

## Robustness and invariance properties of image classifiers

Présentée le 18 août 2022

Faculté des sciences et techniques de l'ingénieur  
Laboratoire de traitement des signaux 4  
Programme doctoral en génie électrique

pour l'obtention du grade de Docteur ès Sciences

par

### Apostolos MODAS

Acceptée sur proposition du jury

Prof. A. M. Alahi, président du jury  
Prof. P. Frossard, directeur de thèse  
Dr J.-H. Jacobsen, rapporteur  
Prof. T. Goldstein, rapporteur  
Prof. N. Flammarion, rapporteur



Philosophers have hitherto only interpreted the world in various ways.  
The point is to change it.  
— Karl Marx

To my beloved family...





# Acknowledgements

First of all, I would like to express my sincere gratitude to my supervisor Prof. Pascal Frossard. Not only he trusted me and gave me the opportunity to pursue my PhD with him, but he also guided and supported me with patience even during the hardest moments of my life. I will always be grateful!

I would also like to thank the members of my thesis committee, Prof. Alexandre Alahi, Prof. Nicolas Flammarion, Prof. Tom Goldstein, and Dr. Jörn-Henrik Jacobsen for their fruitful discussions and feedbacks.

I am extremely glad – and lucky – to have met and worked with Seyed and Guillermo. I definitely enjoyed every single minute of our collaboration. Their fascinating knowledge and intelligence helped me to evolve both as a researcher and as a person, and I am not sure how much of this thesis would exist without their continued support.

I would like to extend my sincere thanks to Mathieu Sinn and Beat Buesser for their hospitality during my internship at IBM Ireland. I also want to thank Prof. Andrea Cavallaro from QMUL for our 3-year collaboration. I really enjoyed our many discussions, and his constant motivation has always been a great support.

I also sincerely thank all the current and former LTS4 labmates Ádám, Ahmet, Alessandro, Arun, Bastien, Beril, Clément, Clémentine, Dorina, Eda, Ersi, Guille, Harshitha, Hermina, Isabela, Javier, Jelena, Mattia, Mireille, Nikos, Ortal, Renata, Roberto, Sahar, Seyed, Stefano, William, Yamin. My special thanks go to Mireille for sharing the office with me during a big part of my PhD years. It was really a pleasure to have such a kind and understanding friend as an officemate. I also want to thank our administrative Anne for her genuine help and for her incredible efficiency.

A very big thank you to my good friends Christos, Dimitris, George, Haris, Irene, Kostas, Maksym, Manos, Marios, Nikos, Sotiris, Thanos, Vaggelis, Victoria and Vlas for giving me strength during all these years away from home.

Finally, I am indebted to my family for their unconditional love and support, and for always being there for me since the beginning. And I am especially grateful to Meli for her support and sacrifices during all these years. To all of you, I know that I cannot give back all those things you have offered me.

*Lausanne, June 1, 2022*

Apostolos Modas



# Abstract

Deep neural networks have achieved impressive results in many image classification tasks. However, since their performance is usually measured in controlled settings, it is important to ensure that their decisions remain correct when deployed in noisy environments. In fact, deep networks are not robust to a large variety of semantic-preserving image modifications, even to imperceptible image changes – known as adversarial perturbations – that can arbitrarily flip the prediction of a classifier. The poor robustness of image classifiers to small data distribution shifts raises serious concerns regarding their trustworthiness. To build reliable machine learning models, we must design principled methods to analyze and understand the mechanisms that shape robustness and invariance. This is exactly the focus of this thesis.

First, we study the problem of computing sparse adversarial perturbations, and exploit the geometry of the decision boundaries of image classifiers for computing sparse perturbations very fast. We evaluate the robustness of deep networks to sparse adversarial perturbations in high-dimensional datasets, and reveal a qualitative correlation between the location of the perturbed pixels and the semantic features of the images. Such correlation suggests a deep connection between adversarial examples and the data features that image classifiers learn.

To better understand this connection, we provide a geometric framework that connects the distance of data samples to the decision boundary, with the features existing in the data. We show that deep classifiers have a strong inductive bias towards invariance to non-discriminative features, and that adversarial training exploits this property to confer robustness. We demonstrate that the invariances of robust classifiers are useful in data-scarce domains, while the improved understanding of the data influence on the inductive bias of deep networks can be exploited to design more robust classifiers.

Finally, we focus on the challenging problem of generalization to unforeseen corruptions of the data, and we propose a novel data augmentation scheme that relies on simple families of max-entropy image transformations to confer robustness to common corruptions. We analyze our method and demonstrate the importance of the mixing strategy on synthesizing corrupted images, and we reveal the robustness-accuracy trade-offs arising in the context of common corruptions. The controllable nature of our method permits to easily adapt it to other tasks and achieve robustness to distribution shifts in data-scarce applications.

Overall, our results contribute to the understanding of the fundamental mechanisms

## Abstract

---

of deep image classifiers, and pave the way for building more reliable machine learning systems that can be deployed in real-world environments.

**Keywords:** image classification, robustness, invariance, adversarial examples, distribution shifts, sparse perturbations, image transformations, data augmentation, decision boundary, deep learning, convolutional neural networks.

# Résumé

Les réseaux de neurones profonds ont obtenu des résultats impressionnants dans de nombreuses tâches de classification d'images. Cependant, comme leurs performances sont généralement mesurées dans des environnements contrôlés, il est important de s'assurer que leurs décisions restent correctes lorsqu'ils sont déployés dans des environnements bruyants. En fait, les réseaux profonds ne sont pas robustes à une grande variété de modifications d'images préservant la sémantique, même à des changements d'images imperceptibles – connus sous le nom de perturbations adverses – qui peuvent arbitrairement faire basculer la prédiction d'un classificateur. La faible robustesse des classificateurs d'images aux petits changements de distribution des données soulève de sérieuses inquiétudes quant à leur fiabilité. Pour construire des modèles d'apprentissage automatique fiables, nous devons concevoir des méthodes fondées sur l'analyse et la compréhension des mécanismes qui façonnent la robustesse et l'invariance. C'est l'objet de cette thèse. Tout d'abord, nous étudions le problème du calcul de perturbations adverses éparses et nous exploitons la géométrie des limites de décision des classificateurs d'images pour calculer très rapidement des perturbations éparses. Nous évaluons la robustesse des réseaux profonds aux perturbations adverses dispersées dans des ensembles de données à haute dimension, et nous révélons une corrélation qualitative entre l'emplacement des pixels perturbés et les caractéristiques sémantiques des images. Cette corrélation suggère une connexion profonde entre les exemples adverses et les caractéristiques des données que les classifieurs d'images apprennent.

Pour mieux comprendre cette connexion, nous fournissons un cadre géométrique qui fait le lien entre la distance des échantillons de données à la frontière de décision et les caractéristiques existant dans les données. Nous montrons que les classifieurs profonds ont un fort biais inductif vers l'invariance des caractéristiques non-discriminatives, et que l'entraînement contradictoire exploite cette propriété pour conférer de la robustesse. Nous démontrons que les invariances des classifieurs robustes sont utiles dans les domaines où les données sont rares, tandis que la meilleure compréhension de l'influence des données sur le biais inductif des réseaux profonds peut être exploitée pour concevoir des classifieurs plus robustes.

Enfin, nous nous concentrons sur le problème difficile de la généralisation aux corruptions imprévues des données, et nous proposons un nouveau schéma d'augmentation des données qui s'appuie sur des familles simples de transformations d'images à entropie maximale pour conférer de la robustesse aux corruptions courantes. Nous analysons notre

## Résumé

---

méthode et démontrons l'importance de la stratégie de mélange pour synthétiser les images corrompues, et nous révélons les compromis entre robustesse et précision dans le contexte des corruptions courantes. La nature contrôlable de notre méthode permet de l'adapter facilement à d'autres tâches et d'atteindre la robustesse aux changements de distribution dans les applications où les données sont rares.

Dans l'ensemble, nos résultats contribuent à la compréhension des mécanismes fondamentaux des classifieurs d'images profonds et ouvrent la voie à la construction de systèmes d'apprentissage automatique plus fiables pouvant être déployés dans des environnements réels.

**Mots clés :** classification d'images, robustesse, invariance, exemples adverses, changements de distribution, perturbations éparses, transformations d'images, augmentation des données, limite de décision, apprentissage profond, réseaux de neurones convolutifs.

# Contents

<b>Acknowledgements</b>	<b>i</b>
<b>Abstract (English/Français)</b>	<b>iii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Robustness of deep networks . . . . .	2
1.2 Adversarial robustness . . . . .	3
1.3 Towards better understanding of deep networks . . . . .	5
1.4 Thesis outline . . . . .	7
<b>2 Related work</b>	<b>9</b>
2.1 Robustness to distribution shifts . . . . .	9
2.1.1 Evaluating robustness to distribution shifts . . . . .	9
2.1.2 Improving robustness to distribution shifts . . . . .	11
2.2 Robustness to adversarial perturbations . . . . .	12
2.2.1 Evaluating robustness to adversarial perturbations . . . . .	12
2.2.2 Improving robustness to adversarial perturbations . . . . .	14
2.3 Understanding deep learning through robustness . . . . .	15
2.4 Summary . . . . .	16
<b>3 Sparse adversarial perturbations and image features</b>	<b>19</b>
3.1 Introduction . . . . .	19
3.2 Minimal sparse adversarial perturbations . . . . .	20
3.2.1 Sparsity constraints . . . . .	20
3.2.2 Linearization and boundary approximation . . . . .	21
3.2.3 SparseFool . . . . .	24
3.3 Experimental evaluation . . . . .	26
3.3.1 Performance analysis . . . . .	27
3.3.2 Perceptibility . . . . .	28
3.4 Analysis of sparse perturbations . . . . .	30
3.4.1 Shared semantic features . . . . .	30
3.4.2 Exclusiveness of adversarial robustness . . . . .	32
3.5 Conclusions . . . . .	33
	vii

<b>4</b>	<b>Analysis of learned features using adversarial proxies</b>	<b>35</b>
4.1	Introduction . . . . .	35
4.2	Decision boundary and discriminative features . . . . .	37
4.2.1	Proposed framework . . . . .	37
4.2.2	Evidence on synthetic examples . . . . .	38
4.3	Discriminative features of real datasets . . . . .	39
4.3.1	Boundary adaptation to data representation . . . . .	42
4.3.2	Invariance and elasticity of decision boundary . . . . .	42
4.3.3	Connections to catastrophic forgetting . . . . .	44
4.3.4	Discussion . . . . .	45
4.4	Sensitivity to position of training samples . . . . .	46
4.4.1	Evidence on synthetic examples . . . . .	46
4.4.2	Connections to adversarial training . . . . .	48
4.5	Implications in data-scarce applications . . . . .	50
4.5.1	An off-the-shelf task: filling level classification . . . . .	51
4.5.2	Adversarial invariance prevents overfitting . . . . .	54
4.5.3	Discussion . . . . .	57
4.6	Conclusions . . . . .	57
<b>5</b>	<b>Robustness to non-adversarial distribution shifts</b>	<b>59</b>
5.1	Introduction . . . . .	59
5.2	Towards robustness to common corruptions . . . . .	61
5.2.1	Invariance by removing features . . . . .	61
5.2.2	General model of visual corruptions . . . . .	63
5.3	PRIME data augmentations . . . . .	65
5.3.1	Instantiating the general model of visual corruptions . . . . .	65
5.3.2	Performance on common corruptions . . . . .	67
5.3.3	Unsupervised domain adaptation . . . . .	70
5.4	Robustness insights . . . . .	70
5.4.1	Contribution of transformations . . . . .	71
5.4.2	The importance of mixing . . . . .	72
5.4.3	Robustness vs Accuracy trade-off . . . . .	74
5.4.4	Sample complexity . . . . .	75
5.5	Improving filling level classification with PRIME . . . . .	76
5.5.1	Distribution shifts and PRIME augmentations . . . . .	77
5.5.2	PRIME transformation parameters . . . . .	79
5.5.3	Mixing parameters in PRIME . . . . .	79
5.5.4	Experimental validation . . . . .	81
5.6	Discussion . . . . .	82
5.7	Conclusions . . . . .	84
<b>6</b>	<b>Conclusion</b>	<b>87</b>
6.1	Summary . . . . .	87



6.2	Future directions . . . . .	89
<b>A</b>	<b>Appendix of Chapter 4</b>	<b>91</b>
A.1	Margin distribution of a linear classifier . . . . .	91
A.2	Training parameters . . . . .	91
A.3	Cross-dataset performance . . . . .	93
A.4	Adversarial training parameters . . . . .	94
A.5	Spectral decomposition on frequency “flipped” data . . . . .	94
<b>B</b>	<b>Appendix of Chapter 6</b>	<b>97</b>
B.1	Maximum entropy transformations . . . . .	97
B.1.1	Spectral domain . . . . .	98
B.1.2	Spatial domain . . . . .	98
B.1.3	Color domain . . . . .	99
B.2	PRIME implementation details . . . . .	99
B.2.1	Additional transformed examples . . . . .	99
B.2.2	Transformation parameters . . . . .	101
B.3	Performance per corruption . . . . .	101
B.4	Performance per severity level . . . . .	102
	<b>Bibliography</b>	<b>105</b>
	<b>Curriculum Vitae</b>	



# 1 Introduction

*“It ain’t what you don’t know that gets you into trouble.  
It’s what you know for sure that just ain’t so.”*

— Mark Twain

In recent years, deep neural networks have become the state-of-the-art in most machine learning benchmarks with the emergence of deep learning. Driven by the vast amounts of available data, deep learning systems have achieved outstanding performance in a wide range of applications, especially in the field of image classification. A standard way of assessing the “outstanding performance” of a classifier, is through its generalization. In practice, the generalization of a classifier is usually determined by its test accuracy, which is the performance on some held-out (test) data that have never been observed during training, but are typically assumed to come from the same distribution as the training data.

When deploying deep learning models in the real world, though, we expect to face very different environments than those of the training data, which sometimes can be noisy or even hostile. Building classifiers that are able to generalize to such conditions is crucial, especially for safety-critical applications like autonomous driving or biomedical imaging. Therefore, it is of utmost importance that the decisions of the classifiers remain robust, even in the presence of worst-case scenarios.

However, deep classifiers are actually brittle and far from robust, and their generalization performance significantly degrades when evaluated in conditions that are slightly different from the training ones. In particular, the decisions of deep classifiers can easily change through small semantic-preserving modifications of their images, even if such changes

would not affect the human perception in general.

Building robust deep networks that are invariant to small modifications of their data is a challenging problem, and a large body of research has focused on developing a variety of techniques for improving the robustness of image classifiers. Nevertheless, the problem of creating classifiers that are able to robustly generalize under different settings is far from being solved, while the exact mechanisms behind the profound vulnerability of deep networks are still not well understood. In this thesis, we propose novel algorithms for evaluating and understanding the robustness and invariance properties of deep networks, and build upon our insights to design methodologies for improving the robustness of image classifiers to distribution shifts and nuisances of their data.

### 1.1 Robustness of deep networks

In general, the low robustness of deep networks to different types of data distribution shifts can be observed in various scenarios and settings. For instance, deep classifiers struggle to generalize to images that have been slightly modified by common types of distortions that may occur during the acquisition or processing of the images (i.e., blur, color jitter etc) [1]. In the spatial domain, carefully crafted shifts on the image pixels are enough to cause significant performance drops [2]. In the spectral domain, the predictions of image classifiers are more sensitive to small perturbations in the low-frequency part of the image spectrum rather than the high-frequency one [3]. Furthermore, image classifiers tend to rely on spurious features of the image background, and hence when image objects appear on backgrounds that are typically presented in different object categories (i.e., a fish on a grass field), is enough to force the network to an erroneous decision (i.e., rabbit instead of fish) [4]. Finally, image classifiers do not exhibit low robustness only in cases where the images are somehow manipulated or distorted, but they can also struggle with completely new images, even if they are collected in the exact same way and from a similar distribution as the one used for training [5, 6].

One way to improve the robustness of the classifiers to distribution shifts or distortions, is by increasing the variability of the training data. The hope is that, the more samples the classifier observes, the better knowledge it obtains about the data distribution; hence, it is more likely to generalize to new samples. In this regard, an obvious approach for improving the robustness of deep classifiers is to increase the amount of training samples [7]. However, this technique might be impractical, since training on huge datasets requires a lot of computational power, and for many tasks the available training data can be quite scarce. In practice, one of the most common techniques for artificially increasing the amount and variability of training data is the so-called data augmentation, where one expects that the classifier becomes invariant – to some extent – to the transformations used to generate the augmented data during training [8]. That is, the decision of the classifier does not vary when test images are manipulated with the transformations

used during training. Note here that the choice of transformations used during data augmentation might be crucial for the overall robustness of the model. For instance, if we want the classifier to be robust to image rotations, then applying a random rotation in  $(0^\circ, 10^\circ]$  on the training images might not be sufficient to achieve robustness to rotations that are larger than  $10^\circ$ .

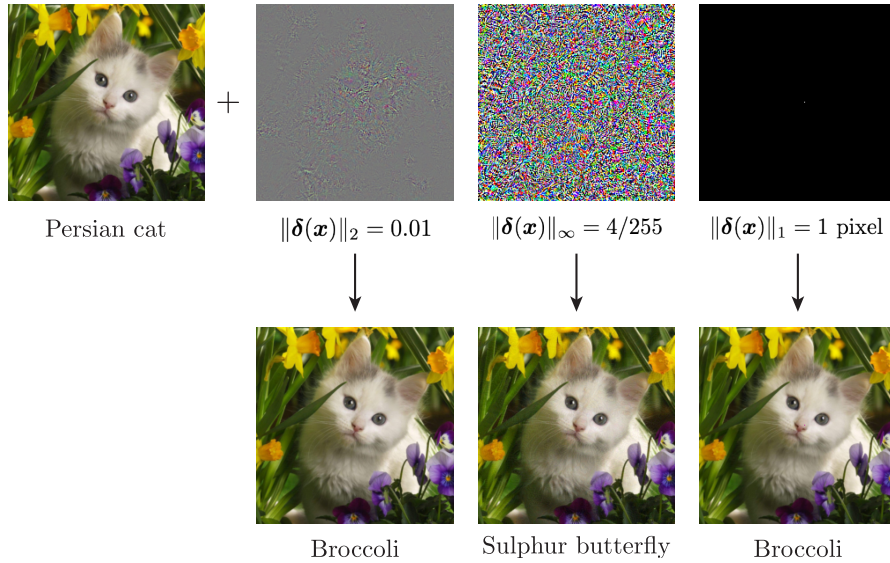
In general, although data augmentation or the use of additional data might work in practice for specific tasks, these methods are not definitive solutions for building robust classifiers. Ideally, given access to a finite amount of training data, we would still want a robust classifier to maintain its predictions when new data come from a slightly shifted distribution, or in the presence of distortions that do not change the semantics of the image. This implies that, for achieving robust and invariant image classification, the classifiers should ideally learn the underlying “concepts” that characterize the class of an image, rather than learning spurious image features that provide generalization only to specific test sets.

## 1.2 Adversarial robustness

Building classifiers with the desired invariance properties is however far from trivial. In fact, the robustness of deep neural networks is heavily challenged by the existence of the so-called *adversarial examples* [9]. These are data samples that have been modified by carefully-crafted, semantic-preserving perturbations that are often imperceptible to the human eye, but can change the prediction of the classifier to any arbitrary class. Adversarial examples represent one of the most intriguing phenomenon in the robustness literature, since their existence is not only constrained to image classifiers [10, 11, 12], but extends to virtually any application of deep learning such as autonomous driving [13], malware detection [14], natural language processing [15], speech recognition [16] and reinforcement learning [17].

Formalizing adversarial examples is an ill-posed problem, since there can exist multiple ways for defining and computing them depending on the data they are applied to, while the notion of perceptibility is subjective and can be arbitrary. Nevertheless, it is common in the literature to compute the  $\ell_p$ -norm of the perturbations as a perceptibility proxy. The goal then is to find adversarial perturbations of a small  $\ell_p$ -norm. An illustration of adversarial examples generated with adversarial perturbations of different  $\ell_p$  norms are shown in Fig. 1.1.

The simplest and most common adversarial perturbations are the additive ones. Let  $\mathbf{x} \in \mathbb{R}^D$  be a  $D$ -dimensional input, and  $f : \mathbb{R}^D \rightarrow \mathbb{R}^K$  be the final layer of a neural network (i.e., logits), such that, for any input  $\mathbf{x}$ ,  $F(\mathbf{x}) = \operatorname{argmax}_k f_k(\mathbf{x})$  represents the decision of that network, where  $f_k(\mathbf{x})$  denotes the component of  $f(\mathbf{x})$  that corresponds to the  $k$ th class. Formally, for a given data sample  $\mathbf{x}$  along with its associated label  $y$ , an



$$\begin{aligned} & \underset{\boldsymbol{\delta}}{\operatorname{argmax}} \mathcal{L}_{\theta}(\boldsymbol{x} + \boldsymbol{\delta}, y) \\ & \text{s.t. } \|\boldsymbol{\delta}\|_p \leq \varepsilon \\ & \quad \boldsymbol{\delta} \in \mathcal{C}, \end{aligned} \tag{1.1}$$

$$\begin{aligned} & \underset{\boldsymbol{\delta}}{\operatorname{argmin}} \quad \|\boldsymbol{\delta}\|_p \\ & \text{s.t.} \quad F(\mathbf{x}) \neq F(\mathbf{x} + \boldsymbol{\delta}) \\ & \quad \boldsymbol{\delta} \in \mathcal{C}. \end{aligned} \tag{1.2}$$

adversarial perturbations are not restricted to the additive  $\ell_p$ -norm definitions of Eq. (1.1) and Eq. (1.2), but can rather extend for instance to geometric [20, 21], functional [22], or even physical perturbations [23, 24].

Building classifiers that are robust against adversarial perturbations is still an open problem. Surprisingly, performing standard data augmentation does not improve adversarial robustness, even when it boosts robustness to specific types of distortions and transformations. Instead, one might need to adapt data augmentation in the worst-case settings and perform *adversarial training* [11, 25]. Adversarial training is a data augmentation method that replaces the clean images with their adversarial examples during training, and has been empirically shown to consistently result into more robust classifiers. Nevertheless, performing adversarial training can be computationally expensive, and it also comes at the cost of building classifiers that may perform worse on their standard test sets [26, 27]. Finally, adversarial training typically results into classifiers that are robust mostly on the perturbations they were trained on (e.g., perturbations of a specific  $\ell_p$  norm) [28], while its contribution to the robustness to non-adversarial distortions might not be significant [29].

### 1.3 Towards better understanding of deep networks

The profound gap between human and machine perception, along with the low robustness of deep neural networks, raises serious concerns regarding the implications on the security, safety, and fairness of deep learning systems. In order to integrate deep learning in sensitive tasks, such robustness vulnerabilities should be addressed. Therefore, it is imperative that we focus on understanding the mechanisms that govern the robustness properties of deep networks, and eventually design principled methods for building accurate and robust classifiers.

The main theme of this thesis is to exploit the connection between adversarial robustness and the geometry of deep networks [19, 30] as a proxy for understanding some fundamental properties of deep neural networks, and build upon our intuitions to design novel methodologies that confer robustness to distribution shifts of the data. Our main contribution is a better understanding of the role of data in the inductive biases of (robust) deep neural networks, which enables us to design new methods for improving the invariance properties of image classifiers and achieve better robustness to different types of distribution shifts. In what follows, we provide more details regarding the individual contributions of this thesis.

First, we focus on the computation of sparse adversarial perturbations. In general, evaluating the robustness of deep networks to perturbations in non-standard  $\ell_p$  regimes, for instance sparse perturbations with  $p \in \{0, 1\}$ , can be useful for identifying salient image features, due to the localized structure of the perturbations. However, generating sparse

adversarial perturbations can be computationally expensive, and existing methods [31, 32, 33] are impractical in large datasets. As a first contribution of this thesis, we provide a very fast, geometry-inspired sparse attack that exploits the low mean curvature of the decision boundaries to generate adversarial perturbations. This efficient algorithm enables us to thoroughly analyze the transferability and spatial properties of sparse perturbations, and provides empirical insights on the correlation between the data features and adversarial perturbations.

As a second step, we study the relationship between adversarial perturbations and the features of the data. The features of real datasets are not known a priori, and hence current literature [34, 35] relies either on experimental evaluations or on synthetic examples in order to investigate the connections between adversarial perturbations and data features. It is still not established which mechanisms are responsible for creating these connections, and it is unclear if this phenomenon depends on the data, the network architecture, or the learning algorithm. Propelled again by the geometric properties of adversarial examples, we propose a new geometric framework that connects the norm of minimal adversarial perturbations (i.e., distance to the classifier decision boundary) with the data features. By carefully manipulating the input data, we demonstrate that deep classifiers have a strong inductive bias towards invariance to non-discriminative features. In fact, we show that adversarial training actually exploits this property to build more robust classifiers. Furthermore, our insights explain why some methods for crafting adversarial examples are more efficient if they constrain the perturbations to the low-frequency subspaces of the Fourier basis: the reason is that the discriminative features of standard image datasets are aligned with low-frequency directions. Finally, we also demonstrate that the invariance properties of robust classifiers are beneficial in tasks where the available training data are scarce.

Last, we extend our focus to the more general case of robustness to non-adversarial perturbations of the data. In particular, we deal with the problem of robustness to common corruptions, a term that generally refers to typical image distortions (i.e., blur, color jitter, brightness variations, random noise etc.) that can occur during acquisition, storage, or processing of the images. In these settings, achieving robustness is more complex than merely inducing invariance through simple data feature manipulations. At the same time, defining such “common corruptions” is an ill-posed problem. For achieving robustness to common corruptions, prior works [36, 37, 38] have built complex data augmentation strategies, combining multiple methods to enrich the training data. These works typically introduce intricate design choices or heuristics, and it is hard to understand which elements of the methods are indeed crucial for improving robustness. We formulate a set of primitive image transformations in the spatial, color and spectral domains, and propose a systematic data augmentation scheme for improving the robustness of deep classifiers to common corruptions of their data. Our method achieves state-of-the-art robustness in multiple benchmarks, while its simplicity permits to perform an in-depth analysis of robustness in the context of common corruptions. In particular, we highlight



the importance of deploying a mixing strategy during the generation of the augmentation instances, and analyze the potential robustness-accuracy trade-offs and the benefits of generating the data augmentations during training (on-line). Finally, we demonstrate that our method can serve as an off-the-shelf solution for achieving robustness to distribution shifts that extend beyond the concept of common corruptions.

It should be noted that, although our focus in this thesis is on images, our geometric framework for connecting the data features to the decision boundaries, as well as the basic principles of our novel augmentation strategy can be easily extended to other modalities of data.

## 1.4 Thesis outline

The rest of thesis is organized as follows:

In Chapter 2, we review some of the prior works related to the problem of (i) evaluating the robustness of image classifiers, (ii) improving the robustness of classifiers to adversarial examples and distribution shifts, and (iii) studying the connections between adversarial robustness and the features learned by deep classifiers.

In Chapter 3, we study the problem of efficiently generating sparse adversarial perturbations. In particular, we design a geometry-inspired algorithm that is by orders of magnitude faster than existing methods and efficiently scales to high-dimensional datasets. Our empirical analysis sheds new light on the transferability of sparse adversarial perturbations, and on the connections between the image semantics and the features that deep networks learn.

In Chapter 4, we develop a novel methodology to characterize the relationship between the distance of a set of samples to the decision boundary, and the discriminative features of the dataset that are used by a classifier. We demonstrate that convolutional neural networks are invariant to non-discriminative features of a dataset. We further show that the decision boundary is very sensitive to the position of the training samples, and that adversarial training exploits this sensitivity and invariance bias to build robust classifiers. Finally, we demonstrate that the invariance properties of robust classifiers prevent overfitting in the non-classical task of estimating the amount of content within a container in scarce data regimes.

In Chapter 5, we study the general problem of robustness to common corruptions of the images. We formulate a new model for semantically-preserving image corruptions, and build on basic concepts to characterize the notions of transformation strength and diversity using a few transformation primitives. We propose a general data augmentation scheme that relies on simple yet rich families of max-entropy image transformations. Our method tops the current baselines on different common corruption datasets, while its

## Chapter 1. Introduction

---

simplicity makes it an effective tool for understanding common corruption robustness, and build classifiers with improved out-of-distribution generalization properties. In particular, we demonstrate that our method can easily be tuned for the context of classifying the filling level within a container, and generate augmentations with properties that resemble those of test-time distribution shifts.

Finally, in Chapter 6 we summarize the main outcomes of this thesis, and outline some of the potential future research directions.

## 2 Related work

In this chapter, we review some of the relevant works from the literature that are linked to the problems studied in this thesis. In particular, in Section 2.1 we summarize the methods for evaluating and improving the robustness of deep neural networks to distribution shifts, while in Section 2.2 we focus on robustness under adversarial settings. Finally, in Section 2.3 we review related works that connect the robustness of deep networks to different other properties of deep learning.

### 2.1 Robustness to distribution shifts

#### 2.1.1 Evaluating robustness to distribution shifts

**Natural distribution shifts** Reliable classification under distribution shifts has received a growing amount of interest in the field of machine learning [39]. Especially in the context of image classification [40], measuring robustness to natural distribution shifts (i.e., not synthetically induced through perturbations or transformations) that can arise in the real world is very important for a broad deployment of machine learning models. In fact, the authors in [5, 41] showed that classifiers that achieve state-of-the-art performance on standard benchmarks practically overfit to the given test sets of such benchmarks. In particular, they exhibit a significant accuracy drop on newly collected unseen images (new test set), despite being sampled from a distribution that is very similar to the one of the original test set. This phenomenon can be even more pronounced, as shown in [6], if the new test set is explicitly selected (adversarially) to cause extreme accuracy drops (up to 90%). Having a unified approach for evaluating the robustness of classifiers to natural distribution shifts is very important. To this end, the authors in [7] recently defined a very large testbed, and demonstrated that classifiers trained on more diverse data typically achieve better robustness. At the same time, they also demonstrated that robustness to synthetic shifts (i.e., through artificial perturbations or transformations) does not necessarily imply natural robustness.

**Common corruptions and nuisances** One of the most broadly studied distribution shifts are those resulting from the different types of visual distortions, or nuisances, that can synthetically/artificially happen on the images (i.e., during acquisition or processing of the images, or through artificial occlusions). Note, that, it is important to focus on cases where the distortions do not alter the semantic information of the images. For instance, although deep networks are relatively robust to small random noise [42, 43], they can still change their decision for larger random noise regimes, despite the fact that the image semantics are preserved. In particular, as shown in [44, 45], although the human visual system can be quite robust to different visual distortions (e.g., strong Gaussian noise or blurs), deep networks are not. And even if the classifiers are trained (i.e., through data augmentation) to be robust to some specific distribution shift (i.e., distortions introduced with Gaussian noise), they tend to overfit to that specific shift and perform poorly on other type of distortion types [45].

Apart from the different types of additive noise or blur [46, 47], the generalization of deep classifiers can also be affected by the existence of different types of nuisances in the data distribution. For instance, deep classifiers can exhibit a very low robustness to slight combinations of translations and rotations of the images [2, 48]. In general, measuring the robustness of a classifier on every possible corruption or nuisance is an ill-posed problem that lacks a formal description. For this reason, the research community has developed multiple standardized benchmarks to measure the robustness of classifiers to different distribution shifts, such as (i) common corruptions (e.g., noises, blurs, weather effects, digital transforms, spectral transforms) [1, 49], (ii) naturally captured blurry images [37], or (iii) visual artistic renditions [37].

**Inductive and distribution biases** Beyond standard corruptions or nuisances, deep networks might have problems generalizing to other types of distribution shifts, where one should not consider what the networks “have not seen”, but rather what they “have already seen” or how they learn to generalize. Hence, we have to focus on the different types of (inductive) biases that deep classifiers have, or inherit from the data. For instance, the large accuracy drop on newly collected unseen images (new test set) that was observed in [5] has been later found to be mainly caused by some statistical bias introduced during the collection/replication of the new dataset [50]. Another example of bias that can cause a generalization drop is the image resolution discrepancy between train and test time [51]: if random resize and crop applied during training generate images where the objects are larger, then the classifier might not perfectly generalize to test images of smaller objects.

Apart from data- or user-induced biases, the inductive bias of the architecture and/or the learning algorithm can also cause problems in generalizing to different distribution shifts. The authors in [52] showed that CNNs are more biased towards the texture rather than the shape of visual objects, and that they exhibit a significant accuracy drop when the texture changes or is totally absent (i.e., edge-based images or silhouettes).

Furthermore, the authors in [4] showed that deep networks learn spurious background features rather than focusing on the actual object, hence irrelevant background changes result in quite low accuracy. Additionally, the authors in [3] showed that Gaussian data augmentation and adversarial training, bias the model towards low-frequency information, which increases its robustness to high-frequency changes but makes it more vulnerable to low-frequency distortions. Furthermore, the authors in [53] and [54] showed that deep classifiers have a very strong inductive bias on fitting the data along specific directions, and that they cannot generalize to data where the information is not aligned with these directions. Finally, a recent study in [55] has assigned the low robustness properties of deep networks to the problem of shortcut learning, where the classifier learns some “easy” features (i.e., background information) for specific objects, but those features are not representative and generalizable for the actual object class.

### 2.1.2 Improving robustness to distribution shifts

**Additional training data** One straightforward approach to improve the robustness of deep networks to natural distribution shifts, consists in increasing the size of the training set by collecting new samples. However, this technique might be impractical, since training on huge datasets requires a lot of computational power, while for many tasks the amount of available data can be quite scarce. In particular, the authors in [56] showed that exploiting unlabelled data in combination with knowledge distillation can improve the generalization of a classifier to natural distribution shifts and corruptions. Furthermore, the authors in [7] demonstrated through a big testbed that using more data was the only technique that consistently led to better robustness on multiple distribution shifts. Interestingly, though, they observed that some models did not have any robustness benefits, since they are limited in exploiting further information from additional data (i.e., due to limited capacity).

**Data augmentation** The most common technique for improving the robustness to image distortions/transformations and nuisances, is to artificially increase the size and variability of the training set through data augmentation. The most standard augmentation is to apply a random horizontal flip and crop on the training image, in order to train classifiers so that they are invariant to the horizontal orientation and location of the object, which can also be beneficial for handling occlusions. For the latter, typical methods randomly mask part of the input [57, 58] in order to force the network to focus on different image features and become robust to occlusions. Other methods that prevent the network from overfitting to specific image features replace parts of the image with crops of another image [59, 60] or synthesize a new image as an interpolation of two images [61, 62]. Note here that such techniques necessarily require the use of soft labels during classification. Also, the texture-shape bias introduced in Sec. 2.1.1 can be avoided by training the classifier on a “stylized” version of the dataset, such that the classifier becomes invariant to the actual texture of the object [52].

Another line of research has managed to increase the generalization performance of image classifiers by exploiting during training a set of different spatial (e.g., translations, rotations), color (e.g., brightness), and spectral (e.g., sharpness) image transformations [63, 64]. When properly combined with a mixing strategy, such transformations during training have been later shown to be quite beneficial in improving the robustness of classifiers to different corruption benchmarks [36]; especially if the mixing is performed in an adversarial (worst-case) way [65]. Nevertheless, such transformations work well on small benchmarks (i.e., CIFAR-10) but do not perform equally well on higher-dimensional datasets. For this reason, the research community has proposed to address these limitations (i) either with more complex methods that use large autoencoders for generating more diverse augmentations [37, 38] or “denoise” the data [66], or (ii) with the creation of multi-view networks in the spectral domain for learning invariant representations of the data [67]. Nevertheless, such methods can be conceptually – and computationally – complex, which prevents from pinpointing their elements that actually contribute to the overall robust.

Adversarial training has also been used as a data augmentation method for improving the robustness of image classifiers to different non-adversarial corruptions. The authors in [68] proposed an adversarial training scheme that uses noise generated from uncorrelated distributions that maximize the classification loss. The authors in [69] studied theoretically and empirically the settings in which adversarial training improves robustness to out-of-distribution samples, while the authors in [29] analyzed the effectiveness of standard  $\ell_p$  adversarial training against common image corruptions, and proposed a relaxation of adversarial training in the embedding space [70].

**Architecture choice** Finally, it is important to mention that the deep network architecture plays an important role in the robustness to distribution shifts. For instance, compressing through pruning a standard CNN can significantly improve its robustness to common corruptions [71]. In addition, Vision Transformers [72] have been shown to be more robust to common corruptions of their data compared to standard CNNs [73, 74], when trained on very large data regimes, i.e., ImageNet-21K [75]. This is due, in part, to their different inductive bias [74]. Besides, performing properly tailored variants of adversarial training on ViTs can further boost their robustness to common corruptions [76, 77].

## 2.2 Robustness to adversarial perturbations

### 2.2.1 Evaluating robustness to adversarial perturbations

**Additive adversarial perturbations** The most widely studied way of measuring the adversarial robustness of deep networks is through the use of  $\ell_p$ -norm additive adversarial perturbations. These perturbations were first introduced in [9] and were computed using

a box-constrained L-BFGS algorithm, which is a simple method but not scalable to high dimensional image classification tasks. Nevertheless, computing adversarial examples is as fast and easy as slightly moving the image along the direction of the gradient of the loss function. That was initially shown with the FGSM algorithm [10], which inspired the creation of multiple gradient-based adversarial attacks that either solve Eq. (1.1) for a given perturbation budget  $\varepsilon$  [25, 78, 79], or compute the minimal adversarial perturbation of Eq. (1.2) that changes the decision of the classifier [11, 32, 80, 81]. Note here that one quite particular case is the so-called *universal perturbation*, which is a single perturbation that can be applied to every image and still change the classifier’s decision with high probability [12]. Another interesting case is the construction of  $\varepsilon$ -constrained adversarial perturbations (input space) that minimize the distance between specific internal representations of the network [82].

In some applications, evaluating the robustness of classifiers to perturbations with certain properties, such as sparsity, may be required. Sparse perturbations represent a very special case, since they are constrained to non-standard  $\ell_p$  norms, such as  $p = 0$ . Interestingly, the vulnerability of deep classifiers is extreme, since they can rather easily “break” by perturbing just a single or a few pixels of the image [31, 32, 33, 83, 84]. However, finding such perturbations is an NP-hard problem and computationally expensive, and some of the existing methods cannot scale in very high-dimensional datasets.

Another interesting type of structured perturbations is that of the subspace-constrained perturbations, which were first studied in [42]. Such perturbations can be constrained to any desired subspace, i.e., different frequency bands defined by the Fourier basis, and can shed light onto multiple spectral properties of deep networks. In particular, it has been shown that deep classifiers are more sensitive to low-frequency perturbations, compared to their high-frequency counterparts [85]. Furthermore, low-frequency perturbations can be more effective even against adversarially trained models [86, 87], while they have also been exploited to design more efficient black-box attacks [88, 89, 90]. That is, adversarial attacks where the adversary has access only to the output of the classifier.

**Non-additive adversarial perturbations** Beyond additive perturbations, one can think of more sophisticated ways to construct adversarial examples. For instance, the robustness of deep classifiers to adversarial geometric transformations is studied in [2, 21, 48], where image classifiers are shown to be very vulnerable to small rotations, translations and affine transformations. Furthermore, some works have evaluated the robustness of deep classifiers to other adversarial perturbation regimes such as color transformations [22, 91, 92], occlusions [23, 93, 94], and deformations [20, 94].

**Black-box settings** Finally, for the sake of completeness, note that the robustness of deep networks can be evaluated in black-box settings, where the only available information during the construction of the adversarial examples is represented by the predictions or the class probabilities of the network. This scenario is quite realistic and many methods

have been proposed for computing black-box adversarial perturbations [95, 96, 97, 98, 99]. However, this is mainly a security concern, which is not the main scope of this thesis.

### 2.2.2 Improving robustness to adversarial perturbations

**Adversarial training** The most standard way for building classifiers that are robust to adversarial examples is the so-called *adversarial training*, which can be seen as a type of data augmentation that replaces the clean images with their adversarial examples during training. Adversarial training was early introduced together with the first methods for computing adversarial examples [9, 10, 11]. One of the risks of using adversarial training is that it can cause the classifier to overfit to specific types of adversarial examples (i.e., generated with FGSM). One can avoid this issue if the algorithm that generates adversarial perturbations is selected properly. In this sense, the scheme that seems to consistently result into more robust classifiers against a variety of adversarial attacks is the one proposed in [25], which generates adversarial examples using the PGD algorithm [100]. This algorithm maximizes the classifier’s loss within a ball of specific radius around the data samples. An alternative, as shown recently [70], would be to perform adversarial training with adversarial examples that maximize the classifier’s loss, but at the same time constrain the adversarial representations of the classifier to be close to the original representations. This technique has been shown to build classifiers that are more robust to a variety of adversarial attacks that also extend beyond  $\ell_p$ -norm settings. Nevertheless, the main problem with adversarial training is that it is computationally expensive, and the main focus of current adversarial training schemes is to achieve similar or higher robustness, but at a lower cost [101, 102].

**Regularization** In order to avoid the computational cost of adversarial training, many works have focused on increasing the stability of the classifiers through different types of regularization. The authors in [103] proposed to smooth the norm of the gradient at each layer of the classifier, while other works have attempted to improve the robustness by regularizing the input gradient or the full Jacobian [104, 105, 106]. Furthermore, it has been shown that, second-order regularization techniques that penalize the curvature of the input loss function [107, 108, 109] result into classifiers that exhibit robustness similar to adversarial training. Lately, some works have focused on penalizing the curvature of the loss landscape in the weight space to improve robustness [110], and others that improve the stability of adversarial training by imposing constraints between the clean and noisy input gradients [111].

Note here, that, reducing the curvature of the input loss landscape to create more robust models can also be exploited to achieve robustness with guarantees, i.e., theoretical certificates that demonstrate a specific level of robustness for a given neural network. Indeed, certifiable adversarial defenses, like randomized smoothing [112, 113, 114], also implicitly regularize curvature by averaging the decision of a classifier on randomly



perturbed samples. This way, one effectively convolves the loss landscape of a classifier with the probability density function of the perturbation distribution, hence, reducing the mean curvature of the loss landscape and smoothing the input geometry of the classifier.

## 2.3 Understanding deep learning through robustness

**Geometric insights** Although the lack of robustness in deep networks raises serious concerns regarding their security and trustworthiness, the whole process of evaluating and improving their robustness has revealed many important properties of deep learning. The first benefit of adversarial perturbations is that they enable us to study the local geometry of the decision boundaries [18] and obtain multiple insights regarding the topology and the geometry of the decision regions [115]. Furthermore, it has been shown that adversarial perturbations span a low-dimensional subspace of the input space [116, 117], and that such subspaces of different networks are also aligned, which can justify the transferability properties of adversarial examples across neural networks [116]. In addition, it has been observed that adversarial perturbations can exist in multiple directions [42, 115, 116] and by studying these different perturbations one can reveal geometric properties such as flatness and curvature of the decision boundaries. In fact, adversarial directions are mostly assigned to curved decision boundaries, and universal adversarial perturbations correspond to those shared curved directions [117]. Moreover, it has been observed that adversarial training creates decision boundaries that lie further away from most data samples [11] and with lower curvature compared to standard models [107], which justifies why regularization methods improve the robustness of deep networks. Nevertheless, the fact that the decision boundaries exhibit a low mean curvature might in some cases increase the vulnerability of deep networks, since this “flatness” property can be exploited as a prior for designing better adversarial attacks [89, 90, 118].

**Connection to data features** In deep learning, the networks are supposed to find “good” features of the training data, for a given learning task. It thus means that the key to the success of deep learning is the choice of features exploited by a neural network, and adversarial examples are actually correlated with such features. In particular, the authors in [34] showed that adversarial perturbations span a low-dimension but highly discriminative subspace of the input, and that deep networks exploit simple and brittle features of the dataset, i.e., non-robust features that are aligned with adversarial perturbations. Furthermore, this was also supported by the authors in [35]. They showed that training on adversarial examples and with the corresponding adversarial labels, results into classifiers that achieve non-trivial accuracy on the original unmodified dataset. This means that, the only way that the network trained on the adversarial samples can generalize to the unmodified test set, is by exploiting the non-human-aligned features introduced by the adversarial perturbations themselves, i.e., the non-robust features. On the other hand, it has been shown that adversarially trained (robust) classifiers learn features that correlate better with semantically meaningful features of the input

images [27, 119, 120, 121]. In this sense, one can actually use the (robust) representations of these models as effective primitives for semantic image manipulations, in order to perform complex tasks such as image generation or inpainting [119].

**Generalization** Adversarial training filters out the non-robust spurious features which, however, would typically be used by the network to achieve good accuracy on the test set. The effect of this is reflected on the so-called robustness/accuracy tradeoff [26, 27, 122], while it also justifies the empirical observation that one needs more data to generalize when adversarial training is used [123]. In practice, it has also been shown that adding more data into the adversarial training process can improve robustness and decrease the generalization gap [124, 125, 126]. On the contrary, humans are not susceptible to adversarial perturbations, suggesting that they do not exploit non-robust features [127]. Hence, it is argued that deep networks should also be able to achieve good generalization by using only robust features [128, 129]. Some theoretical results indicate that there exist at least some synthetic distributions in which adversarial robustness and accuracy are positively correlated [130, 131]. In fact, it has been recently shown that adversarial robustness and generalization are tightly close, and that the gap between the data samples and the decision boundary (minimal adversarial perturbations) can be used to predict the test accuracy [132, 133].

**Dynamics of learning and inductive bias** Finally, by tracking the evolution of adversarial perturbations during training, one can reveal different inductive biases and geometric properties of deep networks. For instance, the inductive bias of deep networks towards invariance has recently been argued to be prejudicial for classification as it can decrease the alignment between our human perception and the network’s decisions [134, 135]. This happens because adversarial training forces the networks to latch onto overly-robust features of the training set that are not human-aligned. Furthermore, in the deep learning community it is a common belief that neural networks are relatively immune to overfitting [136]. Indeed, the train and the validation losses of a neural network during training are clearly decreasing. Nevertheless, with adversarial training, it has recently been shown that the best robustness in the validation set is consistently found at the middle stages of training [137]. This confirms that adversarially trained neural networks have a tendency to overfit to the adversarial examples observed during training. This phenomenon is known as robust overfitting, and it has recently been shown to be one reason for some of the reported differences between different adversarial defenses. In fact, the state-of-the-art adversarial training technique uses early-stopping to obtain the best robustness results [138].

## 2.4 Summary

We summarize the main points of this chapter, in the light of the contributions of this thesis and upcoming challenges:

- Deep image classifiers are extremely vulnerable to adversarial manipulations of their input samples. Different methods have been developed to assess their robustness properties, even in very challenging settings where the perturbations are governed by sparsity constraints. However, such methods are computationally expensive, thus rendering them impractical for high-dimensional datasets. One of the goals of this thesis is to provide a fast and scalable method for evaluating the robustness of deep networks to sparse adversarial perturbations. This will allow to analyze potential correlations between the spatial location of the perturbed sparse pixels and the semantic features of the images.
- The role of data in the generalization and robustness properties of deep classifiers, and the mechanisms that drive the networks to learn or ignore specific image features are not fully understood yet. We here provide a framework that connects the local geometry of the decision boundaries with the features of the dataset, and demonstrate that adversarial training exploits the invariance bias of deep networks and their sensitivity to the position of the training samples for building robust classifiers.
- Deep networks are not only vulnerable to adversarial manipulations, but they also exhibit poor robustness to common corruptions of their data. For improving this vulnerability, prior works have mostly focused on increasing the complexity of their training pipelines in the name of diversity, making it hard to pinpoint which elements of these methods meaningfully contribute to the overall robustness. In this thesis, we formulate a model for characterizing semantically-preserving image corruptions, and propose a principled and simple approach that is based on a mixture of few transformation primitives to confer robustness to common corruptions.



# 3 Sparse adversarial perturbations and image features

*“It is the little bits of things that fret and worry us.  
We can dodge an elephant, but we can’t dodge a fly.”*

— Josh Billings

## 3.1 Introduction

Most of the existing methods in the adversarial robustness literature compute  $\ell_p$ -norm adversarial perturbations for  $p \in \{2, \infty\}$ . However, understanding the vulnerabilities of deep neural networks in non-standard  $\ell_p$  regimes is also important. In particular, sparse perturbations for  $p \in \{0, 1\}$  are quite interesting, since their localized nature can reveal important parts of the image that such perturbations exploit [139]. Prior works on sparse perturbations change the pixels either based on their saliency score [31], or using evolutionary algorithms [33], or with greedy local search algorithms [83]. In general though, computing sparse adversarial perturbations with minimal  $\ell_0$  norm is an NP-hard problem, and current algorithms are all characterized by high complexity and can hardly scale to high-dimensional datasets. Hence, a fast and accurate method for computing sparse perturbations is still needed to easily analyze different robustness properties of image classifiers.

In this chapter, we propose SparseFool, a geometry inspired algorithm that exploits the low mean curvature of the decision boundaries to linearize the sparsity constraints, and thus compute adversarial perturbations efficiently. We show through extensive evaluations that (i) our method computes sparse perturbations much faster than the

---

Part of this chapter has been published in  
“SparseFool: A few pixels make a big difference”. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019 [140].

existing methods, and (ii) it can scale efficiently to high dimensional datasets. We further propose a method to control the magnitude of the perturbation applied on every pixel – and hence the perceptibility of the resulting perturbation –, while retaining the levels of sparsity and complexity. We analyze visually the image features affected by our attack, and show the existence of some shared semantic information across different images and networks, which suggests a strong correlation between adversarial examples and the semantic features of the images. Finally, we show that classifiers that are adversarially trained with  $\ell_\infty$  perturbations are not robust to sparse perturbations, which indicates that the image features that are related to  $\ell_1$  perturbations are different from those of  $\ell_\infty$  perturbations.

The rest of the chapter is organized as follows: in Sec. 3.2, we describe the challenges for computing minimal sparse adversarial perturbations, and provide an efficient method that linearizes the initial optimization problem to obtain an approximate solution. In Sec. 3.3 we evaluate our algorithm on multiple datasets and networks, and perform comparisons with other state-of-the-art methods. Finally, in Sec. 3.4 we analyze empirically the resulting perturbations and demonstrate visual correlations between the perturbed pixels and the semantic information of the images.

## 3.2 Minimal sparse adversarial perturbations

### 3.2.1 Sparsity constraints

Recall from Sec. 1.2 that minimal adversarial perturbations are defined as

$$\begin{aligned} & \underset{\boldsymbol{\delta}}{\operatorname{argmin}} \|\boldsymbol{\delta}\|_p \\ & \text{s.t. } F(\mathbf{x}) \neq F(\mathbf{x} + \boldsymbol{\delta}) \\ & \quad \boldsymbol{\delta} \in \mathcal{C}, \end{aligned} \tag{3.1}$$

where  $f : \mathbb{R}^D \rightarrow \mathbb{R}^K$  is the final layer of a neural network (i.e., logits), such that, for any input  $\mathbf{x} \in \mathbb{R}^D$ ,  $F(\mathbf{x}) = \operatorname{argmax}_k f_k(\mathbf{x})$  represents the decision of that network, with  $f_k(\mathbf{x})$  denoting the  $k$ th component of  $f(\mathbf{x})$  that corresponds to the  $k$ th class.

Most of the existing adversarial algorithms solve Eq. (3.1) for  $p = 2$  or  $\infty$ , resulting in dense but imperceptible perturbations. For the case of sparse perturbations, the goal is to minimize the number of perturbed pixels required to fool the network, which corresponds to minimizing  $\|\boldsymbol{\delta}\|_0$ . Unfortunately, this is an NP-hard problem, for which reaching a global minimum cannot be guaranteed in general [141, 142, 143]. There exist different methods [143, 144] to avoid the computational burden of this problem, with the  $\ell_1$  relaxation being the most common: the minimization of  $\|\boldsymbol{\delta}\|_0$  under linear constraints

can be approximated by solving the corresponding convex  $\ell_1$  problem [145, 146, 147]<sup>1</sup>.

DeepFool [11] is an algorithm that exploits such a relaxation, by adopting an iterative procedure that includes a linearization of the classifier at each iteration, in order to estimate the minimal adversarial perturbation  $\delta$ . Specifically, at each iteration  $i$ , the classifier  $f$  is linearized around the current point  $\mathbf{x}^{(i)}$ , the minimal perturbation  $\delta^{(i)}$  (in an  $\ell_2$  sense) is computed as the projection of  $\mathbf{x}^{(i)}$  onto the linearized hyperplane, and the next iterate  $\mathbf{x}^{(i+1)}$  is updated. Such a linearization procedure could actually be used to solve Eq. (3.1) for  $p = 1$ , so as to obtain an approximation to the  $\ell_0$  solution. In fact, by generalizing the projection to any  $\ell_p$  norm,  $\ell_1$ -DeepFool provides an efficient way for computing sparse adversarial perturbations using the  $\ell_1$  projection.

Although the  $\ell_1$ -DeepFool efficiently computes sparse perturbations, it does not explicitly respect the validity of the adversarial image values; that is  $\mathcal{C} = \{\delta : \mathbf{x} + \delta \in [0, 255]^D\}$  in Eq. (3.1). For  $\ell_2$  and  $\ell_\infty$  perturbations, almost every pixel of the image is typically distorted with noise of small magnitude. Hence, one can practically “ignore” such constraints [10, 11] since it is unlikely that many pixels will be out of their valid range; and even then, clipping the invalid values after the computation of such adversarial images could have a minor impact. This, however, is not the case for sparse perturbations, which typically result in a few distorted pixels but of high-magnitude perturbation, and clipping the values after computing the adversarial image can have a significant impact on the success of the attack.

We investigate the effect of such clipping operation on the quality of adversarial perturbations generated by  $\ell_1$ -DeepFool. For example, with perturbations computed for a VGG-16 [151] trained on ImageNet [152], we observed that  $\ell_1$ -DeepFool achieves almost 100% of fooling rate by perturbing only 0.037% of the pixels on average. However, clipping the pixel values of adversarial images to  $[0, 255]$  results in a fooling rate of merely 13%. Furthermore, incorporating the clipping operator inside the iterative procedure of the algorithm does not improve the results. In other words,  $\ell_1$ -DeepFool fails to properly compute sparse perturbations. This underlies the need for an improved attack algorithm that natively takes into account the validity of generated adversarial images, as proposed in the next sections.

### 3.2.2 Linearization and boundary approximation

Based on the above discussion, minimal sparse perturbations should be obtained as

$$\begin{aligned} & \underset{\delta}{\operatorname{argmin}} \|\delta\|_0 \\ & \text{s.t. } F(\mathbf{x}) \neq F(\mathbf{x} + \delta) \\ & \quad \mathbf{l} \preceq \mathbf{x} + \delta \preceq \mathbf{u}, \end{aligned} \tag{3.2}$$

---

<sup>1</sup>Under some conditions, the solution of such approximation is indeed optimal [148, 149, 150].




Figure 3.1: The approximated decision boundary  $\mathcal{B}$  in the vicinity of the datapoint  $\mathbf{x}$  that belongs to class  $F(\mathbf{x}) = 1$ .  $\mathcal{B}$  can be seen as a one-vs-all linear classifier for class 1.

where  $\mathbf{l}, \mathbf{u} \in \mathbb{R}^D$  denote the lower and upper bounds of the values of  $\mathbf{x} + \boldsymbol{\delta}$ , such that  $l_i \leq x_i + \delta_i \leq u_i$ ,  $i = 1 \dots D$ .

To find an efficient relaxation to Eq. (3.2), we focus on the geometric characteristics of the decision boundary, and specifically on its curvature. It has been shown [18, 34, 115] that the decision boundaries of state-of-the-art deep networks have a quite low mean curvature in the neighborhood of data samples. In other words, for a datapoint  $\mathbf{x}$  and its corresponding minimal  $\ell_2$  adversarial perturbation  $\mathbf{v}$ , the decision boundary at the vicinity of  $\mathbf{x}$  can be locally approximated by a hyperplane passing through the datapoint  $\mathbf{x}_B = \mathbf{x} + \mathbf{v}$ , and a normal vector  $\mathbf{w}$  (see Fig. 3.1 for an illustration). Hence, by exploiting this property we can relax Eq. (3.2) so that sparse adversarial perturbations are computed by solving the following  $\ell_1$  box-constrained optimization problem

$$\begin{aligned} & \underset{\boldsymbol{\delta}}{\operatorname{argmin}} \|\boldsymbol{\delta}\|_1 \\ & \text{s.t. } \mathbf{w}^T((\mathbf{x} + \boldsymbol{\delta}) - \mathbf{x}_B) = 0 \\ & \quad \mathbf{l} \preceq \mathbf{x} + \boldsymbol{\delta} \preceq \mathbf{u}. \end{aligned} \tag{3.3}$$

For solving Eq. (3.3), simply computing the  $\ell_1$  projection of  $\mathbf{x}$  onto the approximated hyperplane does not guarantee a solution. For a perturbed image, consider the case where some of its values exceed the bounds defined by  $\mathbf{l}$  and  $\mathbf{u}$ . Thus, by readjusting the invalid values to match the constraints, the resulted adversarial image may eventually not lie onto the approximated hyperplane. For this reason, we propose an iterative procedure, where at each iteration we project only towards one single coordinate of the normal vector  $\mathbf{w}$  at a time. If projecting  $\mathbf{x}$  towards a specific direction does not provide a solution, then the perturbed image at this coordinate has reached its extrema value. Therefore, at the next iteration this direction should be ignored, since it cannot contribute any further to finding a better solution.

Formally, let  $S$  be a set containing all the directions of  $\mathbf{w}$  that cannot contribute to the



---

**Algorithm 1:** LinearSolver

---

**Input:** image  $\mathbf{x}$ , normal  $\mathbf{w}$ , boundary point  $\mathbf{x}_B$ , projection operator  $Q$ .  
**Output:** perturbed point  $\mathbf{x}^{(i)}$

- 1 Initialize:  $\mathbf{x}^{(0)} \leftarrow \mathbf{x}$ ,  $i \leftarrow 0$ ,  $S = \{\}$
- 2 **while**  $\mathbf{w}^T(\mathbf{x}^{(i)} - \mathbf{x}_B) \neq 0$  **do**
- 3      $\boldsymbol{\delta} \leftarrow \mathbf{0}$
- 4      $d \leftarrow \operatorname{argmax}_{j \notin S} |w_j|$
- 5      $\delta_d \leftarrow \frac{|\mathbf{w}^T(\mathbf{x}^{(i)} - \mathbf{x}_B)|}{|w_d|} \cdot \operatorname{sign}(w_d)$
- 6      $\mathbf{x}^{(i+1)} \leftarrow Q(\mathbf{x}^{(i)} + \boldsymbol{\delta})$
- 7      $S \leftarrow S \cup \{d\}$
- 8      $i \leftarrow i + 1$
- 9 **end**
- 10 **return**  $\mathbf{x}^{(i)}$

---

minimal perturbation anymore. Then, the perturbation  $\boldsymbol{\delta}$  is updated through the  $\ell_1$  projection of the current perturbed iterate  $\mathbf{x}^{(i)}$  onto the estimated hyperplane as

$$\delta_d \leftarrow \frac{|\mathbf{w}^T(\mathbf{x}^{(i)} - \mathbf{x}_B)|}{|w_d|} \cdot \operatorname{sign}(w_d), \quad (3.4)$$

where  $d$  is the index of the maximum absolute value of  $\mathbf{w}$  that has not already been used

$$d \leftarrow \operatorname{argmax}_{j \notin S} |w_j|. \quad (3.5)$$

Before proceeding to the next iteration, we must ensure the validity of the values of the next iterate  $\mathbf{x}^{(i+1)}$ . For this reason, we use a projection operator  $Q(\cdot)$  that readjusts the values of the updated point that are out of bounds, by projecting  $\mathbf{x}^{(i)} + \boldsymbol{\delta}$  onto the box-constraints defined by  $\mathbf{l}$  and  $\mathbf{u}$ . Hence, the new iterate  $\mathbf{x}^{(i+1)}$  is updated as  $\mathbf{x}^{(i+1)} \leftarrow Q(\mathbf{x}^{(i)} + \boldsymbol{\delta})$ . Note here that the bounds  $\mathbf{l}$ ,  $\mathbf{u}$  are not limited to only represent the dynamic range of an image, but can be generalized to satisfy any similar restriction. For example, as we will describe later in Sec. 3.3, they can be used to control the perceptibility of the computed adversarial images.

The next step is to check if the new iterate  $\mathbf{x}^{(i+1)}$  has reached the approximated hyperplane. Otherwise, it means that the perturbed image at the coordinate  $d$  has reached its extrema value, and thus we cannot change it any further; perturbing towards the corresponding direction will have no effect. Thus, we reduce the search space by adding the direction  $d$  to the forbidden set  $S$ , and repeat the procedure until we reach the approximated hyperplane. The algorithm for solving the linearized problem is summarized in Algorithm 1.

Finally, in order to complete our solution we focus on the linear approximation of the

decision boundary. Recall that we need to find a boundary point  $\mathbf{x}_B$ , along with the corresponding normal vector  $\mathbf{w}$ . Finding  $\mathbf{x}_B$  is analogous to computing (in a  $\ell_2$  sense) the minimal adversarial example of  $\mathbf{x}$ . Recall that DeepFool iteratively moves  $\mathbf{x}$  towards the decision boundary, and stops as soon as the perturbed data point reaches the other side of the boundary. Therefore, the final point usually lies very close to the decision boundary, and thus,  $\mathbf{x}_B$  can be very well approximated by  $\mathbf{x} + \boldsymbol{\delta}_{\text{DF}}$ , with  $\boldsymbol{\delta}_{\text{DF}}$  being the  $\ell_2$ -DeepFool perturbation. Let us describe the decision boundary between the class assigned to  $\mathbf{x}_B$ ,  $F(\mathbf{x}_B)$ , and any other class  $F(\mathbf{x})$ , by considering the zero level set of  $f$

$$\mathcal{B} = \left\{ \mathbf{x} : f_{F(\mathbf{x}_B)}(\mathbf{x}) - f_{F(\mathbf{x})}(\mathbf{x}) = 0 \right\}.$$

Using a first-order Taylor expansion at  $\mathbf{x}_B$ , the decision boundary can be expressed as

$$f_{F(\mathbf{x}_B)}(\mathbf{x}_B) + \nabla f_{F(\mathbf{x}_B)}(\mathbf{x}_B)^T (\mathbf{x} - \mathbf{x}_B) - f_{F(\mathbf{x})}(\mathbf{x}_B) - \nabla f_{F(\mathbf{x})}(\mathbf{x}_B)^T (\mathbf{x} - \mathbf{x}_B) = 0.$$

Since  $\mathbf{x}_B$  lies very close to the decision boundary, then  $f_{F(\mathbf{x}_B)}(\mathbf{x}_B) \approx f_{F(\mathbf{x})}(\mathbf{x}_B)$  and

$$\left( \nabla f_{F(\mathbf{x}_B)}(\mathbf{x}_B)^T - \nabla f_{F(\mathbf{x})}(\mathbf{x}_B)^T \right) (\mathbf{x} - \mathbf{x}_B) = 0.$$

We can then define the estimated normal vector  $\mathbf{w}$  to the decision boundary as

$$\mathbf{w} := \nabla f_{F(\mathbf{x}_B)}(\mathbf{x}_B) - \nabla f_{F(\mathbf{x})}(\mathbf{x}_B), \quad (3.6)$$

Hence, the decision boundary can now be approximated through the affine hyperplane  $\mathcal{B} \triangleq \{ \mathbf{x} : \mathbf{w}^T (\mathbf{x} - \mathbf{x}_B) = 0 \}$ , and sparse adversarial perturbations are computed by applying Algorithm 1.

#### 3.2.3 SparseFool

However, although we expected a single-step solution, in many cases the algorithm did not fool the classifier. This is due to the fact that the decision boundaries of the networks are only *locally flat*. Thus, if the  $\ell_1$  perturbation moves the datapoint  $\mathbf{x}$  away from the flat area, then the perturbed point will not reach the other side of the decision boundary.

We mitigate this convergence issue with an iterative method, namely *SparseFool*, where each iteration includes the linear approximation of the decision boundary. Specifically, at iteration  $i$ , the boundary point  $\mathbf{x}_B^{(i)}$  and the normal vector  $\mathbf{w}^{(i)}$  are estimated using  $\ell_2$ -DeepFool based on the current iterate  $\mathbf{x}^{(i)}$ . Then, the next iterate  $\mathbf{x}^{(i+1)}$  is updated through the solution of Algorithm 1, having though  $\mathbf{x}^{(i)}$  as the initial point. The algorithm terminates when  $\mathbf{x}^{(i)}$  changes the label of the network. An illustration of SparseFool is given in Fig. 3.2, and the algorithm is summarized in Algorithm 2.

In our experiments, we observed that instead of using the boundary point  $\mathbf{x}_B^{(i)}$  at the step

### 3.2 Minimal sparse adversarial perturbations




Figure 3.2: Illustration of SparseFool algorithm. With green we denote the  $\ell_2$ -DeepFool adversarial perturbations computed at each iteration. In this example, the algorithm converges after 2 iterations, and the total perturbation is  $\delta = \mathbf{x}^{(2)} - \mathbf{x}^{(0)}$ .

---

#### Algorithm 2: SparseFool

---

**Input:** image  $\mathbf{x}$ , projection operator  $Q$ , classifier  $f$ .  
**Output:** perturbation  $\delta$

- 1 Initialize:  $\mathbf{x}^{(0)} \leftarrow \mathbf{x}$ ,  $i \leftarrow 0$
- 2 **while**  $F(\mathbf{x}^{(i)}) = F(\mathbf{x}^{(0)})$  **do**
- 3      $\delta_{\text{DF}} = \text{DeepFool}(\mathbf{x}^{(i)})$
- 4      $\mathbf{x}_B^{(i)} = \mathbf{x}^{(i)} + \delta_{\text{DF}}$
- 5      $\mathbf{w}^{(i)} = \nabla f_{F(\mathbf{x}_B^{(i)})}(\mathbf{x}_B^{(i)}) - \nabla f_{F(\mathbf{x}^{(i)})}(\mathbf{x}_B^{(i)})$
- 6      $\mathbf{x}^{(i+1)} = \text{LinearSolver}(\mathbf{x}^{(i)}, \mathbf{w}^{(i)}, \mathbf{x}_B^{(i)}, Q)$
- 7      $i \leftarrow i + 1$
- 8 **end**
- 9 **return**  $\delta = \mathbf{x}^{(i)} - \mathbf{x}^{(0)}$

---

6 of SparseFool, better convergence can be achieved by going further into the other side of the boundary, and find a solution for the hyperplane passing through the datapoint  $\mathbf{x}^{(i)} + \lambda(\mathbf{x}_B^{(i)} - \mathbf{x}^{(i)})$ , where  $\lambda \geq 1$ . Specifically, as shown in Fig. 3.3, this over-shooting parameter is used to control the trade-off between the fooling rate, the sparsity, and the complexity. Values close to 1, lead to sparser perturbations, but also to lower fooling rate and increased complexity. On the contrary, higher values of  $\lambda$  lead to fast convergence – even one step solutions –, but the resulted perturbations are less sparse. Since  $\lambda$  is the *only parameter* of the algorithm, it can be easily adjusted to meet the corresponding needs in terms of fooling rate, sparsity, and complexity.

Finally, note that our approach is different from DeepFool’s. DeepFool approximates the decision boundary by linearizing the classifier, while we rather linearize the approximated decision boundary  $\mathcal{B}$ . Furthermore,  $\mathcal{B}$  is the boundary between the adversarial and




Figure 3.3: Fooling rate, sparsity of the perturbations, and average iterations of SparseFool for different values of  $\lambda$ , on 4000 images from ImageNet using an Inception-v3 [153].

the estimated true class, and thus it can be seen as an affine binary classifier. Since at each iteration the adversarial class is computed as the closest (in an  $\ell_2$  sense) to the true one, we can say that SparseFool operates as an untargeted attack. Nevertheless, it can be easily transformed to a targeted one, by simply computing at each iteration the adversarial example – and thus approximating the decision boundary – of a target class.

### 3.3 Experimental evaluation

We evaluate SparseFool on deep convolutional neural network architectures with 10000 images of the MNIST [154] test set, 10000 images of the CIFAR-10 [155] test set, and 4000 randomly selected images from the ImageNet ILSVRC2012 validation set. In order to evaluate our algorithm and compare with related works, we compute the fooling rate, the median perturbation percentage, and the average execution time. Given a dataset  $\mathcal{D}$ , the fooling rate measures the efficiency of the algorithm based on the formula

$$\frac{|\{x \in \mathcal{D} : F(x + \delta_x) \neq F(x)\}|}{|\mathcal{D}|}, \quad (3.7)$$

where  $\delta_x$  is the perturbation of the image  $x$ , while the median perturbation percentage is the median (across all images) percentage of pixels that are perturbed.

We compare SparseFool with JSMA [31]. Since JSMA is a targeted attack, we use its “untargeted” version, where the target is chosen at random. We also make a modification at the success condition; instead of checking if the predicted class is equal to the randomly selected target, we simply check if it is different from the initial class. Note that JSMA is not evaluated on ImageNet, due to its huge computational cost for searching over all pairs of candidates, as also mentioned in [32]. We also compare SparseFool with “One-pixel attack” (1-PA) [33]. Since 1-PA perturbs exactly  $k$  pixels, for every image we start with  $k = 1$  and increase it till 1-PA finds an adversarial example. Again, we do not evaluate 1-PA on ImageNet, due to its computational cost in high dimensional images.

### 3.3 Experimental evaluation

Dataset	Network	Acc. (%)	Fooling rate (%)			Perturbation (%)			Time (sec)		
			SF	JSMA	1-PA	SF	JSMA	1-PA	SF	JSMA	1-PA
MNIST	LeNet [156]	99.14	99.93	95.73	100	1.66	4.85	9.43	0.14	0.66	310.2
CIFAR-10	VGG-19	92.71	100	98.12	100	1.07	2.25	0.15	0.34	6.28	102.7
	ResNet18 [157]	92.74	100	100	100	1.27	3.91	0.2	0.69	8.73	167.4

Table 3.1: The performance of SparseFool (SF), JSMA [31], and “One-pixel attack” (1-PA) [33] on MNIST and CIFAR-10. Due to its high complexity, 1-PA is evaluated on only 100 samples. All the experiments conducted on a GTX TITAN X.

Network	Acc. (%)	Fooling rate (%)	Pert. (%)	Time (sec)
VGG-16	71.59	100	0.18	5.09
ResNet-101	77.37	100	0.23	8.07
DenseNet-161	77.65	100	0.29	10.07
Inception-v3	77.45	100	0.14	4.94

Table 3.2: The performance of SparseFool on the ImageNet dataset, using the pre-trained models provided by PyTorch [158]. All experiments were conducted on a GTX TITAN X.

#### 3.3.1 Performance analysis

We first evaluate the performance of SparseFool, JSMA, and 1-PA on MNIST and CIFAR-10. The control parameter  $\lambda$  in SparseFool was set to 1 and 3 for the MNIST and CIFAR-10 datasets respectively. We observe in Tab. 3.1 that SparseFool computes 2.9x sparser perturbations, and is 4.7x faster compared to JSMA for the MNIST dataset. This behavior remains similar for the CIFAR-10 dataset, where SparseFool computes on average perturbations of 2.4x higher sparsity, and is 15.5x faster. Notice here the difference in the execution time: JSMA becomes much slower as the dimension of the input data increases, while SparseFool’s time complexity remains at very low levels.

Compared to 1-PA, SparseFool computes 5.5x sparser perturbations on MNIST, and is more than 3 orders of magnitude faster. On CIFAR-10, SparseFool still finds very sparse perturbations, but less so than the 1-PA in this case. The reason is that our method does not solve the  $\ell_0$  optimization problem, but it rather computes sparse perturbations through the  $\ell_1$  relaxation. The solution is often sub-optimal, and may be optimal when the image is very close to the boundary, where the linear approximation is more accurate. However, solving the linearized problem is fast and enables our method to efficiently scale to high dimensional data, which is not the case for 1-PA. Considering the tradeoff between the sparsity and complexity, we choose to sacrifice the former. In fact, our method is able to compute sparse perturbations 270x faster than 1-PA.

Finally, due to the computational cost of JSMA and 1-PA, we do not evaluate them on ImageNet. Instead, we compare SparseFool with an algorithm that randomly selects a

subset of elements from each color channel, and replaces their intensity with a random value from  $V = \{0, 255\}$ . The cardinality of each channel subset is constrained to match SparseFool’s per-channel median number of perturbed elements: for each channel, we select as many elements as the median, across all images, of SparseFool’s perturbed elements for this channel. The performance of SparseFool on ImageNet is reported in Tab. 3.2, while the corresponding fooling rates of the random algorithm are 18.2%, 13.2%, 14.5%, and 9.6% respectively. The fooling rates obtained by the random algorithm are far from comparable to SparseFool’s, indicating that our algorithm cleverly finds sparse solutions. Our method is consistent among different architectures, perturbing on average 0.21% of the pixels, with an average execution time of 7 seconds per sample.

To the best of our knowledge, we are the first to provide an adequate sparse attack that efficiently achieves such fooling rates and sparsity, and at the same time scales to high dimensional data. 1-PA does not necessarily find good solutions for all the studied datasets, however, SparseFool – as it relies on the high dimensional geometry of the classifiers – successfully computes sparse enough perturbations for all three datasets.

### 3.3.2 Perceptibility

We now illustrate some adversarial examples generated by SparseFool for three different levels of sparsity: highly sparse perturbations, sparse perturbations, and somewhere in the middle. For MNIST and CIFAR-10 (Fig. 3.4a and Fig. 3.4b respectively), we observe that for highly sparse cases, the perturbation is either imperceptible or can be easily ignored. However, as the number of perturbed pixels increases, the distortion becomes even more perceptible, and in some cases the noise is detectable and far from imperceptible. A similar behavior is also observed for the ImageNet dataset (Fig. 3.5).

To eliminate this perceptibility effect, we focus on the lower/upper bounds of the values of




Figure 3.4: SparseFool adversarial examples for (a) MNIST and (b) CIFAR-10. Each column corresponds to different level of perturbed pixels.



Figure 3.5: SparseFool adversarial examples for ImageNet. Each column corresponds to different level of perturbed pixels. The fooling labels are shown below the images.

the adversarial image  $\hat{\mathbf{x}} = \mathbf{x} + \delta$ . Recall from Sec. 3.2.2 that the bounds  $\mathbf{l}$ ,  $\mathbf{u}$  are defined such that  $l_i \leq \hat{x}_i \leq u_i$ ,  $i = 1 \dots D$ . If these bounds represent the dynamic range of the image, then  $\hat{x}_i$  can take every possible value from this range, and the magnitude of the noise at the element  $i$  can reach visible levels. However, if the perturbed values lie close to the original values  $x_i$ , then we might prevent the magnitude from reaching very high levels. Hence, assuming a dynamic range of  $[0, 255]$ , we explicitly constrain the values of  $\hat{x}_i$  to lie in a small interval  $\pm\alpha$  around  $x_i$ , such that  $0 \leq x_i - \alpha \leq \hat{x}_i \leq x_i + \alpha \leq 255$ .

The sparsity for different values of  $\alpha$  is shown in Fig. 3.6. Higher values give more freedom to the perturbations, but after  $\alpha \approx 25$  the sparsity levels remain almost constant,




Figure 3.6: The resulted sparsity of SparseFool perturbations for  $\pm\alpha$  around the values of  $\mathbf{x}$ , for 100 samples from ImageNet on a ResNet-101 architecture.

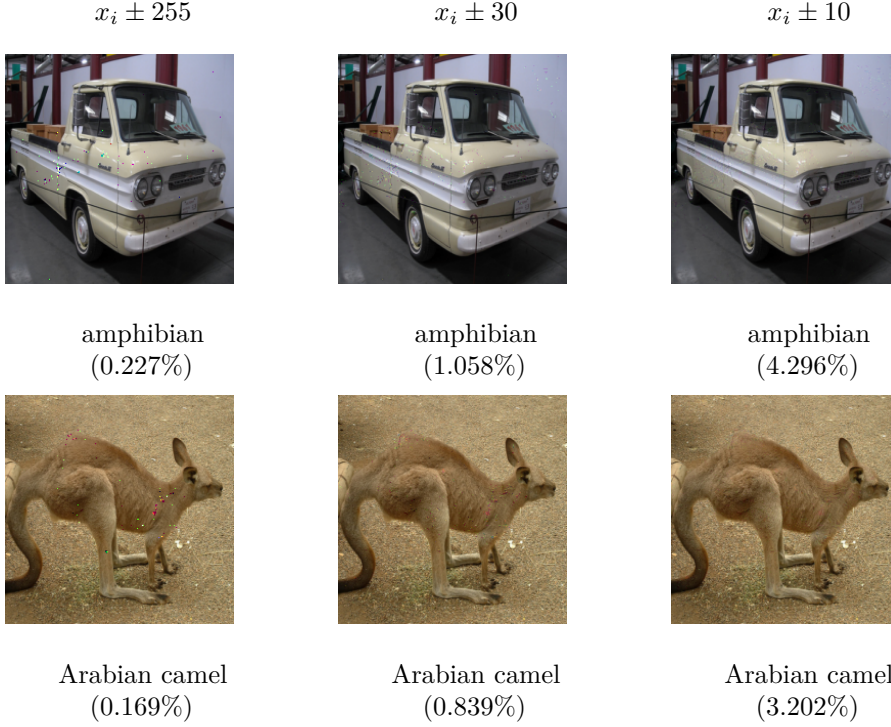


Figure 3.7: The effect of  $\alpha$  on the perceptibility and the sparsity of SparseFool perturbations. The values of  $\alpha$  are shown on top of each column, while the fooling label and the percentage of perturbed pixels are written below each image.

suggesting that we do not need the whole dynamic range. Furthermore, we observed that the average execution time per sample of SparseFool from this value onward remains constant as well, while the fooling rate is 100% regardless  $\alpha$ . Thus, by properly selecting  $\alpha$ , we can control the perceptibility of the perturbations and retain sparsity at a sufficient level. The influence of  $\alpha$  on the perceptibility and sparsity is demonstrated in Fig. 3.7.

## 3.4 Analysis of sparse perturbations

### 3.4.1 Shared semantic features

We now analyze different properties of the perturbations generated with SparseFool. We first investigate if the perturbations transfer across different architectures. For the VGG-16, ResNet-101, and DenseNet-161 architectures, we report in Table 3.3 the fooling rate of each model when fed with adversarial examples generated for another one. We observe that sparse perturbations can generalize only to some extent, and that they are more transferable from larger to smaller architectures. This indicates that there should be some shared semantic information between different architectures that SparseFool exploits, but the exact structure of the perturbations is mostly network dependent.



### 3.4 Analysis of sparse perturbations

	VGG16	ResNet101	DenseNet161
VGG16	100%	10.8%	8.2%
ResNet101	25.3%	100%	12.1%
DenseNet161	28.2%	17.5%	100%

Table 3.3: Fooling rates of SparseFool perturbations between pairs of models for 4000 samples from ImageNet. Row/column denote the source/target model respectively.

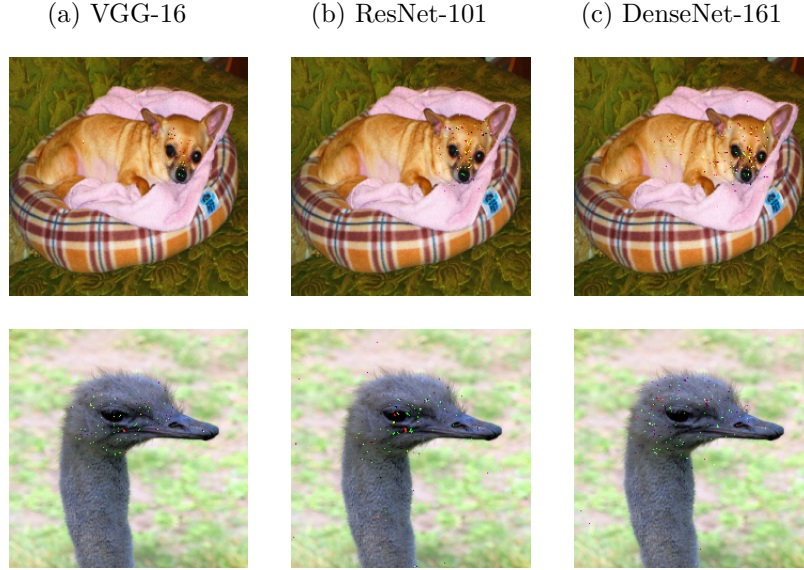


Figure 3.8: Shared information of SparseFool perturbations across three different networks. For all networks, the first row image was classified as “Chihuahua” and misclassified as “French Bulldog”, and second row image as “Ostrich” and “Crane” respectively.

We inspect some animal categories of the ImageNet dataset and examine if semantic information is shared across different architectures (Fig. 3.8). For all networks, the perturbation consistently lies around the important – from a human perspective – areas (i.e., head) of the image, but the way it concentrates or spreads differs for each network. One could say that this is different from dense ( $p \in \{2, \infty\}$ ) perturbations, which exploit more shared input directions and hence have higher transferability. In this sense, we can say that sparse perturbations are identifying directions that align well with human intuition, but at the same time are more unique to each architecture.

For the CIFAR-10 dataset, we observe that in many cases of animal classes, SparseFool tends to perturb some common features around the area of the head (i.e., eyes, ears, nose, mouth etc.), as shown in Fig. 3.9a. Furthermore, we tried to understand if there is a correlation between the perturbed pixels and the fooling class. Interestingly, as shown in Fig. 3.9b, we observe that in many cases the algorithm perturbs those regions of the image that correspond to important features of the fooling class, i.e., when changing a



Figure 3.9: Semantic information of SparseFool perturbations for the CIFAR-10 dataset, on a ResNet-18. Observe that the perturbation is concentrated (a) on some features around the area of the face, or (b) on areas that are important for the fooling class.

“bird” label to a “plane”, where the perturbation seems to resemble some parts of the plane (i.e., wings, tail, nose, turbine). This behavior becomes even more evident when the fooling label is a “deer”, where the noise lies mostly around the area of the head in a way that resembles the antlers.

### 3.4.2 Exclusiveness of adversarial robustness

Finally, we want to understand to what extent sparse perturbations can be “covered” by dense perturbations of larger  $\ell_p$  regimes, when building robust classifiers. To this end, we perform adversarial training on a ResNet-18 on the CIFAR-10 dataset, using  $\ell_\infty$  perturbations of  $\varepsilon = 8/255$  crafted with the PGD attack, as described in [25]. Note that these are the most commonly used settings for building more robust classifiers. The accuracy of this more robust model on the CIFAR-10 test set is 82.17%. Then, we compute the adversarial examples for this model using SparseFool and we measure the fooling rate, median perturbation percentage, and average execution time.

Compared to the results of Tab. 3.1, the fooling rate is still 100%. Regarding the average execution time, it dropped from 0.69 to 0.3 sec, which means that it is “easier” to fool the more robust model compared to the original one. This behaviour can be explained by the fact that adversarially trained classifiers have decision boundaries of very small curvature [107], hence the linear approximation of SparseFool is better.

On the other hand, although the perturbation percentage increased from 1.27% to 2.44%, SparseFool is still able to compute very sparse perturbations compared to  $\ell_\infty$  ones, which perturb almost 100% of the pixels. This result indicates that the effect of  $\ell_1$  and  $\ell_\infty$

perturbations on the robustness of the networks is somehow “mutually exclusive”. This also suggests that the image features exploited by  $\ell_1$  perturbations are different from those of  $\ell_\infty$  ones, and that building classifiers that are robust to a specific  $\ell_p$  regime does not guarantee robustness to other regimes.

### 3.5 Conclusions

In this chapter, we proposed a novel geometry-inspired algorithm to compute sparse adversarial perturbations. In order to avoid the NP-hardness of minimizing the  $\ell_0$  norm to compute sparse perturbations, we focused on finding an efficient relaxation. To this end, we exploited the low mean curvature of the decision boundaries in the vicinity of the data samples and designed an iterative method that we coin SparseFool. At each iteration SparseFool performs a linear approximation of the decision boundary and solves the simpler  $\ell_1$  box-constrained problem to compute sparse adversarial perturbations.

We experimentally demonstrated that SparseFool computes very sparse perturbations, is by orders of magnitude faster than existing methods, and can easily scale to high-dimensional datasets. Furthermore, it incorporates a simple technique to improve the perceptibility of the perturbations, without sacrificing either sparsity or complexity.

By visually inspecting the generated adversarial examples, we observed that SparseFool alters features that are shared among different images, and that in many cases the perturbations resemble image features that are correlated with the fooling class. Finally, we demonstrated that adversarial training with  $\ell_\infty$  perturbations does not build classifiers that are invariant to sparse perturbations, suggesting that the image features exploited by  $\ell_1$  perturbations are different from those of  $\ell_\infty$  ones.

One intriguing observation comes from the visual inspection of the adversarial examples. It actually hints that adversarial perturbations might not necessarily be a “hole” in the system, but they might actually reflect some strong connection/correlation between the features of the dataset and the features that the networks use for taking their decisions. In fact, the authors in [34] showed that adversarial perturbations span a low-dimension but highly discriminative subspace of the input, and that deep networks exploit simple and brittle features of the dataset, i.e., non-robust features that are aligned with adversarial perturbations. However, the potential explanations behind these connections is still an open question. Hence, in the next chapter we investigate this question and provide a novel framework for connecting the image features with the geometry of the decision boundaries and the inductive biases of deep learning.



## 4 Analysis of learned features using adversarial proxies

*“An idea is always a generalization, and generalization is a property of thinking.  
To generalize means to think”*

— Georg Wilhelm Friedrich Hegel

### 4.1 Introduction

The existence of adversarial perturbations implies that the decision boundaries of deep classifiers lie very close to any input sample, and in the previous chapter we demonstrated that the local geometry of the decision boundaries can be used to design a novel sparse attack. This unintuitive behaviour contradicts the common belief that classifiers should be invariant to non-discriminative information of the data. However, using our sparse perturbations we also revealed interesting correlations between the adversarial perturbations and the semantic features of the images. This suggests that adversarial examples might actually be something more than just superficial.

In fact, recent works have established that such perturbations are indeed not irrelevant signals, but rather discriminative features of the training set [34, 35]. This has led to the conjecture that, in most datasets, there exist both robust and non-robust features that neural networks exploit to construct their decision boundaries. Besides, it has also been argued that the excessive invariance in the decision boundaries introduced by

---

Part of this chapter has been published in

“Hold me tight! Influence of discriminative features on deep network boundaries”. In *Neural Information Processing Systems (NeurIPS)*, 2020 [159].

“Redundant features can hurt robustness to distribution shift”. In *Uncertainty & Robustness in Deep Learning Workshop (ICML)*, 2020 [160].

“Improving filling level classification with adversarial training”. In *IEEE International Conference on Image Processing (ICIP)*, 2021 [161].

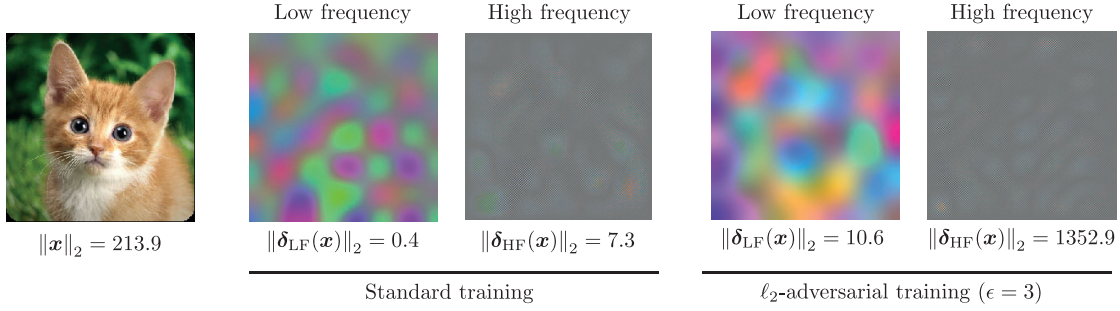


Figure 4.1: Minimal adversarial perturbations constrained in different DCT frequency bands ( $8 \times 8$  subspaces taken from the top left and bottom right of the  $224 \times 224$  DCT matrix) for a ResNet-50 trained (**left**), and adversarially trained (**right**) on ImageNet.

adversarial training can be harmful for standard accuracy, since this invariance causes the classifier to rely on overly-robust features [135]. But, what exactly are these features, and how do networks construct these boundaries? This is still unclear. In this chapter, we propose a novel geometric framework that connects the discriminative features of the training dataset to the norm of the adversarial perturbations. We shed light on these phenomena by describing (i) the strong inductive bias of the networks towards invariance to non-discriminative features, and (ii) the sensitivity of training to small perturbations.

Explaining the mechanisms that construct the decision boundaries is key to understand the dynamics of adversarial training [25], which only differs from standard training in that it slightly perturbs the training samples during optimization. However, these small perturbations can utterly change the geometric properties of these classifiers [107]. An example of such change can be seen in Fig. 4.1, which shows the minimal perturbations – constrained to lie on a low and a high frequency subspace – required to flip the decision of a network. The norm of the perturbations measures the distance (margin) to the decision boundary in these subspaces. Clearly, reaching the boundary using high frequency perturbations requires much more energy (larger  $\ell_2$  norm) than using low frequency ones [3, 86, 162]. But, surprisingly enough, when the network is adversarially trained [25], the largest increase in margin happens in the high frequency subspace. Note that, on the standard network, this distance is already much greater than the size of the training perturbations. Based on this observation, we pose the following questions:

1. *How is the margin in different directions related to the features in the training data?*
2. *How can very small perturbations significantly change the geometry of deep networks?*

In this chapter, we propose a novel approach to answer these questions, and provide a new perspective on the relationship between the distance of a set of samples to the boundary, and the discriminative features used by a network. In particular, we develop a new methodology to construct a local summary of the decision boundary from margin

observations along a sequence of orthogonal directions. This framework permits to carefully tweak the properties of the training samples and measure the induced changes on the boundaries of deep classifiers trained on synthetic and large-scale vision datasets. Via a series of carefully designed experiments, we rigorously confirm the “common belief” that CNNs tend to behave as ideal classifiers and are approximately invariant to non-discriminative features of a dataset. Furthermore, we show that the construction of the decision boundary is very sensitive to the position of the training samples, such that very small perturbations in certain directions can utterly change the decision boundaries in these directions. In fact, we show that adversarial training exploits this training sensitivity and invariance bias to build robust classifiers. Finally, we demonstrate that the invariance properties of robust classifiers can be very beneficial for downstream applications where the available training data are scarce.

The rest of the chapter is organized as follows: In Sec. 4.2 we define our framework and demonstrate on synthetic data how it can be used to associate small margin directions with discriminative features. Then, in Sec. 4.3 we deploy our framework on real datasets, and demonstrate the strong inductive bias of deep networks towards invariance to non-discriminative features. In Sec. 4.4 we use our framework to explain how adversarial training exploits the sensitivity of the network to the position of the training samples, in order to build robust models. Finally, in Sec. 4.5 we demonstrate that the invariances of robust classifiers improve the performance on the off-the-shelf task of estimating the filling level within containers, where the available training data are scarce.

## 4.2 Decision boundary and discriminative features

### 4.2.1 Proposed framework

Recall from Sec. 1.2 that we denote as  $f : \mathbb{R}^D \rightarrow \mathbb{R}^K$  the final layer of a neural network (i.e., logits), such that, for any input  $\mathbf{x} \in \mathbb{R}^D$ ,  $F(\mathbf{x}) = \operatorname{argmax}_k f_k(\mathbf{x})$  represents the decision function of that network, where  $f_k(\mathbf{x})$  denotes the  $k$ th component of  $f(\mathbf{x})$  that corresponds to the  $k$ th class. The decision boundary between classes  $k$  and  $\ell$  is the set  $\mathcal{B}_{k,\ell}(f) = \{\mathbf{x} \in \mathbb{R}^D : f_k(\mathbf{x}) - f_\ell(\mathbf{x}) = 0\}$  (in general, we will omit the dependency with  $k, \ell$  for simplicity). Unless stated otherwise, we assume that all networks are trained using a cross-entropy loss function and some variant of (stochastic) gradient descent. We also assume that training has been conducted for many epochs, and that it has approximately converged to a local minimum of the loss, achieving 100% accuracy on the training data [136]. In general, all our experimental details are listed in App. A.

In this work, we study the role that the training set  $\mathcal{T} = \{(\mathbf{x}^{(i)}, y^{(i)})\}_{i=0}^{N-1}$  has on the boundary  $\mathcal{B}(f)$ . Specifically, we propose to use adversarial proxies to measure the distribution of distances to the decision boundary along a sequence of well defined subspaces. The main quantities of interest are:

**Definition 1** (Subspace-constrained minimal adversarial perturbations). *Based on the definition of Eq. (1.2), for a decision function  $F$ , a sample  $\mathbf{x} \in \mathbb{R}^D$ , and a sub-region of the input space  $\mathcal{S} \subseteq \mathbb{R}^D$ , we define the  $(\ell_2)$  minimal adversarial perturbation of  $\mathbf{x}$  in the subspace  $\mathcal{S}$  as*

$$\begin{aligned} \delta_{\mathcal{S}}(\mathbf{x}) &= \underset{\delta}{\operatorname{argmin}} \|\delta\|_2 \\ \text{s.t. } &F(\mathbf{x}) \neq F(\mathbf{x} + \delta) \\ &\delta \in \mathcal{S}. \end{aligned}$$

In general, we will use  $\delta(\mathbf{x})$  to refer to  $\delta_{\mathbb{R}^D}(\mathbf{x})$ .

**Definition 2** (Margin). *The magnitude  $\|\delta_{\mathcal{S}}(\mathbf{x})\|_2$  is the margin of  $\mathbf{x}$  in  $\mathcal{S}$ .*

Our main objective is to obtain a local summary of  $\mathcal{B}(f)$  around a set of samples, by measuring their margin in a sequence of distinct subspaces  $\{\mathcal{S}_j\}_{j=0}^{R-1}$ . In practice, we use a subspace-constrained version of DeepFool [11]<sup>1</sup> to approximate the margins in each  $\mathcal{S}_j$ .

DeepFool is regarded as one of the most efficient methods to identify minimal adversarial perturbations. Since we want to measure margin, norm-constrained attacks like PGD [25] are not suitable for our study. Besides, more complex attacks like C&W [32], or, even, using unconstrained gradient descent in the input space, are computationally more demanding and harder to tune than DeepFool. Since they in general find very similar adversarial perturbations as DeepFool, we decided to opt for DeepFool in our works.

### 4.2.2 Evidence on synthetic examples

In general, the distance from a sample to the boundary of a neural network can greatly vary depending on the search direction [18]. This behaviour is typically translated into classifiers with small margins along some directions, and large margins along others. We now investigate if neural networks only construct boundaries along discriminative directions, and remain invariant in every other direction<sup>2</sup>.

To this end, we generate a balanced training set  $\mathcal{T}_1(\varepsilon, \sigma)$  by independently sampling  $N$  points  $\mathbf{x}^{(i)} = \mathbf{U}(\mathbf{x}_1^{(i)} \oplus \mathbf{x}_2^{(i)})$ , with  $\mathbf{x}_1^{(i)} = \varepsilon \mathbf{y}^{(i)}$  and  $\mathbf{x}_2^{(i)} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_{D-1})$ , where  $\oplus$  is the concatenation operator,  $\varepsilon > 0$  the feature size, and  $D = 100$ . The labels  $y^{(i)}$  are uniformly sampled from  $\{-1, +1\}$ . The multiplication by a random orthonormal matrix  $\mathbf{U} \in \text{SO}(D)$  is performed to avoid possible biases of the classifier towards the canonical basis.  $\mathcal{T}_1$  is a linearly separable dataset with a single discriminative feature parallel to  $\mathbf{u}_1$  (i.e., first row of  $\mathbf{U}$ ), and all other dimensions filled with non-discriminative noise.

<sup>1</sup>We do not enforce the  $[0, 1]^D$  box constraints on the adversarial images, as we are not interested in finding “plausible” adversarial perturbations, but in measuring the distance to  $\mathcal{B}(f)$ .

<sup>2</sup>This is indeed a desired property for any classification method, but note that for neural networks the existence of adversarial examples contests the idea of it being a reasonable assumption.



### 4.3 Discriminative features of real datasets

	$\mathbf{u}_1$	$\text{span}\{\mathbf{u}_1\}^\perp$	$\mathcal{S}_{\text{orth}}$	$\mathcal{S}_{\text{rand}}$
5-perc.	1.74	4.85	30.68	17.21
Median	2.50	12.36	102.0	27.90
95-perc.	3.22	31.60	229.5	80.61

Table 4.1: Margin statistics of an MLP trained on  $\mathcal{T}_1(\varepsilon = 5, \sigma = 1)$  along different directions ( $N = 10,000$ ,  $M = 1,000$ ,  $S = 3$ ).

To evaluate our hypothesis, we train an overparameterized multi-layer perceptron (MLP) with 10 hidden layers of 500 neurons using SGD (test accuracy: 100%). Table 4.1 shows the margin statistics on the linearly separable direction  $\mathbf{u}_1$ ; its orthogonal complement  $\text{span}\{\mathbf{u}_1\}^\perp$ ; a fixed random subspace of dimension  $S$ ,  $\mathcal{S}_{\text{rand}} \subset \mathbb{R}^D$ ; and a fixed random subspace of the same dimensionality, but orthogonal to  $\mathbf{u}_1$ ,  $\mathcal{S}_{\text{orth}} \subset \text{span}\{\mathbf{u}_1\}^\perp$ . From these values we can see that along the direction where the discriminative feature lies, the margin is much smaller than in any other direction. Therefore, we can see that the classification function of this network is only creating a boundary in  $\mathbf{u}_1$  with median margin  $\varepsilon/2$ , and that it is approximately invariant in  $\text{span}\{\mathbf{u}_1\}^\perp$ .

Comparing the margin values for  $\mathcal{S}_{\text{orth}}$  and  $\mathcal{S}_{\text{rand}}$  we see that, if the observation basis is not aligned with the features exploited by the network, the margin measurements might not be able to separate the small and large margin directions. Indeed, since  $\mathcal{S}_{\text{orth}}$  is orthogonal to the only discriminative direction  $\mathbf{u}_1$ , we see that the margin values reported in this region are much higher than those reported in  $\mathcal{S}_{\text{rand}}$ . The reason for this is that the margin required to flip the label of a classifier in a randomly selected subspace is of the order of  $\sqrt{S/D}$  with high probability [42], and hence the non-trivial correlation of a random subspace with the discriminative features will always hide the differences between small and large margin directions.

Finally, the margin fluctuations and the fact that the classifier is not completely invariant to  $\text{span}\{\mathbf{u}_1\}^\perp$  might indicate that the network has built a complex boundary. However, in App. A.1 we show that similar fluctuations and finite values in  $\text{span}\{\mathbf{u}_1\}^\perp$  are observed even if the model is linear by construction and separates the training data perfectly.

### 4.3 Discriminative features of real datasets

In contrast to the synthetic data, where the discriminative features are known by construction, the exact description of the features presented in *real* datasets is usually not known. In order to identify these features and understand their connection to the local construction of the decision boundaries, we apply the proposed framework on standard computer vision datasets, and investigate if deep networks trained on real data also present high invariance along the non-discriminative directions of the dataset.

In our study, we train multiple networks on MNIST [154] and CIFAR-10 [155], and for ImageNet [152] we use several of the pretrained networks provided by PyTorch [158]<sup>3</sup>. Let  $W, H, C$  denote the width, height, and number of channels of the images in those datasets, respectively. In our experiments we use the 2-dimensional discrete cosine transform (2D-DCT) [163] basis of size  $H \times W$  to generate the observation subspaces. In particular, let  $\mathcal{D} \in \mathbb{R}^{H \times W \times H \times W}$  denote the 2D-DCT generating tensor, such that  $\text{vec}(\mathcal{D}(i, j, :, :) \otimes \mathbf{I}_C)$  represents one basis element of the image space. We generate the subspaces by sampling (see Fig. 4.2)  $K \times K$  blocks from the diagonal of the DCT tensor using a sliding window with step-size  $T$ :  $\mathcal{S}_j = \text{span}\{\text{vec}(\mathcal{D}(j \cdot T + k, j \cdot T + k, :, :) \otimes \mathbf{I}_C) \mid k = 0, \dots, K - 1\}$ .




Figure 4.2: Illustration of sampling from the diagonal of the DCT basis.

The sliding window on the diagonal of the DCT gives a good trade-off between visualization abilities in simple one-dimensional plots, and a diverse sampling of the spatial spectrum of natural images, with a well-defined gradient flowing from low to high frequencies<sup>4</sup>. The DCT has a long application tradition in image processing due to its good approximation of the decorrelation transform (KLT) [164]. Furthermore, in previous studies on the robustness of deep networks to different frequencies, the DCT was also the basis of choice [86] because it avoids dealing with complex subspaces.

The margin distribution of the evaluated test samples is presented in the top of Fig. 4.3. For MNIST and ImageNet, the networks present a strong invariance along high frequency directions and small margin along low frequency ones. We will later show that this is related to the fact that these networks mainly exploit discriminative features in the low frequencies of these datasets. Notice, however, that for CIFAR-10 the margin values are more uniformly distributed; an indication that the network exploits discriminative features across the full spectrum as opposed to the human vision system [165].

We observe in practice that the DCT basis is also quite aligned to the features of these

<sup>3</sup>Experiments on more CNNs (with similar findings) are presented in Appendix I of [159].

<sup>4</sup>A similar analysis including off-diagonal subspaces is presented in Appendix I of [159].

### 4.3 Discriminative features of real datasets

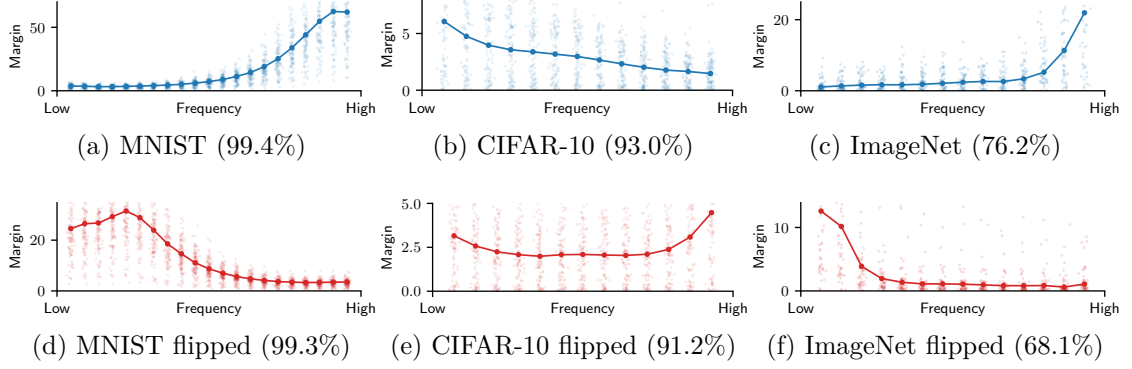


Figure 4.3: Margin distribution of test samples in subspaces taken from the diagonal of the DCT (low to high frequencies). Thick lines indicate the median values of the margin, and shaded points represent its distribution. Parentheses contain the test accuracy. **Top:** (a) MNIST (LeNet), (b) CIFAR-10 (DenseNet-121) and (c) ImageNet (ResNet-50) **Bottom:** (d) MNIST, (e) CIFAR-10 and (f) ImageNet trained on frequency-“flipped” versions of the standard datasets.

datasets, and hence it can give precise information about the discriminative features exploited by the networks. A more aligned basis with respect to the discriminative features would probably show a sharper transition between low and high margins. However, finding such network-agnostic bases is a challenging task without knowing the features *a priori*. The DCT is not perfectly feature-aligned, but it seems to be a good choice for comparing different architectures, especially if we compare its results to those obtained using a random orthonormal basis where differences in margin cannot be identified. Indeed, as shown in Fig. 4.4 it is clear that a random basis is not valid for this task as the margin in any random subspace is of the same order with high probability [42].

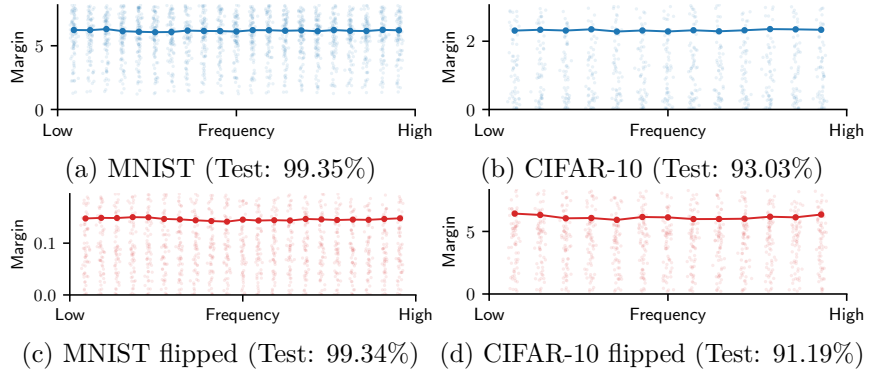


Figure 4.4: Margin distribution of test samples in subspaces taken from a random orthonormal matrix arranged as a tensor of the same dimensionality as the DCT tensor. Subspaces are taken from the diagonal with the same parameters as in DCT. **Top:** (a) MNIST (LeNet), (b) CIFAR-10 (DenseNet-121) **Bottom:** (d) MNIST (LeNet) and (e) CIFAR-10 (DenseNet-121) trained on frequency “flipped” versions of the datasets.

### 4.3.1 Boundary adaptation to data representation

Towards verifying that the proposed framework can capture the relation between the data features and the local construction of the decision boundaries, we must first ensure that the direction of the observed invariance (large margin) is related to the features presented in the dataset, rather than being just an effect of the network itself.

Based on our observation that the margin tends to be small in low frequency directions and large in high frequency ones, we carefully tweak the representation of the data such that the low frequencies are swapped with the high frequencies. In practice, if  $\mathfrak{D}$  denotes the forward DCT transform operator, the new image representation  $\mathbf{x}'$  is expressed as  $\mathbf{x}' = \mathfrak{D}^{-1}(\text{flip}(\mathfrak{D}(\mathbf{x})))$ , where flip corresponds to one horizontal and one vertical flip of the DCT transformed image. Some “flipped” examples are shown in Fig. 4.5.




Figure 4.5: “Flipped” image examples from ImageNet. **Top:** original. **Bottom:** “flipped”.

Thus, if the direction of the resulting margin is strongly related to the data features, the constructed decision boundaries should also adapt to this new data representation, and the margin along the invariant directions (high frequencies) should swap with the margin of the discriminative ones (low frequencies). Informally speaking, the margin distribution should “flip”.

We apply our framework on multiple networks trained on the “flipped” datasets, and the margin distribution is depicted at the bottom of Fig. 4.3. For both MNIST and ImageNet, the directions of the decision boundaries indeed *follow* the new data representation – although they are not an exact mirroring of the original representation. This indicates that the margin strongly depends on the data distribution, and it is not solely an effect of the network architecture. Note again that for CIFAR-10 the effect is not as obvious, due to the quite uniform distribution of the margin.

### 4.3.2 Invariance and elasticity of decision boundary

The second property we need to verify is that the small margins reported in Fig. 4.3 do indeed correspond to directions containing discriminative features in the training set.

### 4.3 Discriminative features of real datasets

For doing so, we use the insights of Fig. 4.3b on CIFAR-10 – where, opposed to the other datasets, we assume that there are exploited discriminative features in the whole spectrum – and show that, by explicitly modifying its features, we can induce a high margin response in the measured curve in a set of selected directions.

In particular, we create a low-pass filtered version of CIFAR-10 ( $\mathcal{T}_{LP}$ ), where we retain only the frequency components in a  $16 \times 16$  square at the top left of the diagonal of the DCT-transformed images. This way we ensure that no training image has any energy/information outside of this frequency subspace. Examples of  $\mathcal{T}_{LP}$  images are shown in the second row of Fig. 4.6.




Figure 4.6: Examples of filtered CIFAR-10 images. **Top** original images, **middle** low-pass and **bottom** high-pass versions.

The median margin<sup>5</sup> of CIFAR-10 test samples for a network trained on  $\mathcal{T}_{LP}$  is illustrated in Fig. 4.7. Indeed, by eliminating the high frequency content, we have forced the network to become invariant along these directions. This clearly demonstrates that there existed discriminative features in the high frequency spectrum of CIFAR-10, and that by removing these from all the samples, the inductive bias of training pushes the network to become invariant to them.




Figure 4.7: Median margin of test samples from CIFAR-10 for a DenseNet-121 (i) trained on CIFAR-10 and fine-tuned on  $\mathcal{T}_{LP}$  (90.79%), and (ii) trained on  $\mathcal{T}_{LP}$  from scratch (89.67%). Parentheses contain the test accuracy.

<sup>5</sup>We do not plot the full distribution to avoid clutter. The 5-percentile of the margin in the last subspace is 5.05.




Figure 4.8: Median margin of  $\mathcal{T}_{LP}$  test samples for a DenseNet-121. **Blue:** trained on  $\mathcal{T}_{LP \cup HP}$ ; **Red:** after forgetting  $\mathcal{T}_{HP}$ ; **Green:** after recovering  $\mathcal{T}_{HP}$ .

Moreover, this effect can *also* be triggered during training. To show this, we start with the CIFAR-10 trained network studied in Fig. 4.3b and continue training it for a few more epochs with a small learning rate using only  $\mathcal{T}_{LP}$ . Figure 4.7 shows the new median margins of this network. The fine-tuned network is again invariant to the high frequencies.

Finally, note that by training with only low frequency data, the test accuracy of the network on the original CIFAR-10 only drops around 3%<sup>6</sup>. Because  $\mathcal{T}_{LP}$  has no high frequency energy, a network trained on it will uniformly extend its boundaries in this part of the spectrum and no high frequency perturbation will be able to flip the network’s output. In contrast, testing  $\mathcal{T}_{LP}$  data on a CIFAR-10 trained network only achieves 27.45% test accuracy. This is because networks trained on CIFAR-10 do have boundaries in the high frequencies, and hence showing them original samples perturbed in this frequency range (i.e.,  $\mathcal{T}_{LP}$ ) can greatly change their decisions.

### 4.3.3 Connections to catastrophic forgetting

The elasticity to the modification of features during training gives a new perspective to the theory of catastrophic forgetting [166], as it confirms that the decision boundaries of a neural network can only exist for as long as the classifier is trained with the samples (features) that hold them together. In particular, we demonstrate this by adding and removing points from a dataset such that its discriminative features are modified during training, and hence artificially causing an elastic response on the network.

To this end, we train a DenseNet-121 on a new dataset  $\mathcal{T}_{LP \cup HP} = \mathcal{T}_{LP} \cup \mathcal{T}_{HP}$  formed by the union of two filtered variants of CIFAR-10:  $\mathcal{T}_{LP}$  is constructed by retaining only the frequency components in a  $16 \times 16$  square at the top-left of the DCT-transformed CIFAR-10 images (low-pass), while for  $\mathcal{T}_{HP}$  only the frequency components in a  $16 \times 16$  square at the bottom-right of the DCT (high-pass; cf. third row of Fig. 4.6). This classifier has a test accuracy of 86.59% and 57.29% on  $\mathcal{T}_{LP}$  and  $\mathcal{T}_{HP}$ , respectively. The

<sup>6</sup>Similar effect was shown on ImageNet [3], although the network was only tested on filtered data. For MNIST, training on low-pass data yields no accuracy drop, since MNIST trained networks exploit mostly low frequencies and already have large margins in the high frequencies. There might exist discriminative information in the high frequencies, but the network does not exploit it.

median margin of 1,000  $\mathcal{T}_{LP}$  test samples along different frequencies for this classifier is shown in blue in Fig. 4.8. As expected, the classifier has picked features across the whole spectrum with the low frequency ones probably belonging to boundaries separating samples in  $\mathcal{T}_{LP}$ , and the high frequency ones separating samples from  $\mathcal{T}_{LP}$  and  $\mathcal{T}_{HP}$ <sup>7</sup>.

After this, we continue training the network with a linearly decaying learning rate (max.  $\alpha = 0.05$ ) for another 30 epochs, but using only  $\mathcal{T}_{LP}$ , achieving a final test accuracy of 87.81% and 10.01% on  $\mathcal{T}_{LP}$  and  $\mathcal{T}_{HP}$ , respectively. Again, Fig. 4.8 shows in red the median margin along different frequencies on test samples from  $\mathcal{T}_{LP}$ . The new median margin is clearly invariant on the high frequencies – where  $\mathcal{T}_{LP}$  has no discriminative features – and the classifier has completely *erased* the boundaries that it previously had in these regions, regardless of the fact that those boundaries did not harm the classification accuracy on  $\mathcal{T}_{LP}$ .

Finally, we investigate if the network is able to recover the forgotten decision boundaries that were used to classify  $\mathcal{T}_{HP}$ . We continue training the network (“forgotten”  $\mathcal{T}_{HP}$ ) for another 30 epochs, but this time by using the whole  $\mathcal{T}_{LP \cup HP}$ . Now this classifier achieves a final test accuracy of 86.1% and 59.11% on  $\mathcal{T}_{LP}$  and  $\mathcal{T}_{HP}$  respectively, which are very close to the corresponding accuracies of the initial network trained from scratch on  $\mathcal{T}_{LP \cup HP}$  (recall: 86.59% and 57.29%). The new median margin for this classifier is shown in green in Fig. 4.8. As we can see by comparing the green to the blue curve, the decision boundaries along the high-frequency directions can be recovered quite successfully.

#### 4.3.4 Discussion

The main claim in Sec. 4.2 and Sec. 4.3 is that deep neural networks *only* create decision boundaries in regions where they identify discriminative features in the training data. As a result, there is a big relative difference in the large margin along the invariant directions, and the smaller margin in the discriminative directions.

The main difficulty for establishing causation in this idea is the fact that the discriminative features of real datasets are not known. Hence, determining their role on the geometry of a trained neural network can only be done by artificially manipulating the data. In particular, there are two main confounding factors that might alternatively explain our results: the network architecture or the training algorithm. However, the experiments in Sec. 4.3 are precisely designed to rule out their influence in this phenomenon.

Specifically, in the flipping experiments, flipping the data – *ceteris paribus* – also flips the margin distribution, thus demonstrating that the margins are necessarily caused by the information present in the data. The other interventions we do on the samples (e.g., low-pass experiments) confirm that, in the absence of information in a certain discriminative

---

<sup>7</sup> $\mathcal{T}_{LP}$  and  $\mathcal{T}_{HP}$  have discriminative features only in the low and high frequency part, respectively.

subspace, the network becomes invariant along this discriminative subspace. Therefore, we believe that there is indeed a causal connection between the features of the data and the measured margins in these neural networks. In fact, parallel theoretical studies have demonstrated that the ability of neural networks to distinguish between discriminative and non-discriminative noise subspaces in a dataset is one of the main advantages of deep learning over kernel methods [167].

## 4.4 Sensitivity to position of training samples

Our novel framework to relate boundary geometry and data features can help track the dynamics of learning. In this section, we use it to explain how training with a slightly perturbed version of the training samples can greatly alter the network geometry. We further analyze how adversarial training can be so successful in removing features with small margin to increase the network’s robustness.

### 4.4.1 Evidence on synthetic examples

We train multiple times an MLP with the same setup as Sec. 4.2.2, but this time using slightly perturbed versions of the same synthetic dataset. In particular, we use a family of training sets  $\mathcal{T}_2(\rho, \varepsilon, \sigma, K)$  consisting in  $N = 10,000$  independent  $D = 100$ -dimensional samples  $\mathbf{x}^{(i)} = \mathbf{U}(\mathbf{x}_1^{(i)} \oplus \mathbf{x}_2^{(i)} \oplus \mathbf{x}_3^{(i)})$  such that  $\mathbf{x}_1^{(i)} = \varepsilon y^{(i)}$ ;  $\mathbf{x}_2^{(i)} = \rho \cdot \mathbf{k}$  when  $y^{(i)} = +1$ , and  $\mathbf{x}_2^{(i)} = \rho \cdot (\mathbf{k} + \frac{1}{2})$  when  $y^{(i)} = -1$ , where  $\mathbf{k}$  is sampled from a discrete uniform distribution with values  $\{-K, \dots, K-1\}$ ; and  $\mathbf{x}_3^{(i)} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_{D-2})$  (see Fig. 4.9). Here,  $\varepsilon, \rho \geq 0$  denote the feature sizes. Again, the multiplication by a random orthonormal matrix  $\mathbf{U} \in \text{SO}(D)$  avoids any possible bias of the network towards the canonical basis. Note that for  $\varepsilon > 0$  this training set will always be linearly separable using  $\mathbf{u}_1$ , but without necessarily yielding a maximum margin classifier. Especially when  $\rho \gg \varepsilon$ .




Figure 4.9: Cross-section of an MLP trained on  $\mathcal{T}_2(\rho = 20, \varepsilon, \sigma = 1, K = 3)$  with  $\varepsilon = 1$  (**top**) and  $\varepsilon = 0$  (**bottom**). Axes scaled differently.

Figure 4.10 shows the median margin of  $M = 1,000$  observation samples for an MLP trained on different versions of  $\mathcal{T}_2(\rho, \varepsilon, \sigma, K)$  with a fixed  $\rho = 20$ , but a varying small  $\varepsilon$ . Based on this plot, it is clear that for very small  $\varepsilon$  the neural network predominantly uses the information contained in  $\mathbf{u}_2$  to separate the different classes. Indeed, for  $\varepsilon < 0.2$ , the



#### 4.4 Sensitivity to position of training samples




Figure 4.10: Median margin values along  $\mathbf{u}_1$  and  $\mathbf{u}_2$  for MLPs (test: 100% always) trained on  $\mathcal{T}_2$  for different values of  $\epsilon$  and  $\rho = 20$ .

network is almost invariant in  $\mathbf{u}_1$ , and it uses a non-linear alternating pattern in  $\mathbf{u}_2$  to separate the data<sup>8</sup> (see bottom row of Fig. 4.9). On the contrary, at  $\epsilon > 0.5$  we notice a sharp transition in which we see that the neural network suddenly changes its behaviour and starts to linearly separate the different points using only  $\mathbf{u}_1$  (see top row of Fig. 4.9).

We conjecture that this phenomenon is rooted on the strong inductive bias of the learning algorithm to build connected decision regions whenever geometrically and topologically possible, as empirically validated in [115]. Here, we go one step further and hypothesize that the inductive bias of the learning algorithm has a tendency to build classifiers in which every pair of training samples with the same label belongs to the same decision region. If possible, connected by a straight path.

We see Fig. 4.10 as a validation of this hypothesis. For small values of  $\epsilon$ , it is hard for the algorithm to find solutions that connect points from the same class with a straight path, as this is very aligned with  $\mathbf{u}_2$ . However, there is a precise moment (i.e.,  $\epsilon = 0.5$ ) in which finding such a solution becomes much easier, and then the algorithm suddenly starts to converge to the linearly separating solution.

At this stage it is important to highlight that repeating the same experiment with a different random seed, or for a fixed initialization, does not affect the results. Furthermore, overfitting cannot be the cause of these results, as the MLP always achieves 100% test accuracy for  $\epsilon < 0.5$ , as well. Finally, adding a small weight decay (i.e.,  $10^{-3}$ ) does not help the network find the linearly separable solution for  $\epsilon < 0.5$ ; it rather hinders its convergence (i.e., final train accuracy is 50%).

It remains unclear whether this inductive bias is the only mechanism that can trigger a sharp transition in the type of learned decision boundaries, or if there are other types of biases that can cause the same effect. In any case, we believe that the significant difference in the type of function that the algorithm learns when trained with very similar training samples (see Fig. 4.9), is an unambiguous confirmation of the sensitivity of deep learning to the exact position of its training input.

<sup>8</sup>This particular pattern can in principle classify any dataset with  $\rho = 20$ , no matter the value of  $\epsilon$ .




Figure 4.11: Margin distribution of test samples in subspaces taken from the diagonal of the DCT (low to high frequencies). Adversarially trained networks using  $\ell_2$  PGD (a) LeNet (Adv: 76%), (b) DenseNet-121 (Adv: 55%) and (c) ResNet-50 (Adv: 35%).

Concurrent work [168] has also used a similarly constructed dataset to  $\mathcal{T}_2(\rho, \varepsilon, \sigma, K)$  to argue that the simplicity bias of a neural network when trained using standard procedures might be responsible for the selection of non-robust features in the dataset [35].

#### 4.4.2 Connections to adversarial training

Finally, we show that adversarial training exploits the type of phenomena described in Sec. 4.4.1 to reshape the boundaries of a neural network. In this regard, Fig. 4.11 shows the margin distribution across the DCT spectrum of a few adversarially trained networks<sup>9</sup>. As expected, the margins of the adversarially trained networks are significantly higher than those in Fig. 4.3.

Surprisingly, though, the largest increase can be noticed in the high frequencies for all datasets. Considering that adversarial training only differs from standard training in that it slightly moves the training samples, it is imperative that deep networks converge to very different solutions under such small modifications. The next experiments on CIFAR-10 shed light on the dynamics of this process.

**Very small adversarial perturbations can trigger large invariance** Slightly perturbing the training samples can remove features in an unpredictable manner. Figure 4.12 shows the spectral decomposition of the adversarial perturbations crafted during adversarial training of CIFAR-10. The energy of the perturbations during training is always concentrated in the low frequencies, and has hardly any high frequency content. However, the greatest effect on margin is seen on the high frequency directions (see Fig. 4.11). This is similar to what is seen in Fig. 4.9, where slightly perturbing the training samples along  $\mathbf{u}_2$  drastically affects the margin along  $\mathbf{u}_2$ .

Overall, we see that adversarial training exploits the sensitivity of the network to small changes in the training samples to hide some discriminative features from the model. This is especially clear when we compare the CIFAR-10 values in Fig. 4.11b and Fig. 4.3b, where it becomes evident that some previously used discriminative features in the high

<sup>9</sup>The analogous effect for the frequency-“flipped” datasets is detailed in Appendix M of [159]




Figure 4.12: Energy of adversarial perturbations on subspaces of the DCT during adversarial training of CIFAR-10 (DenseNet-121). Plot shows 95-percentile.

frequencies are completely overlooked by the adversarially trained network. In the following example, we show that, in practice, it is not actually necessary to change the position of all training points to induce a large invariance reaction.

**Invariance can be triggered by just a few samples** Modifying the position of just a *minimal* number of training samples is enough to locally introduce excessive invariance on a classifier. To demonstrate this, we take a ResNet-18 (test: 90%) trained on CIFAR-10, and randomly select a set of 100 training samples  $\mathcal{P} \subset \mathcal{T}$ . We fine-tune this classifier replacing those 100 samples with  $(\mathbf{x} + \delta^o(\mathbf{x}), y)$  in  $\mathcal{P}$  (test: 90%), where  $\delta^o$  and  $\delta^f$  represent the adversarial perturbations for the original and fine-tuned network, respectively.




Figure 4.13: Margin distribution in different directions of a ResNet-18 trained on CIFAR-10 and fine-tuned on 100 DeepFool examples.

Figure 4.13 shows the magnitude of these perturbations both for the 100 adversarially perturbed points  $\mathcal{P} \subset \mathcal{T}$  and for a subset of 1,000 unmodified samples  $\mathcal{U} \subset \mathcal{T}$ . Here, we can clearly see that, after fine-tuning, the boundaries around  $\mathcal{P}$  have been completely modified, showing a large increase in the distance to the boundary in the direction of the original adversarial perturbation  $\Delta_o^f(\mathbf{x})$  for  $(\mathbf{x}, y) \in \mathcal{P}$ . Meanwhile, the boundaries around  $\mathcal{U}$  have not seen such a dramatic change.

This means that modifying the position of *only* a small fraction of the training samples can induce a large change in the shape of the boundary. Note that this dependency on a few samples resembles the one of support vector machines [169], whose decision boundaries are defined by the position of a few supporting vectors. However, in contrast to SVMs, deep neural networks are not guaranteed to maximize margin in the input space (see Fig. 4.10), and the points that support their boundaries need not be the ones closest to them, hence rendering their identification much harder.

### 4.5 Implications in data-scarce applications

Interestingly, the invariances obtained by robust classifiers that we analyzed in the previous chapters, have been recently shown to be quite beneficial when performing transfer learning [170, 171]. In this section, we will demonstrate that the robust properties of deep networks can be exploited to improve the generalization performance in the off-the-shelf task of estimating the filling level within a container, where the available training data are typically very scarce.

Consider a real-world example of collaborative interactions between humans and robots. In such scenario, estimating through vision the physical properties of objects manipulated by humans is important for performing *accurate* and *safe* grasps of objects handed over by humans [172]. For achieving successful grasps, one important property that should be estimated is the weight of the object based on the shape of the container [173], the type of content inside the container, and the amount of that content.

In particular, estimating the amount of content (filling level) within an *unknown* container is a quite challenging problem due to distribution shifts occurred at test time related to, e.g., (i) occlusions caused by the hand holding the container, (ii) the transparencies of both the container and the filling (e.g., depth estimation may be highly inaccurate for transparent objects [174]), and (iii) by the differences in the shape of the containers. Typically, the few approaches designed to tackle this problem use RGB [175], thermal [176], or a combination of RGB and depth data [177, 178], and usually observe the action of pouring content in a container over multiple frames [176, 177, 178, 179].

By approaching the problem as a classification task, the authors in [175] showed that transfer learning [180] was the best-performing strategy: self-collected data were used as task-specific dataset, the *target domain*, to fine-tune the parameters of a CNN pre-trained on the much larger ImageNet dataset [152], the *source domain*. In the context of filling level estimation, the available training data are usually rather scarce, hence transfer learning introduces useful knowledge when fine-tuning on the target domain.

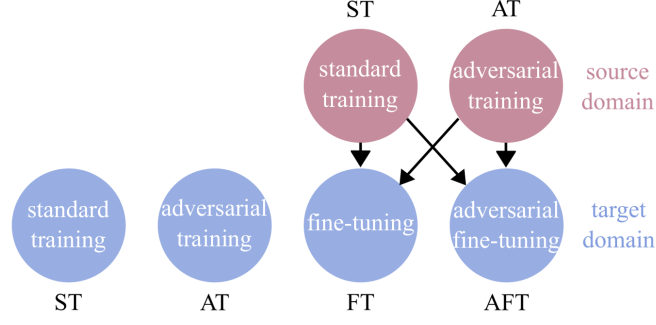


Figure 4.14: The six training strategies analyzed in our experiments: independent standard training (ST) and adversarial training (AT) on the target domain, and four transfer learning strategies from source to target domain via fine-tuning (FT).

#### 4.5.1 An off-the-shelf task: filling level classification

##### Task description and training strategies

We approach the problem of estimating the filling level,  $y$ , of a container captured in an image  $\mathbf{x}$ , as a classification task. We express the filling level as a percentage of the container’s capacity:  $y \in \{0\%, 50\%, 90\%, \text{unknown}\}$ , where the *unknown* class helps handling cases with opaque or translucent containers for which the filling level cannot be estimated through direct vision. Given a train set of image-label pairs  $\mathcal{T} = \{(\mathbf{x}^i, y^i)\}_{i=1}^N$ , the goal is to find a classifier that minimizes a suitable loss function  $\mathcal{L}_\theta(\mathbf{x}, y)$  such that  $F$  correctly predicts  $y$  for  $\mathbf{x} \sim \mathcal{D}$  but  $\mathbf{x} \notin \mathcal{T}$  (generalization).

We refer to the common strategy for training a classifier on a train set,  $\mathcal{T}$ , as Standard Training (ST). A good generalization may be achieved if the number of image-label pairs in  $\mathcal{T}$  is very large, e.g.,  $N \approx 1.2$  millions in ImageNet. However, for the target task of classifying the filling level such amount of data is not available. Transfer learning helps to overcome this limitation by using an additional training set  $\mathcal{S}$ , with  $|\mathcal{S}| = M \gg N$ , that may not be related to the target task. Transfer learning pre-trains the parameters of  $f$  on  $\mathcal{S}$  (source domain) and then refines them on  $\mathcal{T}$  (target domain) via fine-tuning (FT). We refer to this strategy as ST→FT. With ST→FT, the parameters of some layers in the pre-trained model are fixed and FT only refines those of the remaining layers. We will denote with  $L$  the number of layers whose parameters are fixed.

Recall that with Adversarial Training (AT), the resulting models learn features that correlate better with features of the classes of interest [27, 119, 120, 121]. Hence,  $f$  is expected to learn more task-relevant features with AT. We aim to evaluate AT on the filling-level classification task, and to compare it against five other strategies. As training strategies we consider ST→FT [175]; ST on the target domain; AT on the target domain; and three combinations of AT with transfer learning, namely AT on the *source* domain (AT→FT), AT on the *target* domain (ST→AFT), and AT on *both* domains (AT→AFT).

Similarly to what was observed in [170, 171], we expect that the performance of fine-tuning on  $\mathcal{T}$  will further improve if we use a model trained on  $\mathcal{S}$  with AT instead of a model trained with ST, even if the classification performance of the robust model on  $\mathcal{S}$  is worse than the performance of the model trained with ST. The exact reason behind this improvement is still an open question, but it is related to the differences in the learned features between standard and robust models. Also, this improvement depends on the value of  $\varepsilon$  used during AT, and the value that leads to better accuracy may differ across tasks and domains. Smaller values for  $\varepsilon$  generally lead to better performance [171], but its value will be selected empirically.

Finally, the last two training combinations apply AT either on the *target* domain (ST→AFT) via FT or on *both* domains (AT→AFT). Considering the effect of AT on the features learned by a classifier, we will investigate how  $f$  is affected when the transferred learned features from  $\mathcal{S}$  are further filtered by AT on  $\mathcal{T}$ . Fig. 4.14 summarizes the training strategies under analysis, which will be compared in the next section.

### A novel dataset

Since the task is to classify the filling level from a single RGB image, the CCM dataset [181] is a suitable choice due to its large variability in terms of capturing conditions. CCM comprises of four views capturing under different backgrounds and illumination conditions cups and drinking glasses. The containers are transparent, translucent or opaque. The content is transparent (water) or opaque (pasta, rice). Each container stands upright on a surface or is being manipulated by a person. We only consider data of the public CCM repository, namely 4 cups and 4 drinking glasses.

From the CCM video data, we automatically sampled and then visually verified 10,269 frames of containers for which a pouring action was completed. To increase the variability in the sampled data, we selected frames considering that the container is completely visible or occluded by the person’s hand, and under different backgrounds.

For each frame, the final image is extracted by cropping only the region with the container using Mask R-CNN [182], followed by visual verification. Each crop is associated to an annotation of filling type and filling level (empty or filled at 50% or 90% of the capacity of the container), hand occlusion, and transparency of the container. We call this image dataset Crop-CCM or *C-CCM*. Sample C-CCM images<sup>10</sup> are shown in Fig. 4.15.

Finally, in order to evaluate the robustness of the classifier to distribution shifts of the test data, we focus mainly on the shape of the containers, but also in their color, texture, transparency, and size. To investigate these aspects, we split C-CCM into train and test sets under three configurations. The first configuration ( $S_1$ ) considers a champagne flute

---

<sup>10</sup>Sampled images can be found at <https://corsmal.eecs.qmul.ac.uk/filling.html>



Figure 4.15: Sample images (resized crops) from the CORSMAL Containers Manipulation dataset [181]. Each column shows different filling types and levels, and each row shows different backgrounds and hand occlusions.

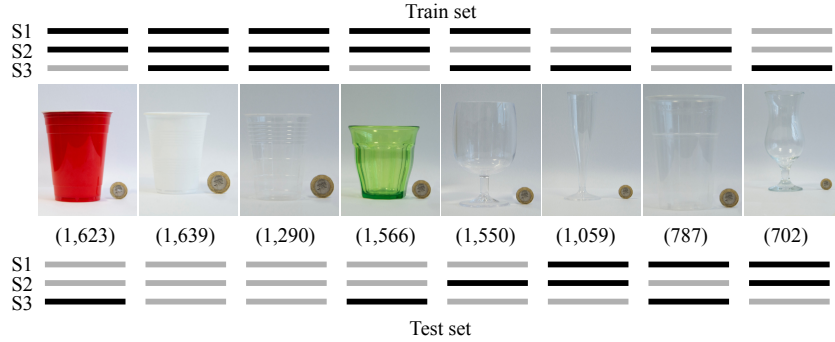


Figure 4.16: Comparison of three train and test splits (S1, S2, S3) of the public containers from CCM for the shape analysis in the experiments. Black lines mean that the set of images belonging to that container are part of the train (test) set in the data split. The number of images for each container are shown in parentheses. Note the diversity in shape, color, texture, transparency, and size.

in the test set to further increase the shape variability of containers not previously seen in the train set. The second configuration ( $S_2$ ) swaps a beer cup with a wine glass to analyze the influence of the stem of the wine glass. The last configuration ( $S_3$ ) places all the containers with a stem in the train set, and the test set contains only cups without stem, as well as a red cup and a green transparent cup, which are characterized by differences in color, texture and transparency with respect to the training data. Fig. 4.16 shows the three configurations and the number of samples for each container type.

### 4.5.2 Adversarial invariance prevents overfitting

In this section we analyze the effect of the transfer learning parameters, and then evaluate the generalization performance of the six different training strategies on the C-CCM dataset. We use as classifier a ResNet-18 [157]. Note that we also conducted experiments using a ResNet-50 and a WideResNet-50 [183], and the findings are similar to the ones of ResNet-18. We will focus on ResNet-18 as it is the least complex network among the three. With ST we train the classifier on C-CCM, whereas with AT we train the classifier on images modified with  $\ell_2$  adversarial perturbations crafted with the 10-iteration PGD [25]. With the transfer learning strategies we fine-tune the available pre-trained models on C-CCM: for ST→FT and ST→AFT we use the pre-trained model provided by PyTorch [158], whereas for AT→FT and AT→AFT we use the robust models provided by [170].

For each strategy, we train or fine-tune the classifier for 30 epochs, using a cross-entropy loss and stochastic gradient descent. The learning rate for updating the weights is set to 0.1 when training directly on C-CCM, and 0.005 when performing transfer learning. The learning rate decays linearly during training. Note that the models we evaluate are the ones obtained at the end of the training epochs (no early-stopping), while for dealing with class imbalances, the training images in a batch are randomly sampled with probabilities that are inversely proportional to the number of images of each class.

#### Sensitivity analysis

We perform a sensitivity analysis on the number of fixed layers ( $L$ ) in fine-tuning with ST→FT, ST→AFT, AT→FT and AT→AFT; and to select the size of the bound for crafting the adversarial perturbation for AT, ST→AFT, AT→FT and AT→AFT. Note that we differentiate  $\varepsilon$  for the source,  $\varepsilon^s$ , and target,  $\varepsilon^t$ , domain. Specifically, we perform the sensitivity analysis only for  $\varepsilon^s$  with AT→FT, and for each data split configuration we select the  $\varepsilon^s$  for which AT→FT achieves the highest accuracy. Then, based on these values of  $\varepsilon^s$ , for each data configuration we set  $\varepsilon^t = \varepsilon^s$ : since we use 10-iteration  $\ell_2$ -PGD, performing a sensitivity analysis or a grid search on  $\varepsilon^t$  is computationally inefficient, as it is analogous to increasing almost  $10\times$  the training epochs.

We first analyze the classification accuracy on the test sets of the three dataset splits when varying the number of fixed layers for ST→FT as  $L = \{0, 1, 2, 3, 4\}$ . Here,  $L = 0$  denotes that no layer remains fixed during fine-tuning (the full network is updated). Note that for a ResNet-18 classifier, a layer is a ResNet block of convolutions and batch normalization (see the original ResNet paper [157]). Since the target dataset is small, it is reasonable to fix the first layer ( $L = 1$ ) in order to prevent the classifier from a possible overfitting [184]. Indeed, Fig. 4.17 (left) shows that the accuracy on the test set of all configurations (S1, S2, S3) is consistently higher for  $L = 1$  (78.34%, 65.63%, 82.32%),






Figure 4.17: Sensitivity analysis for the number of fixed layers  $L$  with ST→FT (left) and for the maximum amount of perturbation bound,  $\epsilon^s$ , with AT→FT on test set of the three dataset splits: first split  $S_1$  (■), second split  $S_2$  (●), third split  $S_3$  (▲). Red indicates the highest achieved accuracy. Note the different scale of the y-axis, and the logarithmic scale for the x-axis (right).

while it gradually decays as  $L$  grows. This is also expected [184], since we allow fewer layers to be fine-tuned on the target datasets, and the classifiers then mostly use fixed features from ImageNet. Therefore, we set  $L = 1$  for ST→FT as well as for ST→AFT, AT→FT, and AT→AFT.

Then, we set  $L = 1$  and analyze the classification accuracy of AT→FT when varying the perturbation size in the source domain,  $\epsilon^s$ . Fig. 4.17 (right) shows that the highest achieved accuracy is different for each dataset configuration: 80.97% for  $S_1$  with  $\epsilon^s = 0.05$ , 73.27% for  $S_2$  with  $\epsilon^s = 1$ , and 88.23% for  $S_3$  with  $\epsilon^s = 0.5$ . As mentioned previously, we use these values of  $\epsilon^s$  also for  $\epsilon^t$  when performing AT, ST→AFT, and AT→AFT. However, we observed that the model trained with ST→AFT is unable to converge (train accuracy around 45%) on  $S_2$  for  $\epsilon^t = 1$  and on  $S_3$  for  $\epsilon^t = 0.5$ , while it successfully converges on  $S_1$  for the smaller  $\epsilon^t = 0.05$ . We believe that this might be caused by the fact that AT with larger  $\epsilon^t$  values eliminates many non-robust, yet useful, features transferred from ImageNet, and prevents the model from fitting the remaining features. Hence, we set  $\epsilon^t = 0.05$  for ST→AFT across all dataset configurations for the rest of the experiments, since with this value the network converges for all dataset configurations.

### Generalization to unseen data

Since all the parameters have been decided, we can now proceed with evaluating the generalization performance of every training strategy on each container in the test set individually. Note that, in this way, we can practically measure the robustness of the classifier to the test-time distribution shifts.

Figure 4.18 shows the filling level classification performance. Constrained by the amount, and hence by the diversity, of training images, the differently trained classifiers could potentially develop biases or overfit to some features, such as the shape of a container.

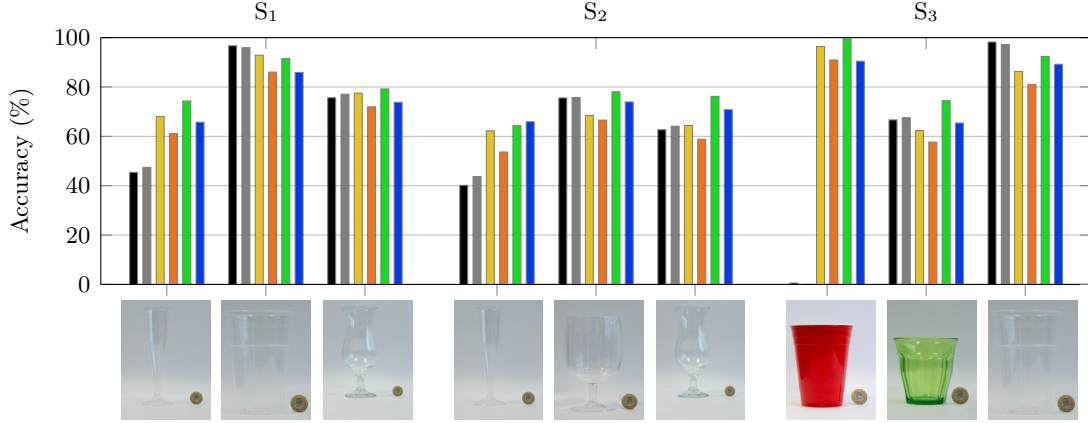


Figure 4.18: Comparison of the per-container filling level classification accuracy (%) for the six training strategies. Note the different containers in the test set for each dataset split (see Fig. 4.16 for the train set of each split). Legend: — ST, — AT, — ST→FT, — ST→AFT, — AT→FT, — AT→AFT.

AT→FT achieved superior performance most of the times. With transfer learning, the features introduced from ImageNet appear to decrease such biases, and enable the classifiers to identify features in the train set that are more generalizable. When combining transfer learning with AT at the source domain, the biases are modulated with the transferred features that are also filtered by AT, and the generalization of the classifier further increases. These results confirm that adversarial training improves transfer learning, even in the context of the challenging filling level classification task.

Overall, whenever the performance of ST is low, all transfer learning strategies lead to a significant improvement. On the contrary, whenever ST performs well, the contribution of transfer learning is insignificant, and sometimes it even decreases the final performance. Furthermore, applying AT on the *target* domain, either alone or combined with transfer learning, may even be harmful for the classifier.

For S<sub>1</sub>, the accuracy of ST on the beer cup (middle) is already very high – due to the similar shape of the small transparent cup in the training set – and the other training strategies do not further improve it. The accuracy on the cocktail glass (right) is similar for all strategies, with AT→FT performing slightly better. As for the champagne flute (left), the performance of ST and AT is quite low ( $\sim 46\%$ ), which might be caused by the unique shape of the flute with respect to the shapes in the training set. However, the accuracy significantly improves with transfer learning, and AT→FT outperforms all the other strategies by  $\sim 30$  percentage points (pp).

For S<sub>2</sub>, the accuracy of all strategies on the champagne flute (left) is similar to the one achieved on S<sub>1</sub>. The accuracy on the cocktail glass (right) is much lower for most strategies ( $\sim 10$ pp less compared to the performance on S<sub>1</sub>), except AT→FT, which drops

only by 3pp and again outperforms the rest of the strategies. The drop of the other strategies could be caused by the lack of a container with a stem in the training set. Finally, the performance on the wine glass (middle) is similar for most strategies, with AT→FT being again slightly better than the rest. Compared to the cocktail glass, the higher accuracy of all strategies on the wine glass could be caused by the similarity of its shape above the stem with the other transparent cups in the training set, despite the fact that no container with a stem is presented in the training set.

For  $S_3$ , the accuracy of ST on the beer cup (right) is high and the other training strategies do not improve it. Instead, the accuracy of ST on the green glass (middle) is lower and reaches an accuracy of 66%. Although ST→FT does not improve the accuracy, AT→FT significantly increases it (almost 10pp). The red cup (left) obtains the most interesting improvement compared to the 0.005% accuracy of ST: all transfer learning techniques achieve an accuracy above 90%, with AT→FT achieving 99.5% classification accuracy. By inspecting the predictions of ST and AT, the classifier assigned the label full (filling level: 90%) almost 99% of the times. In fact, predicting the *unknown* class is conceptually different from estimating the filling level, and it is more related to classifying non-transparent containers. In this sense, the features learned for transparent objects that are full with rice or pasta might be correlated with the features of the red cup.

### 4.5.3 Discussion

In this section, we found out that using adversarial training on the source dataset (ImageNet) followed by transfer learning on the target dataset (C-CCM), permits to consistently improve generalization to unseen containers. Our analysis demonstrates the possibilities of exploiting adversarial robustness for tasks that extend beyond classical image classification settings.

However, adversarial robustness is not necessarily the best way to improve the robustness to certain distribution shifts, such as the one introduced by common corruptions of the data. This is gonna be the main focus of the next chapter, where we will design an efficient data augmentation scheme for conferring robustness to such shifts. In fact, this augmentation scheme can also be applied in the context of filling level estimation. As we will see, transfer learning from adversarially trained models – and transfer learning in general – is not necessary, since one can properly tune our data augmentation scheme of the next section in order to achieve similar or even better generalization performance.

## 4.6 Conclusions

In this chapter, we proposed a new geometric framework that permits to relate the features of a dataset with the distance to the decision boundary along specific directions.

Through a series of carefully designed experiments, both in synthetic and real image datasets, we explained how the inductive bias of the learning algorithm shapes the decision boundaries of neural networks, by creating boundaries that are invariant to non-discriminative directions.

Furthermore, we demonstrated that the decision boundaries are very sensitive to the position of the training samples, and that very small changes along specific directions can cause large and sudden changes in these directions. In fact, adversarial training exploits this sensitivity of the decision boundaries, as well as their inductive bias towards invariance to non-discriminative features, in order to build more robust classifiers. Interestingly, modifying the position of just a *minimal* number of training samples during adversarial training is enough to locally introduce excessive invariance.

Finally, we studied the implications that the invariance properties of robust models have in the downstream application of classifying the filling level within containers. In particular, we showed that, during transfer learning, using robust models in the source domain permits to consistently improve generalization to unseen containers in the target domain that come from a shifted distribution.

We believe that our new framework can be used in future research to investigate the connections between training features and the macroscopic geometry of deep models. This can serve as a tool to obtain new insights on the intriguing properties of deep networks, as we demonstrated for their catastrophic forgetting. On the practical side, there are some important applications that could benefit from our findings. In terms of robustness, identifying the small subspace of discriminative features of a network can lead to faster black box-attacks by restricting the search space of the perturbations. In fact our analysis answers why using low-frequency perturbations improves the query efficiency in recent attacks [86, 87, 90]. Simultaneously, the dependency of boundaries to just a few training samples can be exploited to design faster adversarial training schemes, and is a clear avenue for future research in active learning. Finally, having a better understanding about the mechanisms that lead to excessive invariance [135] after adversarial training could help boost the standard accuracy of robust models.

# 5 Robustness to non-adversarial distribution shifts

*“Your favourite virtue... Simplicity!”*

— Karl Marx

## 5.1 Introduction

In the previous chapters, we mainly exploited adversarial proxies for evaluating the robustness of image classifiers, and for developing methodologies that enabled us to understand, analyze and further explore multiple properties of deep networks. Furthermore, we also saw that, in some tasks, the invariances that adversarially trained classifiers obtain can boost the robustness of deep networks to specific distribution shifts of the data. In this chapter, we steer our focus from the adversarial to the more general setting, and investigate the robustness of deep networks to common distortions of the images.

In general, deep image classifiers do not work well in the presence of various types of distribution shifts [7, 44, 45]. Most notably, their performance can severely drop when the input images are affected by common distortions that are not contained in the training data, such as digital artefacts, low contrast, or blurs [1, 49].

“Common corruptions” is an umbrella term used to describe the set of distortions that can happen to natural images during their acquisition, storage, and processing lifetime, which can be very diverse. Nevertheless, while the space of possible perturbations is huge, the term “common corruptions” is generally used to refer to typical image transformations

---

Part of this chapter appears in

“PRIME: A few primitives can boost robustness to Common Corruptions”, 2021 [185]. Under review.

“Data augmentation with mixtures of max-entropy transformations for filling-level classification”. In *European Signal Processing Conference (EUSIPCO)*, 2022 [186].

that, while degrading the quality of the images, still preserve their semantic information.

Building classifiers that are robust to common corruptions is far from trivial. A naive solution is to include data with all sorts of corruptions during training, but the sheer scale of all possible types of typical perturbations that might affect an image is simply too large. Moreover, the problem is per se ill-defined since there exists no formal description of all possible common corruptions.

Due to the lack of such formal description, the research community has recently favoured increasing the “diversity” of the training data via data augmentation schemes [36, 37, 63]. Intuitively, the hope is that showing very diverse augmentations of an image to a network would increase the chance that the latter becomes invariant to some common corruptions. Still, covering the full space of common corruptions is hard. Hence, current literature has mostly resorted to increasing the diversity of augmentations by designing intricate data augmentation pipelines, e.g., introducing DNNs for generating varied augmentations [37, 38], or coalescing multiple techniques [65], and thus achieve good performance on different benchmarks. This strategy, though, leaves a big range of unintuitive design choices, making it hard to pinpoint which elements of these methods meaningfully contribute to the overall robustness. Meanwhile, the high complexity of recent methods [38, 65] makes them impractical for large-scale tasks. Whereas, some methods are tailored to particular datasets and might not be general enough. Nonetheless, the problem of building robust classifiers is far from completely solved, and the gap between robust and standard accuracy is still large.

In this chapter, we take a step back and provide a systematic way for designing a simple, yet effective data augmentation scheme. By focusing on first principles, we formulate a new model for semantically-preserving corruptions, and build on basic concepts to characterize the notions of transformation strength and diversity using a few transformation primitives. Relying on this model, we propose *PRIME*, a data augmentation scheme that draws transformations from a max-entropy distribution to efficiently sample from a large space of possible distortions. The performance of *PRIME*, alone, already tops the current baselines on different common corruption datasets, whilst it can also be combined with other methods to further boost their performance. Moreover, the simplicity and flexibility of *PRIME* allows to easily understand how each of its components contributes to improving robustness. Finally, we demonstrate that *PRIME* provides a ready-to-use recipe for the problem of filling level classification.

The rest of the chapter is organized as follows: In Sec. 5.2 we demonstrate that simple pre-processing techniques might not be enough for inducing invariance to all possible corruptions, and we formulate a new mathematical model for semantically-preserving corruptions. Then, in Sec. 5.3 we provide a data augmentation scheme that utilizes the proposed model of visual distortions, and evaluate its performance on multiple corruption benchmarks. Furthermore, in Sec. 5.4 we use our method to investigate different aspects

behind robustness to common corruptions, while, in Sec. 5.5 we tune PRIME to confer robustness to distributions shifts in filling level classification. Finally, in Sec. 5.6 we provide a discussion on open challenges and potential extensions of our method.

## 5.2 Towards robustness to common corruptions

### 5.2.1 Invariance by removing features

Data augmentation is the most common technique for improving the robustness of classifiers to common distortions of the images. However, designing efficient data augmentation schemes can be challenging, while at the same time the methods can be computationally intense. To this end, it is reasonable to explore if the ideas introduced in Chapter 4 can also be beneficial. That is, apply some simple pre-processing operations for inducing specific invariances that can improve the robustness to common corruptions, instead of generating augmentations at each training iteration.

In order to evaluate this approach, we will use the widely adopted *Common Corruptions* [1] benchmark. This benchmark consists of 15 natural image distortions that can naturally occur during acquisition, processing or storing of an image, each applied with 5 different severity levels. At a higher level, these corruptions can be grouped into four categories, namely “noise”, “blur”, “weather” and “digital”.

In previous works [3, 29] it was observed that some of these common corruptions affect the high frequency part of the image frequency spectrum (i.e., high-frequency distortions). Hence, following the exact same procedure as in Sec. 4.3.2, we can train the classifier on a low-pass filtered version of the data, such that the network will be invariant to any high-frequency changes introduced by the corruptions. Furthermore, color is another important image space that can be affected by the aforementioned common corruptions. In this sense, we might also want to introduce to the network some invariance to color changes, similarly to what can be achieved by projecting the images onto the low-frequency part of the spectrum to achieve invariance to the high-frequency one. Hence, we can remove the color information from the data by training the classifier on grayscale images. Note, however, that for retaining the architecture of the network such that it can still be able to process standard RGB images at test time, we repeat the grayscale training images along 3 channels.

In order to evaluate the aforementioned operations we conduct a small experiment. Specifically, we train a ResNet-18 for 30 epochs on CIFAR-10 and its pre-processed versions, using SGD and a cross-entropy loss function. For comparison, we also train a network on CIFAR-10 using AugMix [36], a widely-used data augmentation method that achieves very good robustness on the common corruptions benchmark. For simplicity, we do not deploy the Jensen-Shannon Divergence (JSD) consistency loss – as introduced

	Standard	LP	Gray	AugMix	AugMix+LP
Gauss. noise	46.1	<b>57.6</b>	45.6	69.6	<u>70.9</u>
Shot noise	58.2	<b>65.9</b>	56.4	78.3	<u>79.1</u>
Impulse noise	54.5	56.2	<b>57.4</b>	77.5	<u>78.5</u>
Defocus blur	81.5	<b>88.8</b>	71.7	91.8	<u>91.5</u>
Glass blur	51.0	<b>85.3</b>	35.9	64.0	65.4
Motion blur	75.0	<b>84.8</b>	62.9	<u>88.1</u>	88.0
Zoom blur	75.3	<b>87.6</b>	62.0	<u>89.9</u>	89.6
Snow	80.6	<b>85.7</b>	73.0	84.5	85.0
Frost	75.5	<b>86.9</b>	66.5	83.8	84.3
Fog	<b>88.0</b>	81.2	82.3	<u>89.9</u>	89.7
Brightness	<b>92.9</b>	89.0	91.0	92.7	<u>93.0</u>
Contrast	76.1	68.8	<b>78.4</b>	83.6	<u>84.5</u>
Elastic	81.7	<b>86.5</b>	73.1	<u>86.9</u>	86.5
Pixelate	73.9	<b>90.1</b>	63.3	79.5	81.2
JPEG	77.4	<b>83.1</b>	73.0	82.4	<u>83.2</u>
Clean	<u>94.4</u>	91.1	92.4	94.1	94.1
Avg. Corruption	72.5	<b>79.8</b>	66.2	82.8	<u>83.4</u>

Table 5.1: Clean and corruption accuracy of ResNet-18 on CIFAR-10 using pre-processing methods, AugMix, or combining AugMix with low-pass filtering (AugMix+LP). Bold: maximum value among Standard, LP, and Gray. Underline: maximum value per row.

in the original AugMix paper – when training with AugMix. Finally, we evaluate the performance of the resulting networks by measuring their accuracy on the validation set of CIFAR-10 (C-10) and its corruption counterpart, CIFAR-10-C (C-10-C).

The results are presented in Tab. 5.1. Focusing on the pre-processing methods, the first thing to notice is that the low-pass (LP) filtering has a very strong effect on the robustness of the classifier, mainly improving the accuracy on blurs and digital corruptions (e.g., Pixelate). However, as expected from Sec. 4.3.2, there is a drop of around 3% in the clean accuracy of the model. It is quite interesting that such a simple operation of low-complexity can boost the robustness significantly, which verifies the observation that redundant features can hurt robustness to distribution shifts [160]. Nevertheless, this might be just a CIFAR-10 artefact: recall from Sec. 4.3 that for other datasets like ImageNet, the classifiers are already invariant to high-frequency features, and hence a low-pass filtering might not have an effect similar to the one observed here. On the other hand, training on grayscale images seems to hurt the robustness of the classifier, except for the case of impulse noise and contrast. Hence, deploying such pre-processing method to achieve robustness to color changes is not beneficial.

From the results with data augmentation, it is evident that AugMix achieves good results, without really hurting the accuracy on the clean images. Interestingly, though, in many corruption types low-pass filtering achieves better robustness than AugMix. Hence, and in order to investigate any complementary gains between the two methods, we trained the network by combining AugMix with low-pass filtering as well. Such combination can



further increase the overall robustness, achieving better results than simply applying AugMix or LP alone. Nevertheless, it seems that in some cases (i.e., glass blur) the influence of AugMix dominates, and it constrains some beneficial properties of LP.

Overall, from our analysis we can highlight two main insights. First, pre-processing can sometimes be beneficial, but some others can even be hurtful. Beyond the techniques we deployed, there are many other ways to pre-process the data and possibly confer robustness to some corruptions. However, the problem of common corruptions is ill-posed, and it is difficult to introduce pre-processing methods that can cover every possible corruption that might occur. Second, using data augmentation seems to be a sensible direction. Yet, the AugMix+LP experiment hints that existing augmentation practices, such as AugMix, might not properly cover the space of corruptions and can be further improved. Hence, it is important to design a more general data augmentation framework that effectively increases the coverage over the space of possible distortions.

### 5.2.2 General model of visual corruptions

Motivated by the “semantically-preserving” nature of common corruptions, we define a new model of typical distortions. Specifically, we leverage the long tradition of image processing in developing techniques to manipulate images while retaining their semantics and construct a principled framework to characterize a large space of visual corruptions.

Let  $\mathbf{x} : [0, 1]^2 \rightarrow [0, 1]^3$  be a continuous image<sup>1</sup> mapping pixel coordinates  $\mathbf{r} = (r_1, r_2)$  to RGB values. We define our model of common corruptions as the action on  $\mathbf{x}$  of the following additive subgroup of the near-ring of transformations [187]

$$\mathcal{T}_{\mathbf{x}} = \left\{ \sum_{i=1}^n \lambda_i g_1^i \circ \cdots \circ g_m^i(\mathbf{x}) : g_j^i \in \{\omega, \tau, \gamma\}, \lambda_i \in \mathbb{R} \right\}, \quad (5.1)$$

where  $\omega, \tau$  and  $\gamma$  are random primitive transformations which distort  $\mathbf{x}$  along the spectral ( $\omega$ ), spatial ( $\tau$ ), and color ( $\gamma$ ) domains. As we will see, defining each of these primitives in a principled and coherent fashion will be enough to construct a set of perturbations which covers most types of visual corruptions.

To guarantee as much diversity as possible in our model, we follow the principle of maximum entropy to define our distributions of transformations [188]. Note that using a set of augmentations that guarantees maximum entropy comes naturally when trying to optimize the sample complexity derived from certain information-theoretic generalization bounds, both in the clean [189] and corrupted settings [190]. Specifically, the principle of maximum entropy postulates favoring those distributions that are as unbiased as possible given the set of constraints that define a family of distributions. In our case, these

<sup>1</sup>We define our model of common corruptions in the continuous domain for simplicity. However, as is common in image processing, in practice we will work with discrete images on a regular grid.

constraints are given in the form of an expected strength  $\sigma^2$ , some boundary conditions, e.g., the displacement field must be zero at the borders of an image, and finally the desired smoothness level  $K$ . The principle of smoothness helps formalize the notion of physical plausibility, as most naturally occurring processes are smooth.

Formally, let  $\mathcal{I}$  denote the space of all images, and let  $f : \mathcal{I} \rightarrow \mathcal{I}$  be a random image transformation distributed according to the law  $\mu$ . Further, let us define a set of constraints  $\mathcal{C} \subseteq \mathcal{F}$ , which restricts the domain of applicability of  $f$ , i.e.,  $f \in \mathcal{C}$ , and where  $\mathcal{F}$  denotes the space of functions  $\mathcal{I} \rightarrow \mathcal{I}$ . The principle of maximum entropy postulates using the distribution  $\mu$  which has maximum entropy given the constraints:

$$\begin{aligned} & \underset{\mu}{\text{maximize}} && H(\mu) = - \int_{\mathcal{F}} d\mu(f) \log(\mu(f)) \\ & \text{subject to} && f \in \mathcal{C} \quad \forall f \in \text{supp}(\mu), \end{aligned} \tag{5.2}$$

where  $H(\mu)$  represents the entropy of the distribution  $\mu$  [188]. In its general form, solving Eq. (5.2) for any set of constraints  $\mathcal{C}$  is intractable. However, as we show in App. B.1, for the distributions of each of our family of transformations we can derive an analytical expression in closed form, by leveraging results from statistical physics [191].

In what follows, we describe the analytical solutions to Eq. (5.2) for each of our basic primitives. In general, these distributions are governed by two parameters:  $K$  to control smoothness, and  $\sigma^2$  to control strength. These transformations fall back to identity mappings when  $\sigma^2 = 0$ , independently of  $K$ .

**Spectral domain** We parameterize the distribution of random spectral transformations using random filters  $\omega(\mathbf{r})$ , such that the output of the transformation follows

$$\omega(\mathbf{x})(\mathbf{r}) = (\mathbf{x} * (\delta + \omega'))(\mathbf{r}), \tag{5.3}$$

where,  $*$  is the convolution operator,  $\delta(\mathbf{r})$  represents a Dirac delta, i.e., identity filter, and  $\omega'(\mathbf{r})$  is implemented in the discrete grid as a finite impulse response (FIR) filter of size  $K_\omega \times K_\omega$  with i.i.d random entries distributed according to  $\mathcal{N}(0, \sigma_\omega^2)$ . Here,  $\sigma_\omega^2$  governs the transformation strength, while larger  $K_\omega$  yields filters of higher spectral resolution. The bias  $\delta(\mathbf{r})$  makes the output close to the original image.

**Spatial domain** We define our distribution of random spatial transformations, which apply random perturbations over the coordinates of an image, using the following model

$$\tau(\mathbf{x})(\mathbf{r}) = \mathbf{x}(\mathbf{r} + \boldsymbol{\tau}'(\mathbf{r})). \tag{5.4}$$

This model has been recently proposed by Petrini *et al.* [192] to define a distribution of random smooth diffeomorphisms in order to study the stability of neural networks to small spatial transformations. To guarantee smoothness but preserve maximum entropy,

Petrini *et al.* propose to parameterize the vector field  $\boldsymbol{\tau}'$  as

$$\boldsymbol{\tau}'(\mathbf{r}) = \sum_{i^2+j^2 \leq K_\tau^2} \beta_{i,j} \sin(\pi i \mathbf{r}_1) \sin(\pi j \mathbf{r}_2), \quad (5.5)$$

where  $\beta_{i,j} \sim \mathcal{N}(0, \sigma_\tau^2 / (i^2 + j^2))$ . This choice of values guarantees that the resulting mapping is smooth according to the cut frequency  $K_\tau$ , while  $\sigma_\tau^2$  determines its strength.

**Color domain** We follow a similar approach to define the distribution of random color transformations. That is, we build random mappings  $\gamma$  between color spaces such that

$$\gamma(\mathbf{x})(\mathbf{r}) = \mathbf{x}(\mathbf{r}) + \sum_{n=0}^{K_\gamma} \boldsymbol{\beta}_n \odot \sin(\pi n \mathbf{x}(\mathbf{r})), \quad (5.6)$$

where  $\boldsymbol{\beta}_n \sim \mathcal{N}(0, \sigma_\gamma^2 \mathbf{I}_3)$ , with  $\odot$  denoting elementwise multiplication. Again,  $K_\gamma$  controls the smoothness of the transformations and  $\sigma_\gamma^2$  their strength. Note however that, compared to Eq. (5.5), the coefficients in Eq. (5.6) are not weighted by the inverse of the frequency, and have constant variance. In practice, we observe that reducing the variance of the coefficients for higher frequencies creates color mappings that are too smooth and almost imperceptible, so we decided to drop this dependency in our model.

Finally, we note that our model is very flexible with respect to its core primitives. In particular, it can be easily extended to include other distributions of maximum entropy transformations that suit an objective task. For example, one might add the distribution of maximum entropy additive perturbations given by  $\eta(\mathbf{x})(\mathbf{r}) = \mathbf{x}(\mathbf{r}) + \boldsymbol{\eta}'(\mathbf{r})$ , where  $\boldsymbol{\eta}'(\mathbf{r}) \sim \mathcal{N}(0, \sigma_\eta^2)$ . Nonetheless, since most benchmarks of visual corruptions disallow the use of additive perturbations during training [1], our model does not include an additive perturbation category.

Overall, based on the results in Secs. 5.3.2 and 5.4.2, our model is flexible and covers a large part of the semantic-preserving distortions. It also allows to easily control the strength and style of the transformations with just a few parameters. Moreover, changing the transformation strength enables to control the trade-off between corruption robustness and standard accuracy, as shown in Sec. 5.4.3. In what follows, we use our model to design an efficient augmentation scheme to build classifiers robust to common corruptions.

## 5.3 PRIME data augmentations

### 5.3.1 Instantiating the general model of visual corruptions

We now introduce PRIME, a simple yet efficient augmentation scheme that uses our **PR**imitives of **MA**ximum **EN**tropy to confer robustness against common corruptions.

---

**Algorithm 3: PRIME**

---

**Input:** Image  $x$ , primitives  $\mathcal{G} = \{\text{Id}, \omega, \tau, \gamma\}$ , where  $\text{Id}$  is the identity operator  
**Output:** Augmented image  $\tilde{x}$

```

1  $\tilde{x}_0 \leftarrow x$ 
2 for  $i \in \{1, \dots, n\}$  do
3    $\tilde{x}_i \leftarrow x$ 
4   for  $j \in \{1, \dots, m\}$  do
5      $g \sim \mathcal{U}(\mathcal{G})$  ▷ Strength  $\sigma \sim \mathcal{U}(\sigma_{\min}, \sigma_{\max})$ 
6      $\tilde{x}_i \leftarrow g(\tilde{x}_i)$ 
7   end
8 end
9  $\lambda \sim \text{Dir}(\mathbf{1})$  ▷ Random Dirichlet convex coefficients
10  $\tilde{x} \leftarrow \sum_{i=0}^n \lambda_i \tilde{x}_i$ 

```

---

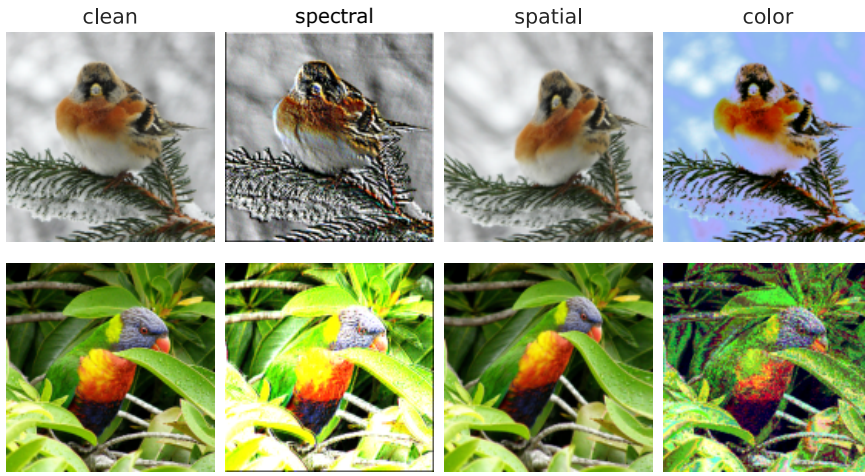


Figure 5.1: Example generated with the transformations of our common corruptions model. Despite the perceptibility of the distortion, the image semantics are preserved.

The pseudo-code of PRIME is given in Algorithm 3, which draws a random sample from Eq. (5.1) using a convex combination of a composition of basic primitives. Below we describe the main implementation details of our algorithm.

**Parameter selection** It is important to ensure that the semantic information of an image is preserved after it goes through PRIME. As measuring semantic preservation quantitatively is not simple, we subjectively select each primitive’s parameters based on visual inspection, ensuring maximum permissible distortion while retaining the semantic content of the image<sup>2</sup>. To avoid relying on a specific strength for each transformation, PRIME stochastically generates augmentations of different strengths by sampling  $\sigma$  from a uniform distribution, with different minimum and maximum values for each primitive.

For the color primitive, we observed that fairly large values for  $K_\gamma$  (in the order of 500) are important for covering a large space of visual distortions. Unfortunately, implementing

---

<sup>2</sup>All the selected transformation parameters of PRIME are provided in App. B.2.




Figure 5.2: Images generated with PRIME, a simple method that uses primitive families of max-entropy transformations in different visual domains to create diverse data augmentations.

such a transformation can be memory inefficient. To avoid this issue, PRIME uses a slight modification of Eq. (5.6) and combines a fixed number  $\Delta$  of consecutive frequencies randomly chosen in the range  $[0, K_\gamma]$ . Finally, as some of our transformations can push the images outside of their color range, we always clip the output of each transformation so that it lies on  $[0, 1]^3$ .

**Mixing transformations** The concept of mixing has been a recurring theme in the augmentation literature [36, 60, 61, 65] and PRIME follows the same trend. In particular, Algorithm 3 uses a convex combination of  $n$  basic augmentations consisting of the composition of  $m$  of our primitive transformations (see Fig. 5.2 for examples generated with PRIME). In our experiments, we fix the total number of generated transformed images (width) to be  $n = 3$ . As for the composition of the transformations (depth), we follow a stochastic approach such that, on every iteration  $i \in \{1, \dots, n\}$ , only  $\hat{m} \in [1, m]$  compositions are performed, with  $m = 3$ . In fact, in Algorithm 3 we do not explicitly select randomly a new  $\hat{m}$  for every  $i$  but we provide the identity operator  $\text{Id}$  instead. This guarantees that, in some cases, no transformation is performed.

In general, the convex mixing procedure (i) broadens the set of possible training augmentations, and (ii) ensures that the augmented image stay close to the original one. We later provide empirical results which underline the efficacy of mixing in Sec. 5.4.2. Finally, note that, the basic skeleton of PRIME is similar to that of AugMix. However, as we will see next, incorporating our maximum entropy transformations leads to significant gains in common corruptions robustness over AugMix.

### 5.3.2 Performance on common corruptions

We compare the classification performance of PRIME on the common corruption datasets [1], with that of two current approaches: AugMix [36] and DeepAugment

Dataset	Method	Clean Acc ( $\uparrow$ )	Common Acc ( $\uparrow$ )	Corruption mCE ( $\downarrow$ )
C-10	Standard	95.0	74.0	24.0
	AugMix	95.2	88.6	11.4
	PRIME	94.2	<b>89.8</b>	<b>10.2</b>
C-100	Standard	76.7	51.9	48.1
	AugMix	78.2	64.9	35.1
	PRIME	78.4	<b>68.2</b>	<b>31.8</b>
IN-100	Standard	88.0	49.7	100.0
	AugMix	88.7	60.7	79.1
	DA	86.3	67.7	68.1
	PRIME	85.9	<b>71.6</b>	<b>61.0</b>
	DA+AugMix	86.5	73.1	57.3
	DA+PRIME	84.9	<b>74.9</b>	<b>54.6</b>
IN	Standard*	76.1	38.1	76.7
	AugMix*	77.5	48.3	65.3
	DA*	76.7	52.6	60.4
	PRIME $^\dagger$	77.0	<b>55.0</b>	<b>57.5</b>
	DA+AugMix	75.8	58.1	53.6
	DA+PRIME $^\dagger$	75.5	<b>59.9</b>	<b>51.3</b>

Table 5.2: Clean and corruption accuracy, and mean corruption error (mCE) for different methods with ResNet-18 on C-10, C-100, IN-100 and ResNet-50 on IN. mCE is the mean corruption error on common corruptions un-normalized for C-10 and C-100; normalized relative to standard model on IN-100 and IN.  $^\dagger$  indicates that JSD consistency loss is not used. \*Models taken from [193].

(DA) [37]. Regarding the training setup, we consider four datasets: CIFAR-10 (C-10), CIFAR-100 (C-100), ImageNet-100 (IN-100) and ImageNet (IN) [152]. IN-100 is a 100-class subset of IN obtained by selecting every 10<sup>th</sup> class in WordNet ID order. We train a ResNet-18 [157] on C-10, C-100 and IN-100; and ResNet-50 on IN. Following AugMix, and for a complete comparison, we also integrate the Jensen-Shannon divergence (JSD)-based consistency loss with PRIME which compels the network to learn similar representations for differently augmented versions of the same input image. Note that all the models are trained for 100 epochs.

Regarding the experimental details, all models are implemented in PyTorch [158] and are trained for 100 epochs using a cyclic learning rate schedule with cosine annealing and a maximum learning rate of 0.2 unless stated otherwise. For IN, we fine-tune a PyTorch regularly pretrained network with a maximum learning rate of 0.01 following Hendrycks *et al.* [37]. We use SGD optimizer with momentum factor 0.9 and Nesterov momentum. On C-10 & C-100, we set the batch size to 128 and use a weight decay of 0.0005. On IN-100 and IN, the batch size is 256 and weight decay is 0.0001. We employ ResNet-18 [157] on C-10, C-100 and IN-100; and use ResNet-50 for IN. The

augmentation hyperparameters for AugMix and DeepAugment are the same as in their original implementations.

We evaluate our models on the corrupted versions (C-10-C, C-100-C, IN-100-C, IN-C) of the aforementioned datasets, and the results are summarized in Tab. 5.2. Amongst individual methods, PRIME yields superior results compared to those obtained by AugMix and DeepAugment alone and advances the baseline performance on the corrupted counterparts of the four datasets. As listed, PRIME pushes the corruption accuracy by 1.2% and 3.3% on C-10-C and C-100-C respectively over AugMix.

On IN-100-C, a more complicated dataset, we observe significant improvements wherein PRIME outperforms AugMix by 10.9%. In fact, this increase in performance hints that our primitive transformations are actually able to cover a larger space of image corruptions, compared to the restricted set of AugMix. Interestingly, the random transformations in PRIME also lead to a 3.9% boost in corruptions accuracy over DeepAugment despite the fact that DeepAugment leverages additional knowledge to augment the training data via its use of pre-trained architectures. Moreover, PRIME provides cumulative gains when combined with DeepAugment, reducing the mean corruption error (mCE) of prior art (DA+AugMix) by 2.7% on IN-100-C.

Lastly, we evaluate the performance of PRIME on full IN-C, but we do not use JSD with PRIME to reduce the computational complexity. Yet, even without JSD, PRIME outperforms, in terms of corruption accuracy, both AugMix (with JSD) and DeepAugment by 6.7% and 2.4% respectively, while the mCE is reduced by 7.8% and 2.9%. Also, when PRIME is combined with DeepAugment, it also surpasses the performance of DA+AugMix (with JSD), reaching a corruption accuracy of almost 60% and an mCE of 51.3%.

Note here, that, not only PRIME achieves superior robustness, but it does so efficiently. Compared to standard training on IN-100, AugMix requires 1.20x time and PRIME requires 1.27x. In contrast, DA is tedious and we do not measure its runtime since it also requires the training of two large image-to-image networks for producing augmentations, and can only be applied offline.

Additionally, given the nuances amongst individual corruption types in common corruptions, we perform a fine-grained analysis with PRIME on IN-100-C to ensure that our method leads to general improvements against all corruption types. As evident from the comparison in Tab. 5.3, PRIME alone, even without the JSD term, improves robustness over current techniques to almost every corruption type except blur. Further, incorporating the JSD term with PRIME attains the best results on all the corruption categories. Relative to the previous best by DeepAugment on IN-100-C, PRIME improves by 4.3% on noises, 2% on blurs, 3.4% on weather changes and 5.8% against digital distortions. This validates that PRIME helps against all common corruption types in IN-100-C, underlining the generality of our model of common corruptions.

Method	IN-100-C	Noise	Blur	Weather	Digital	IN-100
AugMix <sup>†</sup>	55.2	38.9	56.8	57.0	64.2	88.0
AugMix	60.7	44.8	63.1	60.7	70.3	<b>88.7</b>
DA	67.7	75.9	62.5	63.6	70.9	86.3
PRIME <sup>†</sup>	68.8	78.8	58.3	66.0	74.8	87.1
PRIME	<b>71.6</b>	<b>80.2</b>	<b>64.5</b>	<b>67.0</b>	<b>76.7</b>	85.9

Table 5.3: Classification accuracy ( $\uparrow$ ) of various methods on the different corruption types contained in IN-100-C. <sup>†</sup> indicates that JSD consistency loss is not used. Network used: ResNet-18.

### 5.3.3 Unsupervised domain adaptation

Recently, robustness to common corruptions has also been of significant interest in the field of unsupervised domain adaptation [194, 195]. The main difference is that, in domain adaptation, one exploits the limited access to test-time corrupted samples to adjust certain network parameters. Hence, it would be interesting to investigate the utility of PRIME under the setting of domain adaption.

To that end, we combine our method with the adaption trick by Schneider *et al.* [195]. Specifically, we adjust the batch normalization (BN) statistics of our models using a few corrupted samples. Suppose  $z_s \in \{\mu_s, \sigma_s\}$  are the BN mean and variance estimated from the training data, and  $z_t \in \{\mu_t, \sigma_t\}$  are the corresponding statistics computed from  $n$  unlabelled, corrupted test samples, then we re-estimate the BN statistics as follows:

$$\hat{z} = \frac{N}{N+n} z_s + \frac{n}{N+n} z_t \quad (5.7)$$

We consider three adaptation scenarios: single sample ( $n = 1, N = 16$ ), partial ( $n = 8, N = 16$ ) and full ( $n = 400, N = 0$ ) adaptation. Here, we do not perform parameter tuning for  $N$ . As shown in Tab. 5.4, simply correcting BN statistics using as little as 8 corrupted samples pushes the corruption accuracy of PRIME from 71.6% to 75.3%. In general, PRIME yields cumulative gains in combination with adaptation and has the best IN-100-C accuracy.

## 5.4 Robustness insights

In this section, we exploit the simplicity of PRIME to investigate different aspects behind robustness to common corruptions. We first analyze how each transformation domain contributes to the overall robustness of the network. Then, we empirically locate and justify the benefits of mixing the transformations of each domain. Moreover, we demonstrate the existence of a robustness-accuracy trade-off, and, finally, we comment on the low-complexity benefits of PRIME in different data augmentation settings.



Method	IN-100-C accuracy ( $\uparrow$ )				IN-100 ( $\uparrow$ )
	w/o adapt	single adapt	partial adapt	full adapt	single adapt
Standard	49.7	53.8	62.0	63.9	88.1
AugMix	60.7	65.5	71.3	73.0	<b>88.3</b>
DA	67.7	70.2	72.7	74.6	86.3
PRIME	<b>71.6</b>	<b>73.5</b>	<b>75.3</b>	<b>76.6</b>	85.7

Table 5.4: Performance when in concert with domain adaptation on IN-100. Partial adaptation uses 8 samples; full adaptation uses 400 corrupted samples. Network used: ResNet-18.

#### 5.4.1 Contribution of transformations

We want to understand how the transformations in each domain of Eq. (5.1) contribute to the overall robustness. To that end, we conduct an ablation study on IN-100-C by training a ResNet-18 with the max-entropy transformations of PRIME individually or in combination. As shown in Tab. 5.5, spectral transformations mainly help against blur, weather and digital corruptions. Spatial operations also improve on blurs, but on elastic transforms as well (digital). On the contrary, color transformations excel on noises and certain high frequency digital distortions, *e.g.* pixelate and JPEG artefacts, and have a minor effect on weather changes. Besides, incrementally combining these transformations lead to cumulative gains *e.g.* spatial+color help on both noises and blurs. Yet, for obtaining the best results, the combination of all transformations is required. This means that each transformation increases the coverage over the space of possible distortions and the increase in robustness comes from their cumulative contribution.

Transform	IN-100-C	Noise	Blur	Weather	Digital	IN-100
None	49.7	27.3	48.6	54.8	62.6	<b>88.0</b>
$\omega$	64.1	60.7	55.4	66.6	72.9	87.3
$\tau$	53.8	30.1	56.2	57.6	65.4	87.0
$\gamma$	59.9	67.4	52.6	54.4	67.1	86.9
$\omega+\tau$	64.5	58.5	57.3	<b>66.8</b>	73.9	87.7
$\omega+\gamma$	67.5	77.2	55.7	65.3	74.2	87.1
$\tau+\gamma$	63.3	74.7	57.4	56.2	67.8	86.2
$\omega+\tau+\gamma$	<b>68.8</b>	<b>78.8</b>	<b>58.3</b>	66.0	<b>74.8</b>	87.1

Table 5.5: Impact of the different primitives of max-entropy ( $\omega$ : spectral,  $\gamma$ : color,  $\tau$ : spatial) from PRIME on common corruption accuracy ( $\uparrow$ ). All the transformations are essential for the performance of PRIME. All models are trained without JSD loss. Network used: ResNet-18.

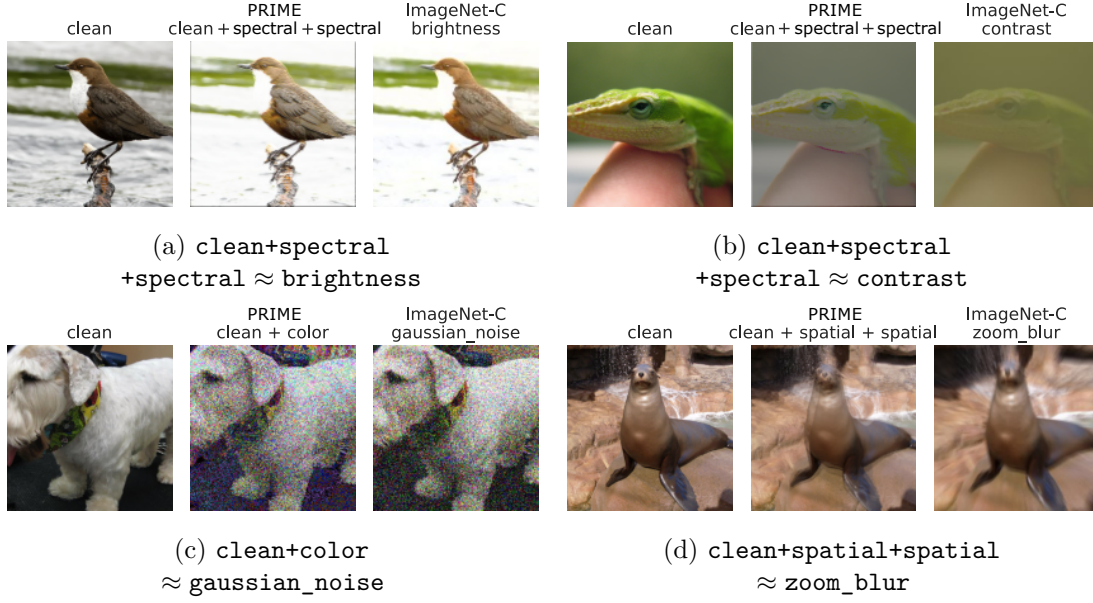


Figure 5.3: The mixing procedure creates distorted images that look visually similar to the test-time corruptions. In each example, we show the clean image, the PRIME image and the corresponding common corruption that resembles the image produced by mixing. We also report the mixing combination used for recreating the corruption. Additional examples can be found in App. D of [185].

### 5.4.2 The importance of mixing

In most data augmentation methods, besides the importance of the transformations themselves, mixing has been claimed as an essential module for increasing diversity in the training process [36, 60, 61, 65]. In our attempt to provide insights on the role of mixing in the context of common corruptions, we found out that it is capable of constructing augmented images that look perceptually similar to their corrupted counterparts. In fact, the improvements on specific corruption types observed in Tab. 5.5 can be largely attributed to mixing. As exemplified in Figs. 5.3a and 5.3b, careful combinations of spectral transformations with the clean image introduce brightness and contrast-like artefacts that look similar to the corresponding corruptions in IN-C. Also, combining spatial transformations creates blur-like artefacts that look identical to zoom blur in IN-C (Fig. 5.3d). Finally, notice in Fig. 5.3c how mixing color transformations helps fabricate corruptions of the “noise” category. This means that the max-entropy color model of PRIME enables robustness to different types of noise without explicitly adding any during training. This might explain the significant improvement over the “noise” category in Tab. 5.3.

Note that one of the main goals of data augmentation is to achieve maximum coverage of the space of possible distortions using a limited transformation budget, i.e., within a few training epochs. The principle of max-entropy guarantees this within each primitive,

Method	Min. cosine distance ( $\times 10^{-3}$ )	
	Avg. ( $\downarrow$ )	Median ( $\downarrow$ )
None (clean)	25.38	6.44
AugMix (w/o mix)	20.57	3.56
PRIME (w/o mix)	<b>10.61</b>	<b>1.88</b>
AugMix	17.48	2.61
PRIME	<b>7.71</b>	<b>1.61</b>

Table 5.6: Minimum cosine distances in the ResNet-50 SimCLRv2 embedding space between 100 augmented samples from 1000 ImageNet images, and their corresponding common corruptions.

but the effect of mixing on the overall space is harder to quantify. In this regard, we can use the distance in the embedding space  $\phi$  of a SimCLRv2 [196] model, as a proxy for visual similarity [197, 198]. We are interested in measuring how mixing the base transformations changes the likelihood that an augmentation scheme generates some sample during training that is visually similar to some of the common corruptions. To that end, we randomly select  $N = 1000$  training images  $\{\mathbf{x}_n\}_{n=1}^N$  from IN, along with their  $C = 75$  (15 corruptions of 5 severity levels) associated common corruptions  $\{\hat{\mathbf{x}}_n^c\}_{c=1}^C$ , and generate for each of the clean images another  $T = 100$  transformed samples  $\{\tilde{\mathbf{x}}_n^t\}_{t=1}^T$  using each augmentation scheme. Moreover, for each corruption  $\hat{\mathbf{x}}_n^c$  we find its closest neighbor  $\tilde{\mathbf{x}}_n^t$  from the set of generated samples using the cosine distance in the embedding space<sup>3</sup>. Our overall measure of fitness is

$$\frac{1}{NC} \sum_{n=1}^N \sum_{c=1}^C \min_t \left\{ 1 - \left( \frac{\phi(\hat{\mathbf{x}}_n^c)^\top \phi(\tilde{\mathbf{x}}_n^t)}{\|\phi(\hat{\mathbf{x}}_n^c)\|_2 \|\phi(\tilde{\mathbf{x}}_n^t)\|_2} \right) \right\}. \quad (5.8)$$

Table 5.6 shows the values of this measure applied to AugMix and PRIME, with and without mixing. For reference, we also report the values of the clean (no transform) images  $\{\mathbf{x}_n\}_{n=1}^N$ . Clearly, mixing helps reducing the distance between the common corruptions and the augmented samples from both methods. We also observe that PRIME, even with only 100 augmentations per image – in the order of number of training epochs – can generate samples which are twice as close to the common corruptions as AugMix. In fact, the feature similarity between training augmentations and test corruptions was also studied in [49], with an attempt to justify the good performance of AugMix on C-10. Yet, we see that the fundamental transformations of AugMix are not enough to span a broad space warranting high perceptual similarity to IN-C. The significant difference in terms of perceptual similarity in Tab. 5.6 between AugMix and PRIME may explain the superior performance of PRIME on IN-100-C and IN-C (cf. Tab. 5.2).

<sup>3</sup>Examples of nearest neighbors can be found in App. E of [185].




Figure 5.4: Robustness vs. accuracy of a ResNet-18 (w/o JSD) on CIFAR-10 (left) and ImageNet-100 (right), when trained multiple times with PRIME. For each training instance, the transformation strength is scaled by  $\alpha$ . Note the different scale in axes.

### 5.4.3 Robustness vs Accuracy trade-off

An important phenomenon observed in the literature of adversarial robustness is the so-called robustness-accuracy trade-off [26, 27, 126], where technically adversarial training [25] with smaller perturbations (typically smaller  $\epsilon$ ) results in models with higher standard but lower adversarial accuracy, and vice versa. In this sense, we want to understand if the strength of the image transformations introduced through data augmentation can also cause such phenomenon in the context of robustness to common corruptions. As described in Sec. 5.2.2, each of the transformations of PRIME has a strength parameter  $\sigma$ , which can be seen as the analogue of  $\epsilon$  in adversarial robustness. Hence, we can easily reduce or increase the strength of the transformations by setting  $\hat{\sigma} = \alpha\sigma$ , where  $\alpha \in \mathbb{R}^+$ . Then, by training a network for different values of  $\alpha$  we can monitor its accuracy on the clean and the corrupted datasets.

We train a ResNet-18 on C-10 and IN-100 using the setup of Sec. 5.3.2. For reducing complexity, we do not use the JSD loss and we train for 30 epochs. This could cause some performance drop compared to the results of Tab. 5.2, but we expect the overall trends in terms of accuracy and robustness to be preserved. Regarding the scaling of the parameters' strength, for C-10 we set  $\alpha \in [10^{-3}, 10^2]$  and sample 100 values spaced evenly on a log-scale, while for IN-100 we set  $\alpha \in [10^{-2}, 10^2]$  and we sample 20 values.

The results are presented in Fig. 5.4. For both C-10 and IN-100, there is a sweet spot for the scale around  $\alpha = 0.2$  and  $\alpha = 1$  respectively, where the accuracy on common corruptions reaches its maximum. For  $\alpha$  smaller than these values, we observe a clear trade-off between validation and robust accuracy. While the robustness to common corruptions increases, the validation accuracy decays. However, for  $\alpha$  greater than the sweet-spot values, we observe that the trade-off ceases to exist since both the validation and robust accuracy present similar behaviour (slight decay). In fact, these observations indicate that robust and validation accuracies are not always positively correlated, and that one might have to slightly sacrifice validation accuracy in order to achieve robustness.




Figure 5.5: Accuracy of ResNet-18 (w/o JSD) on CIFAR-10 (left) and ImageNet-100 (right) when augmenting the training sets with additional PRIME counterparts off-line. Dashed lines denote the accuracy when training under the same setup, but generating the augmentations during training (on-line). Validation accuracy is omitted because it is rather constant: around 93.4% for CIFAR-10 and around 87% for ImageNet-100.

#### 5.4.4 Sample complexity

Finally, we investigate the necessity of performing augmentation during training (on-line augmentation), compared to statically augmenting the dataset before training (off-line augmentation). On the one hand, on-line augmentation is useful when the dataset is huge and storing augmented versions requires a lot of memory. Besides, there are cases where offline augmentation is not feasible as it relies on pre-trained or generative models which are unavailable in certain scenarios, *e.g.* DeepAugment [37] or AdA [38] cannot be applied on C-100. On the other hand, off-line augmentation may be necessary to avoid the computational cost of generating augmentations during training.

To this end, for each of the C-10 and IN-100 training sets, we augment them off-line with  $k = 1, 2, \dots, 10$  i.i.d. PRIME transformed versions. Afterwards, for different values of  $k$ , we train a ResNet-18 on the corresponding augmented dataset and report the accuracy on the validation set and the common corruptions. For the training setup, we follow the settings of Sec. 5.3.2, but without JSD loss. Also, since we increase the size of the training set by  $(k + 1)$ , we also divide the number of training epochs by the same factor, in order to keep the same overall number of gradient updates.

The performance on common corruptions is presented in Fig. 5.5. The first thing to notice is that, even for  $k = 1$ , the obtained robustness to common corruptions is already quite good. In fact, for IN-100 the accuracy (65%) is already better than the best achievable result of on-line AugMix (60.7% with JSD loss cf. Tab. 5.3). Regarding C-10 we observe that for  $k = 4$  the actual difference with respect to the on-line augmentation is almost negligible (88.8% vs. 89.3%), especially considering the overhead of transforming the data at every epoch. Technically, this means that augmenting C-10 with 4 PRIME counterparts is enough for achieving good robustness to common corruptions.

Finally, we also see in Fig. 5.5 that the corruption accuracy on IN-100 presents a very slow improvement after  $k = 4$ . Comparing the accuracy at this point (67.2%) to the one obtained with on-line augmentation and without JSD (68.8% cf. Tab. 5.3) we observe a gap of 1.6%. Hence, given the cost of on-line augmentation on such large scale datasets, simply augmenting the training with 4 extra PRIME samples presents a good compromise for achieving competitive robustness. Nevertheless, the increase of 1.6% introduced by on-line augmentation is rather significant, hinting that generating transformed samples during training might be necessary for maximizing performance. In this regard, the lower computational complexity of PRIME allows it to easily achieve this +1.6% through on-line augmentation, since it only requires  $1.27\times$  additional training time compared to standard training, and only  $1.06\times$  compared to AugMix, but with much better performance. This can be a significant advantage with respect to complex methods, like DeepAugment, that cannot be even applied on-line (require heavy pretraining).

## 5.5 Improving filling level classification with PRIME

So far we demonstrated that data augmentation with PRIME improves the robustness of classifiers to common corruptions of the data. In this section, we will see that PRIME can also be efficiently applied to off-the-self tasks, and improve robustness to other types of distribution shifts. In particular, we will focus again on the problem of estimating the filling level of a container introduced in Sec. 4.5, and we will show that PRIME can be easily tuned to effectively replace current transfer learning approaches, since it can significantly improve the generalization of the classifier.

The main limitation of transfer learning is that it requires the overhead of pre-training large models on very big datasets, while this overhead can further explode if these models are trained adversarially [170, 171]. At the same time, there is actually no guarantee that the transferred features are relevant for the target task, while the exact reason why transfer learning is expected to work is rather obscure. An alternative for increasing the variability of the training data is to perform data augmentation. However, generating additional samples that resemble the previously unseen properties of the test-time data can be quite challenging, since imposing such properties (i.e., changing shape) might require more complex and sophisticated operations, i.e, composition of transformations or mixing strategies [36]

In what follows, we will show that data augmentation using PRIME can improve significantly improve the filling level classification. We will demonstrate that PRIME can be easily tuned to generate augmentations that are tailored for learning classifiers that generalize on test-time images with containers of previously unseen shape, color, and spectral content. Using our scheme, we prevent the underlying classifier from overfitting to undesired features, achieving a filling level classification accuracy that is on-par or better than the one obtained with transfer learning. Yet, through a constructive approach

and without the need of pre-training on large datasets. Finally, we also show that the performance of the classifier may further increase when our data augmentation scheme is used in concert with transfer learning itself.

### 5.5.1 Distribution shifts and PRIME augmentations

Recall from Sec. 4.5 that in an attempt to empirically understand the dataset features that the classifier relies on, we introduced C-CCM, a dataset of cups and glasses that can be empty or filled with water, pasta or rice. The containers can vary a lot in terms of shape and transparency, while they can be captured under different illuminations, backgrounds and occlusions. We further defined three different training and validation splits, where a distribution shift is introduced in the validation set, such that some validation containers always have some property that does not exist in the training set (i.e., a unique shape, a stem, or color).

By measuring the network performance, we observed that the accuracy on the “shifted” containers was systematically lower than the one on containers that share similarities with those in the training set (overfitting). This is probably due to some biases related to the shape, color, and spectral content of the images, and that, with transfer learning, the target model becomes more invariant to features related to these properties. When the source model is also adversarially trained, then its stronger invariance to unnecessary features [159, 160] enables the target model to avoid irrelevant correlations and to rather identify features related to the actual task (i.e., filling level rather than shape).

To avoid both the computational cost of training robust models and the obscurity of transfer learning, it is worth investigating if the variability of the training set can be increased through data augmentation. When performing data augmentation, the choice of the basic transformations used to compose the augmentations is very important. They must be general and diverse enough to cover the desired changes to be induced. In the context of filling level classification, and for the overfitting phenomena discussed in Sec. 5.5.1, one should seek for image transformations that are sufficient to, i.e., change the shape of the container or its color/frequency content.

To this end, PRIME can generate diverse augmentations using a set of primitive max-entropy transformations on the spatial ( $\tau$ ), color ( $\gamma$ ) and spectral ( $\omega$ ) domain. For filling level classification, we expect these transformations to relate to the changes we want to introduce during training: container shape through  $\tau$ , container color through  $\gamma$ , and illumination and texture through  $\omega$ .

Recall from Sec. 5.3 that each RIME transformation has two control parameters:  $K$  for the smoothness and  $\sigma^2$  for the strength. PRIME synthesizes a transformed image  $\tilde{x}$  through a convex combination of  $n$  basic augmentations (width) consisting of the composition of  $m$  of its max-entropy transformations (depth). We make a small modification on the last




Figure 5.6: Sample transformed images using diffeomorphisms with varying smoothness level  $K_\tau$  (top row, from left to right: original image, transformed image with  $K_\tau = 2, 5, 10$ ; bottom row:  $K_\tau = 20, 40, 100, 300$ ).




Figure 5.7: Sample transformed images using color jittering with varying smoothness level  $K_\gamma$  (top row, from left to right: original image, transformed image with  $K_\gamma = 2, 5, 10$ ; bottom row:  $K_\gamma = 20, 40, 100, 300$ ).

step of Algorithm 3, such that the final image  $\hat{\mathbf{x}}$  is synthesized as a linear combination (mixing) of the original image  $\mathbf{x}$  and the transformed image  $\tilde{\mathbf{x}} = \sum_{i=1}^n \lambda_i \tilde{\mathbf{x}}_i$ , with the coefficients of the linear combination drawn from a Beta distribution

$$\hat{\mathbf{x}} = (1 - p) \mathbf{x} + p \tilde{\mathbf{x}} \quad \text{with} \quad p \sim \text{Beta}(\alpha, \beta). \quad (5.9)$$

Note that, for the shape parameters  $\alpha, \beta$  of the Beta distribution, when  $\alpha > \beta$  more importance is given to the pixels of the transformed image  $\tilde{\mathbf{x}}$ , while when  $\alpha < \beta$  more importance is given to the pixels of the original image  $\mathbf{x}$ .

In the next section, we will investigate if transfer learning can be replaced by a more controlled data augmentation strategy tailored for filling level classification. We will explore how to tune PRIME parameters for tackling the dataset-specific distribution shifts and improve the performance of the network on estimating the filling level.





Figure 5.8: Sample transformed images using spectral filtering with varying kernel size (smoothness level)  $K_\omega \times K_\omega$  (top row, from left to right: original image, transformed image with  $K_\omega = 3, 5, 7$ ; bottom row:  $K_\omega = 9, 11, 13, 15$ ).

### 5.5.2 PRIME transformation parameters

We must first identify the proper parameter values, given the dataset-specific shifts that arise for each validation split. We focus first on the splits  $S_1$  and  $S_2$  of the C-CCM dataset, where the shifts are mostly related to the shape of the containers. Hence, the most relevant transformation in PRIME is the one in the spatial domain. Intuitively, we would like to enforce smooth, yet strong, diffeomorphisms that are able to alter the shape of the whole container so it becomes as narrow as a champagne flute, or just a part of it so it resembles the stem of a cocktail glass (see Fig. 5.6). Recall, that, for a fixed value of smoothness  $K_\tau$  the authors in [192] propose to randomly sample the strength  $\sigma_\tau^2$  from a specific interval, such that the resulting diffeomorphism remains bijective. In practice, for smaller values of  $K_\tau$  (smoother), larger values of  $\sigma_\tau^2$  are allowed to be sampled. Hence, we decided to set  $K_\tau = 10$  and let  $\sigma_\tau^2$  to be properly sampled during training. In practice we observed that  $K_\tau \in [10, 20]$  still leads to good results.

For the split  $S_3$ , since we mostly deal with shifts related to the color and frequency content of the containers (i.e., red and green glass, which can lead to different textures and reflections), we focus on the color (see Fig. 5.7) and spectral (see Fig. 5.8) transforms in PRIME. For the smoothness parameter  $K$ , we keep the values proposed in Sec. 5.3.1:  $K_\gamma = 500$  and  $K_\omega = 3$  for the color and the spectral domain respectively. As for the parameter strength, since very strong changes could potentially destroy the information in the images, we decide to only slightly manipulate the color of the pixels and the frequency information of the images, and hence we set  $\sigma_\gamma^2 = 0.001$  and  $\sigma_\omega^2 = 0.01$  respectively.

### 5.5.3 Mixing parameters in PRIME

Regarding the mixing strategy, the width  $n$  specifies the number of transformed instances to be used in the convex combination for synthesizing the transformed image  $\tilde{\mathbf{x}}$ . It is reasonable to assume that we must set  $n > 1$  in order to increase the diversity of the

$m$	$S_1$	$S_2$	$S_3$
1	82.69	<b>73.42</b>	67.91
2	83.16	70.95	65.90
3	<b>84.93</b>	68.92	<b>75.03</b>

Table 5.7: Validation accuracy of a ResNet-18 on the three C-CCM dataset splits ( $S_1, S_2, S_3$ ), when the composition depth  $m$  of PRIME increases. Note that the transformation width is fixed to  $n = 3$ .




Figure 5.9: Effect of the Beta( $\alpha, \beta$ ) distribution on the validation accuracy of a ResNet-18, trained with PRIME, on each split of C-CCM. Note that, during the mixing step of PRIME,  $\alpha > \beta$  imposes more importance to the pixels of the transformed image, while  $\alpha < \beta$  to the pixels of the original image.

generated transformed instances, and hence we decide to use the default value ( $n = 3$ ).

The depth  $m$  specifies how many transformations will be sequentially applied on an image. In general, it is not always possible to determine what will be the exact outcome of such composition, and its impact on the overall performance. Hence, we let the mixing coefficient  $p$  of Eq. (5.9) to be uniformly sampled ( $\alpha = \beta = 1$ ) and perform a sensitivity analysis on the values of  $m$  (see Sec. 5.5.4 for training details). The performance of a ResNet-18 on each dataset split is shown in Tab. 5.7. We observe that for  $S_1$  and  $S_3$  increasing the depth significantly improves the performance. For  $S_2$  though we observe the opposite effect, indicating that applying multiple transformations on the image degrades some important information.

Finally, for the best performing values of  $m$  in Tab. 5.7, we explore the effect of the mixing coefficient  $p$  in equation Eq. (5.9). Specifically, we focus on the Beta distribution parameters, which control the relative importance of the pixels of  $\mathbf{x}$  or  $\tilde{\mathbf{x}}$ . Intuitively, since the classifier overfits to the training data we expect that more importance on  $\tilde{\mathbf{x}}$  is necessary. To that end, we perform a sensitivity analysis on  $\alpha$  and  $\beta$ , by measuring the performance of the network on their different combinations. Based on the results in Fig. 5.9, our initial intuition was right: on every dataset split, the highest validation accuracy is achieved when more importance is given to the pixels of the transformed image. In particular, on  $S_1$  and  $S_2$  the best performance (86.73% and 75.66% respectively) is achieved for Beta(5, 1), while on  $S_3$  (84.21%) it is achieved for Beta(6, 2).

#### 5.5.4 Experimental validation

We conduct experiments on C-CCM [161] using a ResNet-18 [157]. From the C-CCM pre-trained models provided by [161], we evaluate the ones trained with ST (baseline), ST→FT, and AT→FT, with the latter currently being the best one for classifying C-CCM. Furthermore, we train a model directly on C-CCM using data augmentations generated with PRIME ( $\text{DA}_{\text{PRIME}}$ ), and, finally, we also explore the combination of fine-tuning an adversarially trained model [170] with  $\text{DA}_{\text{PRIME}}$ . We denote this strategy as AT→ $\text{DA}_{\text{PRIME}}$ . We evaluate and compare the different methods on the different splits of C-CCM. Whenever PRIME is used, we use the parameters specified in Secs. 5.5.2 and 5.5.3. All model definitions and training procedures are implemented in PyTorch [158].

For  $\text{DA}_{\text{PRIME}}$  and AT→ $\text{DA}_{\text{PRIME}}$  strategies we train or fine-tune the classifier for 50 epochs, using a cross-entropy loss and stochastic gradient descent. The maximum learning rate for updating the weights is set to 0.05 and 0.005 when performing  $\text{DA}_{\text{PRIME}}$  and AT→ $\text{DA}_{\text{PRIME}}$  respectively. The learning rate decays linearly during training. Note that the models we evaluate are the ones that achieve the highest validation accuracy (early-stopping), while for dealing with class imbalances, the training images in a batch are randomly sampled with probabilities that are inversely proportional to the number of images of each class.

#### Classification results

We now evaluate and compare the different methods on the different splits of C-CCM. For the case of AT→ $\text{DA}_{\text{PRIME}}$ , since there are multiple source models adversarially trained with perturbations of different strength  $\varepsilon$ , we decided to choose those that lead to the highest validation accuracy. Hence, for  $S_1$  we select a network trained with  $\varepsilon = 0.05$ , while for  $S_2$  and  $S_3$  a network trained with  $\varepsilon = 0.5$ . Recall that, for the AT→FT models used in [161] the selected values of  $\varepsilon$  were 0.05, 1 and 0.5 for each dataset split respectively.

**Overall performance** Figure 5.10 shows the classification performance of different strategies on the three configurations,  $S_1$ ,  $S_2$  and  $S_3$ . The results indicate that pre-training might not be necessary: properly tuning data augmentation to compensate for the dataset-specific distribution shifts, improves performance with a lower computational cost than using transfer learning.  $\text{DA}_{\text{PRIME}}$  requires only  $1.2\times$  additional training time compared to ST on C-CCM, which is many orders of magnitude lower than (adversarially) training a model on ImageNet for using transfer learning. When training time is not an issue, transfer learning with AT at the source domain combined with  $\text{DA}_{\text{PRIME}}$  (AT→ $\text{DA}_{\text{PRIME}}$ ) generally improves performance. Note that when the performance of ST is low, all strategies lead to significant improvements; whereas when ST performs well, AT→FT has an insignificant contribution or decreases the final performance.

**Detailed analysis for each split** For  $S_1$ , the low performance of ST on the cham-

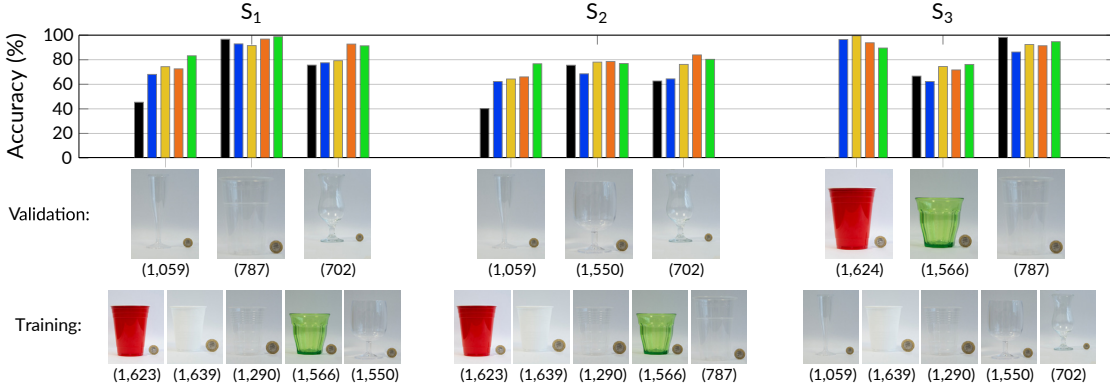


Figure 5.10: Per-container filling level classification accuracy (top) on the three dataset splits (bottom) of C-CCM. Parentheses denote the number of images for each type of container. Legend: — ST, — ST→FT, — AT→FT, —  $DA_{PRIME}$ , — AT→ $DA_{PRIME}$ .

pagne flute (left) is improved by both AT→FT and  $DA_{PRIME}$ , and even more so by AT→ $DA_{PRIME}$ , suggesting that diffeomorphisms compensate for the unique narrow shape of the flute. The accuracy of ST on the beer cup (middle) is high, due to the shape similarity of the small transparent cup in the training set. AT→FT causes a small accuracy drop, whereas  $DA_{PRIME}$  retains the performance and AT→ $DA_{PRIME}$  improves it. The accuracy of ST on the cocktail glass (right) is slightly improved with AT→FT and considerably improved by  $DA_{PRIME}$  and AT→ $DA_{PRIME}$ . Although there is another container with a stem in the training set (wine glass), it seems that the diffeomorphisms better compensate for the different shape above the stem of the cocktail glass.

For  $S_2$ , the accuracy of all strategies on the champagne flute (left) and the cocktail glass (right) is somehow similar in trend to that on  $S_1$ . Note that there are no containers with a stem in the training set. Yet, the performance on the wine glass (middle) is similar for most strategies, which might be due to the similarity of its shape above the stem with the other transparent cups in the training set.

For  $S_3$ , there is no colored container in the training set. ST is unable to generalize for the red cup (left), unlike AT→FT,  $DA_{PRIME}$  and AT→ $DA_{PRIME}$ . Still, the accuracy with data augmentation is not on the same level as with AT→FT, which sets this specific container case as an example of the benefits of adversarially pre-training the network on a large and diverse source dataset. As for the green glass (middle) AT→FT increases on ST, similarly to AT→ $DA_{PRIME}$ . Finally, the accuracy of ST on the beer cup (right) is high and the other strategies cannot reach that level, with AT→ $DA_{PRIME}$  featuring the lowest performance drop.

## 5.6 Discussion

**Coverage over the space of corruptions** In many parts of this chapter we implied that, for conferring robustness to common corruptions, a good augmentation method

should generate augmentations that cover a large space of possible corruptions. In general, formally identifying the space of semantic-preserving corruptions, and providing solid guarantees, is an utmost challenge. However, as with other problems in computer vision, we can rely on empirical proxies to gauge the coverage of an augmentation method. Specifically for PRIME, (i) its superior performance on multiple benchmarks (Tab. 5.2) and (ii) the quantitative study of Sec. 5.4.2 along with the SimCLR embedding distances of Tab. 5.6 suggest that PRIME achieves a broader coverage of the space of common corruptions than other methods. Additionally, the max-entropy principle formally guarantees a good coverage over the space of each of the three primitive transformations. Note here, that similar transformations to our three primitives are commonly used to model many types of image corruptions (e.g., color jitters, or lens artifacts).

The SimCLR embedding distances of Tab. 5.6 can be seen as a measure of “fitness” for the common corruption benchmark; that is, how similar/representative are the generated augmentations with respect to the actual corrupted images in the benchmark. Another measure of interest would be to investigate how diverse are the generated augmentations; that is, how large is the variance of the augmentations as measured in a given space (i.e., embedding space of a network). Hence, to qualitatively compare the diversity the augmentations of PRIME with respect to other methods, we can follow the procedure in [65]. We randomly select 3 images from ImageNet, each one belonging to a different class. For each image, we generate 100 transformed instances using AugMix and PRIME, while with DeepAugment we can only use the original images and the 2 transformed instances that are pre-generated with the EDSR and CAE image-to-image networks that DeepAugment uses. Then, we pass the transformed instances of each method through a ResNet-50 pre-trained on ImageNet and extract the features of its embedding space. On the features extracted for each method, we perform PCA and then visualize the projection of the features onto the first two principal components. We visualize the projected augmented space in Fig. 5.11, which suggests that PRIME, not only generates augmentations that fit better the benchmark of common corruptions (cf. SimCLR embedding distances), but that are also more diverse than AugMix and DeepAugment.

Furthermore, for understanding the potential of PRIME in covering a large space of the common corruptions, one can measure the capability of PRIME to create augmented images that are very similar to those in the benchmark. Since the transformations of PRIME are all differentiable, apart from the SimCLR embedding space distances, one can try to optimize the parameters of the transformations to minimize  $\|\mathbf{x}^t - \mathbf{x}^c\|_2$ ; that is, the distance between a PRIME-transformed image  $\mathbf{x}^t$  and the corresponding image from the corruption benchmark  $\mathbf{x}^c$ . Alternatively, one can focus on the mixing parameters and follow the same approach as in [65]: to find the worst-case (adversarial) mixing coefficients  $\lambda$  such that the mixed augmented images reach a part of the space that causes the most significant change on the classifier.




Figure 5.11: Projections of augmentations generated by different methods on the embedding space of a ResNet-50.

**Extensions** The general model of common corruptions of Eq. (5.1) can be extended with more transformations, as long as they follow the same principles as the ones we introduced (i.e., sampled from a max-entropy distribution). We already discussed the use of additive random noise  $\eta$ , which we do not include in our experiments due to its overlap with corruptions that exist in the evaluation benchmarks. Another possibility is to consider the transformations in the spatial domain. We chose to use diffeomorphisms, but one can also introduce random affine or projective transformations instead. Finally, our model can also easily accept other types of modifications, such as occlusions, or transformations based on task-specific priors.

## 5.7 Conclusions

In this chapter, we took a systematic approach to formulate a universal model that captures a wide variety of semantic-preserving image transformations. In particular, we defined a set of primitive max-entropy transformations in the spatial, color and spectral image domain, which guarantee that we provide the network with the maximum additional information possible about the domain of the transformation. Relying on this model, we proposed a novel data augmentation scheme called *PRIME*, which instantiates our model of corruptions, to confer robustness against common corruptions.

From a practical perspective, our method is principled yet efficient and can be conveniently incorporated into existing training procedures. Moreover, it yields a strong baseline on existing corruption benchmarks outperforming current standalone methods. Additionally, our thorough ablations demonstrate that diversity among basic augmentations (primitives) – which AugMix and other approaches lack – is essential, and that mixing plays a crucial role in the success of both prior methods and PRIME. In general, while complicated methods like DeepAugment perform well, it is difficult to understand, ablate and apply these online. Instead, we show that a simple model-based stance with a few guiding

principles can be used to build an efficient augmentation scheme that can be easily understood, ablated and tuned.

We believe that our insights and PRIME pave the way for building robust models in real-life scenarios. For instance, we demonstrated that PRIME already provides a ready-to-use recipe for conferring robustness in the problem of filling level classification. We believe that our method can be easily adapted to other for data-scarce domains such as medical imaging.





# 6 Conclusion

## 6.1 Summary

In this thesis, we provided novel algorithmic tools, a large a set of new insights and experimental evidences, as well as constructive solutions for understanding and improving the robustness and invariance properties of image classifiers. Our tools allowed us to efficiently evaluate the robustness of deep networks to non-standard adversarial perturbation regimes, as well as to establish strong connections between the data features, the geometric properties of deep classifiers, and their inductive biases. Furthermore, our insights enabled us to further design principled methods for analyzing and improving the robustness of deep classifiers to different types of distribution shifts in multiple tasks.

We first studied methods for measuring the robustness of deep classifiers to sparse additive adversarial perturbations. For avoiding the NP-hardness of minimizing the  $\ell_0$  norm in computing sparse perturbations, we focused on finding an efficient relaxation method. To this end, we exploited the low mean curvature of the decision boundaries in the vicinity of the data samples and designed an iterative method that we coin SparseFool. At each iteration SparseFool performs a linear approximation of the decision boundary and solves the simpler  $\ell_1$  box-constrained problem. Our method computes very sparse perturbations, is by orders of magnitude faster than existing methods, and can easily scale to high-dimensional datasets. By visually inspecting the generated adversarial examples, we observed that SparseFool altered features that are shared among different images, and that, in many cases, the perturbations resembled image features that are correlated with the fooling class. Such observation suggested that adversarial perturbations might not necessarily be a “hole” in the system, but they might actually reflect some strong connection/correlation between the features of the dataset and the features that the networks use for taking their decisions.

Then, we studied this connection in-depth, and we proposed a new geometric framework that permits to relate the features of a dataset with the distance to the decision boundary

along specific directions. Through a series of carefully designed experiments, both in synthetic and real image datasets, we explained how the inductive bias of the learning algorithm shapes the decision boundaries of neural networks by creating boundaries that are invariant to non-discriminative directions. Furthermore, we demonstrated that the decision boundaries are very sensitive to the position of the training samples, and that small changes along specific directions can cause large and sudden changes in orthogonal ones. In fact, adversarial training exploits this sensitivity, as well as the inductive bias towards invariance, in order to shape the boundaries and build robust classifiers. Interestingly, modifying the position of just a minimal number of training samples during adversarial training is enough to locally introduce excessive invariance. In general, our framework can be used to identify parts of the input space that are important for the classifier, to understand intriguing properties of deep networks, such as their catastrophic forgetting, and to design stronger defenses and black-box attacks that exploit small discriminative subspaces. Furthermore, we studied the implications that the invariance properties of robust classifiers have in the downstream, data-scarce application of classifying the filling level within containers. In particular, we showed that, during transfer learning, using robust models in the source domain permits to consistently improve generalization to unseen containers in the target domain that come from a shifted distribution, e.g., containers of unseen shape or color.

However, adversarial robustness is not necessarily the best way to improve the robustness of classifiers to certain distribution shifts, such as common corruptions of the data. Hence, we took a systematic approach to understand the notion of common corruptions and formulated a universal model that captures a wide variety of semantic preserving, common image transformations. In particular, we defined a set of primitive max-entropy transformations in the spatial, color and spectral image domains, in order to guarantee that, with each new augmentation, we provide the network with the maximum additional information possible about the domain of the transformation. Relying on this model, we proposed a novel data augmentation scheme called *PRIME*, which instantiates our semantic-preserving transformations, to confer robustness against common corruptions. From a practical perspective, our method is principled yet efficient and can be conveniently incorporated into existing training procedures. Moreover, it yields a strong baseline on existing corruption benchmarks and outperforms current similar methods. Additionally, our thorough ablations demonstrate that diversity among basic augmentations (primitives) – which AugMix and other approaches lack – is essential, and that mixing plays a crucial role in the success of both prior methods and PRIME. Our findings also underlined the benefits of using PRIME in on-line augmentation settings, due to its scalability and efficiency. Finally, we demonstrated that PRIME already provides a ready-to-use recipe for conferring robustness in the data-scarce problem of filling level classification. PRIME can be easily tuned to generate samples, which are tailored for building classifiers that generalize on images with containers of unknown shape, color, and spectral content.

## 6.2 Future directions

First, we mostly studied the geometric and invariance characteristics of the decision boundaries through the lens of the data features. However, we did not investigate other factors that are responsible for shaping the decision boundaries, such as the architecture. From our experiments we can observe that, when the frequency representation of the data is flipped, the margin distribution is not an exact mirroring of the original one, and that the margin along low-frequency directions cannot reach excessive levels. In fact, this triggered some recent research, which demonstrated that the architecture itself has a very strong inductive bias towards specific input directions, which are typically aligned with low-frequency directions. Such directions are coined Neural Anisotropy Directions (NADS) [53], and can shed new light onto the different types of inductive biases that deep networks have, such as the one towards simple solutions [54]. Developing methodologies for revealing, understanding, and controlling the inductive bias is very important, since it can have strong implications in the generalization and robustness properties of deep networks [199, 200].

On the practical side, the invariance insights of our geometric analysis have already been exploited in [201] to control the excessive margin caused by adversarial training, and eventually balance the robustness/accuracy tradeoff. Furthermore, in our experiments we used the DCT basis. However, such basis is independent of the data and the training process. Incorporating other task- or data-related priors to measure the geometric properties or manipulate the data is quite important. This could have significant application in fairness-sensitive tasks for identifying (or eliminating) the undesired bias from the model. In addition, our framework could also be extended beyond image classification, for example towards natural language processing tasks.

Then, we have seen that formally identifying the space of common corruptions is very difficult and, probably, ill-posed. Currently, one can only follow qualitative or empirical proxies to measure the efficacy of a method, i.e., through accuracy or through the features of the embedding space (i.e., measure similarity or visualize projections on principal feature directions). Nevertheless, it is important to analytically characterize this space, even partially, since it will enable the development of more fundamental methods that provide a good coverage over that space. Additionally, the same principles for identifying the space of corruptions can be followed for characterizing other types of manipulations or distribution shifts [37, 202].

In this regard, and from a practical perspective, it is worth investigating how the primitives of PRIME can be extended for distributions shifts beyond common corruptions, but also to downstream tasks and applications. For instance, it is interesting to explore the applicability of PRIME on semantic segmentation and object detection, or even speech recognition, but also in safety-critical applications, i.e., autonomous driving, or data-scarce domains, i.e., medical imaging.



# A Appendix of Chapter 4

## A.1 Margin distribution of a linear classifier

In this section we demonstrate that even for linear classifiers trained on  $\mathcal{T}_1(\varepsilon, \rho, N)$  the distribution of margins along non-discriminative directions will never be infinite. We demonstrate this effect in practice by repeating the experiment of Sec. 4.2.2, where instead of an MLP we use a simple logistic regression (see Tab. A.1). Clearly, although the values along  $\text{span}\{\mathbf{u}_1\}^\perp$  are quite large, they are still finite. This demonstrates that due to the finiteness of the training set and its high-dimensionality the influence of the non-discriminative directions in the final solution is significant.

	$\mathbf{u}_1$	$\text{span}\{\mathbf{u}_1\}^\perp$	$\mathcal{S}_{\text{ORTH}}$	$\mathcal{S}_{\text{RAND}}$
5-PERC.	2.39	36.7	184.95	11.57
MEDIAN	2.49	38.3	192.98	12.08
95-PERC.	2.60	39.9	201.16	12.59

Table A.1: Margin statistics of a logistic regressor trained on  $\mathcal{T}_1(\varepsilon = 5, \sigma = 1)$  along different directions ( $N = 10,000$ ,  $M = 1,000$ ,  $S = 3$ ).

## A.2 Training parameters

Table A.2 shows the performance and training parameters of the different networks used in the paper. Note that the hyperparameters of these networks were not optimized in any form during this work. Instead they were selected from a set of best practices from the DAWNBench submissions that have been empirically shown to give a good trade-off in terms of convergence speed and performance. In this sense, especially for the non-standard datasets (e.g., “flipped” datasets), the final performance might not be the best reflection of the highest achievable performance of a given architecture. In fact, since the goal of our experiments is not to achieve the most robust models on such non-standard datasets, but rather investigate how the previously observed trends are

## Appendix A. Appendix of Chapter 4

represented in these new classifiers, no further hyperparameter tuning was applied. All the experiments with synthetic data were trained in the same way, namely using SGD with a linearly decaying learning rate (max lr. 0.1), no explicit regularization, and trained for 500 epochs.

DATASET	NETWORK	TEST ACC.	EPOCHS	LR SCHEDULE	MAX. LR	BATCH
MNIST	LeNET	99.35%	30	TRIANG.	0.21	128
	RESNET-18	99.53%				
MNIST FLIPPED	LeNET	99.34%	30	TRIANG.	0.21	128
	RESNET-18	99.52%				
CIFAR-10	VGG-19	89.39%	50	TRIANG.	0.21	128
	RESNET-18	90.05%				
	DENSENET-121	93.03%				
CIFAR-10 LOW PASS	VGG-19	84.81%	50	TRIANG.	0.21	128
	RESNET-18	84.77%				
	DENSENET-121	88.51%				
CIFAR-10 FLIPPED	VGG-19	87.42%	50	TRIANG.	0.21	128
	RESNET-18	88.67%				
	DENSENET-121	91.19%				
IMAGENET	VGG-16	71.59%	–	–	–	–
	RESNET-50	76.15%				
	DENSENET-121	74.65%				
IMAGENET FLIPPED	RESNET-50	68.12%	90(68)	PIECEWISE CONSTANT	0.1	256

Table A.2: Performance and training parameters of multiple networks trained on different datasets. All networks have been trained using SGD with momentum 0.9 and a weight decay of  $5 \times 10^{-4}$ . For ImageNet, we use the pretrained models from PyTorch. For “flipped” ImageNet, the weight decay was set to  $10^{-4}$ , while for computational reasons the training was executed until the 68<sup>th</sup> epoch.

### A.3 Cross-dataset performance

We now show the performance of different networks trained with different variants of the standard computer vision datasets and tested on the rest.

		MNIST	MNIST FLIPPED	MNIST HIGH PASS
MNIST	LeNET	99.35%	18.73%	44.09%
	RESNET-18	99.53%	11.88%	15.73%
MNIST FLIPPED	LeNET	10.52%	99.34%	9.87%
	RESNET-18	16.59%	99.52%	11.23%
MNIST HIGH PASS	LeNET	96.35%	42.36%	98.65%
	RESNET-18	88.38%	21.48%	98.71%

Table A.3: Networks trained on a specific version of MNIST, but evaluated on different variations of it. Rows denote the dataset that a network is trained on, and columns the dataset they are evaluated on. Values on the diagonal correspond to the same variation.

		CIFAR-10	CIFAR-10 FLIPPED	CIFAR-10 LOW PASS
CIFAR-10	VGG-19	89.39%	10.63%	61.4%
	RESNET-18	90.05%	10%	46.99%
	DENSENET-121	93.03%	10.3%	27.45%
CIFAR-10 FLIPPED	VGG-19	10.77%	87.42%	10.79%
	RESNET-18	9.91%	88.67%	9.97%
	DENSENET-121	9.98%	91.19%	10%
CIFAR-10 LOW PASS	VGG-19	85.16%	10.52%	84.81%
	RESNET-18	85.47%	10.45%	84.77%
	DENSENET-121	89.67%	10.45%	88.51%

Table A.4: Networks trained on a specific version of CIFAR-10, but evaluated on different variations of it. Rows denote the dataset that a network is trained on, and columns the dataset they are evaluated on. Values on the diagonal correspond to the same variation.

		IMAGENET	IMAGENET FLIPPED
IMAGENET	VGG-16	71.59%	0.106%
	RESNET-50	76.15%	0.292%
	DENSENET-121	74.65%	0.22%
IMAGENET FLIPPED	RESNET-50	0.184%	68.12%

Table A.5: Networks trained on a specific version of ImageNet, but evaluated on different variations of it. Rows denote the dataset that a network is trained on, and columns the dataset they are evaluated on. Values on the diagonal correspond to the same variation.

## A.4 Adversarial training parameters

Table A.6 shows the performance and adversarial training parameters of the different networks used in the paper. Note that the hyperparameters of these networks were not optimized in any form during this work. Instead they were selected from a set of best practices from the DAWNBench submissions that have been empirically shown to give a good trade-off in terms of convergence speed and performance. Again, as stated in App. A.2, especially for the non-standard datasets (e.g., “flipped” datasets), the final performance might not be the best reflection of the highest achievable performance or robustness of a given architecture, since no further hyperparameter tuning was applied.

DATASET	NETWORK	STANDARD TEST ACC.	ADV. TEST ACC.	EPOCHS	$\ell_2$ BALL RADIUS	STEPS
MNIST	LENET	98.32%	76.01%	25	2	7
	RESNET-18	98.89%	80.26%			
MNIST FLIPPED	LENET	98.29%	74.68%	25	2	7
	RESNET-18	98.75%	81.97%			
CIFAR-10	VGG-19	73.76%	50.15%	50	1	7
	RESNET-18	82.20%	52.38%			
	DENSENET-121	82.90%	54.86%			
CIFAR-10 FLIPPED	VGG-19	71.39%	35.64%	50	1	7
	RESNET-18	73.64%	37.24%			
	DENSENET-121	78.32%	42.32%			
IMAGENET	RESNET-50	57.90%	35.16	–	3	20

Table A.6: Performance and attack parameters of multiple networks adversarially trained using  $\ell_2$ -PGD. The training parameters are similar to the ones of Tab. A.2. For ImageNet we use the adversarially trained ResNet-50 provided by [203].

## A.5 Spectral decomposition on frequency “flipped” data

Following the results presented in Sec. 4.4.2, we now show in Fig. A.1 the spectral decomposition of the adversarial perturbations crafted during adversarial training for the frequency “flipped” CIFAR-10 dataset on a DenseNet-121 network. In contrast to the spectral decomposition of the perturbations on CIFAR-10 (left), the energy of the frequency “flipped” CIFAR-10 perturbations (right) remains concentrated in the high part of the spectrum during the whole training process, and has hardly any presence in the low frequencies. In other words, the frequency content of the  $\ell_2$ -PGD adversarial perturbations also “flips”.



## A.5 Spectral decomposition on frequency “flipped” data




Figure A.1: Energy decomposition in subspaces of the DCT diagonal of adversarial perturbations used during adversarial training ( $\ell_2$  PGD with  $\varepsilon = 1$ ) on 1,000 (a) CIFAR-10 and (b) frequency “flipped” CIFAR-10 training samples per epoch for a DenseNet-121. The plot shows 95-percentile of energy.



# B Appendix of Chapter 6

## B.1 Maximum entropy transformations

To guarantee as much diversity as possible in our model of common corruptions, we follow the principle of maximum entropy to define our distributions of transformations [188]. Note that using a set of augmentations that guarantees maximum entropy comes naturally when trying to optimize the sample complexity derived from certain information theoretic generalization bounds, both in the clean [189] and corrupted setting [190]. Specifically, the principle of maximum entropy postulates favoring those distributions that are as unbiased as possible given the set of constraints that defines a family of distributions. In our case, these constraints are given in the form of an expected strength, i.e.,  $\sigma^2$ , desired smoothness, i.e.,  $K$ , and/or some boundary conditions, *e.g.*, the displacement field must be zero at the borders of an image.

Let us make this formal. In particular, let  $\mathcal{I}$  denote the space of all images  $\mathbf{x} : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ , and let  $f : \mathcal{I} \rightarrow \mathcal{I}$  denote a random image transformation distributed according to the law  $\mu$ . Further, let us define a set of constraints  $\mathcal{C} \subseteq \mathcal{F}$ , which restrict the domain of applicability of  $f$ , i.e.,  $f \in \mathcal{C}$ , and where  $\mathcal{F}$  denotes the space of functions  $\mathcal{I} \rightarrow \mathcal{I}$ . The principle of maximum entropy postulates using the distribution  $\mu$  which has maximum entropy given the constraints:

$$\begin{aligned} \underset{\mu}{\text{maximize}} \quad & H(\mu) = \int_{\mathcal{F}} d\mu(f) \log(\mu(f)) \\ \text{subject to} \quad & f \in \mathcal{C} \quad \forall f \sim \mu, \end{aligned} \tag{B.1}$$

where  $H(\mu)$  represents the entropy of the distribution  $\mu$  [188]. In its general form, solving Eq. (B.1) for any set of constraints  $\mathcal{C}$  is intractable. However, leveraging results from statistical physics, we will see that for our domains of interest, Eq. (B.1) has a simple solution. In what follows we derive those distributions for each of our family of transformations.

### B.1.1 Spectral domain

As we introduced in Sec. 5.2.2, we propose to parameterize our family of spectral transformations using an FIR filter of size  $K_\omega \times K_\omega$ . That is, we are interested in finding a maximum entropy distribution over the space of spectral transformations with a finite spatial support.

Nevertheless, on top of this smoothness constraint we are also interested in controlling the strength of the transformations. We define the strength of a distribution of random spectral transformations applied to an image  $\mathbf{x}$ , as the expected  $L_2$  norm of the difference between the clean and transformed images, i.e.,

$$\mathbb{E}_\omega \|\mathbf{x} - \omega(\mathbf{x})\|_2^2 = \mathbb{E}_{\omega'} \|\omega' * \mathbf{x}\|_2^2, \quad (\text{B.2})$$

which using Young's convolution inequality is bounded as

$$\mathbb{E}_{\omega'} \|\omega' * \mathbf{x}\|_2^2 \leq \|\mathbf{x}\|_1^2 \mathbb{E}_{\omega'} \|\omega'\|_2^2. \quad (\text{B.3})$$

Indeed, we can see that the strength of a distribution of random smooth spectral transformations is governed by the expected norm of its filter. In the discrete domain, this can be simply computed as

$$\mathbb{E}_{\omega'} \|\omega'\|_2^2 = \sum_{i=1}^{K_\omega} \sum_{j=1}^{K_\omega} \mathbb{E}_{\omega'} \omega_{i,j}'^2. \quad (\text{B.4})$$

Considering this, we should then look for a maximum entropy distribution whose samples satisfy

$$\mathcal{C} = \left\{ \omega' \in \mathbb{R}^{K_\omega \times K_\omega} \wedge \mathbb{E}_{\omega'} \|\omega'\|_2^2 = K_\omega^2 \sigma_\omega^2 \mid \omega \sim \mu_\omega \right\}. \quad (\text{B.5})$$

Now, note that this set is defined by an equality constraint involving a sum of  $K_\omega^2$  quadratic random variables. In this sense, we know that the Equipartition Theorem [191] applies and can be used to identify the distribution of maximum entropy. That is, the solution of Eq. (B.1) in the case that  $\mathcal{C}$  is given by Eq. (B.5), is equal to the distribution of FIR filters whose coefficients are i.i.d. with law  $\mathcal{N}(0, \sigma_\omega^2)$ .

### B.1.2 Spatial domain

The distribution of diffeomorphisms of maximum entropy with a fixed norm was derived by Petrini *et al.* in [192]. The derivation is similar to the spectral domain, but with the additional constraint that the diffeomorphisms produce a null displacement at the borders of the image.

### B.1.3 Color domain

We can follow a very similar route to derive the distribution of maximum entropy among all color transformations, where, specifically, we constraint the transformations to yield  $\gamma(0) = 0$  and  $\gamma(1) = 1$  on every channel independently. Doing so, the derivation of the maximum entropy distribution can follow the same steps as in [192].

## B.2 PRIME implementation details

In this section, we provide additional details regarding the implementation of PRIME described in Sec. 5.3. Since the parameters of the transformations are empirically selected, we first provide more visual examples for different values of smoothness  $K$  and strength  $\sigma$ . Then, we give the exact values of the parameters we use in our experiments supported by additional visual examples.

### B.2.1 Additional transformed examples

Here we provide additional visual examples for each of the primitives of PRIME illustrating the effect of the following two factors: (i) smoothness controlled by parameter  $K$ , and (ii) strength of the transformation  $\sigma$  on the resulting transformed images created by the primitives. Figs. B.1, B.2 and B.3 demonstrate the resulting spectrum of images created by applying spectral, spatial and color transformations while varying the parameters  $K$




Figure B.1: Example images (IN-100) generated with spectral transformations from our common corruptions model. In each row, we enlarge the transformation strength  $\sigma_\omega$  from left to right. From top to bottom, we increase the spectral resolution of the filter  $K_\omega$ .




Figure B.2: Example images (IN-100) generated with spatial transformations from our common corruptions model. In each row, we enlarge the transformation strength  $\sigma_\tau$  from left to right. From top to bottom, we increase the cut frequency  $K_\tau$ .




Figure B.3: Example images (IN-100) generated with color transformations from our common corruptions model. In each row, we enlarge the transformation strength  $\sigma_\gamma$  from left to right. From top to bottom, we increase  $K_\gamma$ .

and  $\sigma$ . Notice how increasing the strength  $\sigma$  of each transformation drifts the augmented image farther away from its clean counterpart, yet produces plausible images when appropriately controlled.

### B.2.2 Transformation parameters

**Spectral transform** Regarding the spectral transform of Eq. (5.3) we found out that, for the FIR filter  $\omega'$ , a size of  $K_\omega = 3$  results into semantically preserving images for CIFAR-10/100 and ImageNet. For the latter, one can stretch the filter size to  $5 \times 5$  or even  $7 \times 7$ , but then slight changes on the strength,  $\sigma_\omega$ , might destroy the image semantics. Eventually, given  $K_\omega = 3$ , we observed that  $\sigma_\omega = 4$  is good enough for CIFAR-10/100 and ImageNet.

**Spatial transform** Concerning the spatial transform of Eq. (5.5), for the cut-off parameter  $K_\tau$  we followed the value regimes proposed by Petrini *et al.* [192] and set  $K_\tau = 100$  for CIFAR-10/100;  $K_\tau = 500$  for ImageNet. Furthermore, for a given  $K_\tau$ , Petrini *et al.* also compute the appropriate bounds for the transformation strength,  $\sigma_{\tau_{\min}}^2 \leq \sigma_\tau^2 \leq \sigma_{\tau_{\max}}^2$ , such that the resulting diffeomorphism remains bijective and the pixel displacement does not destroy the image. In fact, in their original implementation, which can be found at <https://github.com/pcsl-epfl/diffeomorphism>, Petrini *et al.* directly sample  $\sigma_\tau \sim U(\sigma_{\tau_{\min}}, \sigma_{\tau_{\max}})$  instead of explicitly setting the strength. In our implementation, we also follow the same approach.

**Color transform** Regarding the color transform of Eq. (5.6) we found out that for CIFAR-10/100 a cut-off value of  $K_\gamma = 10$  and a strength of  $\sigma_\gamma = 0.01$  result into semantically preserving images for CIFAR-10/100; while for ImageNet, the corresponding values are  $K_\gamma = 500$  and  $\sigma_\gamma = 0.05$ . As for the bandwidth (consecutive frequencies)  $\Delta$  we observed that a value of  $\Delta = 20$  was memory sufficient for ImageNet, but for CIFAR-10/100 we can even afford all the frequencies to be used, e.g.,  $\Delta = K_\gamma$ .

Finally, as mentioned in Sec. 5.3, we randomly sample the strength of the transformations  $\sigma$  from a uniform distribution of given minimum and maximum values. Regarding the maximum, we always set it to be the one we selected through visual inspection, while the minimum is set to 0.

## B.3 Performance per corruption

Beyond the average corruption accuracy that we report in Tab. 5.2, we also provide here the performance of each method on the individual corruptions. The results on

Dataset	Method	Clean	CC	Noise			Blur				Weather				Digital			
				Gauss.	Shot	Impulse	Defoc.	Glass	Motion	Zoom	Snow	Frost	Fog	Bright.	Contr.	Elastic	Pixel.	JPEG
C-10	Std.	95.0	74.0	45.1	58.7	54.9	83.2	53.3	76.9	79.1	83.1	79.3	89.0	93.6	76.3	83.9	75.1	77.9
	AugMix	95.2	88.6	79.3	84.8	85.8	94.1	78.9	92.4	93.4	89.7	89.0	91.9	94.3	90.5	90.5	87.6	87.5
	PRIME	94.2	89.8	86.9	88.1	88.6	92.6	85.3	90.8	92.2	89.3	90.5	89.8	93.7	92.4	90.1	88.1	88.8
C-100	Std.	76.7	51.9	25.3	33.7	26.6	60.8	47.1	55.5	57.6	60.8	56.2	62.5	72.2	53.2	63.4	50.1	52.7
	AugMix	78.2	64.9	46.7	55.1	60.6	76.2	47.3	72.6	74.3	67.4	64.4	69.9	75.5	67.4	69.6	64.9	61.8
	PRIME	78.4	68.2	59.0	62.1	68.1	74.0	58.3	70.5	72.3	68.9	68.5	69.8	76.8	74.4	70.1	65.5	64.4

Table B.1: Per-corruption accuracy of different methods on C-10 and C-100 (ResNet-18).

## Appendix B. Appendix of Chapter 6

Dataset	Method	Clean		Noise			Blur				Weather				Digital			
			CC	Gauss.	Shot	Impulse	Defoc.	Glass	Motion	Zoom	Snow	Frost	Fog	Bright.	Contr.	Elastic	Pixel.	JPEG
IN-100	Standard	88.0	49.7	30.9	29.0	22.0	45.6	44.6	50.4	53.9	43.8	46.2	50.5	78.6	42.9	68.8	68.0	70.6
	AugMix	88.7	60.7	45.2	45.8	43.4	58.7	53.3	69.5	71.0	49.1	52.7	60.2	80.7	59.6	73.3	73.6	74.7
	DA	86.3	67.7	76.3	75.6	75.7	64.2	61.7	61.3	62.7	54.4	62.8	55.7	81.6	49.7	69.9	83.3	80.6
	PRIME	85.9	71.6	80.6	80.0	80.1	57.2	66.3	66.2	68.2	61.5	68.2	57.2	81.2	68.3	73.7	82.9	81.9
	DA+AugMix	86.5	73.1	75.2	75.8	74.9	74.1	68.5	76.0	72.1	59.9	66.8	61.4	82.1	72.4	73.1	83.8	81.1
	DA+PRIME	84.9	74.9	81.1	80.9	81.2	70.5	74.2	72.0	71.5	66.3	73.6	56.6	81.9	72.8	74.8	83.4	82.3
IN	Standard*	76.1	39.2	29.3	27.0	23.8	38.8	26.8	38.7	36.2	32.5	38.1	45.4	68.0	39.0	45.3	44.8	53.4
	AugMix*	77.5	48.3	40.6	41.1	37.7	47.7	34.9	53.5	49.0	39.9	43.8	47.1	69.5	51.1	52.0	57.0	60.3
	DA*	76.7	52.6	56.6	54.9	56.3	51.7	40.1	48.7	39.5	44.2	50.3	52.1	71.1	48.3	50.9	65.5	59.3
	PRIME†	77.0	55.0	61.9	60.6	60.9	47.6	39.0	48.4	46.0	47.4	50.8	54.1	71.7	58.2	56.3	59.5	62.2
	DA+AugMix	75.8	58.1	59.4	59.6	59.1	59.0	46.8	61.1	51.5	49.4	53.3	55.9	70.8	58.7	54.3	68.8	63.3
	DA+PRIME†	75.5	59.9	67.4	67.2	66.8	56.2	47.5	54.3	47.3	52.8	56.4	56.3	71.7	62.3	57.3	70.3	65.1

Table B.2: Per-corruption accuracy of different methods on IN-100 (ResNet-18) and IN (ResNet-50). † indicates that JSD consistency loss is not used. \*Models taken from RobustBench [193].

CIFAR-10/100 and ImageNet/ImageNet-100 are shown on Tabs. B.1 and B.2 respectively. Compared to AugMix on CIFAR-10/100, the improvements from PRIME are mostly observed against Gaussian noise (+7.6%/12.3%), shot noise (+3.3%/7.0%), glass blur (+6.4%/11.0%) and JPEG compression (+1.3%/2.6%). These results show that PRIME can really push the performance against certain corruptions in CIFAR-10/100-C despite the fact that AugMix is already good on these datasets. However, AugMix turns out to be slightly better than PRIME against impulse noise, defocus blur and motion blur modifications; all of which have been shown to be resembled by AugMix created images. With ImageNet-100, PRIME enhances the diversity of augmented images, and leads to general improvements against all corruptions except certain blurs. On ImageNet, we observe that, in comparison to DeepAugment, the supremacy of PRIME is reflected on almost every corruption type, except some blurs and pixelate corruptions where DeepAugment is slightly better. When PRIME is used in conjunction with DeepAugment, compared to AugMix combined with DeepAugment, our method seems to lack behind only on blurs, while on the rest of the corruptions achieves higher robustness.

### B.4 Performance per severity level

We also want to investigate the robustness of each method on different severity levels of the corruptions. The results for CIFAR-10/100 and ImageNet/ImageNet-100 are presented in Tabs. B.3 and B.4 respectively. With CIFAR-10/100, PRIME predominantly helps against corruptions with maximal severity and yields +3.9% and +7.1% gains on CIFAR-10 and CIFAR-100 respectively. Besides on ImageNet-100, PRIME again excels at corruptions with moderate to higher severity. This observations also holds when PRIME is employed in concert with DeepAugment. With ImageNet too this trend continues, and we observe that, compared to DeepAugment, PRIME improves significantly on corruptions of larger severity (+3.4% and +5.5% on severity levels 4 and 5 respectively). Also, this behaviour is consistent even when PRIME is combined with DeepAugment and is compared to DeepAugment+AugMix, where we see that again on levels 4 and 5 there is a significant improvement of +2.1% and +3.7% respectively.



## B.4 Performance per severity level

Dataset	Method	Clean	CC Avg.	Severity				
				1	2	3	4	5
C-10	Standard	95.0	74.0	87.4	81.7	75.7	68.3	56.7
	AugMix	95.2	88.6	93.1	91.8	89.9	86.7	81.7
	PRIME	94.2	89.8	92.8	91.6	90.4	88.6	85.6
C-100	Standard	76.7	51.9	66.7	59.4	52.8	45.0	35.4
	AugMix	78.2	64.9	73.3	70.0	66.6	61.3	53.4
	PRIME	78.4	68.2	74.0	71.6	69.2	65.6	60.5

Table B.3: Average accuracy for each corruption severity level of different methods on C-10 and C-100 (ResNet-18).

Dataset	Method	Clean	CC Avg.	Severity				
				1	2	3	4	5
IN-100	Standard	88.0	49.7	73.5	61.0	49.8	37.2	27.0
	AugMix	88.7	60.7	80.4	71.8	63.8	50.3	37.2
	DA	86.3	67.7	81.2	75.4	69.9	61.2	50.8
	PRIME	85.9	71.6	81.7	77.5	73.4	66.9	58.4
	DA+AugMix	86.5	73.1	82.7	78.0	75.5	69.6	59.9
	DA+PRIME	84.9	74.9	82.0	78.7	76.4	71.8	65.5
IN	Standard*	76.1	39.2	60.6	49.8	39.8	27.7	18.0
	AugMix*	77.5	48.3	66.7	58.3	51.1	39.1	26.5
	DA*	76.7	52.6	69.0	61.7	55.4	44.9	32.1
	PRIME <sup>†</sup>	77.0	55.0	68.9	63.1	56.9	48.3	37.6
	DA+AugMix	75.8	58.1	70.3	64.5	60.5	53.0	42.2
	DA+PRIME <sup>†</sup>	75.5	59.9	70.8	66.3	61.6	55.1	45.9

Table B.4: Average accuracy for each corruption severity level of different methods on IN-100 (ResNet-18) and IN (ResNet-50). <sup>†</sup> indicates that JSD consistency loss is not used. \*Models taken from **RobustBench** [193].



# Bibliography

- [1] D. Hendrycks and T. Dietterich, “Benchmarking neural network robustness to common corruptions and perturbations,” in *International Conference on Learning Representations (ICLR)*, May 2019.
- [2] A. Fawzi and P. Frossard, “Manitest: Are classifiers really invariant?” in *Proceedings of the British Machine Vision Conference (BMVC)*, Sep. 2015.
- [3] D. Yin, R. G. Lopes, J. Shlens, E. Cubuk, and J. Gilmer, “A Fourier perspective on model robustness in computer vision,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2019.
- [4] K. Y. Xiao, L. Engstrom, A. Ilyas, and A. Madry, “Noise or signal: The role of image backgrounds in object recognition,” in *International Conference on Learning Representations (ICLR)*, May 2021.
- [5] B. Recht, R. Roelofs, L. Schmidt, and V. Shankar, “Do ImageNet classifiers generalize to ImageNet?” in *International Conference on Machine Learning (ICML)*, Jun. 2019.
- [6] D. Hendrycks, K. Zhao, S. Basart, J. Steinhardt, and D. Song, “Natural adversarial examples,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2021.
- [7] R. Taori, A. Dave, V. Shankar, N. Carlini, B. Recht, and L. Schmidt, “Measuring robustness to natural distribution shifts in image classification,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2020.
- [8] C. Shorten and T. M. Khoshgoftaar, “A survey on image data augmentation for deep learning,” *Journal of Big Data*, vol. 6, no. 1, 2019.
- [9] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, “Intriguing properties of neural networks,” in *International Conference on Learning Representations (ICLR)*, Apr. 2014.

- [10] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” in *International Conference on Learning Representations (ICLR)*, May 2015.
- [11] S.-M. Moosavi-Dezfooli, A. Fawzi, and P. Frossard, “DeepFool: A simple and accurate method to fool deep neural networks,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2016.
- [12] S.-M. Moosavi-Dezfooli, A. Fawzi, O. Fawzi, and P. Frossard, “Universal Adversarial Perturbations,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jul. 2017.
- [13] A. Modas, R. Sanchez-Matilla, P. Frossard, and A. Cavallaro, “Toward robust sensing for Autonomous Vehicles: An adversarial perspective,” *IEEE Signal Processing Magazine (SPM)*, vol. 37, no. 4, pp. 14–23, Jul. 2020.
- [14] K. Grosse, N. Papernot, P. Manoharan, M. Backes, and P. McDaniel, “Adversarial examples for malware detection,” in *European Symposium on Research in Computer Security (ESORICS)*, Sep. 2017.
- [15] M. Alzantot, Y. Sharma, A. Elgohary, B.-J. Ho, M. Srivastava, and K.-W. Chang, “Generating natural language adversarial examples,” in *Empirical Methods in Natural Language Processing (EMNLP)*, Oct. 2018.
- [16] N. Carlini and D. Wagner, “Audio adversarial examples: Targeted attacks on speech-to-text,” in *IEEE Symposium on Security and Privacy Workshops (S&Pw)*, May 2018.
- [17] S. Huang, N. Papernot, I. Goodfellow, Y. Duan, and P. Abbeel, “Adversarial attacks on neural network policies,” in *International Conference on Learning Representations Workshops (ICLRw)*, Apr. 2017.
- [18] A. Fawzi, S.-M. Moosavi-Dezfooli, and P. Frossard, “The robustness of deep networks: A geometrical perspective,” *IEEE Signal Processing Magazine (SPM)*, vol. 34, no. 6, pp. 50–62, Nov. 2017.
- [19] G. Ortiz-Jimenez, A. Modas, S.-M. Moosavi-Dezfooli, and P. Frossard, “Optimism in the face of adversity: Understanding and improving deep learning through adversarial robustness,” *Proceedings of the IEEE*, vol. 109, no. 5, pp. 635–659, May 2021.
- [20] C. Xiao, J.-Y. Zhu, B. Li, W. He, M. Liu, and D. Song, “Spatially transformed adversarial examples,” in *International Conference on Learning Representations (ICLR)*, Apr. 2018.
- [21] C. Kanbak, S.-M. Moosavi-Dezfooli, and P. Frossard, “Geometric robustness of deep networks: Analysis and improvement,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2018.

- 
- [22] C. Laidlaw and S. Feizi, “Functional adversarial attacks,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2019.
  - [23] K. Eykholt, I. Evtimov, E. Fernandes, B. Li, A. Rahmati, C. Xiao, A. Prakash, T. Kohno, and D. Song, “Robust physical-world attacks on deep learning visual classification,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2018.
  - [24] A. Athalye, L. Engstrom, A. Ilyas, and K. Kwok, “Synthesizing robust adversarial examples,” in *International Conference on Machine Learning (ICML)*, Jul. 2018.
  - [25] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards deep learning models resistant to adversarial attacks,” in *International Conference on Learning Representations (ICLR)*, Apr. 2018.
  - [26] A. Fawzi, O. Fawzi, and P. Frossard, “Analysis of classifiers’ robustness to adversarial perturbations,” *Machine Learning*, vol. 107, no. 3, pp. 481–508, 2018.
  - [27] D. Tsipras, S. Santurkar, L. Engstrom, A. Turner, and A. Madry, “Robustness may be at odds with accuracy,” in *International Conference on Learning Representations (ICLR)*, May 2019.
  - [28] P. Maini, E. Wong, and Z. Kolter, “Adversarial robustness against the union of multiple perturbation models,” in *International Conference on Machine Learning (ICML)*, Jul. 2020.
  - [29] K. Kireev, M. Andriushchenko, and N. Flammarion, “On the effectiveness of adversarial training against common corruptions,” *arXiv:2103.02325*, 2021.
  - [30] S.-M. Moosavi-Dezfooli, “Geometry of adversarial robustness of deep networks: methods and applications,” Ph.D. dissertation, Ecole Polytechnique Fédérale de Lausanne (EPFL), 2019.
  - [31] N. Papernot, P. McDaniel, S. Jha, M. Fredrikson, Z. B. Celik, and A. Swami, “The limitations of deep learning in adversarial settings,” *IEEE European Symposium on Security and Privacy (EuroS&P)*, 2016.
  - [32] N. Carlini and D. Wagner, “Towards evaluating the robustness of neural networks,” in *IEEE Symposium on Security and Privacy (SP)*, 2017.
  - [33] J. Su, D. V. Vargas, and K. Sakurai, “One pixel attack for fooling deep neural networks,” *IEEE Transactions on Evolutionary Computation*, vol. 23, no. 5, pp. 828–841, 2019.
  - [34] S. Jetley, N. Lord, and P. Torr, “With friends like these, who needs adversaries?” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2018.

- [35] A. Ilyas, S. Santurkar, D. Tsipras, L. Engstrom, B. Tran, and A. Madry, “Adversarial examples are not bugs, they are features,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2019.
- [36] D. Hendrycks, N. Mu, E. D. Cubuk, B. Zoph, J. Gilmer, and B. Lakshminarayanan, “AugMix: A simple method to improve robustness and uncertainty under data shift,” in *International Conference on Learning Representations (ICLR)*, Apr. 2020.
- [37] D. Hendrycks, S. Basart, N. Mu, S. Kadavath, F. Wang, E. Dorundo, R. Desai, T. Zhu, S. Parajuli, M. Guo, D. Song, J. Steinhardt, and J. Gilmer, “The many faces of robustness: A critical analysis of out-of-distribution generalization,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2021.
- [38] D. A. Calian, F. Stimberg, O. Wiles, S.-A. Rebuffi, A. György, T. A. Mann, and S. Gowal, “Defending against image corruptions through adversarial augmentations,” in *International Conference on Learning Representations (ICLR)*, Apr. 2022.
- [39] J. Quionero-Candela, M. Sugiyama, A. Schwaighofer, and N. D. Lawrence, *Dataset shift in Machine Learning*. The MIT Press, 2009.
- [40] J. G. Moreno-Torres, T. Raeder, R. Alaiz-Rodríguez, N. V. Chawla, and F. Herrera, “A unifying view on dataset shift in classification,” *Pattern Recognition*, vol. 45, no. 1, pp. 521–530, 2012.
- [41] B. Recht, R. Roelofs, L. Schmidt, and V. Shankar, “Do CIFAR-10 classifiers generalize to CIFAR-10?” *arXiv:1806.00451*, 2018.
- [42] A. Fawzi, S.-M. Moosavi-Dezfooli, and P. Frossard, “Robustness of classifiers: from adversarial to random noise,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2016.
- [43] J.-Y. Franceschi, A. Fawzi, and O. Fawzi, “Robustness of classifiers to uniform  $\ell_p$  and Gaussian noise,” in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, Apr. 2018.
- [44] S. Dodge and L. Karam, “A study and comparison of human and deep learning recognition performance under visual distortions,” in *International Conference on Computer Communications and Networks (ICCCN)*, Sep. 2017.
- [45] R. Geirhos, C. R. M. Temme, J. Rauber, H. H. Schütt, M. Bethge, and F. A. Wichmann, “Generalisation in humans and deep neural networks,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2018.
- [46] S. Dodge and L. Karam, “Understanding how image quality affects deep neural networks,” in *International Conference on Quality of Multimedia Experience (QoMEX)*, 2016.

- 
- [47] I. Vasiljevic, A. Chakrabarti, and G. Shakhnarovich, “Examining the impact of blur in recognition by convolutional networks,” *arXiv:1611.05760*, 2016.
  - [48] L. Engstrom, B. Tran, D. Tsipras, L. Schmidt, and A. Madry, “Exploring the landscape of spatial robustness,” in *International Conference on Machine Learning (ICML)*, Jun. 2019.
  - [49] E. Mintun, A. Kirillov, and S. Xie, “On interaction between augmentations and corruptions in natural corruption robustness,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2021.
  - [50] L. Engstrom, A. Ilyas, S. Santurkar, D. Tsipras, J. Steinhardt, and A. Madry, “Identifying statistical bias in dataset replication,” in *International Conference on Machine Learning (ICML)*, Jul. 2020.
  - [51] H. Touvron, A. Vedaldi, M. Douze, and H. Jegou, “Fixing the train-test resolution discrepancy,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2019.
  - [52] R. Geirhos, P. Rubisch, C. Michaelis, M. Bethge, F. A. Wichmann, and W. Brendel, “ImageNet-trained CNNs are biased towards texture: increasing shape bias improves accuracy and robustness,” in *International Conference on Learning Representations (ICLR)*, May 2019.
  - [53] G. Ortiz-Jimenez, A. Modas, S.-M. Moosavi-Dezfooli, and P. Frossard, “Neural Anisotropy Directions,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2020.
  - [54] G. Ortiz-Jimenez, I. F. Salazar-Reque, A. Modas, S.-M. Moosavi-Dezfooli, and P. Frossard, “A neural anisotropic view of underspecification in deep learning,” in *RobustML workshop of International Conference on Learning Representations (ICLR)*, May 2021.
  - [55] R. Geirhos, J.-H. Jacobsen, C. Michaelis, R. Zemel, W. Brendel, M. Bethge, and F. A. Wichmann, “Shortcut learning in deep neural networks,” *Nature Machine Intelligence*, vol. 2, no. 11, pp. 665–673, 2020.
  - [56] Q. Xie, M.-T. Luong, E. Hovy, and Q. V. Le, “Self-training with noisy student improves ImageNet classification,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2020.
  - [57] T. DeVries and G. W. Taylor, “Improved regularization of convolutional neural networks with Cutout,” *arXiv:1708.04552*, 2017.
  - [58] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang, “Random erasing data augmentation,” in *AAAI Conference on Artificial Intelligence (AAAI)*, Apr. 2020.

- [59] R. Takahashi, T. Matsubara, and K. Uehara, “RICAP: Random image cropping and patching data augmentation for deep CNNs,” in *Asian Conference on Machine Learning (ACML)*, Nov. 2018.
- [60] S. Yun, D. Han, S. Chun, S. J. Oh, Y. Yoo, and J. Choe, “CutMix: Regularization strategy to train strong classifiers with localizable features,” in *IEEE International Conference on Computer Vision (ICCV)*, Oct. 2019.
- [61] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, “MixUp: Beyond empirical risk minimization,” in *International Conference on Learning Representations (ICLR)*, May 2018.
- [62] H. Guo, Y. Mao, and R. Zhang, “Mixup as locally linear out-of-manifold regularization,” in *AAAI Conference on Artificial Intelligence (AAAI)*, Jan. 2019.
- [63] E. D. Cubuk, B. Zoph, D. Mané, V. Vasudevan, and Q. V. Le, “AutoAugment: Learning augmentation strategies from data,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2019.
- [64] E. D. Cubuk, B. Zoph, J. Shlens, and Q. Le, “RandAugment: Practical automated data augmentation with a reduced search space,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2020.
- [65] H. Wang, C. Xiao, J. Kossai, Z. Yu, A. Anandkumar, and Z. Wang, “AugMax: Adversarial composition of random augmentations for robust training,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2021.
- [66] I. Kim, S. Han, J.-w. Baek, S.-J. Park, J.-J. Han, and J. Shin, “Quality-agnostic image recognition via invertible decoder,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
- [67] G. Chen, P. Peng, L. Ma, J. Li, L. Du, and Y. Tian, “Amplitude-phase recombination: Rethinking robustness of convolutional neural networks in frequency domain,” in *IEEE International Conference on Computer Vision (ICCV)*, Jun. 2021.
- [68] E. Rusak, L. Schott, R. Zimmermann, J. Bitterwolf, O. Bringmann, M. Bethge, and W. Brendel, “A simple way to make neural networks robust against diverse image corruptions,” in *European Conference on Computer Vision (ECCV)*, Aug. 2020.
- [69] M. Yi, L. Hou, J. Sun, L. Shang, X. Jiang, Q. Liu, and Z. Ma, “Improved OOD generalization via adversarial training and pretraining,” in *International Conference on Machine Learning (ICML)*, Jul. 2021.
- [70] C. Laidlaw, S. Singla, and S. Feizi, “Perceptual adversarial robustness: Defense against unseen threat models,” in *International Conference on Learning Representations (ICLR)*, May 2021.



- 
- [71] J. Diffenderfer, B. R. Bartoldson, S. Chaganti, J. Zhang, and B. Kailkhura, “A winning hand: Compressing deep networks can improve out-of-distribution robustness,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2021.
  - [72] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, “An image is worth 16x16 words: Transformers for image recognition at scale,” in *International Conference on Learning Representations (ICLR)*, May 2021.
  - [73] S. Bhojanapalli, A. Chakrabarti, D. Glasner, D. Li, T. Unterthiner, and A. Veit, “Understanding robustness of transformers for image classification,” in *IEEE International Conference on Computer Vision (ICCV)*, Oct. 2021.
  - [74] K. Morrison, B. Gilby, C. Lipchak, A. Mattioli, and A. Kovashka, “Exploring corruption robustness: Inductive biases in Vision Transformers and MLP-Mixers,” *Uncertainty & Robustness in Deep Learning Workshop, ICML*, Jul. 2021.
  - [75] A. Kolesnikov, L. Beyer, X. Zhai, J. Puigcerver, J. Yung, S. Gelly, and N. Houlsby, “Big transfer (BiT): General visual representation learning,” in *European Conference on Computer Vision (ECCV)*, Aug. 2020.
  - [76] C. Herrmann, K. Sargent, L. Jiang, R. Zabih, H. Chang, C. Liu, D. Krishnan, and D. Sun, “Pyramid adversarial training improves ViT performance,” *arXiv:2111.15121*, 2021.
  - [77] X. Mao, G. Qi, Y. Chen, X. Li, R. Duan, S. Ye, Y. He, and H. Xue, “Towards robust Vision Transformers,” *arXiv:2105.07926*, 2021.
  - [78] A. Kurakin, I. J. Goodfellow, and S. Bengio, “Adversarial machine learning at scale,” in *International Conference on Learning Representations (ICLR)*, Apr. 2017.
  - [79] F. Tramer, N. Carlini, W. Brendel, and A. Madry, “On adaptive attacks to adversarial example defenses,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2020.
  - [80] W. Brendel, J. Rauber, M. Kümmerer, I. Ustyuzhaninov, and M. Bethge, “Accurate, reliable and fast robustness evaluation,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2019.
  - [81] F. Croce and M. Hein, “Minimally distorted adversarial examples with a fast adaptive boundary attack,” in *International Conference on Machine Learning (ICML)*, Jul. 2020.
  - [82] S. Sabour, Y. Cao, F. Faghri, and D. J. Fleet, “Adversarial manipulation of deep representations,” in *International Conference on Learning Representations (ICLR)*, May 2016.

- [83] N. Narodytska and S. Kasiviswanathan, “Simple black-box adversarial attacks on deep neural networks,” in *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRw)*, Jul. 2017.
- [84] F. Croce and M. Hein, “Sparse and imperceivable adversarial attacks,” in *IEEE International Conference on Computer Vision (ICCV)*, Oct. 2019.
- [85] W. Zhou, X. Hou, Y. Chen, M. Tang, X. Huang, X. Gan, and Y. Yang, “Transferable adversarial perturbations,” in *European Conference on Computer Vision (ECCV)*, Sep. 2018.
- [86] Y. Sharma, G. Ding, and M. A. Brubaker, “On the effectiveness of low frequency perturbations,” in *International Joint Conference on Artificial Intelligence (IJCAI)*, Aug. 2019, pp. 3389–3396.
- [87] Y. Tsuzuku and I. Sato, “On the structural sensitivity of deep convolutional networks to the directions of Fourier basis functions,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2019.
- [88] C. Guo, J. Gardner, Y. You, A. G. Wilson, and K. Weinberger, “Simple black-box adversarial attacks,” in *International Conference on Machine Learning (ICML)*, Jun. 2019.
- [89] Y. Liu, S.-M. Moosavi-Dezfooli, and P. Frossard, “A geometry-inspired decision-based attack,” in *IEEE International Conference on Computer Vision (ICCV)*, Oct. 2019.
- [90] A. Rahmati, S.-M. Moosavi-Dezfooli, P. Frossard, and H. Dai, “GeoDA: a geometric framework for black-box adversarial attacks,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2020.
- [91] H. Hosseini and R. Poovendran, “Semantic adversarial examples,” in *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRw)*, Jun. 2018.
- [92] A. S. Shamsabadi, R. Sanchez-Matilla, and A. Cavallaro, “ColorFool: Semantic adversarial colorization,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2020.
- [93] M. Sharif, S. Bhagavatula, L. Bauer, and M. K. Reiter, “Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition,” in *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, Oct. 2016, pp. 1528–1540.
- [94] A. Fawzi and P. Frossard, “Measuring the effect of nuisance variables on classifiers,” in *Proceedings of the British Machine Vision Conference (BMVC)*, Sep. 2016.

- 
- [95] P.-Y. Chen, H. Zhang, Y. Sharma, J. Yi, and C.-J. Hsieh, *ZOO: Zeroth Order Optimization Based Black-Box Attacks to Deep Neural Networks without Training Substitute Models*. Association for Computing Machinery, 2017, pp. 15–26.
  - [96] W. Brendel, J. Rauber, and M. Bethge, “Decision-based adversarial attacks: Reliable attacks against black-box machine learning models,” in *International Conference on Learning Representations (ICLR)*, Apr. 2018.
  - [97] J. Uesato, B. O’Donoghue, P. Kohli, and A. van den Oord, “Adversarial risk and the dangers of evaluating against weak attacks,” in *International Conference on Machine Learning (ICML)*, Jul. 2018.
  - [98] J. Chen, M. I. Jordan, and M. J. Wainwright, “HopSkipJumpAttack: A query-efficient decision-based attack,” *arXiv:1904.02144*, 2019.
  - [99] M. Andriushchenko, F. Croce, N. Flammarion, and M. Hein, “Square attack: a query-efficient black-box adversarial attack via random search,” in *European Conference on Computer Vision (ECCV)*, Aug. 2020.
  - [100] A. Kurakin, I. Goodfellow, and S. Bengio, “Adversarial examples in the physical world,” *arXiv:1607.02533*, 2016.
  - [101] A. Shafahi, M. Najibi, M. A. Ghiasi, Z. Xu, J. Dickerson, C. Studer, L. S. Davis, G. Taylor, and T. Goldstein, “Adversarial training for free!” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2019.
  - [102] E. Wong, L. Rice, and J. Z. Kolter, “Fast is better than free: Revisiting adversarial training,” in *International Conference on Learning Representations (ICLR)*, Apr. 2020.
  - [103] S. Gu and L. Rigazio, “Towards deep neural network architectures robust to adversarial examples,” *arXiv:1412.5068*, 2014.
  - [104] C. Lyu, K. Huang, and H.-N. Liang, “A unified gradient regularization family for adversarial examples,” in *IEEE International Conference on Data Mining (ICDM)*, Nov. 2015.
  - [105] D. Jakubovitz and R. Giryes, “Improving dnn robustness to adversarial attacks using jacobian regularization,” in *European Conference on Computer Vision (ECCV)*, Sep. 2018.
  - [106] A. S. Ross and F. Doshi-Velez, “Improving the adversarial robustness and interpretability of deep neural networks by regularizing their input gradients,” in *AAAI Conference on Artificial Intelligence (AAAI)*, Feb. 2018.
  - [107] S.-M. Moosavi-Dezfooli, A. Fawzi, J. Uesato, and P. Frossard, “Robustness via curvature regularization, and vice versa,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2019.

## Bibliography

---

- [108] C. Qin, J. Martens, S. Gowal, D. Krishnan, K. Dvijotham, A. Fawzi, S. De, R. Stanforth, and P. Kohli, “Adversarial robustness through local linearization,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2019.
- [109] S. Singla and S. Feizi, “Second-order provable defenses against adversarial attacks,” in *International Conference on Machine Learning (ICML)*, Jul. 2020.
- [110] D. Wu, S.-T. Xia, and Y. Wang, “Adversarial weight perturbation helps robust generalization,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2020.
- [111] M. Andriushchenko and N. Flammarion, “Understanding and improving fast adversarial training,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2020.
- [112] J. Cohen, E. Rosenfeld, and Z. Kolter, “Certified adversarial robustness via randomized smoothing,” in *International Conference on Machine Learning (ICML)*, Jun. 2019.
- [113] G. Yang, T. Duan, J. E. Hu, H. Salman, I. Razenshteyn, and J. Li, “Randomized smoothing of all shapes and sizes,” in *International Conference on Machine Learning (ICML)*, Jul. 2020.
- [114] H. Salman, J. Li, I. Razenshteyn, P. Zhang, H. Zhang, S. Bubeck, and G. Yang, “Provably robust deep learning via adversarially trained smoothed classifiers,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2019.
- [115] A. Fawzi, S.-M. Moosavi-Dezfooli, P. Frossard, and S. Soatto, “Empirical study of the topology and geometry of deep networks,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2018.
- [116] F. Tramèr, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel, “The space of transferable adversarial examples,” *arXiv:1704.03453*, Apr. 2017.
- [117] S.-M. Moosavi-Dezfooli, A. Fawzi, O. Fawzi, P. Frossard, and S. Soatto, “Robustness of classifiers to universal perturbations: A geometric perspective,” in *International Conference on Learning Representations (ICLR)*, Apr. 2018.
- [118] A. Rahmati, S.-M. Moosavi-Dezfooli, and H. Dai, “Adversarial training may be a double-edged sword,” in *RobustML workshop of International Conference on Learning Representations (ICLR)*, May 2021.
- [119] S. Santurkar, A. Ilyas, D. Tsipras, L. Engstrom, B. Tran, and A. Madry, “Image synthesis with a single (robust) classifier,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2019.
- [120] L. Engstrom, A. Ilyas, S. Santurkar, D. Tsipras, B. Tran, and A. Madry, “Adversarial robustness as a prior for learned representations,” *arXiv:1906.00945*, 2019.

- 
- [121] Z. Allen-Zhu and Y. Li, “Feature purification: How adversarial training performs robust deep learning,” *arXiv:2005.10190*, 2020.
  - [122] A. Fawzi, H. Fawzi, and O. Fawzi, “Adversarial vulnerability for any classifier,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2018.
  - [123] L. Schmidt, S. Santurkar, D. Tsipras, K. Talwar, and A. Madry, “Adversarially robust generalization requires more data,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2018.
  - [124] J. Alayrac, J. Uesato, P. Huang, A. Fawzi, R. Stanforth, and P. Kohli, “Are labels required for improving adversarial robustness?” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2019.
  - [125] Y. Carmon, A. Raghunathan, L. Schmidt, J. C. Duchi, and P. S. Liang, “Unlabeled data improves adversarial robustness,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2019.
  - [126] A. Raghunathan, S. M. Xie, F. Yang, J. Duchi, and P. Liang, “Understanding and mitigating the tradeoff between robustness and accuracy,” in *International Conference on Machine Learning (ICML)*, Jul. 2020.
  - [127] T. Serre, “Deep Learning: The Good, the Bad, and the Ugly,” *Annual Review of Vision Science*, vol. 5, no. 1, pp. 399–426, Sep. 2019.
  - [128] C. Xie, M. Tan, B. Gong, J. Wang, A. L. Yuille, and Q. V. Le, “Adversarial examples improve image recognition,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2020.
  - [129] Z. Tang, Y. Gao, L. Karlinsky, P. Sattigeri, R. Feris, and D. Metaxas, “OnlineAugment: Online data augmentation with less domain knowledge,” in *European Conference on Computer Vision (ECCV)*, Aug. 2020.
  - [130] E. Dohmatob, “Generalized no free lunch theorem for adversarial robustness,” in *International Conference on Machine Learning (ICML)*, Jun. 2019.
  - [131] J. Gilmer, L. Metz, F. Faghri, S. S. Schoenholz, M. Raghu, M. Wattenberg, and I. Goodfellow, “Adversarial Spheres,” in *International Conference on Learning Representations Workshops (ICLRw)*, Apr. 2018.
  - [132] Y. Jiang, D. Krishnan, H. Mobahi, and S. Bengio, “Predicting the generalization gap in deep networks with margin distributions,” in *International Conference on Learning Representations (ICLR)*, May 2019.
  - [133] R. Werpachowski, A. György, and C. Szepesvari, “Detecting overfitting via adversarial examples,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2019.

- [134] J.-H. Jacobsen, J. Behrmann, R. Zemel, and M. Bethge, “Excessive invariance causes adversarial vulnerability,” in *International Conference on Learning Representations (ICLR)*, May 2019.
- [135] F. Tramèr, J. Behrmann, N. Carlini, N. Papernot, and J.-H. Jacobsen, “Fundamental tradeoffs between invariance and sensitivity to adversarial perturbations,” in *International Conference on Machine Learning (ICML)*, Jul. 2020.
- [136] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, “Understanding deep learning requires rethinking generalization,” in *International Conference on Learning Representations (ICLR)*, Apr. 2017.
- [137] L. Rice, E. Wong, and J. Z. Kolter, “Overfitting in adversarially robust deep learning,” in *International Conference on Machine Learning (ICML)*, Jul. 2020.
- [138] S. Gowal, C. Qin, J. Uesato, T. Mann, and P. Kohli, “Uncovering the limits of adversarial training against norm-bounded adversarial examples,” *arXiv:2010.03593*, 2020.
- [139] K. Xu, S. Liu, P. Zhao, P.-Y. Chen, H. Zhang, Q. Fan, D. Erdogmus, Y. Wang, and X. Lin, “Structured adversarial attack: Towards general implementation and better interpretability,” in *International Conference on Learning Representations (ICLR)*, May 2019.
- [140] A. Modas, S.-M. Moosavi-Dezfooli, and P. Frossard, “SparseFool: A few pixels make a big difference,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2019.
- [141] T. Blumensath and M. E. Davies, “Iterative thresholding for sparse approximations,” *Journal of Fourier Analysis and Applications*, vol. 14, no. 5, pp. 629–654, 2008.
- [142] M. Nikolova, “Description of the minimizers of least squares regularized with  $\ell_0$ -norm. uniqueness of the global minimizer,” *SIAM Journal on Imaging Sciences*, vol. 6, no. 2, pp. 904–937, 2013.
- [143] A. Patrascu and I. Necoara, “Random coordinate descent methods for  $\ell_0$  regularized convex optimization,” *IEEE Transactions on Automatic Control*, vol. 60, no. 7, pp. 1811–1824, 2015.
- [144] M. Nagahara, D. E. Quevedo, and J. Ostergaard, “Sparse packetized predictive control for networked control over erasure channels,” *IEEE Transactions on Automatic Control*, vol. 59, no. 7, pp. 1899–1905, 2014.
- [145] E. J. Candès and T. Tao, “Decoding by linear programming,” *IEEE Transactions on Information Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.
- [146] D. L. Donoho, “Compressed sensing,” *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

- 
- [147] B. K. Natarajan, “Sparse approximate solutions to linear systems,” *SIAM Journal on Computing*, vol. 24, no. 2, pp. 227–234, 1995.
  - [148] E. Candès, M. Rudelson, T. Tao, and E. Vershynin, “Error correction via linear programming,” in *IEEE Symposium on Foundations of Computer Science (FoCS)*, 2005.
  - [149] D. L. Donoho and M. Elad, “Optimally sparse representation in general (nonorthogonal) dictionaries via  $\ell_1$  minimization,” *National Academy of Sciences*, vol. 100, no. 5, pp. 2197–2202, 2003.
  - [150] R. Gribonval and M. Nielsen, “Sparse representations in unions of bases,” *IEEE Transactions on Information Theory*, vol. 49, no. 12, pp. 3320–3325, 2003.
  - [151] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *International Conference on Learning Representations (ICLR)*, May 2015.
  - [152] J. Deng, W. Dong, R. Socher, L.-J. Li, L. Kai, and F.-F. Li, “Imagenet: A large-scale hierarchical image database,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2009.
  - [153] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the Inception architecture for computer vision,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2016.
  - [154] Y. LeCun and C. Cortes. (2010) Mnist handwritten digits database. [Online]. Available: <http://yann.lecun.com/exdb/mnist/>
  - [155] A. Krizhevsky, “Learning Multiple Layers of Features from Tiny Images,” University of Toronto, Tech. Rep., 2009.
  - [156] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
  - [157] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2016.
  - [158] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “PyTorch: An Imperative Style, High-Performance Deep Learning Library,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2019.

## Bibliography

---

- [159] G. Ortiz-Jimenez, A. Modas, S.-M. Moosavi-Dezfooli, and P. Frossard, “Hold me tight! Influence of discriminative features on deep network boundaries,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2020.
- [160] G. Ortiz-Jimenez, A. Modas, S.-M. Moosavi-Dezfooli, and P. Frossard, “Redundant features can hurt robustness to distribution shift,” *Uncertainty & Robustness in Deep Learning Workshop, ICML*, Jul. 2020.
- [161] A. Modas, A. Xompero, R. Sanchez-Matilla, P. Frossard, and A. Cavallaro, “Improving filling level classification with adversarial training,” in *IEEE International Conference on Image Processing (ICIP)*, Sep. 2021.
- [162] S. R. Maiya, M. Ehrlich, V. Agarwal, S.-N. Lim, T. Goldstein, and A. Shrivastava, “A frequency perspective of adversarial robustness,” *arXiv:2111.00861*, 2021.
- [163] N. Ahmed, T. Natarajan, and K. R. Rao, “Discrete Cosine Transform,” *IEEE Transactions on Computers*, vol. C-23, no. 1, pp. 90–93, 1974.
- [164] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 4th ed. Pearson, 2017.
- [165] F. W. Campbell and J. G. Robson, “Application of fourier analysis to the visibility of gratings,” *The Journal of Physiology*, vol. 197, no. 3, pp. 551–566, 1968.
- [166] M. McCloskey and N. J. Cohen, “Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem,” *Psychology of Learning and Motivation*, vol. 24, pp. 109–165, 1989.
- [167] B. Ghorbani, S. Mei, T. Misiakiewicz, and A. Montanari, “When do neural networks outperform kernel methods?” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2020.
- [168] H. Shah, K. Tamuly, A. Raghunathan, P. Jain, and P. Netrapalli, “The pitfalls of simplicity bias in neural networks,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2020.
- [169] C. Cortes and V. Vapnik, “Support-vector networks,” *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [170] H. Salman, A. Ilyas, L. Engstrom, A. Kapoor, and A. Madry, “Do adversarially robust ImageNet models transfer better?” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2020.
- [171] F. Utrera, E. Kravitz, N. B. Erichson, R. Khanna, and M. W. Mahoney, “Adversarially-trained deep nets transfer better,” in *International Conference on Learning Representations (ICLR)*, May 2021.



- 
- [172] R. Sanchez-Matilla, K. Chatzilygeroudis, A. Modas, N. F. Duarte, A. Xompero, P. Frossard, A. Billard, and A. Cavallaro, “Benchmark for human-to-robot handovers of unseen containers with unknown filling,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 5, no. 2, pp. 1642–1649, 2020.
- [173] A. Xompero, R. Sanchez-Matilla, A. Modas, P. Frossard, and A. Cavallaro, “Multi-view shape estimation of transparent containers,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, May 2020.
- [174] S. S. Sajjan, M. Moore, M. Pan, G. Nagaraja, J. Lee, A. Zeng, and S. Song, “ClearGrasp: 3D shape estimation of transparent objects for manipulation,” in *IEEE International Conference on on Robotics and Automation (ICRA)*, Jun. 2020.
- [175] R. Mottaghi, C. Schenck, D. Fox, and A. Farhadi, “See the glass half full: Reasoning about liquid containers, their volume and content,” in *IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [176] C. Schenck and D. Fox, “Visual closed-loop control for pouring liquids,” in *IEEE International Conference on on Robotics and Automation (ICRA)*, May 2017.
- [177] C. Do, T. Schubert, and W. Burgard, “A probabilistic approach to liquid level detection in cups using an RGB-D camera,” in *IEEE International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2016.
- [178] C. Do and W. Burgard, “Accurate pouring with an autonomous robot using an RGB-D camera,” in *International Conference on Intelligent Autonomous Systems (ICoIAS)*, Jul. 2018.
- [179] C. Schenck and D. Fox, “Reasoning about liquids via closed-loop simulation,” in *Robotics: Science and Systems*, Jul. 2017.
- [180] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu, “A survey on deep transfer learning,” in *International Conference on Artificial Neural Networks (ICANN)*, 2018.
- [181] A. Xompero, R. Sanchez-Matilla, R. Mazzon, and A. Cavallaro, “CORSMAL Containers Manipulation,” 2020, (1.0) [Dataset]. Queen Mary University of London. <https://doi.org/10.17636/101CORSMAL1>. [Online]. Available: [https://corsmal.eecs.qmul.ac.uk/containers\\_manip.html](https://corsmal.eecs.qmul.ac.uk/containers_manip.html)
- [182] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick, “Mask R-CNN,” in *IEEE International Conference on Computer Vision (ICCV)*, Oct. 2017.
- [183] S. Zagoruyko and N. Komodakis, “Wide residual networks,” in *Proceedings of the British Machine Vision Conference (BMVC)*, Sep. 2016.

## Bibliography

---

- [184] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, “How transferable are features in deep neural networks?” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2014.
- [185] A. Modas, R. Rade, G. Ortiz-Jiménez, S.-M. Moosavi-Dezfooli, and P. Frossard, “PRIME: A few primitives can boost robustness to common corruptions,” *arXiv:2112.13547*, 2021.
- [186] A. Modas, A. Cavallaro, and P. Frossard, “Data augmentation with mixtures of max-entropy transformations for filling-level classification,” in *European Signal Processing Conference (EUSIPCO)*, 2022, arXiv:2203.04027.
- [187] F. Binder, E. Aichinger, J. Ecker, C. Nöbauer, and P. Mayr, “Algorithms for near-rings of non-linear transformations,” in *International Symposium on Symbolic and Algebraic Computation, (ISSAC)*, Aug. 2000.
- [188] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley-Interscience, 2006.
- [189] A. Xu and M. Raginsky, “Information-theoretic analysis of generalization capability of learning algorithms,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2017.
- [190] M. S. Masiha, A. Gohari, M. H. Yassaee, and M. R. Aref, “Learning under distribution mismatch and model misspecification,” in *IEEE International Symposium on Information Theory, (ISIT)*, 2021.
- [191] P. Beale, *Statistical Mechanics*. Elsevier, 1996.
- [192] L. Petrini, A. Favero, M. Geiger, and M. Wyart, “Relative stability toward diffeomorphisms indicates performance in deep nets,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2021.
- [193] F. Croce, M. Andriushchenko, V. Schwag, E. Debenedetti, N. Flammarion, M. Chiang, P. Mittal, and M. Hein, “Robustbench: a standardized adversarial robustness benchmark,” in *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*, 2021.
- [194] P. Benz, C. Zhang, A. Karjauv, and I. S. Kweon, “Revisiting batch normalization for improving corruption robustness,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, Jan. 2021.
- [195] S. Schneider, E. Rusak, L. Eck, O. Bringmann, W. Brendel, and M. Bethge, “Improving robustness against common corruptions by covariate shift adaptation,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2020.

- 
- [196] T. Chen, S. Kornblith, K. Swersky, M. Norouzi, and G. E. Hinton, “Big self-supervised models are strong semi-supervised learners,” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2020.
  - [197] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2018.
  - [198] M. Moayeri and S. Feizi, “Sample efficient detection and classification of adversarial attacks via self-supervised embeddings,” in *IEEE International Conference on Computer Vision (ICCV)*, Oct. 2021.
  - [199] G. Ortiz-Jimenez, S.-M. Moosavi-Dezfooli, and P. Frossard, “What can linearized neural networks actually say about generalization?” in *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2021.
  - [200] G. Yüce, G. Ortiz-Jimenez, B. Besbinar, and P. Frossard, “A structured dictionary perspective on implicit neural representations,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2022.
  - [201] R. Rade and S.-M. Moosavi-Dezfooli, “Helper-based adversarial training: Reducing excessive margin to achieve a better accuracy vs. robustness trade-off,” in *Adversarial Machine Learning workshop (ICML)*, 2021.
  - [202] P. W. Koh, S. Sagawa, H. Marklund, S. M. Xie, M. Zhang, A. Balsubramani, W. Hu, M. Yasunaga, R. L. Phillips, I. Gao, T. Lee, E. David, I. Stavness, W. Guo, B. Earnshaw, I. Haque, S. M. Beery, J. Leskovec, A. Kundahe, E. Pierson, S. Levine, C. Finn, and P. Liang, “WILDS: A benchmark of in-the-wild distribution shifts,” in *International Conference on Machine Learning (ICML)*, Jul. 2021.
  - [203] L. Engstrom, A. Ilyas, S. Santurkar, and D. Tsipras, “Robustness (python library),” 2019. [Online]. Available: <https://github.com/MadryLab/robustness>





# APOSTOLOS MODAS

Lausanne, Switzerland

Deep Learning • Robustness • Computer Vision

✉ apostolos.modas@gmail.com  
🐙 github.com/amodas  
🌐 linkedin.com/in/amodas/  
🐦 twitter.com/modasapo



## ABOUT ME

My main research focuses on the **robustness** of **deep neural networks**. I have designed novel frameworks to **evaluate the robustness** and **reveal the biases** of deep image classifiers, in order to build more **robust** and **reliable** deep learning systems. The spectrum of my work ranges from **understanding** deep networks, to designing task-specific **robust solutions**.

Publication topics: Understanding deep networks through adversarial robustness [1], [2], [3], [4]  
Robustness to distribution shifts [5], [6], [7]  
Inductive bias of deep networks [8], [9]  
Computer vision for robotics [10], [11]



## WORK & RESEARCH EXPERIENCE

Apr. 2018  
May 2022



### Research Assistant — EPFL

Supervisor: Prof. Pascal Frossard — LTS4

I developed novel frameworks for evaluating the robustness of image classifiers to adversarial perturbations, for connecting adversarial examples with the features in the datasets, and for achieving state-of-the-art robustness to distribution shifts caused by common corruptions of the data. My work has explained various mechanisms behind the robustness of classifiers, and it has also been adapted in the context of human-to-robot collaborations of the CHIST-ERA CORSMAL project, for improving robustness to real-world data distribution shifts.

Feb. 2021  
June 2021



### Research Intern — IBM Research Dublin/Zurich

Supervisor: Dr. Mathieu Sinn — AI and Machine Learning Research Lab, Dublin

I designed simple adaptive defences that can easily improve the robustness of image classifiers against stochastic adaptive black-box adversarial attacks. My findings led to contributions to the open-source Adversarial Robustness Toolbox (ART).

July 2016  
Sept. 2017



### Research Assistant — University of Athens

Supervisor: Prof. Nancy Alonistioti — SCAN Lab

I worked on the detection of affected crop areas using drones (UAVs). Furthermore, I also drafted a thorough plan of the multi-sensory infrastructure required to convert the city of Argos, Greece, into a Smart City.



## EDUCATION

Apr. 2018  
May 2022



### PhD. Machine Learning — EPFL

Supervisor: Prof. Pascal Frossard

Thesis: "Robustness and invariance properties of image classifiers" (*"Best Thesis Award"* nominee)

Sept. 2015  
June 2017



### MSc. Signal Processing & Machine Learning — University of Athens (Top graduate — GPA: 9.77/10)

Supervisor: Prof. Nancy Alonistioti

Thesis: "Adaptive unmanned vehicle autopiloting using WebRTC video analysis"

Sept. 2009  
Sept. 2015



### BSc. Informatics & Telecommunications — University of Athens

Supervisor: Dr. Alex Eleftheriadis

Thesis: "Multi-user video conferencing using SSRC multiplexing and RTP packet retransmissions"



# APOSTOLOS MODAS

Lausanne, Switzerland

Deep Learning • Robustness • Computer Vision

✉ apostolos.modas@gmail.com  
🐙 github.com/amodas  
🌐 linkedin.com/in/amodas/  
🐦 twitter.com/modasapo



## SOFTWARE SKILLS

**Languages:** Python, Matlab, C, Java, JavaScript

**Frameworks:** PyTorch, Torchvision

**Featured repos:** [github.com/LTS4/SparseFool](https://github.com/LTS4/SparseFool)  
[github.com/LTS4/hold-me-tight](https://github.com/LTS4/hold-me-tight)  
[github.com/amodas/PRIME-augmentations](https://github.com/amodas/PRIME-augmentations)  
[github.com/LTS4/neural-anisotropy-directions](https://github.com/LTS4/neural-anisotropy-directions)



## LANGUAGES

Greek	● ● ● ● ●
English	● ● ● ● ●
German	● ○ ○ ○ ○



## PERSONAL INTERESTS

Basketbal, playing the bass, cooking, theater



## FEATURED PUBLICATIONS

- [1] **A. Modas**, SM. Moosavi-Dezfooli and P. Frossard, “*SparseFool: A few pixels make a big difference*”, in IEEE CVPR, 2019.
- [2] G. Ortiz-Jimenez\*, **A. Modas\***, SM. Moosavi-Dezfooli and P. Frossard, “*Hold me tight! Influence of discriminative features on deep network boundaries*”, in NeurIPS, 2020.
- [3] **A. Modas\***, R. Sanchez-Matilla\*, P. Frossard and A. Cavallaro, “*Towards robust sensing for autonomous vehicles: An adversarial perspective*”, in IEEE SPM, 2020.
- [4] G. Ortiz-Jimenez, **A. Modas**, SM. Moosavi-Dezfooli and P. Frossard, “*Optimism in the face of adversity: Understanding and improving deep learning through adversarial robustness*”, in Proc. of IEEE, 2021.
- [5] **A. Modas\***, R. Rade\*, G. Ortiz-Jimenez, SM. Moosavi-Dezfooli and P. Frossard, “*PRIME: A few primitives can boost robustness to Common Corruptions*”, arXiv pre-print, 2021 (Under review).
- [6] **A. Modas**, A. Xompero, P. Frossard and A. Cavallaro, “*Improving filling-level classification with adversarial training*”, in IEEE ICIP, 2021.
- [7] **A. Modas**, A. Cavallaro and P. Frossard, “*Data augmentation with mixtures of max-entropy transformations for filling-level classification*”, in EUSIPCO 2022.
- [8] G. Ortiz-Jimenez\*, **A. Modas\***, SM. Moosavi-Dezfooli and P. Frossard, “*Neural Anisotropy Directions*”, in NeurIPS, 2020.
- [9] G. Ortiz-Jimenez, I.F. Salazar, **A. Modas**, SM. Moosavi-Dezfooli and P. Frossard, “*A neural anisotropic view of underspecification in deep learning*”, in ICLR RobustML Workshop, 2021.
- [10] R. Sanchez-Matilla, K. Chatzilygeroudis, **A. Modas**, N. Ferreira Duarte, A. Xompero, P. Frossard, A. Billard and A. Cavallaro, “*Benchmark for human-to-robot handovers of unseen containers with unknown filling*”, in IEEE RA-L, 2020.
- [11] A. Xompero, R. Sanchez-Matilla, **A. Modas**, P. Frossard and A. Cavallaro, “*Multi-view shape estimation of transparent containers*”, in IEEE ICASSP, 2020.