn hyperparameters. Using a bilevel optimization approach to train a neural network is challenging, as usually there is no closed-form expression of the function learned in the inner loop (Section 3). To address this, Maclaurin et al. [34] and later Franceschi et al. [15] proposed methods to reverse the forward pass to compute the gradient of the validation loss. However, these methods are applicable only when the number of hyperparameters and the complexity of the models are limited due to the memory needed to save the intermediate steps. Another approach to address the computational hurdle in the inner loop is to calculate an approximation of the gradient like in Pedregosa [40] Luketina et al. [32] or MacKay et al. [33]. Our method differentiates from those by using truncated back propagation to estimate the gradient of the validation loss. Finally, note that hyper-parameter optimization presents some similarities to meta learning as shown in Franceschi et al. [16]. For instance, in MAML [14], a shared model initialization is learned to minimize the validation loss and therefore improve the generalization capabilities of the model.

3. Proposed Data Augmentation Method

Consider a labeled set $\mathcal{X} := \{x_i, y_i\}_{i=1}^N$, where x_i is an input image, y_i the associated class label, N the number of samples and $\hat{\mathcal{X}}$ the set of transformed images. We formulate the problem of identifying effective data augmentation transformations as a bilevel optimization problem. In this setup, the augmenter $\mathcal{A}_{\theta} : \mathcal{X} \to \hat{\mathcal{X}}$ is parametrized by θ and is used to minimize the loss \mathcal{L} on the validation data \mathcal{X}_{val} in the outer loop. In the inner loop, the classifier parameters ω are optimized on the training data \mathcal{X}_{tr} in the standard supervised way. This formulation can be written as:

$$\theta^* = \operatorname*{arg\,min}_{\theta} \mathcal{L}(\mathcal{X}_{val}, \omega^*) \tag{1}$$

s.t.
$$\omega^* = \operatorname*{arg\,min}_{\omega} \mathcal{L}(\mathcal{A}_{\theta}(\mathcal{X}_{tr}), \omega).$$
 (2)

While optimizing a few hyperparameters on the validation data is feasible with black-box approaches such as grid and random search [4] or Bayesian optimization [47], it is not efficient. With bilevel optimization, our aim is to efficiently learn an entire neural network A_{θ} (possibly with thousands of parameters θ) which defines a distribution of transformations that should be applied on the training data to improve generalization.

Gradient descent was shown to be an efficient method for optimizing parameters of large networks. In problems such as architecture search [31], the parameters can be directly optimized with gradient descent (or second order methods) against the training and validation data. However, this is not the case for data augmentation. The reason is that the transformation network \mathcal{A}_{θ} is optimized to maximize the validation score, but applies transformations only on the training set. Therefore, first order methods would not work. The aim of data augmentation is to introduce transformations during the training phase that can make the model invariant or partially invariant to any transformations that can occur at test time. If we optimize the transformation network directly on the validation data, the model will simply select trivial solutions such as the identity transformation. This approach has been used for object localization [25] and it did not improve the model generalization performance as much as data augmentation. To solve this issue, new methods relied on reinforcement learning instead of gradient descent to learn effective data augmentation [9, 23, 29].

In this work, we show that in the case of a differentiable augmenter \mathcal{A}_{θ} , there is a simple, efficient way to find optimal data transformations based on gradient descent that generalize well on validation data. We formulate our problem as an approximation to bilevel optimization by using truncated back-propagation as it allows our method to: i) efficiently estimate a large number of parameters to generate the optimal data augmentation transformations by gradient descent; ii) obtain an online estimation of the optimal data augmentation during the different phases of the training, which can also be beneficial [17]; iii) change the training data to adapt to different validation conditions as in supervised domain adaptation.

Although approximate bilevel optimization has already been proposed for hyperparameter optimization [15, 16, 45], in this paper we show that it can be used for training a large, complex model (the augmenter A_{θ} network) in order to learn an effective distribution of transformations.

3.1. Approximate Online Bilevel Optimization

As shown in Eq. 1 and 2, the problem of finding the optimal data augmentation transformations \mathcal{A}_{θ} can be cast as a bilevel optimization problem. This problem can be solved by iteratively solving Eq. 2 to find the optimal network weight ω^* , given the parameters of the transformation



Figure 2: Computational graph of our model at iteration t = J. *K* is the number of gradient unfolding steps, and J is the number of inner loop iterations after which θ gets updated. The case where K=J=T (T being the iteration of the classifier convergence) is the complete bilevel optimization as in Eq.1 whereas K=J=1 corresponds to updating θ at each mini-batch (K = 1), suing only one step of gradient unfolding (J = 1).

 θ and then updating θ :

$$\theta \leftarrow \theta - \eta_{\theta} \nabla_{\theta} \mathcal{L}(\mathcal{X}_{val}, \omega^*)$$
(3)

where η_{θ} is the learning rate used to train the augmenter network.

However, as the augmentations are to be applied only on the training dataset and not on the validation set, calculating $\frac{\partial \mathcal{L}(\mathcal{X}_{val},\omega^*)}{\partial \theta}$ is not trivial. To enable this calculation, we use the fact that the weights ω of the network are shared between training and validation data and use the chain rule to differentiate the validation loss $\mathcal{L}(\mathcal{X}_{val},\omega^*)$ with respect to the hyperparameters θ . In other words, instead of using a very slow black-box optimization for θ , we can exploit gradient information because the model parameters ω^* are shared between the validation and the training loss.

We define the gradient of the validation loss with respect to θ as follows:

$$\nabla_{\theta} \mathcal{L}(\mathcal{X}_{val}, \omega^{*}) = \frac{\partial \mathcal{L}(\mathcal{X}_{val}, \omega^{*})}{\partial \theta} \\ = \frac{\partial \mathcal{L}(\mathcal{X}_{val}, \omega^{*})}{\partial \omega^{*}} \frac{\partial \omega^{*}}{\partial \theta}$$
(4)

By defining $\mathcal{G}^{(t)}$ as the gradient of the training loss at iteration *t*:

$$\mathcal{G}^{(t)} = \nabla_{\omega} \mathcal{L}(\mathcal{A}_{\theta}(\mathcal{X}_{tr}), \omega^{t})$$
(5)

we can write $\frac{\partial \omega^*}{\partial \theta}$ in Eq. 4 as:

$$\frac{\partial \omega^*}{\partial \theta} = \sum_{i=1}^{T-1} \frac{\partial \omega^{(T)}}{\partial \omega^{(i)}} \frac{\partial \omega^{(i)}}{\partial \mathcal{G}^{(i-1)}} \frac{\partial \mathcal{G}^{(i-1)}}{\partial \theta} \tag{6}$$

where T is the iteration when the classifier converges.

As ω^* represents the model weights at training convergence, they depend on θ for each iteration of gradient descent. Thus, to compute $\frac{\partial \omega^*}{\partial \theta}$, one has to back-propagate

throughout the entire T iterations of the training cycle. An example of this approach is in Maclaurin et al. [34]. This approach is feasible only for small problems due to the large requirements in terms of computation and memory. However, as optimizing ω^* is an iterative process, instead of computing $\frac{\partial \omega}{\partial \theta}$ only at the end of the training loop, we can estimate it at every iteration t:

$$\frac{\partial \omega^*}{\partial \theta} \approx \frac{\partial \omega^{(t)}}{\partial \theta^{(t)}} = \sum_{i=1}^t \frac{\partial \omega^{(t)}}{\partial \omega^{(i)}} \frac{\partial \omega^{(i)}}{\partial \mathcal{G}^{(i-1)}} \frac{\partial \mathcal{G}^{(i-1)}}{\partial \theta^{(i)}},\tag{7}$$

This procedure corresponds to dynamically changing θ during the training iterations (thus it becomes $\theta^{(t)}$) to minimize the current validation loss based on the training history. Although this formulation is different from the original objective function, adapting the data augmentation transformations dynamically with the evolution of the training process can improve generalization performance [17]. This relaxation is often used in constrained optimization for deep models, in which constraints are reformulated as penalties and their gradients are updated online, without waiting for convergence, to save computation [39]. However, in our case, we cannot write the bilevel optimization as a single unconstrained formulation in which the constraint in ω^* is summed with a multiplicative factor that is maximized (i.e., Lagrange multipliers), because the upper level optimization should be performed only on θ , while the lower level optimization should be performed only on ω . Nonetheless, even with this relaxation, estimating $\frac{\partial \omega^*}{\partial \theta}$ still remains a challenge as it does not scale well. Indeed, the computational cost of computing $\frac{\partial \omega^{(t)}}{\partial \theta^{(t)}}$ grows with the number of iterations t as shown in Eq. 7. To make the gradient computation constant at each iteration we use truncated back-propagation similarly to what is commonly used in recurrent neural net-works [51]:

$$\frac{\partial \omega^{(t)}}{\partial \hat{\theta}} \approx \sum_{i=t-K}^{t} \frac{\partial \omega^{(t)}}{\partial \omega^{(i)}} \frac{\partial \omega^{(i)}}{\partial \mathcal{G}^{(i-1)}} \frac{\partial \mathcal{G}^{(i-1)}}{\partial \theta^{(i)}}, \tag{8}$$

where K represents the number of gradient unfolding that we use. Fig. 2b. shows the computational graph used for this computation.

Additionally, as Williams and Peng [51], we consider a second parameter J which defines the number of inner loop training iterations after which θ is updated, in other words how often the computation of the gradients of θ is performed. The situation where K = J = T is the exact bilevel optimization as shown in Eq. 1 while K = J = 1corresponds to updating θ at each iteration, in our case minibatch (K = 1), using only one step of gradient unfolding (J = 1). A theoretical analysis of the convergence of this approach is presented in Shaban et al. [45].

3.2. Augmenter Networks

We present augmenter networks that can perform two types of transformations: geometrical and color.

Geometrical transformation is a good form of data augmentation, because it simulates the fact that in the real world the same object can be located at multiple positions and seen from different viewpoints. We use the transformation model of spatial transformer networks [25], but for data augmentation instead of data alignment. Thus, as illustrated in Fig. 1, the augmenter is composed of a module that generates a set of transformation parameters followed by a module that applies the generated transformation to the original image. In our experiments, we consider scenarios where the augmenter network learn affine transformations as well as scenarios where it learns only translation. In this case, only two values are learned (translation values respectively on x and y axis). Note that the learned transformations are not conditioned on the input image but defined only based on random noise.

Color transformations considered are: hue, saturation, contrast and brightness. In this case, the augmenter receives as input a random noise vector and generate a single value representing the amplitude of each color transformation.

4. Experimental Setup

4.1. Datasets

In our experiments, we consider the five following datasets:

CIFAR10 [27] is a dataset composed of 60,000 32x32 natural color images distributed in 10 different classes (6,000 images per class). This dataset is split into a training set of 50,000 images and a test set of 10,000 images.

CIFAR100 [27] is an extension of CIFAR10 dataset. It contains the same number of images at the same resolution, but they are distributed in 100 classes instead of 10.

ImageNet [44] is a dataset of 1.28 million natural color images in the training set and 50,000 images in the test set. The image size is variable, so in our experiments, we resize them to a resolution of 224x224.

Tiny ImageNet is a subset of ImageNet [44] containing 200 classes and images resized to 64x64. Each class has 500 training images, 50 validation images, and 50 test images. Since the test labels are not available, the validation set is used as test set and 20% of the training set is used for validation.

Finally, BACH [2] is a dataset of 400 breast cancer histology images of resolution 2048 x 1536 distributed in 4 balanced classes of 100 images. As there is no test set publicly available, we use in our experiments 50% of the dataset for training and validation and 50% for test.

4.2. Implementation Details

Our model is composed of a classifier and an augmenter network. To facilitate fair comparison of the results, we use in our experiments the same classifiers as in previous works: BadGAN [11], ResNet18 [20], ResNet50 [20] and Wide-ResNet-28-10 [52]. BadGAN is a simple CNN based architecture composed of 9 convolutional layers with Leaky Re-LUs and a MLP classifier. ResNet18 and ResNet50 are respectively 18 and 50 layers deep neural network with residual connections and WideResNet 28-10 is a ResNet network with 28 layers and a width factor of 10.

The augmenter learning the geometric and color transformations is a MLP that receives a noise vector as input and generates the transformation parameters. We experimented with three sizes. The *small* network has an input and output size of n, n being the number of hyperparameters to optimize (6 for affine transformations, 2 for translation and 1 for each color transformations), and it has two layers with respectively n and 10n neurons. The *medium* one has an input size of 100 and two layers of 64 and 32 neurons. The *large* one has an input size of 100 and four layers of 512, 1024, 124 and 512 neurons. In order to have differentiable affine and color transformations, we use the Kornia [43] library and the affine_grid and grid_samples functions of the torchvision package of pytorch framework.

In all experiments, we use 20% of the training set to form the validation set. Although in principle we usually use a separate validation set for training the augmenter, in practice, we noticed that reusing the training data in a variant of this holdout approach (the training set is randomly split into train and validation at each epoch) yields better results. However, it is important that the batch of samples used to learn the augmenter is different from the one used to train the classifier to ensure that the model learns data augmentation parameters that generalize well. In preliminary experiments, we tried different values for the frequency of updat-

ResNet18 / CIFAR10	Trans.	Affine	Cost
Baseline	88.55	88.55	1
Predefined	95.28	94.59	> 60
Transf. invariant (STN)	92.14	90.31	1.1
Validated magnitude	94.58	93.43	11.5
Our model + HFlip	95.35	95.16	5.3

Table 1: Impact and training cost of different geometric data augmentation strategies on classification accuracy on CIFAR10. Considering only translation and affine transformations, our approach is faster than methods requiring a validation loop and is more efficient than predefined data augmentation, STN and validated magnitude of predefined data augmentation.

ing θ J and the number of steps of back-propagation K, but they did not show relevant improvements. Therefore, for all our experiments, we use K = J = 1. In practice, the classifier is updated after each training mini-batch and the augmenter after each validation mini-batch. The only predefined transformation used by our model is the horizontal flip (vertical for BACH) as it is not differentiable.

5. Results

The goal of our method is to learn data augmentation automatically. Our experiments compare the performance of different classifiers without data augmentation (baselines), the same classifiers with the best-known hyperparameters for data augmentation (predefined), state of the art methods and our method. We experiment with two groups of transformations: geometric and color transformations.

5.1. Geometric Transformations

In this section, we evaluate our model by investigating learned geometric transformations.

In a first experiment, we assess the computational efficiency of our method and the utility of the learned geometric transformations for the classification task. In Tab. 1 we compare the performance of our method on CIFAR10 (with ResNet18) against several methods in terms of accuracy and training cost for translation and affine transformations. We define our baseline as a training without any data augmentation and we consider its training time cost as 1. Predefined represents a classifier trained with the usual standard geometric data augmentation: horizontal flip and random translation between -4 and 4 pixels along x and y axis. To estimate the training cost of this scenario, we consider the general case where the best data augmentation setting is not known and many different values have to be tested using a grid or random search. For instance, for the 6 parameters of an affine transformation, and 2 different values to try for each parameter, the number of models to validate is $2^6 = 64$. In the third case, the augmenter is trained to be transformation invariant similarly to the spatial transformer networks [25]. The transformations generated by the aug-

CIFAR-10	BadGAN		ResNet18		
AugClass.	Tr.	Aff.	Tr.	Aff.	
Small	93.65	93.62	95.35	95.16	
Medium	93.75	93.63	95.25	95.06	
Large	93.65	93.39	95.00	94.83	

Table 2: **Impact of architecture on classification accuracy.** Increasing the classifier size improves the model performance. Increasing the augmenter size has no significant impact on the final classification accuracy.

menter are applied on training as well as on test and the update of the augmenter parameters θ is done on the same data as the update of the classifier parameters ω . This approach has a very low computational cost (1.1, just the overhead of applying the augmenter) and its accuracy is better than using no data augmentation, but far from a model trained with a good data augmentation. Finally, we consider a validated magnitude approach that selects only a single parameter defining the magnitude of the transformation parameters from which an actual transformation is sampled from. This is similar to the strategy used in RandAugment [10]. This performs surprisingly well but is still inferior to our model and with a higher cost for validating the magnitude of the transformations. The last row presents the results of our approach. For both translation only and affine transformations, we obtain better results than the other approaches. In terms of computational cost, our approach is around 5 times slower than a basic training without data augmentation. However, as our approach learns the data augmentation parameters directly and does not need to loop over possible values, it is already 14x faster than the simple case of predefined data augmentation described above where we consider only 2 possible values for each parameter.

In a second experiment, we investigate the influence of the augmenter network and the classifier size on the performance of a model trained on CIFAR10. In Tab. 2, results show that a larger classifier (from BadGAN to ResNet18) improves the performance. However, the size of the augmenter network does not have a significant impact on the accuracy of the classifier. Thus, in the following experiments, we use the *small* augmenter, which is faster to train.

In a third experiment, we investigate the efficiency of learned transformations against heuristically chosen ones. For this, we consider the medical imaging dataset BACH, as possibly useful data augmentation for histological images is not trivial to define as opposed to natural images. The usual heuristically chosen geometric transformations for medical images are vertical flip and affine transformations. In Tab. 3, we compare our model learning affine transformations to a *baseline* trained without data augmentation, with *vertical flip only* and finally with *vertical flip and predefined affine transformations*: random translation between -4 and 4 pixels along x and y axis, random scale with factor between 0.5 and 2 and rotation between -10 and 10 degrees. Results

ResNet18 / BACH	Acc.
Baseline	49.50
Baseline + only VFlip	46.00
Baseline + VFlip + Affine	50.60
Our model(affine) + VFlip	56.00

Table 3: **Impact of geometric data augmentation on classification accuracy on BACH.** Vertical flip alone decreases the model performance. Used in combination with affine transformations, it improves the classification accuracy. Best performances are obtained using learned affine transformations.

ResNet18 / CIFAR10	Acc.
Baseline	88.55
Baseline + Color Jitter	88.63
Baseline + Affine + HFlip	94.59
Baseline + Affine + Color Jitter + HFlip	94.96
Our model(color)	94.18
Our model(color) + Color Jitter	94.63
Our model(affine + color) + HFlip	95.16
Our model(affine + color) + HFlip + Color Jitter	95.18

Table 4: **Impact of color and affine transformations on classification accuracy on CIFAR10.** Transformations in parentheses are learned, others are predefined. For this dataset, both color and affine transformations improve the classification accuracy. Best performances are obtained with a combination of transformations of both types.

show that only using vertical flip, which is a common transformation in medical imaging is reducing the performance of the classifier whereas using a broader range of affine transformations is yielding a better model performance than using no augmentation. Our model, using only vertical flip as predefined augmentation, obtains a much better final accuracy, which shows that the learned transformations are more useful for the classifier than the hand defined ones.

5.2. Color Transformations

In this section, we investigate the impact of color transformations alone and in combination with affine transformations on different datasets.

In a first experiment, we study color transformations alone and in combination with affine transformations on CI-FAR10. For the predefined color jitter, we use the same settings as [9]. We consider 2 versions of our model, the first one learning only color transformations and the second one learning color and affine transformations. In Tab. 4, we can see that in both cases, the learned transformations are yielding better results than predefined ones, which illustrates the efficiency of our approach for color transformations. The best results are obtained when combining color and affine transformations.

In a second experiment, we repeat the same protocol as in the previous experiment but on the BACH dataset. As there is no usual color jitter value for this dataset, we use the same default setting as for CIFAR10. In Tab. 5, we can

ResNet18 / BACH	Acc.
Baseline	49.50
Baseline + Color Jitter	43.90
Baseline + VFlip	46.00
Baseline + VFlip + Color Jitter	44.90
Baseline + Affine + VFlip	50.60
Baseline + Affine + VFlip + Color Jitter	43.00
Our model(color)	54.60
Our model(color) + Color Jitter	52.40
Our model(affine + color) + VFlip	56.50
Our model(affine + color) + VFlip + Color Jitter	49.70

Table 5: **Impact of color and affine transformations on classification accuracy on BACH**. Transformations in parentheses are learned, others are predefined. Heuristically chosen color jitter parameters have a negative impact on the classifier accuracy whereas learned color transformations improve the training. Best performances are obtained with a combination of learned color and affine transformations.



Figure 3: **Qualitative results.** The images of the first column are original images, the following ones are images transformed by our augmenter at different epochs. The first three rows contain images from Tiny ImageNet, the next two rows from BACH, and the last two rows from ImageNet.

see that using a heuristically chosen color jitter in the predefined data augmentation leads to a significant degradation of the classifier performance. This confirms that a good data augmentation strategy in one domain is not always transferable to another and that it is safer to let the model learn the optimal transformations. This also shows that for histological images, classifiers are very sensitive to color modifications. Best results are obtained by learning a combination of color and affine transformations.

	CIFAR10	CIFAR100	Tiny ImageNet	ImageNet	BACH
	ResNet18	ResNet18	ResNet18	ResNet50	ResNet18
Baseline	88.55	68.99	59.69	69.39	49.5
Predefined	94.69	73.61	61.10	76.02	50.6
Ours (affine)	95.16	74.31	62.92	76.10	55,7
Ours (full)	95.42	76.10	63.61	76.20	56.5

Table 6: Accuracy of our model on different datasets. ImageNet results reported are Top1. For all datasets, our model performs better than a classifier trained only with standard predefined data augmentation.

5.3. Evaluation on Different Datasets

We now evaluate our approach on different datasets. In addition to the CIFAR10 and BACH results, we report in Tab. 6 also results on CIFAR100, Tiny ImageNet and ImageNet. Predefined transformations used for CIFAR100 and Tiny Imagenet are the same as defined for CIFAR10 in 3.2. Predefined transformations for Imagenet are a resize to 256x256 followed by a random crop of size 224x224, horizontal flip and color transformations as in [20]. Results show that our model performs better than a classifier trained only with predefined transformations on the five datasets considered already with learned Affine transformations, but performances are even better when adding color transformations (Full). This shows that our approach can be applied to datasets with different characteristics. Note that it is suitable not only for large scale datasets like ImageNet, but also for high resolution images like in BACH (2048x1536).

5.4. Comparison with SotA Methods

In Tab. 7, we compare our model to state-of-the-art methods on CIFAR10, CIFAR100 and ImageNet. Predefined transformations are the same as described in 5.3. Results show that on CIFAR10 and using ResNet18 as classifier, our method obtains a better accuracy than GAN-based automatic data augmentation learning methods. AutoAugment has a slightly better accuracy, but note that our model obtains very close results with a smaller network. On bigger networks like Wide ResNet 28-10 and ResNet 50, our approach performs very close to search-based methods. The performance gap is explained by the fact that the searchbased methods are using more transformations, in particular non-differentiable transformations, to train the end classifier. On the other end, our model requires less prior knowledge as it does not require to define a list of possible transformation and to perform an additional loop to learn the best augmentation policy from this predefined list. Considering this, it represents an interesting trade-off between training speed and accuracy, especially for datasets where potentially useful augmentations are not trivial to define.

On Fig. 3, we show some examples of transformations learned during the training process. The first 3 rows show examples on Tiny ImageNet. What is interesting to note is that at the beginning of the training (left) transformations tend to be strong, while towards the end of the training

	Classifier	CIFAR10	CIFAR100	ImageNet
Baseline	ResNet18	88.55	68.99	-
Predefined	ResNet18	91.18	73.61	-
Bayesian DA [50]	ResNet18	91.00	72.10	-
DAN [36]	BadGAN	93.00	-	-
TANDA [42]	ResNet56	94.40	-	-
AutoAugment [9]	ResNet32	95.50	-	-
Ours	ResNet18	95.42	74.31	-
Baseline	WRN 28-10	94.83	69.90	-
Predefined	WRN 28-10	95.76	81.10	-
AutoAugment	WRN 28-10	97.40	82.90	-
Fast AA	WRN 28-10	97.30	82.70	-
PBA	WRN 28-10	97.40	83.30	-
RandAugment	WRN 28-10	97.30	83.30	-
Our model	WRN 28-10	96.44	81.90	-
Baseline	ResNet50	-	-	69.39/89.41
Predefined	ResNet50	-	-	76.02/92.84
Faster AA	ResNet50	-	-	76.50/93.20
AutoAugment	ResNet50	-	-	77.60/93.80
Fast AA	ResNet50	-	-	77.60/93.70
RandAugment	ResNet50	-	-	77.60/93.80
Our model	ResNet50	-	-	76.20/92.90

Table 7: **Comparison with other models**. ImageNet results reported are Top1/Top5. Our model based on affine and color transformations outperforms previous GANbased models and performs at a level very close to searchbased approaches. Those approaches perform better by considering also non-differentiable transformations but our approach requires less prior knowledge and no policy search loop, which makes it easier to train and more suitable for datasets where predefining data augmentation is not trivial.

(right) they are smaller and tend to approach identity. This behavior can also be seen during training on BACH (row 4 and 5) and ImageNet (row 6 and 7).

6. Conclusion

We have presented a novel approach to automatically learn the transformations needed for effective data augmentation. It is based on an online approximation of the bilevel optimization problem defined by alternating between optimizing the model parameters and the data augmentation hyperparameters. Thus, we can train the classifier network and an augmenter network jointly to generate the right transformations at every epoch. We evaluated the proposed approach with different models against a variety of datasets and transformations. The obtained results were comparable or better than the results obtained from defining handengineered transformations. This approach brings us a step closer to having a fully automated learning system that requires minimal human intervention.

7. Acknowledgment

This research was supported by the National Science and Engineering Research Council of Canada (NSERC), via its Discovery Grant program and MITACS via its Accélération program.

References

- A. Antoniou, A. Storkey, and H. Edwards. Augmenting image classifiers using data augmentation generative adversarial networks. In *International Conference on Artificial Neu*ral Networks, 2018.
- [2] G. Aresta, T. Araújo, S. Kwok, S. S. Chennamsetty, P. MohammedSafwanK., A. Varghese, B. Marami, M. Prastawa, M. Chan, M. J. Donovan, G. Fernandez, J. Zeineh, M. Kohl, C. Walz, F. Ludwig, S. Braunewell, M. Baust, Q. D. Vu, M. N. N. To, E. Kim, J. T. Kwak, S. Galal, V. Sanchez-Freire, N. Brancati, M. Frucci, D. Riccio, Y. Wang, L. Sun, K. Ma, J. Fang, I. Koné, L. Boulmane, A. Campilho, C. Eloy, A. Polónia, and P. Aguiar. Bach: Grand challenge on breast cancer histology images. *Medical image analysis*, 56:122– 139, 2019.
- [3] Y. Bengio. Gradient-based optimization of hyperparameters. *Neural computation*, 12(8), 2000.
- [4] J. Bergstra and Y. Bengio. Random search for hyperparameter optimization. J. Mach. Learn. Res., 13, Feb. 2012.
- [5] J. Bergstra, D. Yamins, and D. D. Cox. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. *JMLR*, 2013.
- [6] J. S. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl. Algorithms for hyper-parameter optimization. In *nips*, 2011.
- [7] L. Chongxuan, T. Xu, J. Zhu, and B. Zhang. Triple generative adversarial nets. In Advances in Neural Information Processing Systems (NeurIPS), 2017.
- [8] B. Colson, P. Marcotte, and G. Savard. An overview of bilevel optimization. *Annals of operations research*, 153(1), 2007.
- [9] E. D. Cubuk, B. Zoph, D. Mane, V. Vasudevan, and Q. V. Le. Autoaugment: Learning augmentation strategies from data. In *Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [10] E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le. Randaugment: Practical data augmentation with no separate search. *ArXiv*, abs/1909.13719, 2019.
- [11] Z. Dai, Z. Yang, F. Yang, W. W. Cohen, and R. R. Salakhutdinov. Good semi-supervised learning that requires a bad gan. In Advances in Neural Information Processing Systems (NeurIPS), 2017.
- [12] T. DeVries and G. W. Taylor. Improved regularization of convolutional neural networks with cutout. arXiv preprint arXiv:1708.04552, 2017.
- [13] J. Domke. Generic methods for optimization-based modeling. In Artificial Intelligence and Statistics, 2012.
- [14] C. Finn, P. Abbeel, and S. Levine. Model-agnostic metalearning for fast adaptation of deep networks. In *International Conference on Machine Learning (ICML)*, 2017.
- [15] L. Franceschi, M. Donini, P. Frasconi, and M. Pontil. Forward and reverse gradient-based hyperparameter optimization. In *International Conference on Machine Learning* (*ICML*), 2017.
- [16] L. Franceschi, P. Frasconi, S. Salzo, R. Grazzi, and M. Pontil. Bilevel programming for hyperparameter optimization and meta-learning. In *International Conference on Machine Learning (ICML)*, 2018.
- [17] A. Golatkar, A. Achille, and S. Soatto. Time matters in reg-

ularizing deep networks: Weight decay and data augmentation affect early learning dynamics, matter little near convergence. In Advances in Neural Information Processing Systems (NeurIPS), 2019.

- [18] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. C. Courville, and Y. Bengio. Generative adversarial networks. *ArXiv*, abs/1406.2661, 2014.
- [19] R. Hataya, J. Zdenek, K. Yoshizoe, and H. Nakayama. Faster autoaugment: Learning augmentation strategies using backpropagation. *ArXiv*, abs/1911.06987, 2019.
- [20] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [21] A. Hernández-García and P. König. Data augmentation instead of explicit regularization. arXiv preprint arXiv:1806.03852, 2018.
- [22] A. Hernández-García and P. König. Do deep nets really need weight decay and dropout? arXiv preprint arXiv:1802.07042, 2018.
- [23] D. Ho, E. Liang, I. Stoica, P. Abbeel, and X. Chen. Population based augmentation: Efficient learning of augmentation policy schedules. In *International Conference on Machine Learning (ICML)*, 2019.
- [24] F. Hutter, H. H. Hoos, and K. Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In *International Conference on Learning and Intelligent Optimization*, 2011.
- [25] M. Jaderberg, K. Simonyan, A. Zisserman, et al. Spatial transformer networks. In Advances in Neural Information Processing Systems (NeurIPS), 2015.
- [26] A. Krause, P. Perona, and R. G. Gomes. Discriminative clustering by regularized information maximization. In Advances in neural information processing systems (NIPS), 2010.
- [27] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- [28] J. Lemley, S. Bazrafkan, and P. Corcoran. Smart augmentation learning an optimal data augmentation strategy. *IEEE Access*, 5:5858–5869, 2017.
- [29] S. Lim, I. Kim, T. Kim, C. Kim, and S. Kim. Fast autoaugment. In Advances in Neural Information Processing Systems (NeurIPS), 2019.
- [30] C. Lin, M. Guo, C. Li, W. Wu, D. Lin, W. Ouyang, and J. Yan. Online hyper-parameter learning for autoaugmentation strategy. 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pages 6578–6587, 2019.
- [31] H. Liu, K. Simonyan, and Y. Yang. Darts: Differentiable architecture search. In *International Conference on Learning Representations (ICLR)*, 2019.
- [32] J. Luketina, M. Berglund, K. Greff, and T. Raiko. Scalable gradient-based tuning of continuous regularization hyperparameters. In *International Conference on Machine Learning* (*ICML*), pages 2952–2960, 2016.
- [33] M. MacKay, P. Vicol, J. Lorraine, D. Duvenaud, and R. Grosse. Self-tuning networks: Bilevel optimization of hyperparameters using structured best-response functions.

In International Conference on Learning Representations (ICLR), 2019.

- [34] D. Maclaurin, D. Duvenaud, and R. Adams. Gradientbased hyperparameter optimization through reversible learning. In *International Conference on Machine Learning* (*ICML*), 2015.
- [35] M. Mirza and S. Osindero. Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784, 2014.
- [36] S. Mounsaveng, D. Vázquez, I. B. Ayed, and M. Pedersoli. Adversarial learning of general transformations for data augmentation. *ArXiv*, abs/1909.09801, 2019.
- [37] S. J. Nowlan and G. E. Hinton. Simplifying neural networks by soft weight-sharing. *Neural Computation*, 4:473–493, 1992.
- [38] A. Odena, C. Olah, and J. Shlens. Conditional image synthesis with auxiliary classifier gans. In *International Conference on Machine Learning (ICML)*, 2017.
- [39] D. Pathak, P. Krähenbühl, and T. Darrell. Constrained convolutional neural networks for weakly supervised segmentation. 2015 IEEE International Conference on Computer Vision (ICCV), pages 1796–1804, 2015.
- [40] F. Pedregosa. Hyperparameter optimization with approximate gradient. In *International Conference on Machine Learning (ICML)*, 2016.
- [41] L. Perez and J. Wang. The effectiveness of data augmentation in image classification using deep learning. arXiv preprint arXiv:1712.04621, 2017.
- [42] A. J. Ratner, H. Ehrenberg, Z. Hussain, J. Dunnmon, and C. Ré. Learning to compose domain-specific transformations for data augmentation. In Advances in Neural Information Processing Systems (NeurIPS), 2017.
- [43] E. Riba, D. Mishkin, D. Ponsa, E. Rublee, and G. Bradski. Kornia: an open source differentiable computer vision library for pytorch, 2019.
- [44] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3), 2015.
- [45] A. Shaban, C.-A. Cheng, N. Hatch, and B. Boots. Truncated backpropogation for bi-level optimization. In *Proceedings of* the 22nd International Conference on Artificial Intelligence and Statistics, 2019.
- [46] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.
- [47] J. Snoek, H. Larochelle, and R. P. Adams. Practical bayesian optimization of machine learning algorithms. In Advances in Neural Information Processing Systems (NeurIPS), 2012.
- [48] J. T. Springenberg. Unsupervised and semi-supervised learning with categorical generative adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2016.
- [49] N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15, 2014.
- [50] T. Tran, T. Pham, G. Carneiro, L. Palmer, and I. Reid. A bayesian data augmentation approach for learning deep mod-

els. In Advances in Neural Information Processing Systems (NeurIPS), 2017.

- [51] R. J. Williams and J. Peng. An efficient gradient-based algorithm for on-line training of recurrent network trajectories. *Neural Computation*, 1990.
- [52] S. Zagoruyko and N. Komodakis. Wide residual networks. *ArXiv*, abs/1605.07146, 2016.
- [53] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations (ICLR)*, 2018.