



HAL
open science

Encouraging Intra-Class Diversity Through a Reverse Contrastive Loss for Better Single-Source Domain Generalization

Thomas Duboudin, Emmanuel Dellandréa, Corentin Abgrall, Gilles Hénaff,
Liming Chen

► **To cite this version:**

Thomas Duboudin, Emmanuel Dellandréa, Corentin Abgrall, Gilles Hénaff, Liming Chen. Encouraging Intra-Class Diversity Through a Reverse Contrastive Loss for Better Single-Source Domain Generalization. 2021. hal-03260124v1

HAL Id: hal-03260124

<https://hal.science/hal-03260124v1>

Preprint submitted on 14 Jun 2021 (v1), last revised 23 Feb 2023 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Encouraging Intra-Class Diversity Through a Reverse Contrastive Loss for Better Single-Source Domain Generalization

Thomas Duboudin
École Centrale de Lyon

thomas.duboudin@ec-lyon.fr

Corentin Abgrall
Thales LAS France

corentin.abgrall@fr.thalesgroup.com

Emmanuel Dellandréa
École Centrale de Lyon

emmanuel.dellandrea@ec-lyon.fr

Gilles Hénaff
Thales LAS France

gilles.henaff@fr.thalesgroup.com

Liming Chen
École Centrale de Lyon
liming.chen@ec-lyon.fr

Abstract

Traditional deep learning algorithms often fail to generalize when they are tested outside of the domain of training data. Because data distributions can change dynamically in real-life applications once a learned model is deployed, in this paper we are interested in single-source domain generalization (SDG) which aims to develop deep learning algorithms able to generalize from a single training domain where no information about the test domain is available at training time. Firstly, we design two simple MNIST-based SDG benchmarks, namely MNIST Color SDG-MP and MNIST Color SDG-UP, which highlight the two different fundamental SDG issues of increasing difficulties: 1) a class-correlated pattern in the training domain is missing (SDG-MP), or 2) uncorrelated with the class (SDG-UP), in the testing data domain. This is in sharp contrast with the current domain generalization (DG) benchmarks which mix up different correlation and variation factors and thereby make hard to disentangle success or failure factors when benchmarking DG algorithms. We further evaluate several state-of-the-art SDG algorithms through our simple benchmark, namely MNIST Color SDG-MP, and show that the issue SDG-MP is largely unsolved despite of a decade of efforts in developing DG algorithms. Finally, we also propose a partially reversed contrastive loss to encourage intra-class diversity and find less strongly correlated patterns, to deal with SDG-MP and show that the proposed approach is very effective on our MNIST Color SDG-MP benchmark.

1. Introduction

While deep neural networks reach or even surpass human-level performance on an increasing number of computer vision tasks, *e.g.*, classification, object detection or semantic segmentation [15, 14], it is found that their performance could drop sharply [17] when mismatch between training and testing data distributions occurs. This is an issue of critical importance for real-world application, where test-time domain shift is common: the data distribution can change over time because of varying factors as diverse as lighting, view angle, sensor, *etc.*, ending up being significantly different from the one used during training. For embedded networks deployed in the wild, that can't be easily retrained (autonomous cars, for instance), geographical distribution change are an issue too. It is impossible to gather enough data to cover all possible shifts, so another solution has to be found.



Figure 1. Samples of the MNIST Color SDG-MP & UP datasets. First row is training data, second row is testing data for SDG-MP, third row is testing data for SDG-UP. All digits have their class color, without exception.

To generalize without having access to any kind of




Figure 2. An illustration of our reverse contrastive approach on 3 anchors. Anchor 1 in the latent space is sufficiently far from its closest negative neighbor (full line) for it to be pushed away from its closest positive neighbor (dotted line). So is anchor 3. Anchor 2, however, is too close to a negative point to be pushed. The arrows show the moving direction of each anchor.

information regarding the target domain is the goal of the research field known as domain generalization (DG) (or *out-of-distribution* generalization). Most of the works in this field assume to have access to data coming from several identified domains during training and aim to make the network’s internal representations invariant on the domains, by finding patterns common to all domains. The idea behind is that patterns shared among several domains are more likely to be found in a new, unseen, domain [26, 27]. This multi-source domain generalization (MDG) setting is not realistic for all tasks and all kind of data: health-related data, for instance, cannot be easily shared and the gathering of data from different sources (such as country) may prove difficult. In this paper, we are concerned with a situation, coined as single-source domain generalization (SDG), where only one single source domain is available during training, either that all data come from a same distribution, or that the data come from several distributions which are not identified differently. Both settings (single and multi-source) are still largely unsolved in the general case: on several complex and realistic datasets, *e.g.*, Domain-Net [31], PACS [25], methods designed to improve the out-of-distribution generalization capability of networks were found to perform as badly as the standard training procedure [13].

In order to shed light into factors leading to success or failure of DG algorithms, we argue that it is important to disentangle issues of different nature often mixed up in the state-of-the-art DG benchmarks to enable isolation of success or failure factors when dealing with each of

these DG issues of increasing difficulties. Specifically, as illustrated in Figure 1, we design two simple MNIST-based benchmarks, namely MNIST Color SDG-MP and MNIST Color SDG-UP, to enable benchmarking of SDG algorithms upon the two different SDG issues of increasing difficulties: when a domain shift occurs, the patterns used by the network to correctly classify an image in the source domain may be either 1) - missing (SDG-MP) or 2) - uncorrelated or anti-correlated (SDG-UP) with the image label in the target domain. By missing, we mean a pattern whose corresponding filter does not activate above its training mean at test-time.

In this paper we focus our attention on the issue of missing patterns at test-time, *i.e.*, the SDG-MP issue. Intuitively, this problem could be solved by having the network focus on the maximum amount of discriminant patterns, which is unfortunately not its natural behavior. Indeed, by learning more patterns in training for its decisions, the network could still make a sound prediction with the remaining ones, in case one particular pattern identified in the learning phase is missing at test-time. This, and its consequences on a network’s ability to generalize outside of its training domain, have been previously identified and partially mitigated in [19, 32]. However, using our MNIST Color SDG-MP benchmark specifically crafted to evaluate the ability of a model to focus on other patterns than the most strongly correlated ones, we show the inability of existing methods, *e.g.*, Jigsaw [4], Spectral Decoupling [32], RSC [19], to correctly handle this issue, even in our simple benchmark: none of the existing works reach satisfying performance on our dataset and some of them can record a performance drop as high as 56 accuracy points when the correlated pattern identified in training data, *i.e.*, color, is missing in the test data.

Given these findings, we posit that, when training for a classification task, naturally learned patterns are learned because they tend to maximize inter-class specificity, while ignoring intra-class variability. These learned patterns are the most effective when it comes to discriminate different classes. As a result, we hypothesize that, to find less efficient patterns, we need to look for class-discriminant patterns with a higher intra-class variability than those learned naturally. That is, patterns enabling the network to better discriminate samples of the same class. Because we don’t have access to additional auxiliary labels to discriminate elements within a class, the idea is to expand the class-wise internal representations of a deep network. We do so by using a partially reversed contrastive loss: in a traditional contrastive approach, same class samples are brought together in the latent space while different class samples are pushed away from each other. In our approach,

same class samples are pushed away from each others until they reach a different class sample (see Figure 2 for an illustration of our method). Our strategy yields state-of-the-art performance on our MNIST Color SDG-MP dataset.

Our contributions are threefold:

- We design two MNIST-based SDG benchmarks, namely MNIST Color SDG-MP and MNIST Color SDG-UP, specifically crafted to highlight the behavior of Single Domain Generalization (SDG) algorithms with respect to two different SDG issues, respectively, *i.e.*, when a class correlated pattern in training is missing (SDG-MP) or even uncorrelated with the class (SDG-UP) in the testing domain;
- We show that several state-of-the-art SDG algorithms fail to handle the SDG-MP issue using our MNIST Color SDG-MP benchmark;
- We propose a reverse contrastive loss (RCL-SDG) to find less strongly correlated patterns by encouraging intra-class diversity to a certain extent.

The remaining of the paper is organized as follows. Sec.2 discusses related works. Sec.3 introduces the proposed MNIST Color SDG benchmarks. Sec.4 presents our proposed RCL-based SDG algorithm. Sec.5 benchmarks our proposed SDG algorithm along with several state-of-the-art algorithms for comparison. Sec.6 concludes the paper and draws some future works.

2. Related Works

2.1. Domain Generalization (DG)

2.1.1 Domain shift robustness through multi source training

Most settings in DG assume to have data coming from several different source domains, which are identified. By learning domain invariant features, through a great diversity of practical approaches, the model is supposedly able to generalize to a new domain. Most recent works use adversarial strategies: in the work of Tzeng *et al.* [36] the features extracted by a feature extractor are fed to a discriminator tasked to find the original domain of the image. The discriminator is trained to minimize the domain classification error, while the feature extractor is trained to maximize it, hence making the features indistinguishable between domains. This family of works were initially developed in the context of domain adaptation (DA) but can be adapted in the multi-source domain generalization setting by employing all the domains at disposal during training, *e.g.*, [5] and [26]. The work of Krueger *et al.* [23] is enforcing invariance at loss level by minimizing the

variance of the loss over all training domains alongside with the usual average of the loss per batch. Self-supervised strategies have also emerged to solve the problem of multi-source domain generalization. Carlucci *et al.* [4] used a self-supervised strategy based on solving jigsaw puzzles: the images are divided into tiles, which are shuffled before being given to the classifier. The classifier has two goals: classify the samples and find the shuffling permutation. By having a supplementary objective, not related in any way to the classification, the network will learn other patterns than those strictly sufficient to the classification task and hence generalize better to a new distribution. So far, it is unclear why features learned with self-supervised tasks are more robust to domain shifts.

In 2019, Gulrajani and Lopez-Par [13] analysed a lot of DG algorithms and benchmarks "in search of lost domain generalization". They detailed ways to carefully and realistically evaluate multi-source DG methods, by using larger models, stronger data augmentation, correct model and hyper-parameters selection (without any use of the target domain), and doing all on more datasets. Their conclusion was that no methods were significantly above the basic expected risk minimization procedure (ERM), where data from all domains is merged into one single dataset with the network trained on it. Their findings are corroborated by Cha *et al.* [6] who look into the latent spaces from differently trained networks, one with DANN [9] and another with the basic ERM. They have shown that ERM naturally tends to yield domain invariant features. These findings lead us to study DG settings of increasing difficulties where different factors for DG are disentangled, and to subdivide the complicated out-of-distribution generalization problem.

2.1.2 Domain shift robustness through stylization

Style transfer as a mean of data augmentation has been recently introduced independently by several authors [28, 39, 22, 33] as a way to promote invariance to texture changes. The images are simply stylized differently, with AdaIN based [18] or GAN-based [12, 41] architectures, before being fed to the main task model. It has been hypothesized that the low ability of deep networks to transfer to unseen domains was related to their texture bias: deep networks grant more importance to the texture than the shape in image recognition tasks [10]. As such, correcting this texture bias is expected to improve out-of-distribution generalization, provided that the distribution shift is mainly due to a texture shift. The stylization as data augmentation exists in different flavors: Nam *et al.* [28] use intra-domain extra-class stylization: the styles used to transfer are coming from the images of the training distribution but from different class samples than the image being stylized. Yue

et al. [39] use several auxiliary style datasets. Jackson *et al.* [20] approximated the style codes distribution of paintings with a multivariate normal, and completely discard the painting dataset once the transfer network is trained. The stylization is a strong baseline in single-source domain generalization (SDG) [39] when the domain shift encountered is a texture shift, such as a synthetic to real transfer, which, unfortunately, does not always hold in the general case.

2.2. Overcoming Biases

The issue of missing patterns is related to both out-of-distribution generalization and bias avoidance in deep networks. Deep neural networks learn correlations between patterns and labels no matter how spurious they seem to a human being and therefore sometimes need to be explicitly told not to use certain patterns for the recognition task at hand [29, 21]. In a domain shift situation, the patterns naturally learned by a network are precisely the "biases" we need to overcome to find other and semantically different patterns. A line of work, *e.g.*, Kim *et al.* [21], is able to create bias invariant features, but requires the bias to be given as an auxiliary label, and are therefore not suitable to the SDG setting where no information on testing data, *e.g.*, bias shift, is assumed available. Another line of work, *e.g.*, [29, 8, 1], focuses on finding counter-examples *i.e.*, samples that do not share the majority biases of their classes, and increase their importance during training. Assuming the necessary existence of counter-examples is an optimistic hypothesis in certain situations, such as with synthetic data where only a few number of synthetic models of interest may be available. Besides, it could lead to overstate the importance of what could be an aberrant outlier, such as an annotation error. As such, there also exists approaches which do not need counter-examples. Representation Self-Challenging [19] (RSC) introduces a dropout-based strategy, where the feature map coefficients most responsible for the prediction are muted, instead of random ones, therefore leading the network to focus on other less correlated patterns. Spectral Decoupling [32], a method introduced by Pezeshki *et al.*, proposes an L2 regularization on the output logits of the network, with a theoretical motivation, to push the network to learn other patterns than the obvious and spurious ones. Both these methods can work without counter-examples.

2.3. Supervised Contrastive Learning

Contrastive learning [35, 16] has been introduced in deep learning to enable class-wise manipulation of a network's latent space with a simple idea: elements of the same class should be close to each other in the latent space, elements of different classes should be away from each other in the latent space. Our approach is to loosen the latent space to identify less correlated patterns in a recognition task. While

this idea has been previously studied in [40] and [38], it was applied with a different goal. Zhang *et al.* [40] proposes to take all the space available on the unit sphere evenly, but the regularization only impacts negative pairs. Xuan *et al.* [38] aims to create a latent space with better generalization to unseen classes using a modified triplet loss that also pushes together elements of different classes.

3. MNIST Color SDG Dataset

We introduce a new toy dataset based on the MNIST dataset [24] for benchmarking single-source domain generalization (SDG) algorithms. For the training data, each digit, coming from the original MNIST training data, is colored with a color depending on its class. There are no counter-examples (digit colored differently than the majority of the other digits in its class), as in [1] or label noise (digit colored with their correct class color, but with an assigned label different from the original one), as in [32]. We do so to study the ability of a network to find more semantically different patterns than what is done naturally, without having any kind of data signals that could help it find those patterns. The validation data is the same for both SDG-MP and SDG-UP and is colored the same way as the training data but come from the original MNIST test split. A sample of the dataset can be found Figure 1. A summary of the synthetic MNIST-based DG datasets is available in Table 1.

Method	SDG/MDG	UP/MP	LN/CE	Classes
IRM [2]	MDG (2)	UP	LN	2
SD [32]	SDG	UP	LN	2
GIP [1]	SDG	UP+MP	CE	10
Ours	SDG	UP+MP	None	10

Table 1. Comparison of MNIST-datasets introduced to study domain generalization strategies. SDG/MDG refers to single or multiple training domains, UP/MP is uncorrelated (or anti-correlated) or missing patterns, LN is label noise, CE is counter-examples, Classes is the number of classes.

3.1. SDG-MP

The SDG-MP dataset aims to benchmark SDG algorithms in presence of missing patterns, *i.e.*, when a class correlated pattern in training, *e.g.*, color, is missing in testing. It is made by coloring each digit with the same color, no matter their class. The images come from the original MNIST test data. The test color and the training colors are not chosen randomly. To precisely confront a deep network with a situation where useful training patterns are missing at test-time, the test color has to produce an activation that is roughly the same for all color specific filters in the network. This way, the network is unable to use color related information to predict the class of a digit and has to make

use of other class correlated patterns, *e.g.*, shape. Because we can't directly abide by this constraint we propose to use the average training color as the test color, provided that no specific training color is closer to the average than any other one and that all training colors are sufficiently different from each other. If the training color of class C is closer to the average color than the other colors, the network will wrongfully believe that test samples are all of class C. If two training colors are close to each other, the network won't be able to easily differentiate samples of the two classes using only the color and will be driven to learn other class related patterns, *e.g.*, shape, at least partially. The colors (a triplet of value in range $[0, 1]$) are therefore chosen to be on a sphere centered in the average (6 of them) and on the vertices of a cube centered in the average (the 4 left). This way no particular color is closer to the average than at least three others. All colors are not on the sphere to avoid colors that are too similar to one another.

3.2. SDG-UP

SDG-UP aims to benchmark SDG algorithms when a class related pattern in training data becomes uncorrelated or anti-correlated in testing data. We make it by permuting the class-color correspondences, on the original MNIST test set. This means that colors and labels correlated in the training set are now anti-correlated. Absence of correlation and anti-correlation are in fact the same from the network's perspective: prediction on test samples is done sample-wise, without considering the overall test-time correlations. While for us humans the colors are obviously spurious and should be ignored, the network should rely on colors to make its predictions. The correlation color - label is stronger than the correlation shape - label (because there is more variability in the shapes of digits) and, as such, it is natural to expect the network to use primarily the colors for classification. Therefore, a network is not expected to naturally reach good classification accuracy on SDG-UP when trained on the colored training set and should instead have an accuracy close to 0%. This is effectively what happens for a standard training. A network able to classify anti-correlated colored samples has to ignore the color information and rely only on the shapes. Thus, a supplementary information is required to obtain this behavior (hypothesis or test-time data).

4. Reverse Contrastive Approach

4.1. The proposed method

Our method starts from the observation that patterns that are learned with a standard training procedure (*e.g.*, stochastic gradient descent with a cross-entropy loss, for a classification task) are learned because they exhibit the maximum inter-class specificity but they ignore intra-class diversity.

To make the network learn semantically different ways to do the task, we propose to look for class-discriminant patterns with a higher intra-class variability than those learned naturally. A higher intra-class variability means that we are able to discriminate elements inside a single class, something that is difficult with the standard training procedure (see Figure 3, a). To find such patterns, we spread the intermediate features (the output of the convolutional features extractor F, before the fully connected classifier layers) of elements of the same class, to a certain extent. The limit in the class-wise repulsion is fixed with elements of the other classes: the intra-class features are repelled from each other until we reach elements of the other classes, with a margin. This approach is akin to a partially reversed contrastive loss where positive and negative pairs are switched, but a margin must be maintained between samples of different classes. We train the network with two objectives: the conventional classification objective for both the features extractor and the classifier, and the following for the features extractor only:

$$\begin{aligned} \min_F \{ & -d(x_a, x_p) \} \\ \text{s.t. } & d(x_a, x_p) < m \times d(x_a, x_n) \end{aligned} \quad (1)$$

x_a is the anchor feature map, x_p a feature map belonging to a point of the same class as x_a *i.e.* the positive sample, x_n a feature map belonging to a point of a different class than x_a , *i.e.* the negative sample, and m is the margin, a scalar to chose between 0 and 1 if we want the inter-class distance to be larger than the intra-class distance. We use a multiplicative margin, instead of an additive margin, contrary to the standard triplet loss [16], to ease the hyperparameter search. The distance d used is the L1-distance. Feature maps are rescaled in the range $[0, 1]$ before distance calculation, by using the maximum and minimum activation values computed batch-wise. This is done to avoid divergence as there are otherwise no lower bounds for this loss. Usually, features are normalized with their L2-norm, putting them on the unit sphere, but we found better results with the min/max strategy. Practically, we use the following reverse contrastive loss (RCL-SDG):

$$L_{rc} = \begin{cases} -d(x_a, x_p) & \text{if } d(x_a, x_p) < m \times d(x_a, x_n) \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

The triplet of features used is not chosen randomly in the batch: we follow an easy-positive hard-negative sampling strategy. The anchors are sampled randomly, the positive is the closest element of the same class, the negative is also the closest element in a different class. It is useless to push away elements of the same class that are already far from each other in the latent space: patterns extracted for these samples are already different. We compare the closest intra-class distance with the closest inter-class distance

Method	Validation Accuracy	Test Accuracy
Standard Training Procedure	99.8 (± 0.006)	26.0 (± 4.7)
Dropout	99.8 (± 0.01)	31.9 (± 3.1)
Dropout & Orthogonality [3]	99.8 (± 0.008)	42.7 (± 2.6)
Dropout & Covariance [7]	99.8 (± 0.002)	42.6 (± 2.0)
Jigsaw Puzzle [4]	99.8 (± 0.01)	43.0 (± 3.0)
(with early stopping @95)	98.5 (± 0.3)	59.9 (± 4.5)
Reconstruction	99.8 (± 0.007)	28.5 (± 4.8)
Spectral Decoupling [32]	99.8 (± 0.004)	47.7 (± 2.9)
(with early stopping @95)	99.4 (± 0.09)	49.8 (± 1.5)
RSC [19]	99.2 (± 0.2)	43.5 (± 5.0)
RCL-SDG (ours)		
$m = 0.2$	99.8 (± 0.007)	12.5 (± 2.6)
$m = 0.4$	99.8 (± 0.007)	19.9 (± 5.5)
$m = 0.6$	99.8 (± 0.009)	36.8 (± 5.6)
$m = 0.8$	99.7 (± 0.03)	55.7 (± 4.7)
$m = 0.9$	99.4 (± 0.08)	68.2 (± 4.0)
(with early stopping @95)	95.84 (± 0.6)	74.7 (± 8.5)
$m = \infty$	96.0 (± 0.3)	89.9 (± 0.9)
Standard Training Procedure on the original MNIST	97.8 (± 0.12)	97.8 (± 0.12)

Table 2. Results of the methods on the MNIST Color SDG-MP dataset. The results are obtained with a best validation selection strategy. The first column is the validation accuracy (colored digits), and the second the test (grey digits, SDG-MP) accuracy. Accuracies reported are averages over 10 runs, with the standard deviation between parenthesis. The last line indicates the accuracy reached by our backbone on the original MNIST dataset without domain shift, it thus gives the highest bound of the accuracy in the test domain for the methods.

to avoid creating a mixed latent space and define precise boundaries in the latent space. The positive points used in the loss are detached from the computation graph so that only the anchor is moved in the latent space: the distance check with negative neighbors is only relevant for the anchor. This loss cannot work on its own without the classification loss: at initialization, all features are clustered in the same area of the latent space and since the inter-class distances are small, we cannot expand the size of the intra-class clusters. The cross entropy loss will push away the features of elements of different classes, allowing the second objective to loosen the intra-class clusters. A visualisation of the proposed method can be found in Figure 2. We found even better results when pushing away features of the same class without any kind of limitation, that is, minimizing $-d(x_a, x_p)$, or having a margin set to infinity. While counter-intuitive, this might be explained by the influence of the classification loss, which prevents a complete scattering of intra-class features and forces the differentiation of inter-class features.

4.2. State-of-the-art SDG baselines

For comparison without any bias which could be introduced by differences in the backbone network architecture, we implement using a same backbone several other state-

of-the-art SDG methods. A first SDG algorithm can simply make use of dropout [34] to find several useful patterns for classification. By zeroing activations inside the network, we force it to look for new patterns. A limitation of dropout is that nothing prevents the network from learning the same patterns through several filters. To avoid this redundancy phenomenon, we further implement two straightforward variants using two different regularizations: orthogonality of filters, used in [3] (more precisely, we apply the double soft orthogonality regularization) and constraint over the covariance matrix of the filters activations, used in [7]. Both regularizations are applied on the same layer as dropout, but before dropout is applied for the calculation. An orthogonality constraint for filters is supposed to prevent a filter to be close to another. The penalisation of the covariance matrix of the activations is based on the idea that filters that generally activate together, or don't activate together, are probably semantically related, even though their weights might be different. This constitutes what we call the naive strategies.

We also implement more elaborated SDG models inspired by DG works reviewed in sec.2. First, we compare our approach with one inspired from the jigsaw puzzle multi-task strategy used in [4]. Our approach is also com-

pared to a reconstruction multi-task strategy (inspired from [11]), where the features obtained by the features extractor are used for a classification task, with a classification head, and for an image reconstruction task, with a decoder. The feature extractor must extract sufficiently complete patterns to reconstruct the input, which are more than what is extracted with a single classification objective. The two strategies tailored precisely for the SDG problem at hand are Representation Self Challenging [19], and Spectral Decoupling [32].

5. Results and Discussion

5.1. Experimental setup

Our experiments are conducted on MNIST Color SDG-MP using a small neural network. The architecture we use was introduced in [4] to study MNIST to SVHN [30] transfer. It is composed of two convolutional layers and three fully connected layers. Max pooling operations are inserted between each convolutional. The non-linearity used is ReLU for all the layers. The convolutional layers define the feature extractor (128 channels), and the fully connected layers the classifier. For all experiments, we use stochastic gradient descent, with a batch size of 128, with nesterov momentum at 0.9, a fixed learning rate of $1e-3$ and an L2 weight decay at $1e-5$. Models are trained for 10 epochs. There is no data augmentation. The jigsaw puzzle strategy uses an additional fully connected layer to the network, as in [4], alongside another fully connected layer used for the classification. The images are divided into 2×2 squared tiles, which are then shuffled, yielding 24 possible permutations. Each batch is used for both labels: class with the original batch sent through the network, permutation with the shuffled batch. The decoder in the reconstruction method is composed of 4 transposed convolution layers, with an hyperbolic tangent as the last activation function. When dropout is used, the dropout rate is chosen at random between 0 and 1 for each iteration: a fixed dropout rate helps the network introduce redundancy in a simple fashion: it only has to create more redundancies than there are dropped channels. Likewise, only full channel dropout is used because of the correlation between spatially close activations in the same channel, which enable the network to recover the color all the time. For RSC, we reuse these dropout hyper-parameters and the batch percentage (the proportion of samples per patch for which RSC is used) is fixed at 100% as we want the network to look beyond the color for every image.

Most strategies use two objectives during training: the classification cross-entropy and another objective (jigsaw puzzle, reverse contrastive loss, reconstruction) or regularization (orthogonality, spectral decoupling). The weight for

the classification loss is always set to 1. The weightings for the supplementary losses (or regularizations) were selected in the list (0.001, 0.01, 0.1, 1, 5, 10) with the following principle: the value selected is the largest value that does not lead to a collapse of the validation accuracy. The idea is that the regularization weighting should have a positive slope for out-of-distribution accuracy, *i.e.* the larger the weight, the better the out-of-distribution accuracy, up until a certain point. It is grounded in the fact that most additional objectives tend to counter the natural behavior of the network. The weight is 1.0 for the orthogonality constraint, 0.01 for the covariance constraint, 10.0 for the Jigsaw Puzzle, 1.0 for the reconstruction, 5.0 for Spectral Decoupling and 1.0 for our reverse contrastive constraint.

During training, we select the model with the highest validation accuracy, and evaluate this model on the testing data (SDG-MP). It has been noted that domain generalization is a setting where the initialization of the network is more important than usual and that results may vary in a greater fashion than with a training and testing dataset coming from the same distribution [23]. Therefore, we average the results over 10 runs, and report the standard deviation alongside the average.

5.2. Results and analysis

Table 2 synthesizes the overall results. As can be seen, our reverse contrastive method yields a significant improvement over the previous works and the naive strategies on our dataset. Dropout compares favorably to normal training, but yields far better results when used together with a regularization to prevent redundancy. This redundancy issue might explain why RSC yields results only marginally above dropout and regularization. During training, we monitored the test accuracy over the epochs and noticed that jigsaw puzzle and spectral decoupling succeed to some extent but suffers from an over-fitting issue: the best test accuracy does not happen for the model with the best validation accuracy. Early stopping is useful in this situation. We employ a simple yet realistic early stopping strategy that does not use the test set: training is stopped as soon as the validation accuracy reaches a satisfying threshold, fixed here at 95%. This only enables us to recover a closer accuracy than the best test accuracy but not go higher. Most methods specifically designed to prevent the network to focus only on a subset of useful features are not effective enough to consider the problem solved on our simple dataset.

Beside the accuracy table, we also illustrate the impact of our reverse contrastive loss in Figure 3 by looking at the latent space extracted from models trained differently, through the T-SNE dimensionality reduction method [37]

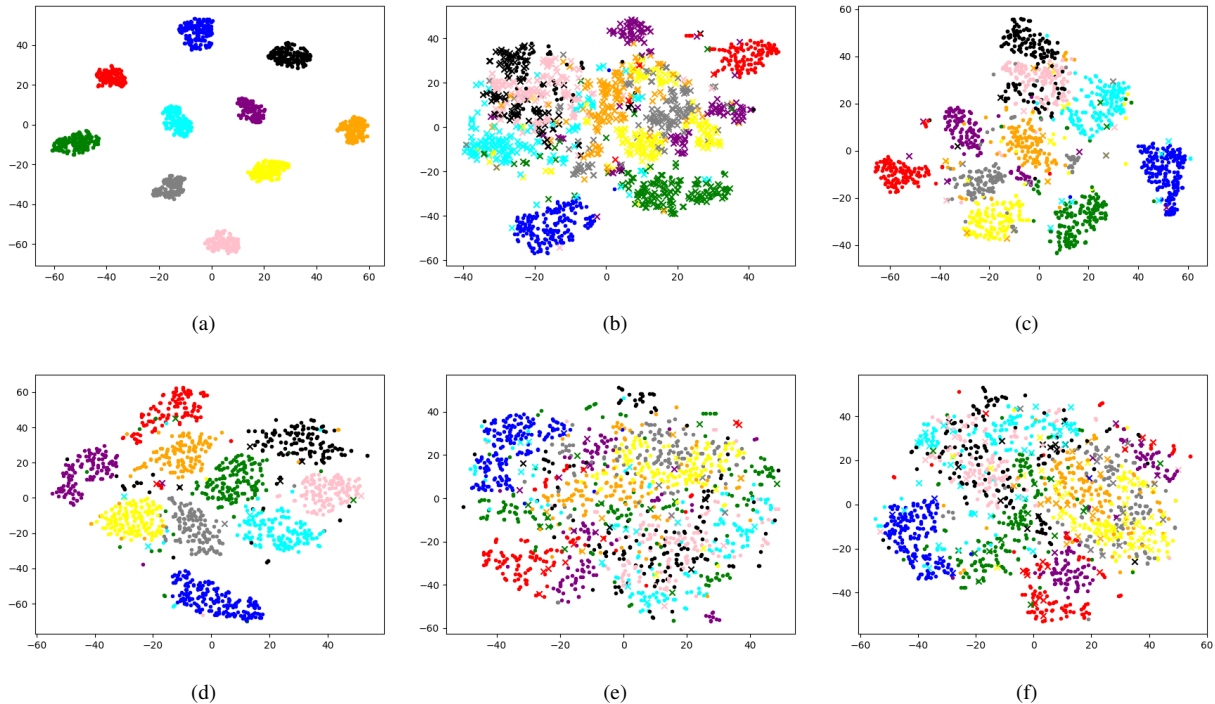


Figure 3. Visualisation of different latent spaces through a 2-dimensional T-SNE [37]. The two axis are the dimensions obtained by the T-SNE. (a) shows the validation (colored digits) latent space, for a model trained normally. (b) shows the test (grey digits) latent space for the same model. (c) shows the validation latent space for a model trained normally, on the original MNIST. (d) shows the validation latent space for a model trained with the RCL-SDG and a margin of 0.9. (e) shows the validation latent space for a model trained with the RCL-SDG without limitation. (f) shows the test latent space for the same model. The color of a dot is its ground truth class (not the same as the colors used for the digits) and its shape represents whether or not the network successfully predicted the correct class: circle if so, cross if not. Best viewed in color.

applied on the features. With a margin of 0.9, the class clusters are still separated but we can see their expansions (d), compared to the standard training situation (a). Without limitation, the clusters are mixed together, and can only be discriminated one another only roughly (e). This explains why the performance in the training domain is lower: the classifier can't perfectly separate the classes, however the performance in the testing domain is higher. This illustrates an underlying trade-off in our method: the more the intra-class clusters are expanded, the more the network will find "hidden" useful patterns but the more it will also learn useless patterns, that are not inter-class discriminant. By comparing with the latent space of the same model trained on the original MNIST (c), we see an inefficiency of the method to find useful patterns. The cluster's sizes needed to obtain an accuracy around 90% are larger than the size of the clusters on the original MNIST, indicating that noise and instance specific patterns have been learned by the network. We hypothesize that on more complex datasets, the RCL-SDG without limitation might not yield

good results due to the images specificities being more prevalent than in MNIST. To tackle this issue is a future work direction.

Our approach tends to qualify the common principle in the deep learning research community that a good latent space, one able to generalize well, is supposed to have tight intra-class clusters with large margins between clusters. While this is true when there is no domain shift, it might not always hold true when so. This can be seen in Figure 3 where a testing latent space corresponding to a large-margin tight-clusters training latent space is completely misunderstood by the network (b). On the opposite, the blurred boundaries training latent space (e) stays similar in the testing domain (f).

6. Conclusion

In this paper we carefully created a benchmark to study one of the reasons deep networks fail to generalize outside

of their training domain: the reliance of a model on only the most obvious discriminant patterns has dramatic consequences if, in the test domain, such patterns are missing. We showed that existing methods only mitigate the damage. Therefore, we proposed a counter-intuitive strategy: instead of aiming for a tight-cluster large-margin latent space, it is beneficial to try to expand the class-wise clusters, as the cluster-size is linked to the diversity of patterns learned. Experiments have shown that our approach performs significantly better than state-of-the-art approaches on our MNIST SDG-MP benchmark. The future work will be dedicated to the study of more real life situations and to the problem of remaining but uncorrelated patterns at test-time.

Acknowledgment

This work was in part supported by the 4D Vision project funded by the Partner University Fund (PUF), a FACE program, as well as Alegoria project and Arès labcom projects, both funded by the French Research Agency, l'Agence Nationale de Recherche (ANR), under grants ANR-17-CE38-0014-03 and ANR-16-LCV2-0012-01, respectively.

References

- [1] Faruk Ahmed, Yoshua Bengio, Harm van Seijen, and Aaron Courville. Systematic generalisation with group invariant predictions. In *International Conference on Learning Representations*, 2021.
- [2] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2020.
- [3] Nitin Bansal, Xiaohan Chen, and Zhangyang Wang. Can we gain more from orthogonality regularizations in training deep networks? In *Advances in Neural Information Processing Systems*, 2018.
- [4] Fabio M Carlucci, Antonio D'Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [5] Fabio Maria Carlucci, Paolo Russo, Tatiana Tommasi, and Barbara Caputo. Hallucinating agnostic images to generalize across domains. In *IEEE/CVF International Conference on Computer Vision Workshop*, 2019.
- [6] Junbum Cha, Hanchchol Cho, Kyungjae Lee, Seunghyun Park, Yunsung Lee, and Sungrae Park. Domain generalization needs stochastic weight averaging for robustness on domain shifts. *arXiv preprint arXiv:2102.08604*, 2021.
- [7] Michael Cogswell, Faruk Ahmed, Ross Girshick, Larry Zitnick, and Dhruv Batra. Reducing overfitting in deep networks by decorrelating representations. *International Conference on Learning Representations*, 2016.
- [8] Nikolay Dagaev, Brett D Roads, Xiaoliang Luo, Daniel N Barry, Kaustubh R Patil, and Bradley C Love. A too-good-to-be-true prior to reduce shortcut reliance. *arXiv preprint arXiv:2102.06406*, 2021.
- [9] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International Conference on Machine Learning*, 2015.
- [10] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*, 2019.
- [11] Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, and David Balduzzi. Domain generalization for object recognition with multi-task autoencoders. In *IEEE/CVF International Conference on Computer Vision*, 2015.
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. In *Advances in Neural Information Processing Systems*, 2014.
- [13] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *International Conference on Learning Representations*, 2021.
- [14] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *IEEE/CVF International Conference on Computer Vision*, 2017.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016.
- [16] Elad Hoffer and Nir Ailon. Deep metric learning using triplet network. In *International Workshop on Similarity-Based Pattern Recognition*, 2015.
- [17] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International Conference on Machine Learning*, 2018.
- [18] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *IEEE/CVF International Conference on Computer Vision*, 2017.
- [19] Zeyi Huang, Haohan Wang, Eric P. Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In *IEEE/CVF European Conference on Computer Vision*, 2020.
- [20] Philip TG Jackson, Amir Atapour Abarghouei, Stephen Bonner, Toby P Breckon, and Boguslaw Obara. Style augmentation: data augmentation via style randomization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2019.
- [21] Byungju Kim, Hyunwoo Kim, Kyungsu Kim, Sungjin Kim, and Junmo Kim. Learning not to learn: Training deep neural networks with biased data. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2019.
- [22] Myeongjin Kim and Hyeran Byun. Learning texture invariant representation for domain adaptation of semantic segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [23] David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (rex). *arXiv preprint arXiv:2003.00688*, 2020.

- [24] Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs*. Available: <http://yann.lecun.com/exdb/mnist>, 2010.
- [25] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Deeper, broader and artier domain generalization. In *IEEE/CVF International Conference on Computer Vision*, 2017.
- [26] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.
- [27] Daniel Moyer, Shuyang Gao, Rob Brekelmans, Aram Galstyan, and Greg Ver Steeg. Invariant representations without adversarial training. In *Advances in Neural Information Processing Systems*, 2018.
- [28] Hyeonseob Nam, HyunJae Lee, Jongchan Park, Wonjun Yoon, and Donggeun Yoo. Reducing domain gap via style-agnostic networks. *arXiv preprint arXiv:1910.11645*, 2019.
- [29] Junhyun Nam, Hyuntak Cha, Sungsoo Ahn, Jaeho Lee, and Jinwoo Shin. Learning from failure: Training debiased classifier from biased classifier. In *Advances in Neural Information Processing Systems*, 2020.
- [30] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisaccho, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In *Advances in Neural Information Processing Systems Workshop*, 2011.
- [31] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *IEEE/CVF International Conference on Computer Vision*, 2019.
- [32] Mohammad Pezeshki, Sékou-Oumar Kaba, Yoshua Bengio, Aaron Courville, Doina Precup, and Guillaume Lajoie. Gradient starvation: A learning proclivity in neural networks. *arXiv preprint arXiv:2011.09468*, 2020.
- [33] Nathan Somavarapu, Chih-Yao Ma, and Zsolt Kira. Frustratingly simple domain generalization via image stylization. *arXiv preprint arXiv:2006.11207*, 2020.
- [34] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 2014.
- [35] Yaniv Taigman, Ming Yang, Marc’ Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2014.
- [36] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017.
- [37] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 2008.
- [38] Hong Xuan, Abby Stylianou, Xiaotong Liu, and Robert Pless. Hard negative examples are hard, but useful. In *IEEE/CVF European Conference on Computer Vision*, 2020.
- [39] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In *IEEE/CVF International Conference on Computer Vision*, 2019.
- [40] Xu Zhang, Felix X. Yu, Sanjiv Kumar, and Shih-Fu Chang. Learning spread-out local feature descriptors. In *IEEE/CVF International Conference on Computer Vision*, 2017.
- [41] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE/CVF International Conference on Computer Vision*, 2017.