



Graphs for deep learning representations

Carlos Eduardo Rosar Kós Lassance

► To cite this version:

Carlos Eduardo Rosar Kós Lassance. Graphs for deep learning representations. Machine Learning [cs.LG]. Ecole nationale supérieure Mines-Télécom Atlantique, 2020. English. NNT : 2020IMTA0204 . tel-03080186

HAL Id: tel-03080186

<https://theses.hal.science/tel-03080186>

Submitted on 17 Dec 2020

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

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

THESE DE DOCTORAT DE

L'ÉCOLE NATIONALE SUPERIEURE MINES-TELECOM ATLANTIQUE
BRETAGNE PAYS DE LA LOIRE - IMT ATLANTIQUE

ECOLE DOCTORALE N° 601
*Mathématiques et Sciences et Technologies
de l'Information et de la Communication*
Spécialité : *Signal, Image, Vision*

Par

Carlos Eduardo ROSAR KOS LASSANCE

Graphs for deep learning representations

Thèse présentée et soutenue à Brest, le 15/10/2020

Unité de recherche : Lab-STICC, UMR CNRS 6285

Thèse N° : 2020IMTA0204

Rapporteurs avant soutenance :

Pierre BORGNAT	Directeur de Recherche - CNRS/ENS de Lyon
Stephan GÜNNEMAN	Professeur - TU Munich

Composition du Jury :

Rapporteurs :	Pierre BORGNAT	Directeur de Recherche - CNRS/ENS de Lyon
	Stephan GÜNNEMAN	Professeur - TU Munich

Examineurs :	Sophie ACHARD	Directrice de recherche - CNRS/Université Grenoble Alpes
	Antonio ORTEGA	Professeur - USC (University of Southern California)
	Vincent GRIPON	Chargé de recherche - IMT Atlantique
Dir. de thèse :	Michel JEZEQUEL	Professeur - IMT Atlantique

Invité(s)

Jian Tang	Professeur Adjoint	- HEC Montréal/Mila
-----------	--------------------	---------------------

Contents

Acknowledgements	vii
Résumé	ix
Abstract	ix
Introduction	x
Chapitre 2: Concepts en apprentissage profond	xix
Chapitre 3: Concepts en traitement de signaux sur graphe	xx
Chapitre 4: Réseaux de neurones profonds pour des signaux sur graphe	xxi
Chapitre 5: Représentations latentes de réseaux profonds sur graphe	xxii
1 Introduction	1
1.1 Context and motivation	2
1.2 Graphs for deep learning latent representations	6
1.3 Document structure	10
2 Deep Neural Networks	13
2.1 Definitions	14
2.2 Deep Learning Layers	27
2.3 Tasks and Datasets	34

2.4	Compression	43
2.5	Robustness	48
2.6	Summary of the chapter	61
3	Graph Signal Processing	65
3.1	Definitions	65
3.2	Graphs for samples of features	74
3.3	Inferring graph topology from signals	78
3.4	Graph filters	86
3.5	Summary of the chapter	100
4	Deep Learning for inputs supported on graphs	103
4.1	Definitions	104
4.2	Supervised classification of graph signals	113
4.3	Semi supervised classification of vertices	121
4.4	Summary of the chapter	128
5	Deep Neural Networks latent spaces supported on graphs	131
5.1	Characterizing DNN behavior via smoothness of intermediate representations	132
5.2	Smoothness as an objective function for DNNs	138
5.3	Controlling DNN smoothness to improve robustness	146
5.4	Using intermediate representation graphs to compress DNNs	163
5.5	Summary of the chapter	172
6	Conclusion	175
6.1	Summary of contributions	176

6.2	Perspectives and future work	178
6.3	Discussion and considerations about the field	181
	Bibliography	183

Acknowledgements

This is for me the most important part of this thesis, as it would not be possible to do this alone. I can only hope to acknowledge everyone that helped me in this pursuit, but for the ones that are not cited here, it is 100% my fault as I have left this part to the very end of my writing and I hope you forgive me.

First, I have to thank my whole family. My parents (Jacqueline ROSAR KOS LASSANCE and Carlos Alberto KOS LASSANCE JUNIOR) that supported me in all my decisions and that greatly invested in my education, both academic/formal but most importantly on my social/informal education. I would not be anywhere near where I am now without you.

The same could be said about my brother (Luiz Carlos BANDEIRA LASSANCE) and sisters (Tatiana BANDEIRA LASSANCE and Patricia BANDEIRA LASSANCE BURNS), without you in my life I would not be able to develop in the person that I am now.

Second, I have to thank my supervisor (Vincent GRIPON) and my thesis director (Michel JEZEQUEL). Vincent taught me more than I could help to understand and both gave me the full support I needed to develop my thesis in the appropriate time. I also need to thank them for their patience, when I was down and thought that I was not doing any work. I can only hope to be able to pay this forward to the next generation.

Third, I have to thank my PUC-Rio professor (Sergio LIFSCHITZ) for all the things he taught me during my time at PUC-Rio, for nudging me into considering Télécom Bretagne (now IMT Atlantique) for my double-degree and for his friendship.

In this vein, I also need to thank all my co-authors for the invaluable discussions and for improving my knowledge of the field. The same has to be said about my lab-mates who had to endure endless hours of me nagging about the most innocuous things.

I also have to thank my international collaborators. First Antonio ORTEGA from

USC, who was almost a second supervisor during my thesis and helped immensely to advance on my thesis and my writing. Second, Jian TANG from Mila and HEC-Montréal, who took me under an one year internship and helped me develop a lot of skills in the graph neural network domain. Third, Yasir LATIF and Ravi GARG from University of Adelaide that taught me a lot on the domain of VBL (Visual-based localization) and allowed me extend my competences to another field of Deep Learning. Finally, Gonzalo MATEOS from University of Rochester, who helped me a lot with graph inference and took me under a one month visit of his lab.

To all my friends, everywhere in the world, who helped me not obsess about the thesis and to be able to arrive here at the end.

I also have to thank all the members of the jury for accepting to be a part of this process. I was eager to read their comments and discuss the work with them and it was with major excitement that I read their reports and answered their questions during my thesis defense.

Finally, I feel that I have to acknowledge Claude BERROU for believing in my capabilities and allowing me to work with him in the context of Deep Learning when I knew almost nothing of the domain. I would have never worked in the field if not for this.

Résumé

Abstract	ix
Introduction	x
Chapitre 2: Concepts en apprentissage profond	xix
Chapitre 3: Concepts en traitement de signaux sur graphe	xx
Chapitre 4: Réseaux de neurones profonds pour des signaux sur graphe	xxi
Chapitre 5: Représentations latentes de réseaux profonds sur graphe	xxii

Abstract

Ces dernières années, les méthodes d'apprentissage profond ont atteint l'état de l'art dans une vaste gamme de tâches d'apprentissage automatique, y compris la classification d'images et la traduction automatique. Ces architectures sont assemblées pour résoudre des tâches d'apprentissage automatique de bout en bout. Afin d'atteindre des performances de haut niveau, ces architectures nécessitent souvent d'un très grand nombre de paramètres. Les conséquences indésirables sont multiples, et pour y remédier, il est souhaitable de pouvoir comprendre ce qui se passe à l'intérieur des architectures d'apprentissage profond. Il est difficile de le faire en raison de: i) la dimension élevée des représentations, and ii) la stochasticité du processus de formation. Dans cette thèse, nous étudions ces architectures en introduisant un formalisme à base de graphes, s'appuyant notamment sur les récents progrès du traitement de signaux sur graphe (TSG). À savoir, nous utilisons des graphes pour représenter les espaces latents des réseaux neuronaux profonds. Nous montrons que ce formalisme des graphes nous permet de répondre à diverses questions, notamment: i) mesurer des capacités de généralisation, ii) réduire la quantité de des choix arbitraires dans la conception du processus d'apprentissage, iii) améliorer la robustesse aux petites

perturbations ajoutées sur les entrées, and iv) réduire la complexité des calculs.

Introduction

Ces dernières années, les réseaux de neurones profonds (DNN) ont explosé en popularité, créant un nouveau domaine appelé «Apprentissage Profond» [39]. Si le concept de réseaux de neurones [134] et les DNN [138] sont tous deux assez anciens, ils n'ont commencé à gagner en popularité que ces dernières années. Ce changement est dû aux deux avancées en matériel, spécialement les cartes graphiques (GPU) [44] et aux premières victoires dans les défis de vision par ordinateur comme AlexNet [82] gagnant le LSRVC 2012-Imagenet [139] et DanNet [20] remportant le «Contest on Mitosis Detection in Breast Cancer Histological Images» [136]. La figure 1 présente des exemples d'images issues de ces concours.

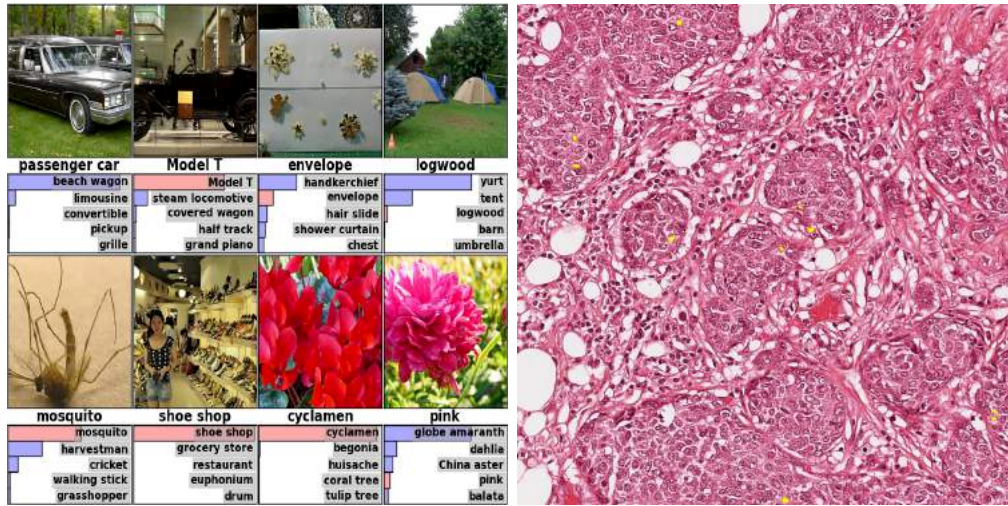


Figure 1: **Gauche:** Huit images du «test set» de ILSVRC-2010 et les cinq étiquettes les plus probables par AlexNet. L'étiquette correcte est écrite au dessous de chaque image, et la probabilité que le réseau a assigné pour la bonne étiquette est montrée par la barre rouge (seulement si la bonne étiquette est dans le top5). Cette image est tirée de [82] ©2012 Neural Information Processing Systems Foundation. **Droite:** Image tirée de [136] pour la détection de mitose, où les parties jaunes sont les parties que le DNN doit détecter

Ces réseaux ont été construits sur la base de deux grands principes : i) un a-priori convolutif, and ii) des représentations hiérarchiques apprises par rétropropagation du gradient d'erreur. Le premier guide la forme de base du réseau, afin d'imposer l'invariance aux translations et le partage des poids. Le second assure que c'est à la méthode d'optimisation d'adapter le réseau à un extracteur de caractéristiques suivi d'un

classificateur, sans contrôle spécifique de l'évolution de ce dernier. En effet, dans [93] les auteurs disent : «A potentially more interesting scheme is to eliminate the feature extractor, feeding the network with “raw” inputs (e.g. normalized images), and to rely on backpropagation to turn the first few layers into an appropriate feature extractor». Par conséquent, nous pouvons considérer le traitement d'une entrée dans un DNN comme la génération d'une séquence de **représentations intermédiaires** qui font partie des **espaces latents** du DNN.

Pour mieux comprendre ce que nous appelons un DNN et les représentations intermédiaires, nous allons illustrer les DNN dans la figure 2 et les représentations intermédiaires dans la figure 3. Notons comment les représentations intermédiaires sont adaptées à diverses résolutions et concepts abstraits [181], par exemple dans la figure 3 nous pouvons dire que la couche 2 est spécialisée pour détecter des coins et des bords tandis que la couche 5 est spécialisée pour des objets entiers.

Contexte et motivation

Comme nous l'avons dit précédemment, les architectures d'apprentissage profond sont capables d'atteindre l'état de l'art dans de nombreux défis dans le domaine de l'apprentissage machine. Elles le font parce qu'elles sont capables d'exploiter la quantité colossale d'informations disponibles. Elles sont souvent présentées comme un cas extrême de méthodes basées sur les données, où il n'a pas de connaissance sur la forme de la fonction à trouver. En tant que telles, elles souffrent de quelques inconvénients:

1. Elles contiennent de nombreux paramètres qui sont réglés à l'aide de routines d'optimisation complexes qui dépendent à la fois de leur initialisation et des données d'entraînement. En conséquence, elles sont souvent déployées comme des boîtes noires associant des entrées à des sorties. Il existe peu de théories capables de fournir des résultats exploitables sur les mécanismes de ces boîtes noires;
2. Il est tout à fait habituel d'observer une optimalité de pareto entre la complexité des modèles et les performances sur les tâches considérées. Autrement dit, pour atteindre l'état de l'art, les modèles nécessitent un grand nombre de paramètres, de calculs et de mémoire [44];
3. Il est juste de dire que les méthodes d'apprentissage profond ont connu un grand succès grâce à leurs performances expérimentales. Les modèles proposés ont connu plusieurs générations de complexification depuis le renouvellement du domaine au début des années 2010. Il existe donc un écart croissant entre ce que la théorie de

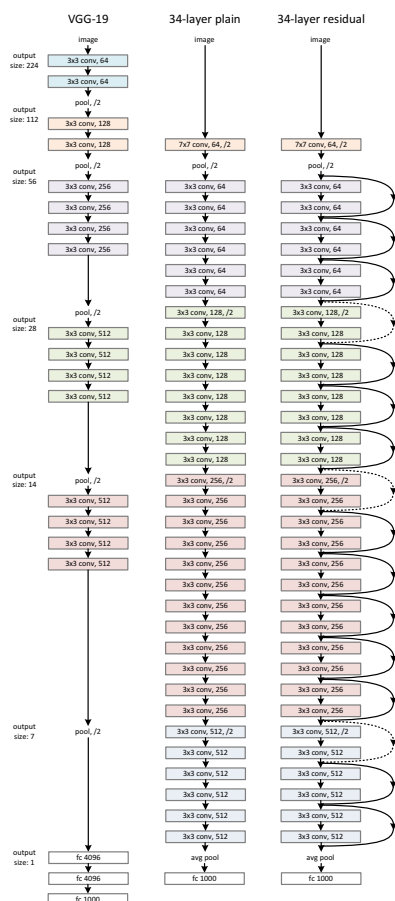


Figure 2: Exemples de architectures de DNNs. **Gauche:** VGG-19 [149] (19.6 milliards de FLOPs). **Centre:** un réseau avec 34 couches (3.6 milliards de FLOPs). **Droite:** un réseau résiduel [50] avec 34 couches (3.6 milliards FLOPs). Image retiré de [50] ©2016 IEEE.

l'apprentissage profond peut expliquer et ce que les solutions pratiques actuelles mettent en œuvre pour résoudre les problèmes.

DNN - une «boîte noire»

Comme nous l'avons présenté au tout début de ce résumé, le changement de paradigme consistant à passer de caractéristiques et de modèles triés sur le volet à des architectures d'apprentissage profond a été le principe directeur des recherches récentes dans le domaine. Si des caractéristiques conçues par un expert humain sont considérées comme bien maîtrisées ou interprétables, les DNN n'ont de leur côté aucun contrôle explicite, ce qui conduit à un très haut degré de liberté et à des solutions qui sont basées à 100% sur des données. De manière empirique, il a été constaté que les DNN ont tendance à dépasser

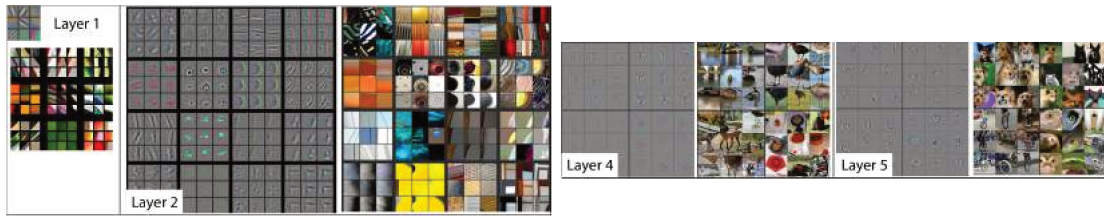


Figure 3: Visualisation des caractéristiques dans un modèle entraîné. Pour les couches (2,4,5), les 9 principales activations sont représentées à partir d'un sous-ensemble aléatoire de cartes d'éléments dans les données de validation, projetées dans l'espace des pixels en utilisant une approche de réseau déconvolutif. Les reconstructions sont des échantillons du modèle : il s'agit de modèles reconstruits à partir de l'ensemble de données de validation qui provoquent des activations élevées dans une carte de caractéristiques donnée. Pour chaque carte de caractéristiques, nous montrons également les patches d'images correspondants. On note : (i) le regroupement important dans chaque carte de caractéristiques, (ii) une plus grande invariance dans les couches supérieures et (iii) l'exagération des parties discriminantes de l'image. Il est préférable de visualiser l'image sous forme numérique. Image et légende adaptées de [181].

les performances des systèmes experts au fur et à mesure que la quantité de données disponibles augmente.

Si cela a conduit à diverses améliorations de l'apprentissage machine et a permis à l'apprentissage profond d'être l'état de l'art de la plupart des tâches d'apprentissage machine, cela entraîne divers inconvénients. Par exemple, dans les domaines où l'interprétabilité est essentielle, comme le domaine médical [110], la précision des modèles d'apprentissage profond pourrait ne pas être un argument assez fort pour permettre son adoption mondiale. Ouvrir la «boîte noire» et comprendre les mécanismes sous-jacents de chaque étape d'une architecture d'apprentissage profond est un problème encore ouvert que nous abordons partiellement dans cette thèse.

Quantité de paramètres et de calculs

Un autre problème qui découle de l'utilisation des DNN est la quantité de paramètres et les exigences en calculs. Les physiciens citent fréquemment la célèbre déclaration de von Neumann : "... with four parameters I can fit an elephant, and with five I can make him wiggle his trunk." [28] pour soutenir que les modèles d'apprentissage machine ont tendance à être sur-paramétrés. D'un autre côté, nous avons maintenant des praticiens de l'apprentissage profond qui dépassent fréquemment les millions et parfois même les

milliards de paramètres [14]. Il est ainsi très compliqué de gérer la liberté dont disposent les modèles d'apprentissage profond et de faire en sorte qu'ils apprennent le comportement souhaité.

En effet, si les éléments constitutifs des architectures d'apprentissage approfondi sont des fonctions très simples, la quantité totale de paramètres rend très difficile l'interprétation exacte de ce qui est traité dans un DNN. En outre, la complexité de calcul de ces modèles s'est accrue, ce qui nécessite non seulement du matériel spécialisé (comme les GPU) mais aussi une grande consommation d'énergie. L'étude de modèles d'apprentissage profond et la réduction de la quantité de paramètres et de calculs nécessaires sont non seulement nécessaires du point de vue des connaissances sous-jacentes (par exemple, l'interprétabilité et la robustesse), mais aussi en raison de problèmes sociétaux plus complexes tels que le coût environnemental du matériel et de l'énergie nécessaires pour développer et utiliser ces modèles, ainsi que pour rendre les systèmes d'apprentissage profond accessibles à la plupart des gens.

Manque de compréhension théorique

Comme la plupart des recherches dans le domaine de l'apprentissage profond sont fortement axées sur les applications, l'expérimentation et les résultats de référence sont bien souvent la partie la plus importante des articles récents. Bien qu'il soit très important de comparer les différentes méthodes utilisées dans la littérature, en particulier dans un domaine qui évolue aussi rapidement que l'apprentissage approfondi, cela peut présenter plusieurs inconvénients:

1. Les améliorations peuvent provenir de petits détails de mise en œuvre plutôt que de la théorie sous-jacente. Considérons par exemple les problèmes de reproductibilité dans le domaine de l'apprentissage du renforcement [52], où il est parfois impossible de reproduire un résultat sans regarder l'implémentation directe du code au lieu de regarder simplement l'article;
2. Les améliorations peuvent ne pas être en accord avec la théorie qui leur a été présentée, par exemple les couches de «batch normalization» (BN) [64] ont été proposées pour traiter le décalage des covariables («covariate shift»), ce que d'autres chercheurs soutiennent qu'elles ne sont pas adaptées pour traiter [23, 140, 185], mais les couches BN sont toujours la pierre fondamentale de diverses architectures;
3. Des comparaisons injustes entre les méthodes peuvent se produire en raison de la combinaison de différentes méthodes ou de la référence elle-même. Par exemple,

dans l'apprentissage des métriques, il a été démontré que les méthodes traditionnelles peuvent être plus performantes que les méthodes plus récentes si elles sont correctement formées (c'est-à-dire dans des conditions d'égalité avec les méthodes récentes) [135].

Graphes pour représenter les espaces latents des réseaux neuronaux profonds

En résumé, les lacunes présentées dans les paragraphes précédents peuvent être considérées comme provenant de la principale force de l'apprentissage profond, à savoir : le fait que les modèles sont capables d'utiliser leur grande liberté pour apprendre des fonctions très complexes, alors que le fonctionnement sous-jacent de chaque étape individuelle n'est pas compris ou explicitement contrôlé. Dans cette thèse, nous proposons d'attaquer ces inconvénients en nous concentrant sur l'étude des représentations intermédiaires dans les DNN. En effet, l'étude des représentations intermédiaires peut être considérée comme une «ouverture de la boîte noire des DNN» et vise à mieux comprendre ce qui est traité à chaque étape. Notez que l'étude des représentations intermédiaires n'augmente pas la complexité du calcul (il faut de toute façon les calculer) et il a déjà été démontré qu'elles contiennent des informations importantes, par exemple la compression des DNN par la distillation des connaissances [55, 133].

L'étude des représentations intermédiaires et de l'effet global de chaque couche intermédiaire devrait permettre une compréhension plus fine des différentes propriétés des DNN, telles que la robustesse [87] et la généralisation globale [43]. Dans cette thèse, nous nous concentrons sur l'étude des représentations latentes/intermédiaires des DNN. Comme nous l'avons vu dans les paragraphes précédents, les principaux inconvénients de l'apprentissage profond viennent de sa force centrale : le fait qu'il est capable d'exploiter pleinement les données disponibles sans aucune contrainte forte. Cela tend à conduire à des réseaux très complexes où il est difficile de comprendre à quoi chaque partie est destinée, ainsi qu'à des difficultés pour évaluer si la fonction apprise est une bonne approximation de celle qui est visée.

Pour contrer cet inconvénient, nous analysons et proposons de nombreuses méthodes qui exploitent les connaissances intrinsèques des représentations intermédiaires des DNN. Dans la suite de ce document, nous présentons les définitions nécessaires pour comprendre les domaines explorés de l'apprentissage profond et les méthodes que nous proposons. En particulier, nous nous concentrons sur trois de ces domaines i) apprentissage de la représentation et du transfert des éléments, ii) compression des architectures, and

iii) surapprentissage (généralisation et robustesse). Afin d'effectuer l'analyse nécessaire et de concevoir les méthodes, nous utilisons le cadre du **traitement de signal sur graphe (TSG)** [148].

Nous avons choisi le cadre du TSG car il étend l'analyse harmonique traditionnelle aux domaines irréguliers représentés par des graphes. Les graphes présentent un avantage unique dans le domaine de l'apprentissage profond car ils exploitent les relations des données elles-mêmes. Cela est très conforme à la philosophie de l'apprentissage profond où les données sont essentielles. Par conséquent, l'utilisation des graphes nous permet d'étudier les DNN et leurs représentations intermédiaires car elle fournit un support pour les relations qui sont générées à chaque représentation intermédiaire. Cela facilite l'étude des représentations intermédiaires des DNN, car nous pouvons examiner les données et leurs relations au lieu de l'espace irrégulier de grande dimension.

Afin d'illustrer et de donner une idée de ce que sont les graphes d'une représentation intermédiaire, nous décrivons et illustrons un exemple dans les paragraphes suivants.

Exemple: Illustration des graphes de représentation intermédiaire

Considérons que nous avons un DNN déjà entraîné sur un ensemble de données. Nous construisons trois graphes de similarité où les sommets correspondent aux échantillons et les arêtes relient les échantillons les plus similaires. Nous le construisons en utilisant un petit sous-ensemble de l'ensemble de données. Le premier graphe utilise les représentations de l'espace initial (l'espace des images) et les deux derniers utilisent les représentations intermédiaires du DNN. Ces représentations proviennent d'une couche intermédiaire et d'une des couches finales. Sur le plan qualitatif, nous nous attendons à ce que les échantillons qui appartiennent à une même classe soient plus faciles à séparer à mesure que nous nous enfonçons dans l'architecture considérée. Ceci serait en accord avec la définition citée de [93]. Nous décrivons cet exemple dans la figure 4. Comme prévu, nous pouvons voir qualitativement la différence de séparation entre l'espace image et les espaces latents, comme on peut le constater par la quantité d'arêtes entre les éléments de classes distinctes (nous gardons le même nombre k de voisins sur chaque graphe) et aussi par la séparation géométrique lorsque nous utilisons des «Laplacian eigenmaps» [6] pour placer les différents échantillons dans un espace 2D régulier.

Notez que ces représentations illustrent clairement le principe de démêlage dans les réseaux neuronaux profonds. En effet, les DNN peuvent être considérés comme une cascade d'opérations qui transforment l'espace d'entrée dans lequel les échantillons d'une même classe sont mélangés avec d'autres, en espaces latents qui sont progressivement

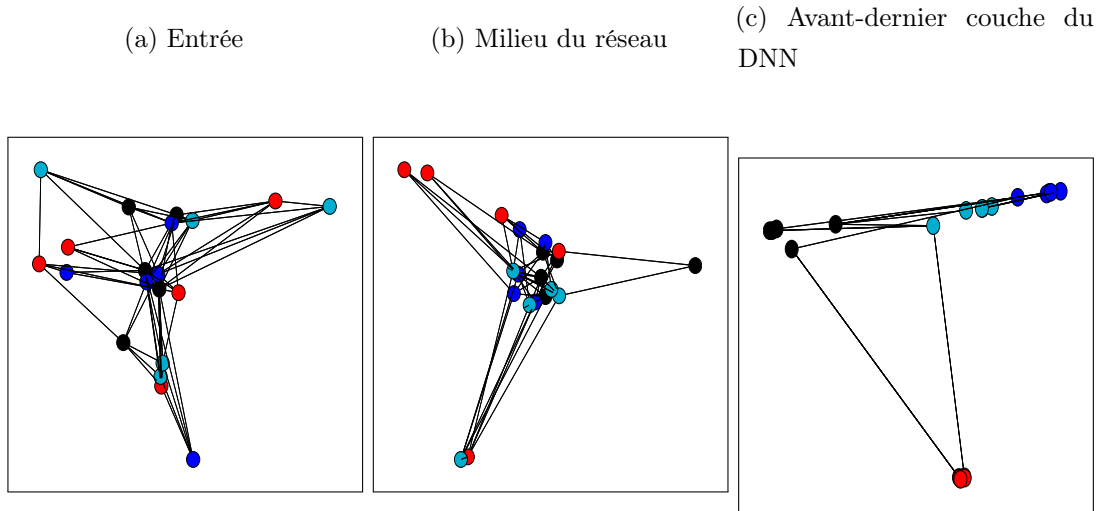


Figure 4: Exemple de représentation sous forme de graphes, de l’espace des entrées (gauche) à l’avant dernière couche du DNN (droite). Les différentes couleurs des sommets représentent la classe de l’objet. Pour faciliter la visualisation, nous ne représentons que les arêtes entre des exemples de classes distinctes. Notez qu’il y a beaucoup plus de bords à l’entrée (a) et que le nombre de bords diminue au fur et à mesure que l’on s’enfonce dans l’architecture (b et c).

mieux alignés sur la tâche considérée à laquelle le DNN a été entraîné.

Contributions

Nous considérons que cette thèse comporte trois types de contributions différentes. Premièrement, nous avons cherché à rendre toutes les productions de la thèse aussi ouvertes/libres que possible. Cela nous est très cher, car le but de cette thèse n’était pas de générer une “application de premier ordre” ou une preuve de concept, mais d’étudier les orientations de recherche que nous pensons être d’intérêt pour la communauté. Pour atteindre cet objectif, nous avons mis à disposition la plupart de notre production textuelle (par exemple des articles) soit sur des sites d’archives bien connus (tels que arxiv.org), soit sur les pages web personnelles des auteurs. Nous avons également mis à disposition, lorsque cela était possible, le code utilisé pour les expériences et les preuves de concepts sur le site de contrôle de version github.com. Nous fournissons une liste des contributions à la fin de cette section.

Deuxièmement, nous avons fait un effort pour communiquer nos résultats et diffuser les connaissances via des présentations, via la conception de cours et via l’enseignement. Nous avons présenté nos résultats lors de diverses conférences et ateliers internationaux

et nationaux, afin de promouvoir et de discuter de nos résultats avec la communauté scientifique. Nous avons également consacré une partie de la thèse à l'enseignement et à la conception de cours dans les domaines de la théorie des graphes et de l'apprentissage automatique, y compris les cours ouverts massifs en ligne (MOOC). En fait, au cours de mon doctorat, j'ai eu l'occasion de contribuer à la création de deux cours. Le premier est un MOOC intitulé «Advanced Algorithmics and Graph Theory with Python», disponible sur le platform EdX et qui a rassemblé plus de 10 000 étudiants de plus de 50 pays depuis son lancement en 2018. Le second est un cours d'introduction à l'IA moderne qui est conçu pour les étudiants de l'IMT Atlantique. Dans les deux cas, j'ai participé à la fois à la conception générale et aux parties techniques des cours.

Enfin, nous nous sommes consacrés à l'étude des technologies qui, selon nous, devraient contribuer à la société dans son ensemble. Prenons par exemple deux des domaines que nous avons étudiés : la compression et la robustesse des DNN. Dans le premier, nous visons à réduire la consommation totale d'énergie (et donc les émissions de carbone) des réseaux neuronaux, indépendamment de la tâche en aval. Dans le second, nous visons à accroître la confiance générale que nous pouvons avoir dans les résultats générés par un DNN.

Dans les paragraphes suivants, nous présentons une liste des contributions de cette thèse :

- Lassance, C. E. R. K., Gripon, V., and Ortega, A. (2018a). Laplacian networks: Bounding indicator function smoothness for neural network robustness. *arXiv preprint arXiv:1805.10133*, under journal review since 01/2020
- Lassance, C. E. R. K., Vialatte, J.-C., and Gripon, V. (2018b). Matching convolutional neural networks without priors about data. In *2018 IEEE Data Science Workshop (DSW)*, pages 234–238. IEEE
- Lassance, C. E. R. K., Gripon, V., and Ortega, A. (2018c). Predicting under and overfitting in deep neural networks using graph smoothness. *2018 Graph Signal Processing Workshop (Non-archival)* available at <https://cadurosar.github.io/papers/GSP2018.pdf>
- Lassance, C., Latif, Y., Garg, R., Gripon, V., and Reid, I. (2019b). Improved visual localization via graph smoothing. *arXiv preprint arXiv:1911.02961*
- Lassance, C., Gripon, V., Tang, J., and Ortega, A. (2019). Structural robustness for deep learning architectures. In *2019 IEEE Data Science Workshop (DSW)*, pages 125–129

- Lassance, C., Gripon, V., and Mateos, G. (2020b). Graph topology inference benchmarks for machine learning. *arXiv preprint arXiv:2007.08216, to appear in 2020 IEEE 30th International Workshop on Machine Learning for Signal Processing (MLSP)*
- Lassance, C., Bontonou, M., Hacene, G. B., Gripon, V., Tang, J., and Ortega, A. (2020a). Deep geometric knowledge distillation with graphs. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8484–8488. IEEE
- Grelier, N., Lassance, C. E. R. K., Dupraz, E., and Gripon, V. (2018). Graph-projected signal processing. In *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 763–767. IEEE
- Bontonou*, M., Lassance*, C., Hacene, G. B., Gripon, V., Tang, J., and Ortega, A. (2019). Introducing graph smoothness loss for training deep learning architectures. In *2019 IEEE Data Science Workshop (DSW)*, pages 160–164, * authors contributed equally
- Bontonou, M., Lassance, C., Vialatte, J.-C., and Gripon, V. (2019). A unified deep learning formalism for processing graph signals. *arXiv preprint arXiv:1905.00496*
- Bontonou, M., Lassance, C., Gripon, V., and Farrugia, N. (2019). Comparing linear structure-based and data-driven latent spatial representations for sequence prediction. In *Wavelets and Sparsity XVIII*, volume 11138, page 111380Z. International Society for Optics and Photonics

Dans le suivi de cette résumé, nous décrivons via un sommaire rapide chacun des chapitres de cette thèse.

Chapitre 2: Concepts en apprentissage profond

Dans ce premier chapitre, nous présentons les réseaux neuronaux profonds (DNN), en nous concentrant plus particulièrement sur les architectures résiduelles. Nous introduisons également le concept de représentations intermédiaires dans les DNN, qui sera un élément central des chapitres suivants de cette thèse. Ces représentations intermédiaires peuvent être utilisées afin d’effectuer un apprentissage par transfert, tel que présenté dans la section 2.1.2 et également afin d’abstraire les DNN en tant qu’extracteurs de caractéristiques suivis de classificateurs.

Nous présentons divers problèmes pour lesquels les DNN sont pertinents. Ces problèmes vont être étudiés plus en détail dans les prochains chapitres et comprennent les tâches suivantes: i) localisation basée sur la vision, ii) classification des images, iii) classification des tâches neurologiques, and iv) classification des documents.

Nous faisons également une revue de littérature sur la compression des réseaux de neurones, notamment les méthodes de distillation et les couches de convolution plus efficaces. Nous présentons SAL, une contribution sur le sujet des couches de convolution efficaces, qui a fait l'objet de l'article de conférence suivant:

- Hacene, G. B., Lassance, C., Gripon, V., Courbariaux, M., and Bengio, Y. (2019). Attention based pruning for shift networks. *arXiv preprint arXiv:1905.12300*, to appear in *25th International Conference on Pattern Recognition (ICPR2020)*

En outre, nous introduisons le concept de robustesse d'un classificateur, et nous démontrons empiriquement comment il peut être lié à la capacité de bien fonctionner en présence d'entrées corrompues. Ce concept de robustesse et ses expériences empiriques ont été publiés dans l'article de conférence suivant:

- Lassance, C., Gripon, V., Tang, J., and Ortega, A. (2019). Structural robustness for deep learning architectures. In *2019 IEEE Data Science Workshop (DSW)*, pages 125–129

Chapitre 3: Concepts en traitement des signaux sur graphe

Dans ce chapitre, nous introduisons les concepts de graphes et de signaux de graphes, ainsi que les outils nécessaires du cadre du traitement des signaux de graphes (TSG). Ces concepts et outils nous permettent d'analyser les représentations latentes profondes et d'en tirer de nouvelles contributions destinées à la communauté d'apprentissage machine et qui seront présentées dans les chapitres suivants.

Parmi les outils présentés dans cette section figurent la transformée de Fourier sur graphe (GFT) et l'analyse de la fluidité des signaux de graphes. Nous abordons également les méthodes permettant de déduire des graphes à partir de données pour lesquelles la structure de support des graphes n'est pas explicitement disponible, y compris une nouvelle contribution :

- Lassance, C., Gripon, V., and Mateos, G. (2020b). Graph topology inference

benchmarks for machine learning. *arXiv preprint arXiv:2007.08216, to appear in 2020 IEEE 30th International Workshop on Machine Learning for Signal Processing (MLSP)*

Nous dérivons ensuite des filtres de graphe, qui émulent les filtres traditionnels de traitement du signal dans le domaine des graphes. Ces filtres de graphe seront utilisés pour relier les couches convolutives et les couches convolutives de graphes dans le prochain chapitre. Nous présentons également deux applications de filtres de graphes qui nous permettent de réduire la quantité de bruit des éléments extraits à l'aide de DNN et d'améliorer les performances des tâches en aval, notamment : l'apprentissage avec peu d'exemples ; la classification des images ; la localisation visuelle (VBL) et l'extraction d'images (IR). L'application de localisation visuelle a fait l'objet d'une contribution :

- Lassance, C., Latif, Y., Garg, R., Gripon, V., and Reid, I. (2019b). Improved visual localization via graph smoothing. *arXiv preprint arXiv:1911.02961*

Chapitre 4: Réseaux de neurones profonds pour des signaux sur graphe

Dans ce chapitre, nous approfondissons le domaine des réseaux neuronaux profonds définis sur des graphes. Nous nous sommes appuyés sur les concepts des chapitres précédents afin de définir les méthodes récentes dans un cadre de filtrage sur graphes unique que nous avons présenté par ordre croissant de complexité dans la section 4.1. Bien que ce cadre ne soit pas exactement nouveau, nous l'avons étendu à d'autres méthodes et avons introduit une discussion sur les inconvénients de ces méthodes.

Nous avons ensuite discuté des applications des DNN définis sur les graphes dans le contexte de la classification supervisée des signaux des graphes dans la section 4.2. Nous avons discuté des contributions récentes qui montrent les inconvénients des approches actuelles dans ce domaine et avons ensuite présenté deux de nos contributions. Leur objectif est de combler l'écart entre les convolutions des graphes et les convolutions 2D/3D classiques. Nos deux contributions introduites ont été publiées dans des conférences :

- Grelier, N., Lassance, C. E. R. K., Dupraz, E., and Gripon, V. (2018). Graph-projected signal processing. In *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 763–767. IEEE

- Lassance, C. E. R. K., Vialatte, J.-C., and Gripon, V. (2018b). Matching convolutional neural networks without priors about data. In *2018 IEEE Data Science Workshop (DSW)*, pages 234–238. IEEE

Enfin, nous discutons des applications dans le contexte de la classification semi-supervisée des sommets d'un graphe. Nous discutons d'abord du problème de l'évaluation équitable des différentes méthodes GNN sur cette tâche. Bien qu'il ne s'agisse pas d'un problème nouveau dans le domaine, les travaux récents présentent encore les deux écueils les plus courants : i) l'utilisation d'une seule répartition train/validité/essai dont il a déjà été démontré qu'elle faussait les résultats [145], and ii) des expériences ne comparant pas équitablement les méthodes, par exemple, la méthode A est plus performante que la méthode B, mais cela est principalement dû à l'ajout du dropout plutôt qu'à la méthode elle-même. Notez que ces problèmes ne sont pas nécessairement dus à une faute de connaissance ou à une malveillance, mais surtout à un simple problème de quantité des calcul qui seraient nécessaires pour tout exécuter correctement. En effet, nous proposons un cadre afin de résoudre à la fois les problèmes i) et ii), mais nous montrons que nous ne pourrions jamais exécuter la version optimale dans un délai raisonnable. Nous avons ajouté un cadre plus souple et présentons nos résultats sur l'ensemble de données de Cora.

Chapitre 5: Représentations latentes de réseaux profonds sur graphe

Dans ce chapitre nous présentons principalement nos contributions dans le domaine de l'utilisation de graphes pour représenter les espaces latents de réseaux de neurones profonds. Bien que ce domaine ne soit pas très développé, nous espérons que nos contributions pourront apporter un éclairage et permettre de le développer davantage, car nous pensons qu'il y a beaucoup de contributions intéressantes à poursuivre.

Nous présentons d'abord le travail qui a été le début de notre intérêt pour le domaine [43], dans lequel les auteurs ont montré qu'il était possible de caractériser différents comportements de DNN en analysant l'évolution de la fluidité d'un signal sur un graphe. Nous nous sommes ensuite appuyés sur ces travaux pour proposer une mesure qui est empiriquement corrélée avec la performance de généralisation des DNN. Cette mesure a fait l'objet d'une contribution à une conférence :

- Lassance, C. E. R. K., Gripon, V., and Ortega, A. (2018c). Predicting under and overfitting in deep neural networks using graph smoothness. *2018 Graph Signal*

Processing Workshop (Non-archival) available at <https://cadurosar.github.io/papers/GSP2018.pdf>

Nous nous sommes ensuite concentrés sur les utilisations possibles de la fluidité du signal sur le graphe pendant la formation des réseaux de neurones. Nous avons d’abord montré que nous sommes capables de former de bons extracteurs de caractéristiques en entraînant le réseau à minimiser la fluidité des signaux des indicateurs d’étiquettes sur les graphes générés par leurs sorties. Cette nouvelle fonction objectif possède trois caractéristiques importantes qui ne sont pas présentes dans l’entropie croisée et nous démontrons à l’aide d’expériences que nous sommes capables d’obtenir des réseaux plus robustes, sans perdre trop de performance de généralisation. Deuxièmement, nous proposons d’utiliser un régularisateur afin de contrôler l’évolution de la fluidité des signaux indicateurs des étiquettes sur les graphes qui sont générés par les représentations intermédiaires des DNN. Nous montrons que ces régulariseurs sont non seulement théoriquement conformes à notre définition de la robustesse (Definition 2.5.2), mais aussi que nous pouvons démontrer empiriquement leur efficacité lorsqu’ils sont comparés (ou ajoutés) à d’autres méthodes dans la littérature. Ces deux utilisations de la fluidité d’un signal sur graphe ont fait l’objet de contributions, l’une à une conférence et l’autre est en cours d’examen dans une revue :

1. Bontonou*, M., Lassance*, C., Hacene, G. B., Gripon, V., Tang, J., and Ortega, A. (2019). Introducing graph smoothness loss for training deep learning architectures. In *2019 IEEE Data Science Workshop (DSW)*, pages 160–164, * authors contributed equally
2. Lassance, C. E. R. K., Gripon, V., and Ortega, A. (2018a). Laplacian networks: Bounding indicator function smoothness for neural network robustness. *arXiv preprint arXiv:1805.10133*, under journal review since 01/2020

Enfin, nous présentons une méthode qui ne s’appuie pas sur le cadre du TSG, mais qui nous permet d’utiliser le cadre du TSG sur des techniques préalablement définies. En d’autres termes, nous avons spécialisé le cadre RKD en GKD, dont nous avons démontré empiriquement et analytiquement qu’il permettait d’améliorer les performances des réseaux compressés. Ce travail d’introduction a été publié lors d’une conférence :

- Lassance, C., Bontonou, M., Hacene, G. B., Gripon, V., Tang, J., and Ortega, A. (2020a). Deep geometric knowledge distillation with graphs. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8484–8488. IEEE

Conclusion

L'idée principale que nous avons poursuivie au cours des trois dernières années était de remédier à certaines lacunes des architectures d'apprentissage profond en examinant leurs représentations intermédiaires. Pour effectuer nos analyses, nous avons utilisé le cadre du traitement du signal des graphes, dans lequel les graphes sont utilisés pour représenter la topologie d'un domaine complexe (ici : les espaces latents des architectures d'apprentissage profond). Nous avons considéré les applications d'apprentissage profond dans trois domaines d'apprentissage machine : i) apprentissage et transfert des représentations, ii) compression des architectures d'apprentissage profond, and iii) étude du sur-apprentissage (generalisation et robustesse).

Chapter 1

Introduction

1.1	Context and motivation	2
1.2	Graphs for deep learning latent representations	6
1.3	Document structure	10

In recent years, Deep Neural Networks (DNN) have exploded in popularity, creating a new domain called "Deep Learning" [39]. While both the concept of neural networks [134] and DNNs [138] are quite old, they only started gaining popularity in recent years. This change was due to both advances in hardware, specially Graphic Processing Unit (GPUs) [44] and with the first victories in computer vision challenges such as AlexNet [82] winning the 2012 CVPR **L**arge **S**cale **V**isual **R**ecognition **C**hallenge (LSRVC-Imagenet) [139] and DanNet [20] winning the Contest on Mitosis Detection in Breast Cancer Histological Images of ICPR 2012 [136]. Figure 1.1 depict example images from these competitions.

These networks were built upon two main principles: i) convolutional priors, and ii) hierarchical backpropagation-learned representations. The former guides the base form of the network, in order to enforce invariance to shifts and weight sharing. The latter informs that the network should receive the input as-is and it is thus the job of the optimization method to adapt the network to a feature extractor followed by a classifier, without any specific control of the network evolution. Indeed, in [93] the authors say: "A potentially more interesting scheme is to eliminate the feature extractor, feeding the network with "raw" inputs (e.g. normalized images), and to rely on backpropagation to turn the first few layers into an appropriate feature extractor". Therefore, we can look at the processing of an input in a DNN as the generation of a sequence of **intermediate**

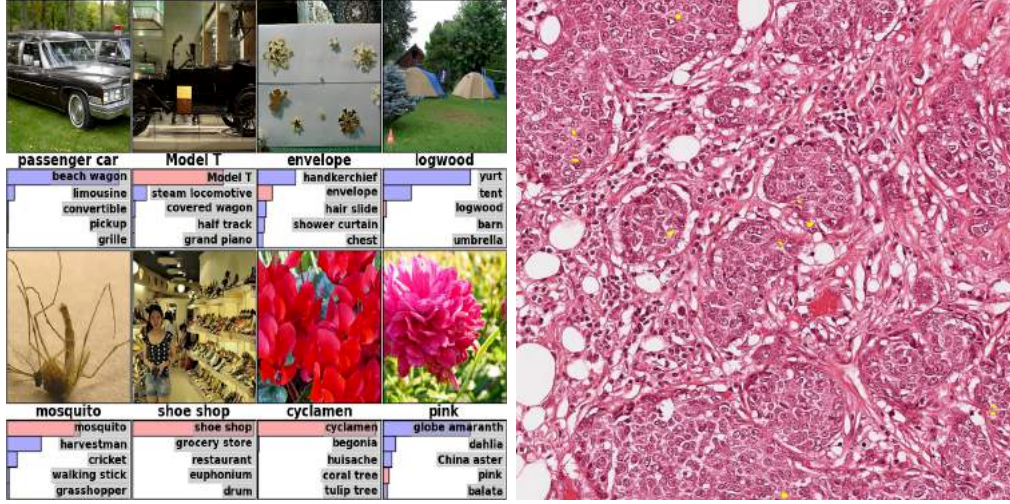


Figure 1.1: **Left:**Eight ILSVRC-2010 test images and the five labels considered most probable by AlexNet. The correct label is written under each image, and the probability assigned to the correct label is also shown with a red bar (if it happens to be in the top 5). Image extracted from [82] @2012 Neural Information Processing Systems Foundation. **Right** Example image from [136] for mitosis detection, where the yellow parts represent the parts that should be detected by the DNN

representations that are part of the **latent spaces** of the DNN.

To better understand what we call a DNN and an intermediate representation, let us illustrate these concepts. We first depict in Figure 1.2 some typical deep neural networks. Note how they tend to follow a mostly sequential structure, with few shortcuts.

Now, in Figure 1.3 we depict the intermediate representation evolution from one layer to the next in the same architecture. Note how the intermediate representations are adapted to various abstract resolutions and concepts [181], for example in Figure 1.3 we can say that layer 2 responds to corners/edges while layer 5 responds to entire objects.

1.1 Context and motivation

As we said in the previous section, deep learning architectures are able to reach state-of-the-art performance in many challenges in the field of machine learning. They do so because they are able to exploit the colossal amount of information contained in the training data. They are often presented as an extreme case of data-driven methods (i.e. a discriminative approach), where very few priors are given about the function to be found.

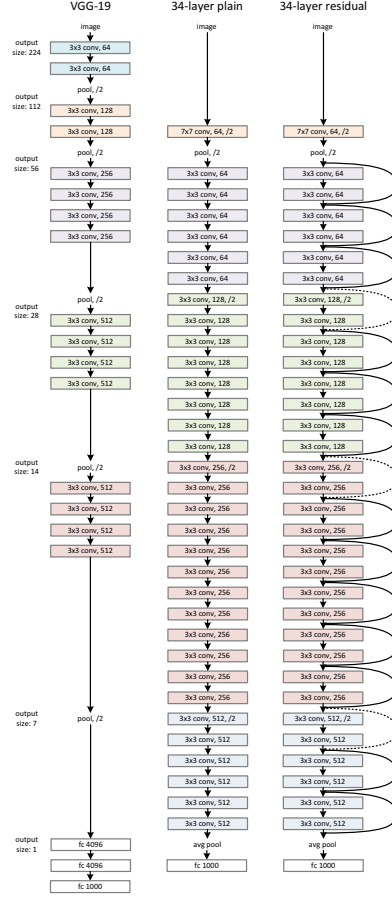


Figure 1.2: Examples of network architectures. **Left:** the VGG-19 model [149] (19.6 billion FLOPs). **Middle:** a plain network with 34 parameter layers (3.6 billion FLOPs). **Right:** a residual network [50] with 34 parameter layers (3.6 billion FLOPs). Image extracted from [50] ©2016 IEEE.

As such, they suffer from the same shortcomings as most discriminative approaches:

1. They contain a lot of parameters that are tuned using complex optimization routines that depend on both their initialization and on the training data. As a consequence, they are often seen as black boxes associating inputs with outputs. There is little theory able to provide exploitable results about the inside of these black boxes;
2. It is quite usual to observe a pareto optimality between complexity of the models and performance on the considered tasks. Said otherwise, in order to reach state-of-the-art accuracy, models require a huge number of parameters, computations and memory [44];
3. It is fair to say that deep learning methods have known a great success thanks to their experimental performance. Proposed models have known several generations

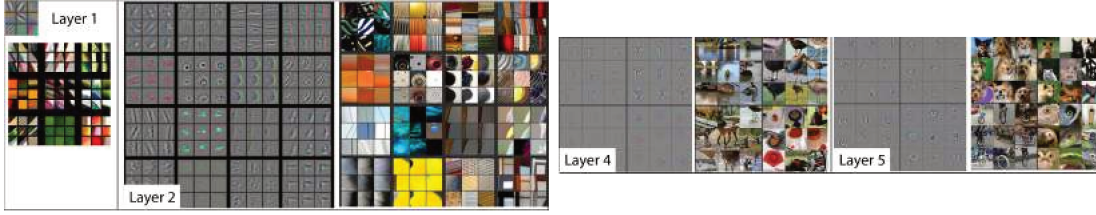


Figure 1.3: Visualization of features in a fully trained model. For layers (2,4,5) the top 9 activations are shown from a random subset of feature maps across the validation data, projected down to pixel space using a deconvolutional network approach. The reconstructions are *not* samples from the model: they are reconstructed patterns from the validation set that cause high activations in a given feature map. For each feature map, we also show the corresponding image patches. Note: (i) the strong grouping within each feature map, (ii) greater invariance at higher layers and (iii) exaggeration of discriminative parts of the image. Best viewed in electronic form. Figure and caption adapted from [181].

of complexifications since the renewal of the domain in the early 2010s. There is therefore an increasing gap between what the theory of deep learning can explain and what current practical solutions implement to solve complex problems.

1.1.1 DNN as a black box

As we have introduced at the very beginning of this document, changing the paradigm from hand-picked features and models to deep learning architectures has been the guiding principle of recent deep learning research. The former is seen as well-behaved or interpretable as the features and models are designed to solve the task in a very specific way. The latter however thrives in the fact that no explicit control is performed¹, which leads to a very high degree of freedom and to solutions that are 100% data based. Empirically it has been found that the latter tends to overtake the former as the amount of available data increases.

While this has led to various improvements in machine learning and has allowed deep learning to be the de-facto state-of-the-art of most machine learning tasks, this leads to various drawbacks. For example, in domains where interpretability is key, such as the medical domain [110], the accuracy of deep learning models might not be an argument strong enough to allow its global adoption. To “open” the black box and understand the underlying mechanisms of each step of a deep learning architecture is still an open problem that we partially tackle in this thesis.

¹safe for some priors such as shift-invariance or temporal connections.

1.1.2 Amount of parameters and computational complexity

Another problem that arises from the use of DNNs is the amount of parameters and the computational requirements of recent DNN models. Physicists frequently cite the famous von Neumann statement “... with four parameters I can fit an elephant, and with five I can make him wiggle his trunk.” [28] to argue that machine learning models tend to be overparametrized. On the other hand we now have deep learning practitioners frequently surpass the millions and sometimes even billions of parameters [14]. Dealing with the amount of freedom that deep learning models have and enforcing that they learn the desired behaviour is very complicated.

Indeed, while the building blocks of deep learning architectures are very simple functions, the total amount of parameters make it very hard to interpret exactly what is being processed inside a DNN. Moreover, the computational complexity of these models have been increasing, which not only requires dedicated hardware (such as GPUs) but a large amount of energy consumption. Studying deep learning models and reducing the amount of parameters and computations needed is not only necessary from an underlying knowledge aspect (e.g. interpretability and robustness) but also due to more complex societal problems such as the environmental cost of material and energy necessary for developing and using those models, as well as making deep learning systems accessible to most.

1.1.3 Lack of theoretical understanding

As most of the research on the domain of deep learning is highly application focused, experimentation and benchmark results are now seen as the most important part of recent papers. While it is very important in order to compare different methods in the literature, especially in a domain that evolves as quickly as deep learning, it may come with several drawbacks:

1. Improvements may come from small implementation details instead of the underlying theory. Consider for example the reproducibility concerns in the domain of reinforcement learning [52], where sometimes it is not possible to reproduce a result without looking at the direct code implementation instead of just looking at the paper;
2. Improvements may not agree with the theory they were presented with, for example batch normalization (BN) [64] layers were proposed to deal with covariate shift,

what other researchers argue that they are not suited to address [23, 140, 185], but the BN layers are still the cornerstone of various architectures;

3. Unfair comparisons between methods may happen due to the combination of different methods or from the benchmark itself. For example, in metric learning it has been shown that traditional methods may outperform more recent methods if they are properly trained (i.e. in equality of conditions with the recent methods) [135].

1.1.4 Limitations of deep discriminative models

In summary, the shortcomings presented in the previous paragraphs can be seen as originating from the main strength of deep learning, that is: the fact that the models are able to use their high amount of freedom to learn very complex functions, while the underlying functioning of each individual step is not understood or explicitly controlled. In this thesis we propose to attack these three drawbacks by concentrating on the study of intermediate representations in DNNs. Indeed, studying the intermediate representations can be seen as “opening” the black box of the DNN and aiming to better understand what is being processed at each step. Note that looking at the intermediate representations does not increase the computational complexity (one has to compute them anyway) and it has already been shown that they contain important information, e.g., compression of DNNs via knowledge distillation [55, 133].

Studying the intermediate representations and the overall effect of each intermediate layer should lead to a more fine-grain understanding of different properties of the DNNs, such as robustness [87] and overall generalization [43]. In the following paragraphs, we describe how we propose to leverage these representations in order to contribute to deep learning research.

1.2 Graphs for deep learning latent representations

In this thesis, we focus on studying the latent/intermediate representations of DNNs. As we have discussed in the previous paragraphs, the main drawbacks of deep learning come from its central strength: the fact that it is able to fully leverage the available data without any strong constraint. This tends to lead to very complex networks where it is difficult to understand what each part is meant for, as well as difficulties in assessing whether the learned function is a good approximation of the targeted one.

To counter this drawback we analyze and propose many different methods that exploit

the intrinsic knowledge from the intermediate representations of DNNs. In the following of this document we present the definitions needed to understand the explored domains of deep learning and the methods we propose. In particular, we focus on three such domains i) representation and transfer learning, ii) compression of architectures, and iii) overfitting (generalization and robustness). In order to perform the needed analysis and design the methods we use the framework of **Graph Signal Processing (GSP)** [148].

We have chosen the GSP framework as it extends traditional harmonic analysis to irregular domains represented by graphs. Graphs have a unique advantage in the deep learning domain as they exploit the relationships from the data itself. This is very inline with the philosophy of deep learning where data is key. Therefore, using graphs allows us to study the DNNs and intermediate representations as it provides a support for the relationships that are generated at each intermediate representation. This facilitates the study of the intermediate representations of DNNs, as we can look at the data and its relationships instead of the highly-dimensional irregular space.

In order to illustrate and to give an idea of what are intermediate representation graphs, we describe and depict an example in the following paragraphs.

1.2.1 Example: Depiction of intermediate representation graphs

Consider that we have a pre-trained DNN on an image dataset. What does it look like if we create our graph representations in such a scenario? In order to create such a depiction, we construct three similarity graphs where vertices correspond to samples and edges connect samples that are the most similar. We build it by using a small subset of the training dataset. The first graph uses the representations from image space and the latter two use the intermediate representations of the DNN. Such representations come from an intermediate layer and one of the final (end) layers. What we expect to see qualitatively is that the samples that belong to a same class will be easier to separate as we go deeper in the considered architecture. This would be in hand with the quoted definition from [93] in the first page of this thesis. We depict this example in Figure 1.4. As expected, we can qualitatively see the difference in separation from the image space to the latent spaces, as can be noted by the amount of edges between elements of distinct classes (we keep the same number k of neighbors at each graph) and also by the geometric separation when using Laplacian eigenmaps [6] to position the different samples in a regular 2D space.

Note that these representations clearly illustrate the principle of disentangling in deep neural networks. Indeed, DNNs can be thought of as a cascade of operations that smoothly transform the input space in which samples from a same class are likely to be

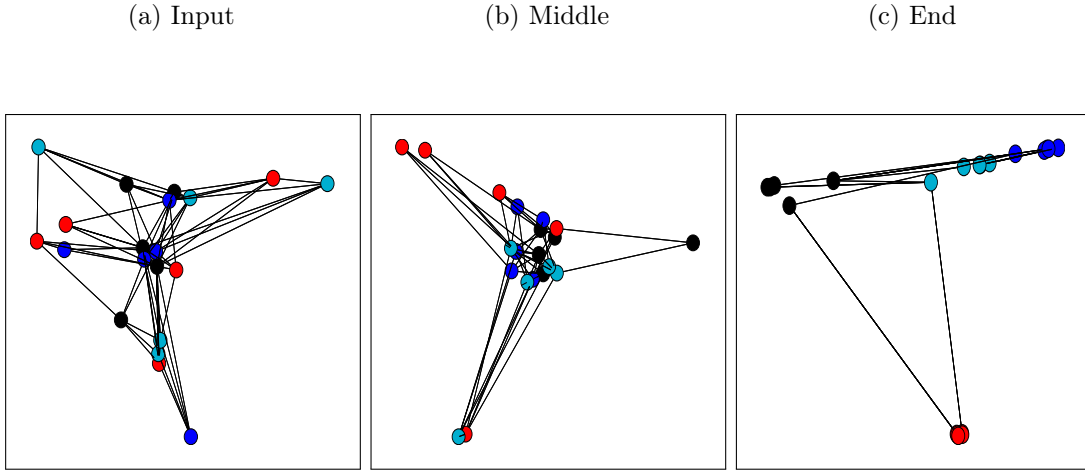


Figure 1.4: Graph representation example, from the input space (left) to the output of the network (right). The different vertex colors represent the class of the object. To help the visualization, we only depict the edges between examples of distinct classes. Note how there are many more edges at the input (a) and how the number of edges decrease as we go deeper in the architecture (b and c).

spread among other ones, to latent spaces that progressively are better aligned with the considered task the network is trained upon.

1.2.2 Contributions

We consider this thesis to have three different types of contributions. First, we have sought to turn all productions of the thesis as open access/free as possible. This is very dear to us, as the goal of this thesis was not to generate a so called “top notch application” or proof-of-concept, but to investigate research directions that we believe to be of interest to the community. To achieve this goal we have made available most of our textual production (e.g. articles) either at well known archival websites (such as arxiv.org) or at the authors personal webpages. We have also made available, when possible, the code responsible for the experiments and proofs of concepts on the version control site github.com. We provide a list of contributions at the end of this subsection.

Second, we have done an effort to communicate our results and disseminate knowledge via presentations, course designs and teaching. We have presented our findings in various international and national conferences and workshops, in order to promote and discuss our results with the scientific community. We have also dedicated a part of the thesis to the teaching and course designs in the domains of graph theory and machine learning,

including open Massive Online Open Courses (MOOCs). As a matter of fact, during my PhD I had the opportunity to contribute to the creation of two courses. The first one is a MOOC entitled “Advanced Algorithmics and Graph Theory with Python” that is available on the EdX platform and that has gathered more than 10k students from 50+ countries since its launch in 2018. The second one is an introductory course to modern AI that is designed for students at IMT Atlantique. In both cases, I have participated to both the overall design and to the technical parts of the courses.

Finally, we have dedicated ourselves to the study of technologies that we believe should contribute to society at large. For example consider two of the domains that we have studied: compression and robustness of DNNs. In the former we aim to reduce the total energy consumption (and therefore carbon emissions) of neural networks, independently of the downstream task. In the latter, we aim to increase the general confidence that we can have on the results generated by a DNNs.

In the following paragraphs, we present a list of the academic contributions of this thesis:

- Lassance, C. E. R. K., Gripon, V., and Ortega, A. (2018a). Laplacian networks: Bounding indicator function smoothness for neural network robustness. *arXiv preprint arXiv:1805.10133*, under journal review since 01/2020
- Lassance, C. E. R. K., Vialatte, J.-C., and Gripon, V. (2018b). Matching convolutional neural networks without priors about data. In *2018 IEEE Data Science Workshop (DSW)*, pages 234–238. IEEE
- Lassance, C. E. R. K., Gripon, V., and Ortega, A. (2018c). Predicting under and overfitting in deep neural networks using graph smoothness. *2018 Graph Signal Processing Workshop (Non-archival)* available at <https://cadurosar.github.io/papers/GSP2018.pdf>
- Lassance, C., Latif, Y., Garg, R., Gripon, V., and Reid, I. (2019b). Improved visual localization via graph smoothing. *arXiv preprint arXiv:1911.02961*
- Lassance, C., Gripon, V., Tang, J., and Ortega, A. (2019). Structural robustness for deep learning architectures. In *2019 IEEE Data Science Workshop (DSW)*, pages 125–129
- Lassance, C., Gripon, V., and Mateos, G. (2020b). Graph topology inference benchmarks for machine learning. *arXiv preprint arXiv:2007.08216*, to appear in *2020 IEEE 30th International Workshop on Machine Learning for Signal Processing (MLSP)*

- Lassance, C., Bontonou, M., Hacene, G. B., Gripon, V., Tang, J., and Ortega, A. (2020a). Deep geometric knowledge distillation with graphs. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8484–8488. IEEE
- Grelier, N., Lassance, C. E. R. K., Dupraz, E., and Gripon, V. (2018). Graph-projected signal processing. In *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 763–767. IEEE
- Bontonou*, M., Lassance*, C., Hacene, G. B., Gripon, V., Tang, J., and Ortega, A. (2019). Introducing graph smoothness loss for training deep learning architectures. In *2019 IEEE Data Science Workshop (DSW)*, pages 160–164, * authors contributed equally
- Bontonou, M., Lassance, C., Vialatte, J.-C., and Gripon, V. (2019). A unified deep learning formalism for processing graph signals. *arXiv preprint arXiv:1905.00496*
- Bontonou, M., Lassance, C., Gripon, V., and Farrugia, N. (2019). Comparing linear structure-based and data-driven latent spatial representations for sequence prediction. In *Wavelets and Sparsity XVIII*, volume 11138, page 111380Z. International Society for Optics and Photonics

In addition to these contributions, we are currently working on a book chapter dedicated to the use of graphs to represent latent spaces of DNNs.

1.3 Document structure

Overall this manuscript is divided into four parts. First, we present an introduction of the two main subjects of this thesis: i) Deep Learning (Chapter 2), and ii) Graph Signal Processing (Chapter 3). In these two chapters we define concepts that were essential to the thesis in our own words and introduce contributions that are directly linked to either deep learning or GSP.

Second we introduce and discuss the domain of “Deep Learning for inputs supported on graphs” in Chapter 4. This domain combines both Deep Learning and GSP concepts and studies the application of deep learning methods to inputs that are defined in the graph domain. As we did not have a specific focus on these types of applications during the thesis, we present it more as an overview using a proposed mathematical framework and discuss some applications.

Third, we introduce the main contribution of this thesis, which is the development of the domain of graph based methods for improving deep learning. We do so by studying the intermediate representations of deep neural networks using the formalism from GSP. In contrast with the previous part, we do not need for the inputs to be defined over a graph domain to be able to deploy our methods. In this part we are actually introducing a new domain of research: the study of general intermediate representations of DNNs using the GSP framework.

Finally, we present a summary in Chapter 6, including a quick recall of our contributions and the research directions that are now open for future work.

Chapter 2

Classification and feature extraction with Deep Neural Networks

2.1	Definitions	14
2.2	Deep Learning Layers	27
2.3	Tasks and Datasets	34
2.4	Compression	43
2.5	Robustness	48
2.6	Summary of the chapter	61

In this chapter, we introduce the concepts of classification and feature extraction using Deep Neural Networks (DNNs). This chapter is organized as follows: first in Section 2.1, we introduce and define neural networks, then in Section 2.2, we introduce the layers used in the scope of this work and in Section 2.3 we introduce the datasets considered in this thesis. Finally, we introduce compression tools in Section 2.4 and robustness definitions in Section 2.5.

2.1 Definitions

Deep Neural Networks (DNNs) contain the term “neural” as they are loosely inspired by the functioning of brain neurons. However, it is fair to say that this inspiration is becoming less important in the recent developments in the field. This is why in this Chapter we adopt a network-based definition of these models.

So let us consider a DNN architecture. Such an architecture is mathematically described by its “network function” f . We call f “deep” as it is obtained through a long cascading sequence of intermediate functions from its input to its output. More precisely, f receives an input tensor \mathbf{x} , which typically represents the pixel values of an image, and outputs a corresponding tensor $f(\mathbf{x})$ which dimensions and interpretation depends on the task for which the network was initially designed.

There exists a lot of ways to obtain deep neural network functions, but the simplest to formalize mathematically consists in a composition of layer functions:

$$f = f^{\ell_{\max}} \circ f^{\ell_{\max}-1} \circ \dots \circ f^1. \quad (2.1)$$

Here, each function f^ℓ , called “layer function”, is highly constrained. Indeed, each f^ℓ is typically defined as a parametrized linear function followed by a non-parametric non-linear function.

In modern literature, it is rare to encounter such constructions of deep neural network functions. Instead, many authors use residual networks [50]. Indeed, residual networks have been demonstrated to reach state-of-the-art performance in many challenges in the context of classification. Residual networks (**Resnet**) are composed of blocks of layers, as depicted in Figure 2.1. Note that even in the case of Resnets, the core idea remains that network functions are built as an assembly of layer functions.

In the literature one can find many types of layers. The most notable ones are i) Fully connected, ii) Convolutional, iii) Pooling, iv) Normalization, v) Graph Convolutional, vi) Recurrent Neural Network, and vii) Long Short-Term Memory. In Section 2.2, we will describe the first four items in more detail; item 5 will be introduced and detailed in Chapter 4; items 6 and 7 are outside the scope of this work.

The outputs of layer functions are called intermediate representations.

Definition 2.1.1 (intermediate representation). We call the output of an intermediate function f^ℓ an intermediate representation. In other words, in the simple case of architectures that can be written using Equation 2.1, \mathbf{x}^ℓ is the intermediate representation

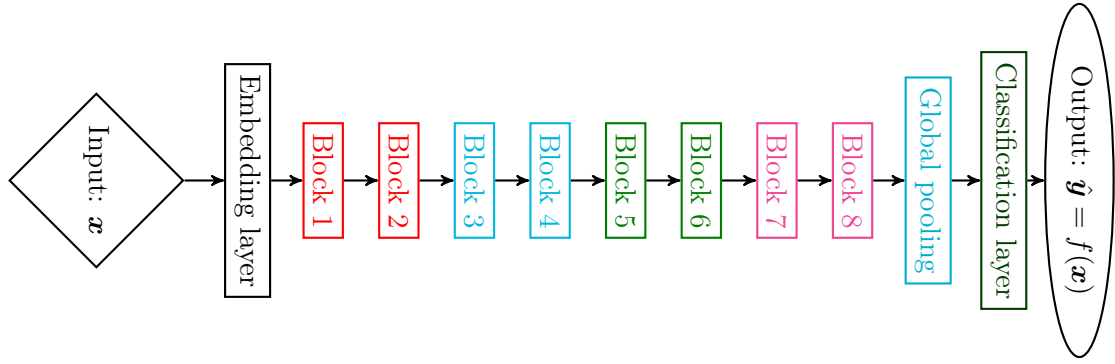


Figure 2.1: Simplified depiction of a Resnet with eight residual blocks, divided into four block groups. We depict the residual blocks in Figure 2.2. The color of each block indicates their group block and that they have the same dimensions.

generated by applying the network function f from f^1 until f^ℓ on \mathbf{x} where ℓ represents the depth in the DNN architecture.

The goal of the intermediate representations is to capture the internal state of the DNN. In simple architectures that respect Equation 2.1, they also obey the Markov property (i.e., they fully capture the actual state of the input traversing the network and are sufficient to compute the output). Note that in more complex architectures this property does not necessarily holds.

Recall the concept of Resnets that are grouped in blocks of layers. These blocks define splits f', f'', \dots . The output of each block is fully characterized by its weights and its inputs. Intermediate representations obtained between blocks are therefore Markovian. However, this property is not valid inside a block, as a residual connection exists. We depict residual network blocks in Figure 2.2.

The function f is characterized using tunable values called **parameters**. Initially, these parameters are randomly sampled in a distribution \mathcal{N} , so that the output of the network function can be interpreted as a random projection of data. In the context of classification, the most used framework is to output a class-wise classification score $\hat{\mathbf{y}}$. Therefore, in order to train the network function to solve the classification task, it is common to use the **label indicator vector** \mathbf{y} associated with the input \mathbf{x} .

Definition 2.1.2 (label indicator vector). A binary vector with as many coordinates as the number of classes in the problem. Only the coordinate corresponding to the class of input \mathbf{x} is set to 1 while all the other coordinates are zeroed.

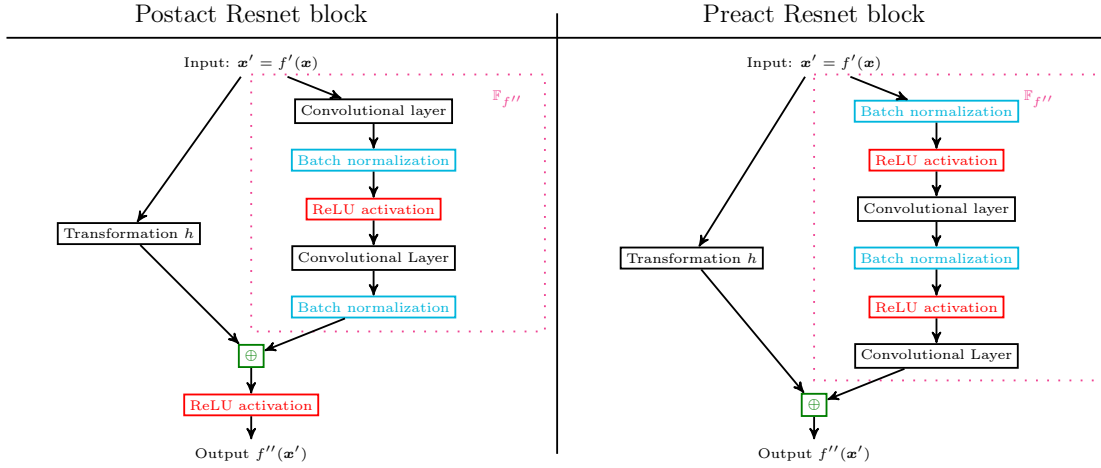


Figure 2.2: Depiction of the preact and postact residual network blocks proposed in [50]. The transformation h ensures that $h(\mathbf{x}')$ and $\mathbb{F}_{f''}(\mathbf{x}')$ have the same dimensions so that we can perform the sum operation between them. h is normally implemented as either the identity operation (if both \mathbf{x}' and $\mathbb{F}_{f''}(\mathbf{x}')$ have the same dimensions) or as a convolution layer if the dimensions differ. In many cases, this convolution has only one parameter for each feature map.

In classification tasks, the goal of a DNN is to correctly classify all the inputs \mathbf{x} from the domain of possible inputs \mathbb{D} . As it is impossible to collect all the possible inputs for most tasks, we use a subset \mathcal{D} that we call dataset. Therefore, the parameters of the DNN are tuned during a learning phase using a dataset \mathcal{D} and an objective function \mathcal{L} that measures the discrepancy between the outputs of the network functions and expected label indicator vectors, i.e., discrepancy between \mathbf{y} and $\hat{\mathbf{y}}$. We present the datasets used in this thesis in Section 2.3.

Usually, the objective function is a loss function, which is minimized over a subset of the dataset that we call “**training set**” ($\mathcal{D}_{\text{train}}$), composed of training examples \mathbb{X} . This optimization is usually performed using variants of the stochastic gradient descent algorithm [12]. As such, the network function f , which is typically composed of a vast number of parameters, is adjusted during the learning phase. We will see in Section 2.4 exactly how vast is this number of parameters and some techniques aiming at reducing the number of computations and memory needed for both training and inference in neural networks.

Note that in the context of vision, it is quite common to introduce a **data augmentation** scheme to go alongside our training.

Definition 2.1.3 (data augmentation). We define data augmentation as the act of

artificially generating new inputs \mathbf{x}_{DA} to increase the size of $\mathcal{D}_{\text{train}}$ by performing a set of transformations \mathbb{H}_{DA} on the inputs $\mathbf{x} \in \mathcal{D}_{\text{train}}$.

Data augmentation is indeed an instrumental technique, as it allows us to increase the size of $\mathcal{D}_{\text{train}}$ without the cost of drawing new labeled samples from \mathbb{D} . There are two main types of data augmentation, which we refer to as **domain-driven** and **data-driven**. The former uses the knowledge one has over the domain \mathbb{D} to design transformations h which are known to generate valid new inputs without compromising the nature of their corresponding class. Typical domain-driven data augmentation in image scenarios include randomly removing a small part of the image (also called *random crop*) and horizontal flipping.




Data-driven data augmentation, on the other hand, uses information from the dataset to generate new inputs. Doing so allows us to have a significant advantage as one does not need to be a specialist on the domain \mathbb{D} to propose the data augmentation scheme. However, it may lead to training on inputs \mathbf{x}_{da} that are outside of the domain \mathbb{D} and possibly inputs \mathbf{x}_{da} that are misclassified.

Two of the most used data-driven techniques are *autoaugment* [22] where one tries various data augmentation schemes at random and keeps the ones that work the best in terms of final accuracy of the model and *mixup* [184] where both the input \mathbf{x}_1 and its desired output \mathbf{y}_1 are interpolated with another example which input is \mathbf{x}_2 and output is \mathbf{y}_2 to generate a new input \mathbf{x}_{da} and its associated output \mathbf{y}_{da} . Note that both data-driven techniques come with drawbacks. Autoaugment still requires some domain knowledge (to design the data augmentation schemes that are tested). On the other hand, mixup may be incompatible with some datasets and requires adapting the objective function. We present some examples of data augmented samples in Table 2.1 (mixup) and in Figure 2.3 (autoaugment).

After the learning phase comes the *inference phase*. During inference, one first fixes the weights of the network and then evaluates its performance. Note that training and inference phases can be performed iteratively, and most of the time alternate. At the end of the training, we obtain the final architecture with its corresponding weights. This architecture obtains a score (most often the score consists in measuring the accuracy of the model on a dataset that is distinct from the training set).

Most of the time, the whole process of obtaining the final score of an architecture is repeated several times for two main reasons: i) to verify that the score is robust against a different initialization of the parameters and sampling of the training examples,, and ii) to search for more efficient **hyperparameters** of the network.

Table 2.1: Examples of mixup based data-augmentation samples. Figures extracted from [179] ©2019, IEEE.

	\mathbf{x}	$\mathbf{x}_{da}[184]$	$\mathbf{x}_{da}[179]$
Image			
Label indicator vector	Dog 1.0	Dog 0.5 Cat 0.5	Dog 0.6 Cat 0.4



















	Original	Sub-policy 1	Sub-policy 2	Sub-policy 3	Sub-policy 4	Sub-policy 5
Batch 1						
Batch 2						
Batch 3						
		Equalize, 0.4, 4 Rotate, 0.8, 8	Solarize, 0.6, 3 Equalize, 0.6, 7	Posterize, 0.8, 5 Equalize, 1.0, 2	Rotate, 0.2, 3 Solarize, 0.6, 8	Equalize, 0.6, 8 Posterize, 0.4, 6

Figure 2.3: Examples of autoaugment based data-augmentation samples. Figure extracted from [22]: ©2019, IEEE.

Definition 2.1.4 (hyperparameters). Hyperparameters are a set of parameters that are fixed during training. For example, the parameters concerning the architecture of the DNN (e.g., the number of layers in the network) and the training methodology are considered to be hyperparameters of a DNN.

During the inference phase there are two types of evaluation we can consider: i) memorization, and ii) generalization. In the former, the goal is to verify that the network can correctly classify the inputs used during the training phase, i.e., the inputs in $\mathcal{D}_{\text{train}}$. *Generalization*, on the other hand, aims at evaluating the capacity of the network to extrapolate from the seen inputs and correctly classify unseen ones, i.e., not presented during the training.

Generalization is a fundamental property for DNNs, as it is not possible to access

the full domain \mathbb{D} . This is why alongside of the training set it also common to define a **validation set** $\mathcal{D}_{\text{valid}}$ and a **test set** $\mathcal{D}_{\text{test}}$. The main goal is that all sets come from the same distribution \mathcal{D} , but do not intersect $\mathcal{D}_{\text{train}} \cap \mathcal{D}_{\text{valid}} = \emptyset$, $\mathcal{D}_{\text{valid}} \cap \mathcal{D}_{\text{test}} = \emptyset$ and $\mathcal{D}_{\text{train}} \cap \mathcal{D}_{\text{test}} = \emptyset$. We illustrate \mathbb{D} and its relationship with the dataset in Figure 2.4.

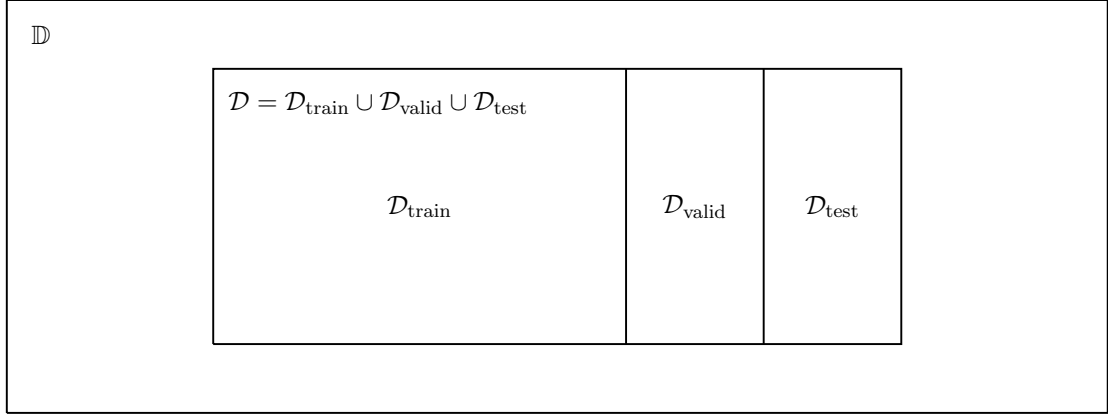


Figure 2.4: Diagram depicting the relationship between \mathbb{D} , \mathcal{D} , $\mathcal{D}_{\text{train}}$, $\mathcal{D}_{\text{valid}}$ and $\mathcal{D}_{\text{test}}$.

The *validation set* is to be used during the iterative phases of training and inference. In this case, the validation set helps one make decisions about the architecture and the training process of the network. On the other hand, the *test set* should not be used to influence the decisions during the iterative phases of the process. The test set should serve only as an a posteriori measure of performance.

This setting with validation and test sets has been proposed to avoid the shortcoming of **overfitting**.

Definition 2.1.5 (overfitting). In this work, we define overfitting as a three-fold phenomenon:

1. *Overfit to $\mathcal{D}_{\text{train}}$* : a network that has an excellent performance on memorization, but an inferior one on generalization is said to be overfitted to $\mathcal{D}_{\text{train}}$. Ideally, we would aim for both evaluations having similar performance.
2. *Overfit to $\mathcal{D}_{\text{valid}}$* : on the other hand, having an excellent generalization performance to the $\mathcal{D}_{\text{valid}}$ is not sufficient. By using $\mathcal{D}_{\text{valid}}$ to define the hyperparameters, the trained network can be biased to $\mathcal{D}_{\text{valid}}$. For this reason, it is recommended to have two distinct sets for evaluating generalization, one used for tuning hyperparameters ($\mathcal{D}_{\text{valid}}$) and a second one as an external measure of performance ($\mathcal{D}_{\text{test}}$).
3. *Overfit to \mathcal{D}* : finally, even if the network shows excellent generalization to both $\mathcal{D}_{\text{valid}}$ and $\mathcal{D}_{\text{test}}$, it could still be overfitted to the dataset, i.e., the network would

not generalize to the rest of the domain \mathbb{D} . Another way to say that a network is overfitted to \mathcal{D} is to say that it is not robust. We delve into more details on this problem in Section 2.5.

Given this definition of overfitting, it is quite surprising that in the recent literature, most works ignore the use of a $\mathcal{D}_{\text{valid}}$. Instead, they mostly use the union of $\mathcal{D}_{\text{train}}$ and $\mathcal{D}_{\text{valid}}$ as the de facto training set, and the $\mathcal{D}_{\text{test}}$ is used to tune the hyperparameters. Fortunately, this seems not to be biasing the DNNs to the $\mathcal{D}_{\text{test}}$ as described in a recent work [132] that draws a new $\mathcal{D}_{\text{test}2}$ from the same \mathbb{D} and shows that the networks that performed the best in $\mathcal{D}_{\text{test}}$ also were the best in $\mathcal{D}_{\text{test}2}$. However, this finding only covers the first two definitions of overfitting.

As the domain \mathbb{D} grows, it is harder to represent it accurately with the subset \mathcal{D} . Therefore it is often the case that the construction of the dataset biases neural networks, i.e., they may disregard features that are not present or underrepresented on \mathcal{D} . We discuss the robustness of neural networks to inputs that are in the domain \mathbb{D} but are not represented in \mathcal{D} in Section 2.5.

2.1.1 Residual networks (Resnet)

There is a vast number of DNN architectures proposed in recent literature [50, 122, 157, 163, 180]. In this work we mostly use residual networks [50], that we introduce in more detail this subsection. This architecture was chosen for three main reasons:

1. Ease of training: as we are going to introduce in the next paragraphs, residual networks are easy to train and to find the correct hyperparameters;
2. Standard of the literature: residual networks are used very often in the literature, therefore choosing this architecture allows us an easier and fairer comparison with recent works;
3. Certified correct implementation: as residual networks are both easy to train and widely used in recent literature, there exist standardized implementations of residual networks for almost all languages and frameworks, removing the possible fail-point that our implementation is incorrect or differs from the works we compare to.

Residual networks are named because they introduce residual connections between layers in the DNN. The main goal of residual connections is to ease the training of deep neural networks. Residual network paths are formed using the original input and a set

of sequential layers. The input goes through the set of sequential layers generating an output. This output is then summed with a simple and direct transformation of the original input. This additive path, called residual path, is also commonly called a block. More formally, we define a residual block f'' that receives an input \mathbf{x}' by:

$$f''(\mathbf{x}') = g(h(\mathbf{x}') + \mathbb{F}_{f''}(\mathbf{x}')), \quad (2.2)$$

where h is a simple transformation that ensures the sum is performed between two tensors with the same dimensions, g is an optional activation function and $\mathbb{F}_{f''}$ is the set of sequential layers that belongs to block f'' .

One key interest of residual networks is that a block can easily implement the identity function, and as such deeper architectures can emulate shallower ones. Among other interests, being able to behave as a shallower network eases the training process. In other words, residual DNN can behave like a shallower network at the start of the training in order to be able to warm up, and when their weights are well-conditioned, they can start to use their entire depth and reap the benefits of deeper networks. In the literature, it is not rare to see residual networks with hundreds of layers.

The property of behaving like a shallower network and its influence in training deep residual networks was demonstrated in recent works, where authors [23, 185] show that one can remove the normalizations of the DNN by starting the network biased for the shallower paths. In Figure 2.5 we depict the two residual network blocks used in this work (*preact* and *postact* blocks) and the residual connection. Note that these two types of blocks are not the only ones used in the literature and that the *preact* block is slightly different from the original presentation in [50] and in Figure 2.2. This choice was made because most of the recent literature shifted to this design.

As shown in Figure 2.1, Resnets start with a first convolutional layer to increase the feature map size of the input to F_{initial} , sometimes called embedding layer. The residual blocks then follow the embedding layer. After the residual blocks, it is common to add a pooling layer. This pooling layer is sometimes called global if the 3D intermediate representation $([F, w, h])$ is downsampled to an 1D representation $([F])$. A fully connected layer then follows this downsampling. This FC layer is sometimes called the classification layer.

To determinate the depth of Resnets and, therefore, their nomenclature, one has to look at the number of blocks of the network and how they are constructed. Note that Resnet blocks do not always treat the same feature map size, and it is quite common to split these blocks into groups, where each block may either end or start with a downsampling operation. We use the nomenclature $\text{Resnet}n\text{-}w_i$ where n is the number of layers, and w_i

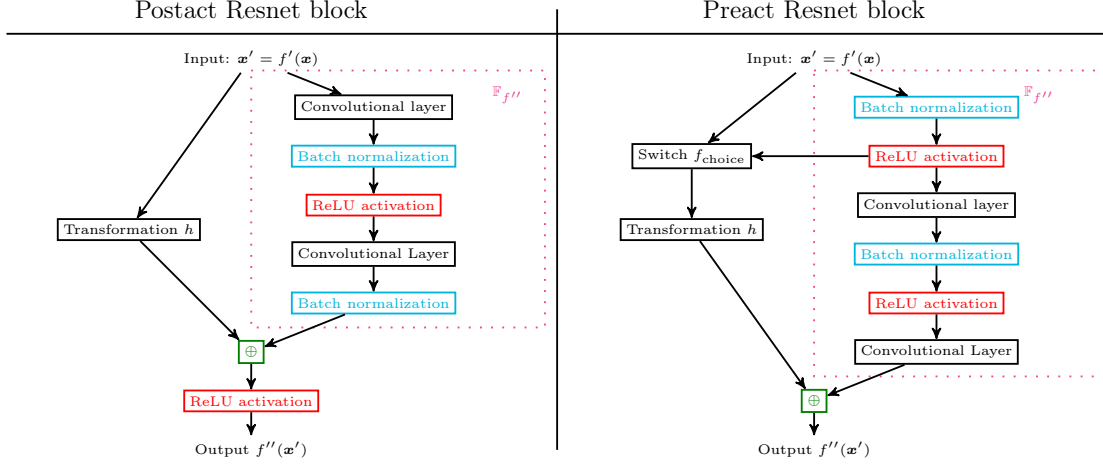


Figure 2.5: Depiction of the preact and postact residual network blocks used in this work. The switch f_{choice} decides the input of h based on the size of the input \mathbf{x}' and the output from $\mathbb{F}_{f''}(\mathbf{x}')$: if both have the same dimensions the direct path is chosen, if not the indirect (the one that starts on the right and then goes to the left) is chosen. We note that the blocks depicted here are different from the ones in Figure 2.2 as most works in recent literature shifted to this design.

is a widen factor first defined in WideResnets [180].

The amount of layers n is characterized by the following equation $n = 2(cd + 1)$, where c is the number of convolutional layers per block, and d is the number of blocks. Unfortunately this notation can be misleading because two Resnet18- w could have very different block configurations (e.g, $[2, 2, 2, 2]$, $[1, 3, 1, 3]$ and $[3, 3, 2]$ would be valid Resnet18- w configurations).

It is also standard to define a common feature map amount for all convolutions of the same block group and to define a feature map scaling that doubles at each downsampling operation. We recall that WideResnets [180] also add a widen factor w_i that is a multiplier to the feature map amounts from the first convolution to the first block, allowing the network to be wider. We note that WideResnets follow a different nomenclature ($n = 2(cd + 2)$). As we try to be consistent on our nomenclature, networks in this work are always presented following $2(cd + 1)$, which may lead to confusion if one is familiarized with WideResnets (e.g., WideResnet28-10 will be presented as WideResnet26-10 in this work).

In more detail, in this work, we only use Resnets with blocks with two convolutions, as presented in Figure 2.5 and groups with an equal amount of blocks. We also limit our configurations to have either 3 or 4 groups of blocks. Finally, the residual networks used in

this work will use either $F_{\text{initial}} = 16$ if there are three groups of blocks and $F_{\text{initial}} = 64$ if there are four groups. We depict Resnet18- w_i in Figure 2.6 and Resnet20- w_i in Figure 2.7 to illustrate these architectures as they will be widely used in our experiments.

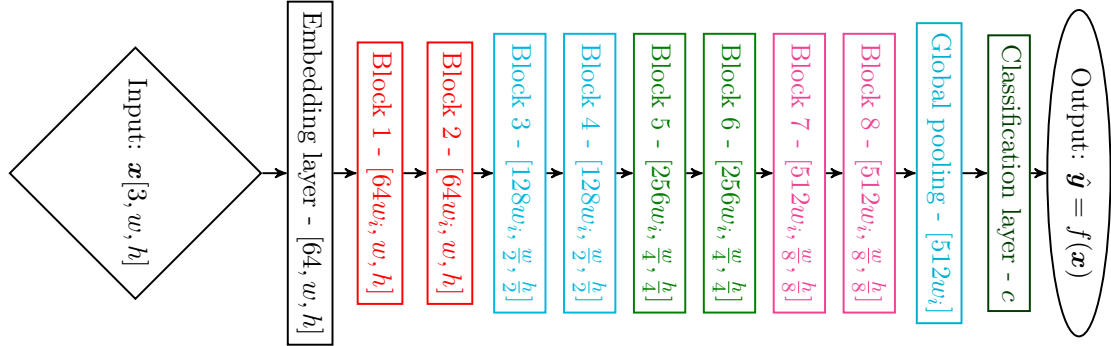


Figure 2.6: Depiction of the Resnet18- w_i used in this work. We write after the layer/block the dimensions of the output. If the output dimensions of the block differ from the input, the first convolution of the block will be the one responsible for the change, by either increasing the number of output feature maps or performing strided convolutions to reduce width and/or height.

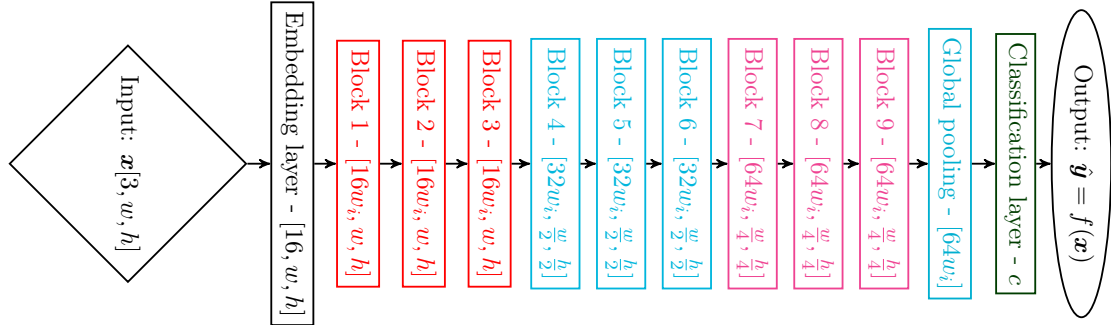


Figure 2.7: Depiction of the Resnet20- w_i used in this work. We write after the layer/block the dimensions of the output. If the output dimensions of the block differ from the input, the first convolution of the block will be the one responsible for the change, by either increasing the number of output feature maps or performing strided convolutions to reduce width and/or height.

2.1.2 Feature extraction with DNNs

One of the most significant advantages of DNNs is that they can extract relevant features from the input, as seen in Figure 1.3. This property of DNNs seems to be in line with the way biological neural classification operates. Indeed, works such as [61] and the more modern [49] show that there is a hierarchical organization of transformation from raw

images to concepts in the brain. These works were used as a first reasoning basis for DNNs and are the historical reason for terms such as neurons and artificial neural networks. In this work, we try to avoid the word neuron, using the network nomenclature instead, i.e., we use node instead of neuron.

In this vein, we can also define the network function f by a feature extractor (\mathcal{F}) followed by a classifier (\mathcal{C}), such that $f(\mathbf{x}) = \mathcal{C}(\mathcal{F}(\mathbf{x}))$.

As a consequence, it is possible to re-use networks in order to perform *transfer learning* [156, 182]. In transfer learning we first train a network f_1 on a first dataset \mathcal{D}_1 . Then when a subsequent dataset $\mathcal{D}_2 \cap \mathcal{D}_1 = \emptyset$ appears, one can either re-use the feature extractor \mathcal{F}_1 and train a new classifier \mathcal{C}_2 or can fine-tune the entire DNN using f_1 as a starting point. The former approach is preferable if $|\mathcal{D}_{\text{train}_2}|$ is small [2, 106], while the latter is used when the $|\mathcal{D}_{\text{train}_2}|$ is sufficiently large, e.g., it is quite common for the state of the art of medium-sized image datasets to be achieved by fine-tuning a network trained using a large image dataset [77].

The decision of training from scratch or performing transfer learning depends on the cardinality of both \mathcal{D}_1 and \mathcal{D}_2 and on the type of classification that is performed. For example, image classification tasks seem to do better with fine-tuning [77, 182], while image retrieval tasks seem to be better with re-using the feature extractor [2, 101].

2.1.3 Classifiers for DNNs

Multiple classifiers can be used in DNNs, depending on the goal/task at hand. In this subsection, we introduce some of the most used classifiers for DNNs.

Please note that one does not need to split \mathcal{F} and \mathcal{C} exactly at the last layer of a DNN, as is usually done in the literature. As a matter of fact, it can be beneficial in some cases to have deep classifiers (e.g., [77, 182]).

Throughout this section, we consider given a feature extractor \mathcal{F} , and we introduce various ways to perform downstream classification. Note that when training a DNN end-to-end the classifier choice will directly impact the training of the entire network.

2.1.3.1 Logistic Regression (LR)

The logistic regression is the most used form of the classifier for DNNs. It applies a linear transformation to the input so that the output is of the same dimension as the number of

classes in the problem, provided we are facing a classification one. Its goal is to optimize an objective function using a logarithmic model of the probabilities for each class.

In neural networks, it is common to generate the pseudo-probabilities of logistic regression using the softmax function:

$$\hat{\mathbf{y}}_i = \text{softmax}(\mathbf{x}')_i = \frac{e^{\frac{x'_i}{T}}}{\sum_{x'_j \in \mathbf{x}'} e^{\frac{x'_j}{T}}}, \quad (2.3)$$

where T is a temperature parameter set to 1 in most cases, $\hat{\mathbf{y}}$ is the output of the network, and \mathbf{x}' is the output of the last layer (commonly called classification layer). The logarithmic model then uses the **cross-entropy loss** as the objective function \mathcal{L} :

$$\mathcal{L}_{\text{cross-entropy}} = - \sum_i \mathbf{y} \log \hat{\mathbf{y}}. \quad (2.4)$$

where \mathbf{y} is the label indicator vector of \mathbf{x} . We depict three logistic regression classifiers under different values of T in Figure 2.8. Note that higher values of temperature create softer decisions and lower values lead to strict decisions.

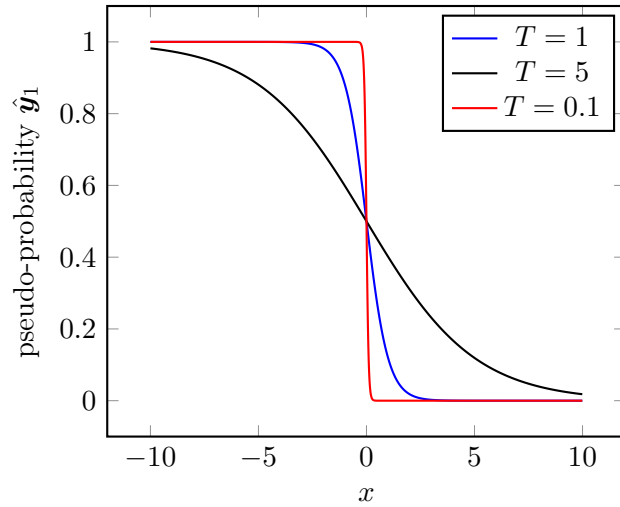


Figure 2.8: Logistic regression of a problem with two classes ($x \leq 0 \in c_1$ and $x > 0 \in c_2$) under different temperature ($T \in \{0.1, 1, 5\}$) conditions. We depict the output \hat{y}_1 , and consider that the decision is performed using the arg max function.

2.1.3.2 Support Vector Machine (SVM)

Another widely used classifier in machine learning is the Support Vector Machine [159]. The primary motivation of a SVM is to generate a set of hyperplanes to separate the

classes. These hyperplanes may be applied directly to the raw data or after the data has been transformed using a kernel k . The latter is also called the kernel-trick, and its main contribution is that it allows one to learn nonlinear classifiers using convex optimization techniques that are guaranteed to converge efficiently [39].

The most used kernel k is the radial basis function (**RBF**) defined as:

$$k(\mathbf{x}_1, \mathbf{x}_2) = e^{-\gamma \|\mathbf{x}_1 - \mathbf{x}_2\|_2^2}, \quad (2.5)$$

where \mathbf{x}_1 and \mathbf{x}_2 are two inputs and γ is an adjustable parameter. The RBF kernel can also be seen as a dot product in an infinite-dimensional space [39]. Note that the output of $k(\mathbf{x}_1, \mathbf{x}_2)$ is a similarity measure between \mathbf{x}_1 and \mathbf{x}_2 .

As the DNN already generates a suitable nonlinear feature extractor, it is more common to apply the hyperplane separators directly. In this case, given $\hat{\mathbf{y}}$ the output of the network, the objective function $\mathcal{L}_{\text{LinearSVM}}$ is defined as:

$$\mathcal{L}_{\text{LinearSVM}} = \max(0, m - \mathbf{y}\hat{\mathbf{y}}) \quad (2.6)$$

where m is the classifier margin, and \mathbf{y} is a modified binary label indicator vector, where one uses -1 to indicate that it does not belong to a class instead of using 0. In other words, the objective of the network is to output at least m for the coordinate corresponding to the class of the input and at most $-m$ for the other classes. Figure 2.9 depicts both linear and RBF kernel SVM classifications.

2.1.3.3 k -Nearest Neighbor classifier (k -nn)

Not all classifiers need to be optimized. Using k -nearest neighbors, the idea is to classify an input by looking at the closest training samples in the output domain of the feature extractor. Usually, the classification of an unseen input, commonly called **query**, is performed using the k closest examples from the training set (also called **support set** in this scenario) and a majority vote. One of the advantages of not having a training phase for the classifier is that one can quickly create ensemble decisions by adding a classifier per layer, as seen in [116]. Also, this technique can be easily deployed in the context of streaming data.

An interesting characteristic of a 1-nn classifier is that it has perfect memorization by definition. On the other hand it will probably have a poor generalization performance. We depict an example of a 1-nn classifier and a 20-nn classifier in Figure 2.10. We note the trade-off between the smoothness of the classifier border and memorization that is

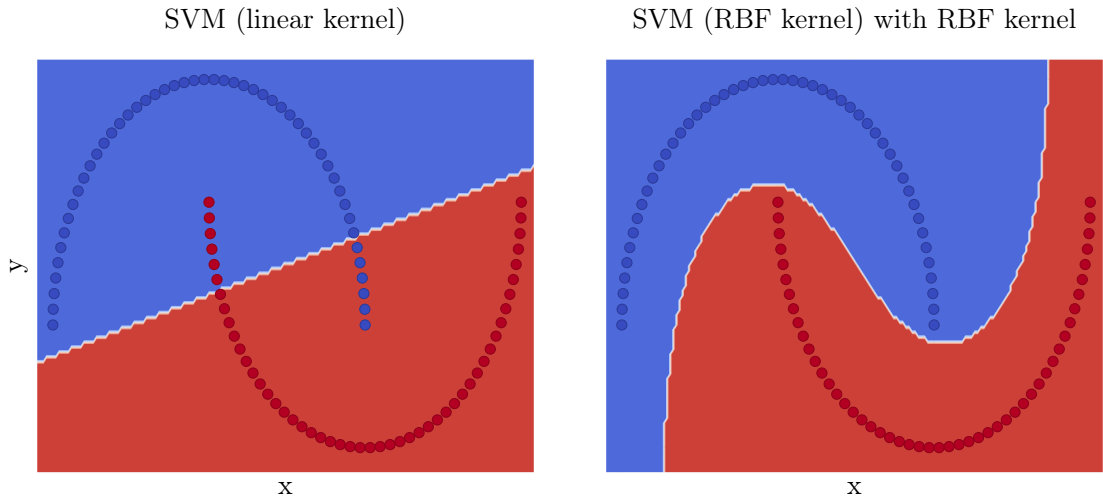


Figure 2.9: SVM classification applied to the “two moons” dataset. The goal is to classify the blue and the red dots. We depict the hyperplanes applied to the raw data (Linear SVM) and using an RBF kernel (RBF SVM). Note how the RBF SVM can completely separate the two classes, while the Linear SVM misclassifies 12 out of the 100 points.

displayed in the image. The 1-nn classifier has perfect memorization but a very rough classification border which can sometimes be a problem for generalizing to unseen data, especially when using inputs that could be mislabeled. On the other hand the 20-nn classifier has a very smooth classification border, but fails to correctly classify some examples of the dataset.

2.1.3.4 Nearest Centroid Mean classifier (NCM)

Another example of a classifier that does not need to be training is the NCM or Rocchio classifier [107]. In this case, one uses the mean representations for each cluster (one or more clusters per class) as the support set for a k -nn classifier. A comparison of the k -nn classifier and the NCM classifier is depicted in Figure 2.11. Note that the classification border of the NCM classifier is even smoother than that of the 20-nn classifier.

2.2 Deep Learning Layers

In this section, we introduce the layers that are used in this work. We recall that layers are the elementary functions on top of which the network function f of a DNN is built. In the remaining of this section, we denote by \mathbf{x}' the input of a layer and by f'' the layer

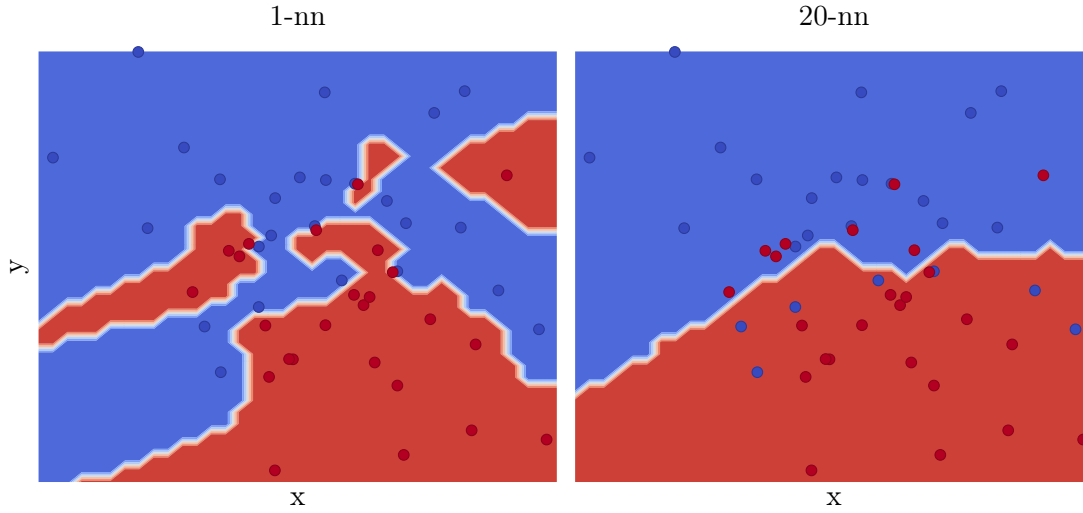


Figure 2.10: Example of 1-nn and 20-nn classification. The goal is to classify the blue and the red dots. Note how the 1-nn can completely memorize the dataset, but does so using a very non smooth border which can be a problem for generalizing to unseen data.

function.

In this work we always consider 4D inputs $[b_{\text{size}}, F, w, h]$, where the first dimension b_{size} indicates the number of samples treated concurrently as a **batch**, and the other three dimensions typically correspond to the number of feature maps, width and height of \mathbf{x}' .

Feature maps typically aggregate various components of the input (e.g., RGB images have 3 feature maps, one for each color: red, green, and blue). Sometimes, the three dimensions F, w, h can be flattened; in this case, we consider an input of dimension $d = whF$. Note that each coordinate of these dimensions is typically referred to as a *node* or *neuron* in the literature.

Layers typically compose a nonlinear activation function g with a linear function. Popular nonlinear activation functions include sigmoids, ReLUs, tanh. . . These functions are essential to ensure the expressivity of the deep learning models. Indeed, as the algebra of matrices is associative, removing the nonlinear functions would mean that deep models would be equivalent to a linear one-layer model. If one-layer models can only solve linearly separable problems, it has been known for a long time that the deep learning models are universal approximators under mild conditions [56].

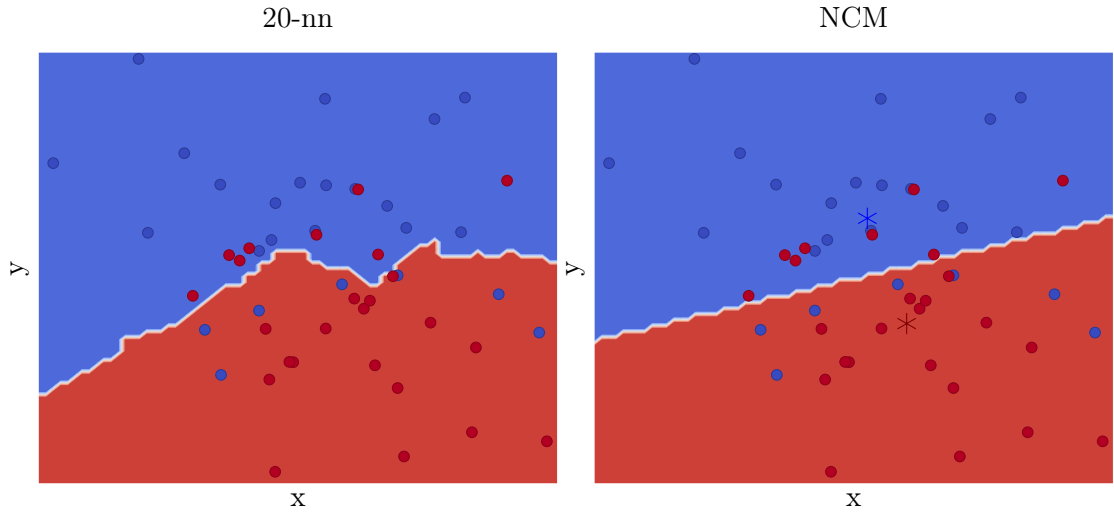


Figure 2.11: Comparison between the k -nn and NCL classifiers. The goal is to classify the blue and the red dots, with the asterisks representing the centroids of each class. Note that the classification border of the NCM classifier is even smoother than that of the 20-nn classifier.

2.2.1 Fully connected layer (FC)

A layer is said to be fully connected if each node i of the input is connected to each node j of the output with a proper weight $W_{i,j}$. In this case we can say that:

$$f''(\mathbf{x}') = g(\mathbf{W}\mathbf{x}' + \mathbf{b}), \quad (2.7)$$

where \mathbf{W} is the trainable weight matrix, \mathbf{b} is a trainable bias added to the layer, and finally, g is the activation function. In other words, we first perform a parametric linear transformation ($\mathbf{W}\mathbf{x}' + \mathbf{b}$) and then we apply a non-parametrized non-linear function g .

The fully connected layer is the most generic layer and the basis of deep learning. However, its use is severely limited when data to be processed is structured. For example, modern architectures for image classification usually only use one FC layer as the final classification layer. Figure 2.6 and Figure 2.7 depict two examples of the architectures used in this work, and both have only one FC layer. This limitation happens because FC layers ignore the intrinsic structure of the data. Indeed, the indexing of the inputs of a FC layer has no consequence on the accuracy of the trained model.

Theoretically, it would be possible for a DNN composed only of FC layers to understand the structure during training and limit its connections to obey this structure. However, it is tough to do so in practice [91]. We will explore this drawback in more detail in Chapter 4.

Examples of data with intrinsic structure range from images to citation networks. The most common solution to this problem is to use convolutional layers that we introduce in the next subsection. Another drawback of FC layers is their huge amount of trainable parameters. Considering that the input has dimension d_{input} and the output has d_{output} dimensions, the amount of trainable parameters of the FC layer is $(d_{\text{input}} + 1)d_{\text{output}}$. We represent the FC layer in Figure 2.12.

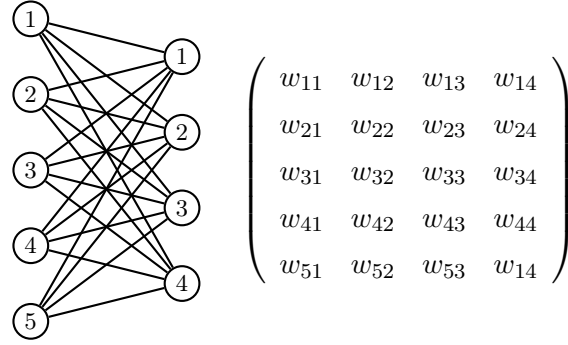


Figure 2.12: Representation of a FC layer with $d_{\text{input}} = 5$ and $d_{\text{output}} = 4$.

2.2.2 Convolutional layer

Convolutional layers [93] can be associated with the principle of filtering from signal processing. Instead of generating one full representation of the data by connecting every node of the input to every node of the output the idea of convolutional layers is to convolve several small filters over the coordinates of the input, generating multiple representations of the data on the output. We depict a convolutional layer in Figure 2.13 (unrolled) and Figure 2.14 (compressed).

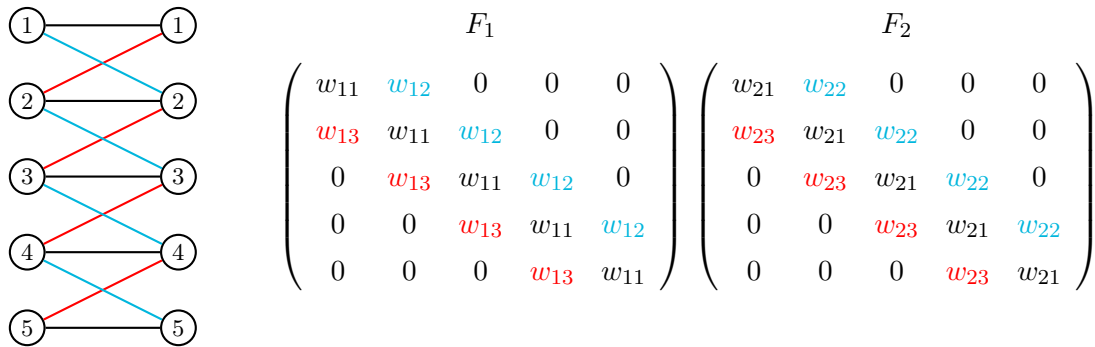


Figure 2.13: Representation of an unrolled 1D convolutional layer with $f_w = 3$, $s = 1$, $F_{\text{input}} = 1$, $F_{\text{output}} = 2$, $d_{\text{input}} = d_{\text{output}} = 5$.

We call **feature map** each one of those representations (f'_{maps}), i.e., each interme-

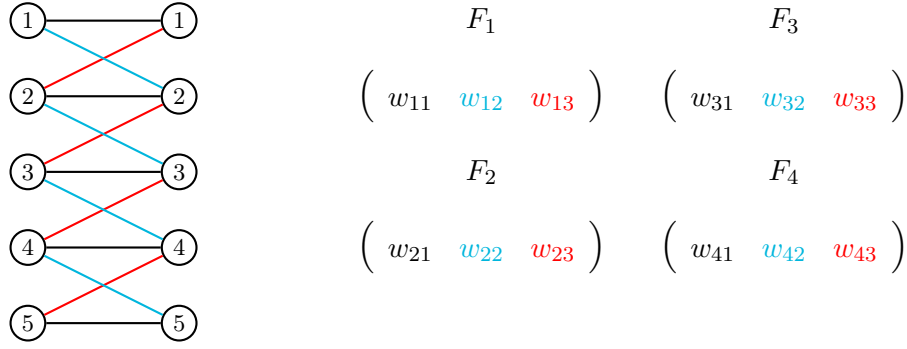


Figure 2.14: Representation of a 1D convolutional layer with $f_w = 3$, $s = 1$, $F_{input} = 1$, $F_{output} = 4$, $d_{input} = d_{output} = 5$.

diagram representation from a convolutional layer is composed of multiple feature maps of dimension d_{output} . In this work we focus on 2D convolutions (except for the illustrative figures 2.13, 2.14, and 2.15), and this will be reflected on our notations, but it should be straightforward to extend our work to any dimensional convolutions. In 2D convolutions the filters can be organized as a 4 dimensional weight tensor with dimensions $[F_{input}, F_{output}, f_w, f_h]$.

As these small (f_w by f_h) filters are convolved on the input, they respect the intrinsic structure of the data and allow the operation to be **translation equivariant**.

Definition 2.2.1 (translation equivariance). A function f' is said to be translation equivariant if applying a translation h on the input \mathbf{x}' is equivalent to applying a consistent mapping M_h on the output of f' . Formally we define translation equivariance by:

$$\forall \mathbf{x} \in \mathbb{D} : f'(h(\mathbf{x})) \approx M_h(f'(\mathbf{x})) \quad (2.8)$$

Being translation equivariant is an essential feature of convolutional layers. It allows for patterns to be recognized even if they are shifted on the data, e.g., in images, the position of the object we want to classify should not change the overall class of the image. Note that translation equivariance and invariance are different features, even if some works use the terms interchangeably. Both translation equivariance and translation invariance are desired features in image classification.

As it was the case of FC layers, we can characterize a convolutional layer by a parametric linear transformation followed by a non-parametrized non-linear function g :

$$f''(\mathbf{x}') = g(\mathbb{H} \otimes \mathbf{x}' + \mathbf{b}), \quad (2.9)$$

where \otimes is the convolution operator and \mathbb{H} is the set of f_w by f_h filters of the convolutional layer. Note that this convolution has a parameter s called *stride* that specifies the gap between the center of each convolution. In this work we consider two values of stride, $s = [1, 2]$. If $s = 1$, then the convolution is performed pixel by pixel, this means that every 2D coordinate is used as the center of every filter. On the other hand, if $s = 2$, then the convolution is performed using a quarter of the pixels, and the 2D representation of each feature map is therefore downsampled to a quarter of its original value. In this work we call convolutions with $s = 2$ *strided convolutions*, even if all convolutions are by definition strided. We depict a strided convolution in Figure 2.15.

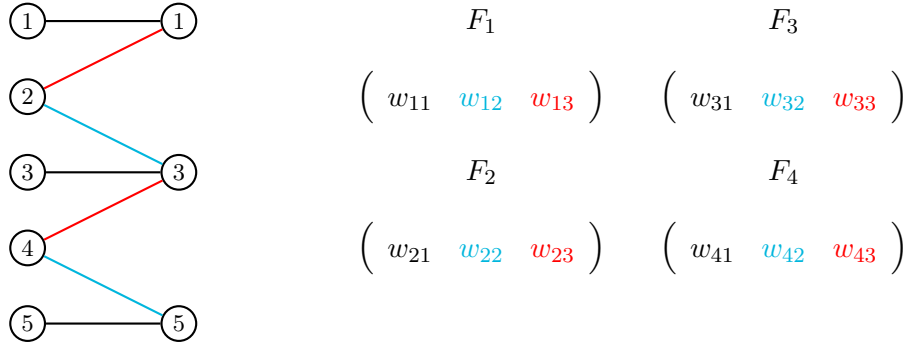


Figure 2.15: Representation of a strided 1D convolutional layer with $f_w = 3$, $s = 2$, $F_{input} = 1$, $F_{output} = 4$, $d_{input} = 5$, $d_{output} = 3$.

We ignore border effects that could be caused by misaligned w and f_w or h and f_h thanks to a zero padding. Zero padding means that if there would be a problem of misalignment between the dimensions we add zeros to the border in order to ensure that either $w_{input} = w_{output}$ or $w_{input} = \frac{w_{output}}{2}$ (in case of $s = 2$). The same is done for h .

Works such as [11, 167] aim to use a similar representation for both FC and convolution layers. This common representation is discussed in more detail in Chapter 4 as this abstraction is fundamental to understanding graph convolutions.

In the case of convolutional layers the amount of parameters is: $(F_{input}f_wf_h+1)F_{output}$, which is typically vastly smaller when compared to the FC layer as $F \ll d$ for both input and output.

2.2.2.1 Pooling layer

Pooling layers have two primary goals.

The first one is to downsample the intermediate representations, allowing to use

more coarse or simplified representations as we go deeper in the DNN. Using downsampled representations increases the real filter size of the network (e.g., a 3x3 filter on a downsample representation has a larger visual receptive field than a 3x3 filter on the full representation). In this work, this downsampling operation is executed with a 2D square of side p and an aggregation function agg . The two most common aggregation functions are the max operator (i.e., the output of the pooling operation is the maximum value inside each 2D square) and the average operator (i.e., the output of the pooling operation is the average value inside each 2D square).

The second goal of pooling layers is to provide a weak translation invariance in the radius of the square p .

2.2.3 Normalization layer

It is common to add normalization layers as the DNN grows deeper. Normalization layers are added in order to ease the training of the network, even if the theoretical reasoning behind these layers is not exactly solid. In this work, we consider only the most used normalization layer that is the **batch normalization** layer.

2.2.3.1 Batch normalization layer (BN)

Batch normalization (BN) layers were first proposed in order to reduce the internal covariate shift [64]. Covariate shift happens when the distribution in a learning system changes [147], which can be a big problem for DNNs as the deeper the network is, the easier it is for a small change in the distribution to affect the deeper layers. The goal is to normalize the intermediate representations \mathbf{x}' by applying f_{BN} so that each feature map has a fixed mean and standard deviation. To do so, we first perform f'' to normalize a batch of inputs \mathbf{X}' to have zero mean and standard deviation equal to one:

$$\mathbf{x}'' = f''(\mathbf{X}') = \frac{\mathbf{X}' - \mu_{\mathbf{X}'}}{\sigma_{\mathbf{X}'}^2} \quad (2.10)$$

where $\mu_{\mathbf{X}'}$ is a vector with the mean value for each feature map over the batch, and $\sigma_{\mathbf{X}'}^2$ is a vector with the standard deviation for each feature map over the batch. Note that during the inference phase, one cannot consider that it will receive a batch of inputs, therefore in this work, we always use the running mean and running standard deviation computed during the training phase. Now that the intermediate representation is normalized, we can fix the mean and standard deviation to trainable parameters $\mu_{f_{\text{BN}}}$ and $\sigma_{f_{\text{BN}}}^2$ as follows:

$$f_{\text{BN}}(\mathbf{X}') = \sigma_{f_{\text{BN}}}^2 f''(\mathbf{X}') + \mu_{f_{\text{BN}}} \quad (2.11)$$

Various recent works contest the covariate shift claim. In this work, we concentrate on three such studies [23, 140, 185], but many others exist in recent literature. The first one is a more theoretical paper that argues that BN helps the network to converge by smoothing the optimization landscape instead.

The latter two are more empirical, and while they have a similar argument, they are more focused on the application in residual networks. The authors explain that the BN layers smooth the optimization landscape by biasing the residual network to follow the shallower paths at the start of the training. The smoothing of the optimization landscape is done by giving less importance (sometimes even ignoring) to most of the blocks in the Resnet. They propose to remove the BN layers and to either change the initialization of the parameters or to add a parameter α to the residual connection in order to mimic this behavior and show similar results to batch normalized networks.

2.3 Tasks and Datasets

In this section, we first introduce the tasks as subsections, and then the datasets used in this thesis for each task as their subsubsections.

2.3.1 Image Classification

Image classification is one of the most common task in computer vision [39]. In this work, we consider only single-labeled classifications, where the goal is to classify the most prominent object in the image, even if more than one object is visible. In this task, we mostly focus on a measure of performance which is called the top- k classification accuracy. Top- k accuracy is the proportion of samples for which the expected output is among the k highest ranked ones in the obtained output.

2.3.1.1 CIFAR-10/100

CIFAR-10 and CIFAR-100 [81] are tiny (32x32 pixels) image datasets extracted from the 80 million tiny images dataset [161]. They are mostly used because they offer a good trade-off between complexity (i.e., trivial solutions and DNNs with only FC layers

do not provide good performance) and training times. This trade-off means that while trivial solutions, such as KNN on the pixel domain and DNNs with only FC layers, do not provide optimal performance, it is still possible to train near state of the art DNNs with even 3 to 4-year-old GPUs in less than a day. These two characteristics allow for quick/low-cost training and idea iteration on a significant problem.

The 10 and 100 after the dataset names specify the number of classes of the problem, but this does not mean that CIFAR-10 is a subset of CIFAR-100. The 100 classes of CIFAR-100 may be divided into 20 supergroups of 5 classes. CIFAR-10 classes, on the other hand, may be divided into 2 supergroups: i) transportation methods, and ii) animals. Note that it is infrequent to use these supergroups in the literature.

The CIFAR datasets come with a standard train/test split, and it is up to the authors to split the $\mathcal{D}_{\text{train}}$ to define the $\mathcal{D}_{\text{valid}}$ set. As we discussed in Section 2.1, most authors opt for directly training on $\mathcal{D}_{\text{train}} \cup \mathcal{D}_{\text{valid}}$ using the $\mathcal{D}_{\text{test}}$ to optimize the hyperparameters, which of course is not completely fair. Both datasets are composed of 60,000 images, being 50,000 images on the training set (5,000 per class for CIFAR-10 and 500 per class for CIFAR-100), and 10,000 images on the test set (1,000 per class for CIFAR-10 and 100 per class on CIFAR-100). We illustrate some of the examples from the datasets in Figure 2.16.



Figure 2.16: Illustrative examples from the CIFAR-10 dataset.

2.3.1.2 SVHN

The Street View House Numbers (SVHN) [113] is a dataset composed of images taken from House numbers, where the goal is to identify the most prominent number on the cropped image, where sometimes more than one number appears on the image. The

dataset was first proposed as a harder alternative to the MNIST dataset [94] on the digit recognition task. SVHN is composed of 10 classes (1 for each digit) and comes with a standard split of 73,257 training images and 26,032 test images. Note that an additional 531,131 extra images are available and can be incorporated into the dataset. We depict some examples of this dataset in Figure 2.17.



Figure 2.17: Illustrative examples from the SVHN dataset, extracted from the dataset website <http://ufldl.stanford.edu/housenumbers/>

2.3.1.3 Imagenet and variants

As introduced in Chapter 1, one of the reasons for the quick expansion of DNNs and deep learning was that these systems won the CVPR2012 Large Scale Visual Recognition (LSVR) challenge [82, 139]. This challenge has been used as the de-facto benchmark since then, and most works (including this one) still use the $\mathcal{D}_{\text{train}}$ and $\mathcal{D}_{\text{valid}}$ from that competition for benchmarking purposes. In other words, while it is common place to call this dataset Imagenet, a more appropriate name would be LSVR2012 Imagenet, as Imagenet is the database where the images were extracted. Note that it is uncommon to

use the $\mathcal{D}_{\text{test}}$ of this dataset.

The Imagenet dataset is one of the most complex challenges in computer vision given not only its high number of classes (1,000) and images (1.2 million images on the training set and 50,000 images on the validation set) but also the high image resolution of the dataset (it varies from 75x56 to 4,288x2,848). Treating images of high resolution takes considerably more time to process, but tends to lead to better results. In this work, due to computational constraints, we use the Imagenet32 variant [19], which downscales all images to the same size of the CIFAR datasets (32x32) allowing us to have a good trade-off between the cost of training and the relevance of the problem. We depict some examples of Imagenet in Figure 2.18 and Figure 1.1.



Figure 2.18: Illustrative examples from the Imagenet dataset.

2.3.2 Image retrieval

Image retrieval differs from image classification as the focus is mostly on the mean average precision (mAP) instead of the top- k accuracy.

Definition 2.3.1 (mean average precision (mAP)). To define mAP, we first need to define average precision (AP). To compute the average precision of an example $\mathbf{x}_q \in \mathcal{D}_{\text{query}}$, we use the sum of the precision at each correctly retrieved image of the support set ($\mathcal{D}_{\text{support}}$) based on its ranking k :

$$AP(\mathbf{x}_q) = \frac{\sum_{k=1}^{|\mathcal{D}_{\text{support}}|} P(k) r(k)}{|\mathcal{D}_{\text{relevant}_q}|} \quad (2.12)$$

where $P(k)$ is the precision at ranking k (the percentage of correctly retrieved images until rank k), $\mathcal{D}_{\text{relevant}_q}$ is the subset of $\mathcal{D}_{\text{support}}$ containing the relevant examples to \mathbf{x}_q and $r(k)$ is an indicator function returning 1 if the item at ranking k is in $\mathcal{D}_{\text{relevant}}$, and 0 otherwise. Given the average precision we can now define the mAP over a query set $\mathcal{D}_{\text{query}}$ as:

$$mAP(\mathcal{D}_{\text{query}}) = \frac{\sum_{\mathbf{x}_q \in \mathcal{D}_{\text{query}}} AP(\mathbf{x}_q)}{|\mathcal{D}_{\text{query}}|} \quad (2.13)$$

In other words, mAP grows when the top ranked retrieved images from the support set ($\mathcal{D}_{\text{support}}$) are from closer location and/or present the same objects as the considered query image $\mathbf{x}_q \in \mathcal{D}_{\text{query}}$. Note that doing simple classification and then outputting all images from the found class would be heavily penalized by this measure in case of misclassification. This is why most methods rely on alternative solutions.

2.3.2.1 Revisited oxford5k and paris6k

The revisited oxford5k (\mathcal{ROxf}) and revisited paris6k (\mathcal{RPar}) were first introduced in [128], in order to better represent the image retrieval task when compared to their original versions [123, 124]. Images are divided into $\mathcal{D}_{\text{support}}$ and $\mathcal{D}_{\text{query}}$. For each object that we want to retrieve (13 for Oxford and 12 for Paris), the images of the support set are divided, depending on the quality of the image and how apparent are the objects we want to retrieve, into three sets: i) easy ($\mathcal{D}_{\text{easy}_q}$), ii) hard ($\mathcal{D}_{\text{hard}_q}$), and iii) unclear ($\mathcal{D}_{\text{unclear}_q}$). The task is then divided into two difficulties:

1. Medium: Easy and hard images have to be retrieved and unclear images are disregarded (i.e., are not taking into account for computing mAP). More formally: $\mathcal{D}_{\text{support}_q} = \mathcal{D}_{\text{support}} - \mathcal{D}_{\text{unclear}_q}$ and $\mathcal{D}_{\text{relevant}_q} = \mathcal{D}_{\text{hard}_q} \cup \mathcal{D}_{\text{easy}_q}$.
2. Hard: Only the hard images have to be retrieved while easy and unclear images are disregarded. More formally $\mathcal{D}_{\text{support}_q} = \mathcal{D}_{\text{support}} - (\mathcal{D}_{\text{unclear}} \cup \mathcal{D}_{\text{easy}})$ and $\mathcal{D}_{\text{relevant}_q} = \mathcal{D}_{\text{hard}_q}$.

The datasets are composed of 70 query images ($\mathcal{D}_{\text{query}}$) and a $\mathcal{D}_{\text{support}}$ of 4,993 images for the Oxford dataset and 6,332 for the Paris one. We depict example images from the datasets in Figure 2.19 (oxford5k) and in Figure 2.20 (paris6k).

2.3.3 Vision-based localization

Vision-Based Localization (VBL) is at the same time closely related to the image retrieval problem as we want to find the landmarks on the image that allows us to recognize the location and to the image classification problem as it could be seen as a regression task where given an image we want to return the physical location of the camera. In a more formal way, VBL refers to the problem of retrieving both the location and orientation (pose) of the camera based on a query image. We refer the readers to [125] for a review on the subject.



Ashmolean museum

Christ church

Magdalen college

Figure 2.19: Illustrative examples from the revisited oxford5k dataset.



Les invalides

Tour Eiffel

Moulin Rouge

Figure 2.20: Illustrative examples from the revisited paris6k dataset.

2.3.3.1 Adelaide and Sydney datasets

In this work, we focus mostly on the datasets we introduced in [86]. These datasets were constructed by collecting images from the Mapillary API¹, which contains data that was publicly sourced over time using dashcams (i.e., cameras mounted on the windshield of vehicles). The images are extracted from videos, meaning that they can be divided into sequences. We extracted two sets of images from the road imagery of Australian cities. The first (Adelaide) covers the Central Business District (CBD) area of Adelaide, Australia. The second set is collected around the Greater Sydney region and covers an area of around 200km². Since the data is publicly sourced, there are some extra difficulties in these datasets (that are kept to simulate real life scenarios better):

1. There are viewpoint, illumination and dynamic changes. This is expected to happen

¹<https://www.mapillary.com/developer/api-documentation/>

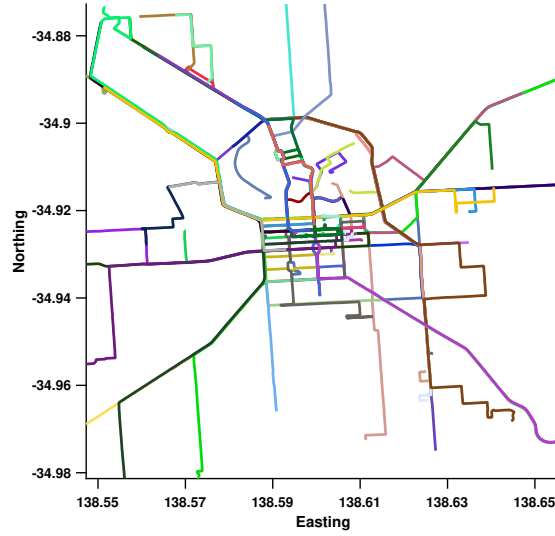


Figure 2.21: GPS Tracks of image sequence collected around Adelaide CBD from Mapillary. Figure and caption extracted from [86]

in real-life scenarios too (e.g., winter images would have snow while summer images may be brighter),

2. In the Sydney dataset, some of the sequences were generated using different equipment (e.g., panoramic cameras) and different positioning (e.g., some of the images are not of dashcams but cameras mounted on the side windows) from the ones used in traditional VBL problems.

In addition to imagery, the collected data provides sequence information and GPS. We depict the GPS tracks in Figure 2.21 and Figure 2.22

The Adelaide dataset is divided into one support set and two query sets ($\mathcal{D}_{\text{valid}}$ and $\mathcal{D}_{\text{test}}$). For the Sydney database, we split the images into one support and two query sets ($\mathcal{D}_{\text{easy}}$ and $\mathcal{D}_{\text{hard}}$). The split into easy and hard was performed due to the greater difficulty of the Sydney set. Note that the definition of $\mathcal{D}_{\text{easy}}$ and $\mathcal{D}_{\text{hard}}$ is not the same as in the case of Image retrieval problems. Statistics for each dataset are presented in Table 2.2.

2.3.4 Neurological task classification

In this work we also consider datasets composed of fMRI (functional Magnetic Resonance Imaging) scans. The main goal of using this type of data is to study architectures that leverage the underlying structure, even though the latter is not as simple as a standard

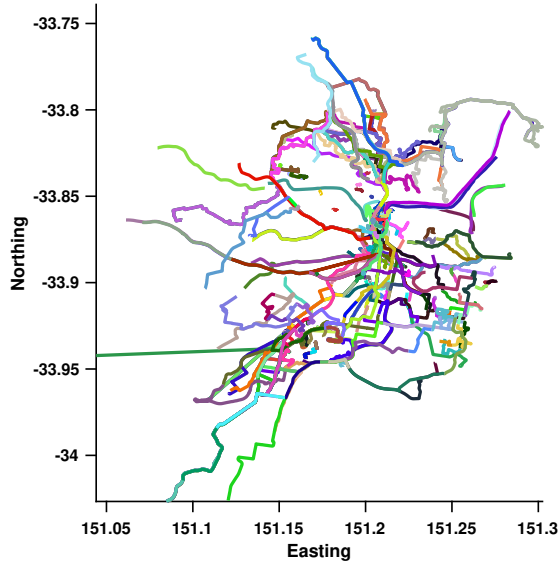


Figure 2.22: GPS Tracks of image sequence collected around Sydney from Mapillary. Figure and caption extracted from [86]

Table 2.2: Summary of the VBL datasets.

City	Adelaide	
	# Sequences	# Images
$\mathcal{D}_{\text{support}}$	44	24,263
$\mathcal{D}_{\text{valid}}$	4	2,141
$\mathcal{D}_{\text{test}}$	5	1,481
	Sydney	
	# Sequences	# Images
$\mathcal{D}_{\text{support}}$	284	117,860
$\mathcal{D}_{\text{easy}}$	5	1,915
$\mathcal{D}_{\text{hard}}$	5	2,285

euclidean 2D space. We will delve into more details on this in Chapter 4.

2.3.4.1 Pines

In [91], we introduce the use of the Pines dataset in the context of deep neural networks for the first time. The task of the PINES dataset is to identify the emotional rating a subject gives to a picture by using the fMRI scan of its brain. There are 182 subjects in the study [17]. To generate the dataset, we fetched first-level statistical maps (beta images) of each individual with minimal and maximal ratings from <https://neurovault>.

`org/collections/1964/`. Final volumes used for classification contain 369 signals per sample distributed on a 16mm cube. The dataset is composed of a $\mathcal{D}_{\text{train}}$ containing 1,949 samples and a $\mathcal{D}_{\text{test}}$ of 1,010 samples. The samples are divided into two classes, one for the minimal rating and the other for the maximal rating.

2.3.5 Document classification on citation networks

In this work, we also consider the problem of classifying scientific papers on a citation network. All datasets are constructed using the bag-of-words [150] method. The bag-of-words approach consists of using indicator vectors over a dictionary as the features of each document (i.e., the amount of times each word of the dictionary appears in the document). In this work, we use binarized versions of the bag-of-words vectors as it is commonplace in recent literature. All the datasets presented for this task are accompanied by a citation graph connecting documents that either cite or are cited by other documents.

The datasets presented in this subsection have no definitive split into $\mathcal{D}_{\text{train}}$ and $\mathcal{D}_{\text{test}}$, but it is common to use the split from [176] for benchmarking. We delve into more detail on this choice and these datasets in Chapter 4. In particular, we will show that they suffer from many limitations and biases, despite being the cornerstone of benchmarking in the field of graph-supported semi-supervised learning [145].

2.3.5.1 cora

A dataset of machine learning papers, composed of 2,708 documents, divided into seven classes, with a dictionary of 1,433 words (i.e., 1,433 features per document). This dataset was first proposed in [143], using articles from the cora database [109].

2.3.5.2 citeseer

The citeseer dataset is composed of 3,312 documents with a dictionary of 3,703 words and divided into six classes. This dataset was also first proposed in [143] using data from the citeseer database [36].

2.3.5.3 pubmed

A dataset of diabetes medicine research [112]. This dataset is composed of 19,717 papers with a dictionary of 500 words and divided into 3 classes depending on the type of diabetes

addressed in the publication.

2.4 Compression

In the previous sections we have introduced DNNs and the tasks and datasets on which we are going to apply them in this work. Now let us take a step back and consider the computational complexity of DNNs. Indeed, to achieve state of the art results, CNNs will often rely on a large number of trainable parameters, and considerable computational complexity. This is why there has been a lot of interest in the past few years towards the compression of CNNs, so that they can be deployed onto embedded systems or in real-time settings. The purpose of this section is to analyze the techniques that allow for the efficient compression of DNNs in order to reduce their computational complexity and memory footprint, while maintaining a high level of accuracy.

Prominent areas in neural network compression include distilling knowledge from a larger teacher network to a smaller (student) network [55, 79, 85, 117, 133], binarizing (or quantifying) weights and activations [15, 59], pruning network connections during or before training [3, 97], changing the way convolutions are performed [46, 172] and many others.

Most of the works in this area focus only on reducing complexity and footprint in the inference phase. Indeed, many of these methods increase the cost of the training phase. This choice is justified because the training of the network is considered to be done in an unconstrained environment, while the inference phase would be run in a highly constrained environment, which is applicable to most (yet not all) practical applications.

In this work, we focus only on distillation and in restricting the possible convolutions in DNNs and refer the reader to [44] for a more detailed review on this subject.

2.4.1 Distillation

Distillation based approaches aim at distilling knowledge from a pre-trained larger network that we call teacher to a smaller yet to be trained network called student. More formally, let T and S denote the architectures of the teacher and student. The goal is to transfer knowledge from T to S . For presentation simplicity, we assume that both architectures always generate the same number of intermediate representations, even if they are not from the same depth in the network architecture.

Distillation methods can be divided into two groups depending on how they perform the distillation: i) Individual Knowledge Distillation (IKD), and ii) Relational Knowledge Distillation (RKD). IKD methods consider each example of the training set separately, while RKD methods use information between sets of examples in order to distill knowledge. More formally, we can define the objective function of the student networks trained with knowledge distillation as:

$$\mathcal{L} = \mathcal{L}_{\text{task}} + \lambda_{\text{KD}} \cdot \mathcal{L}_{\text{KD}} , \quad (2.14)$$

where $\mathcal{L}_{\text{task}}$ is typically the same loss that was used to train the teacher (e.g., cross-entropy), \mathcal{L}_{KD} is the distillation loss and λ_{KD} is a scaling parameter to control the importance of the distillation with respect to that of the task. We now describe some IKD techniques. The first we describe is often called HKD (Hinton Knowledge Distillation) [55], and it is focused on the output $\hat{\mathbf{y}}$ of the networks, forcing S to mimic the output of T . The knowledge acquired by the teacher during the training phase is therefore diffused throughout the student network due to the backpropagation of weights.

Recent works [79, 133] advise that even if the knowledge is diffused through back-propagation, it is better also to force S to mimic the intermediate representations of T in order to improve the performance of S . However, as IKD treats each example individually, it can only do layer-wise mimicking if the student and the teacher have intermediate data representations with the same dimension [79]. This drawback restricts the architecture choice of the student networks. In an attempt to avoid this limitation, the authors of [133] propose to include affine transformations to ensure that the intermediate representations have the same size. The authors introduced extra layers meant to perform distillation during training. These transformations are therefore discarded during the inference of S , which could be a problem as these transformations will encode part of the knowledge that comes from the teacher network. In Figure 2.23, we depict the architecture from [79] where the output of each block is compared.

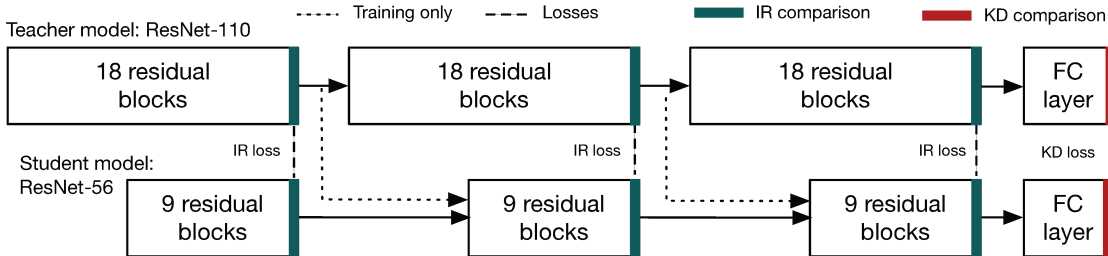


Figure 2.23: Depiction of the LIT architecture, IR comparison refers to a sample-wise L2 loss between the blocks, while KD comparison refers to the loss proposed in [55]. Figure extracted from [79] ©2019 PMLR.

More formally, we denote $\mathbf{X} \in \mathcal{D}_{\text{train}}$ a batch of input examples and \mathbb{X}' the set of intermediate representations generated using \mathbf{X} that are used for inferring knowledge. As previously described in Section 2.1, when an input \mathbf{x} is processed, a series of inner representations $\mathbf{x}', \mathbf{x}'', \dots$ is generated. IKD approaches will directly compare the inner representations of both teacher and student when processing \mathbf{x} . We therefore define the IKD loss from [55, 79, 133] as follows:

$$\mathcal{L}_{\text{IKD}} = \sum_{\mathbf{x}' \in \mathbb{X}'} \mathcal{L}_d(\mathbf{x}'^T, \mathbf{x}'^S), \quad (2.15)$$

where \mathbf{x}'^A is the intermediate representation of architecture A and \mathcal{L}_d is in most cases a measure of the distance between its arguments, which requires that they have the same dimensions.

However, forcing architectures for teacher and student to have intermediate representations with the same dimensions is not always desirable. Indeed, recent recommendations on DNN architecture design [157] show that efficient neural network scaling considers three main aspects: i) network depth (number of layers);, ii) network width (number of F per layer);, and iii) resolution (w by h size of the input and intermediate representations).. The two latter points are directly related to the dimensions of the intermediate representations and are therefore incompatible with IKD techniques.

To mitigate this drawback, recent works such as [85, 117] have introduced distillation that can be performed in both dimension-agnostic fashion and without adding extra transformations to the architecture of the DNNs. These methods are a part of the RKD group, where the focus is in the relative distances between the intermediate representations rather than on their exact positions. We delve into more details in RKD based compression in Section 5.4, where we describe Graph Knowledge Distillation (GKD), a method that we proposed in [85] to improve performance of RKD based approaches.

2.4.2 Shift attention layers (SAL)

In [46], we propose a more memory and computation efficient variation of convolutional layers that we call SAL, which we will detail in this subsection. Recently, the authors of [172] have proposed to replace the convolution operator with the combination of shifts and 1x1 convolutions, an approach they called shift convolutions. In other words, shift convolutions propose to limit the filter construction of convolutions to just one weight per filter, while still keeping the original filter shape. In this first work, all the shifts

were hand-crafted (i.e., decided arbitrarily before training). Note that previous works have shown that these shift convolutions are well suited for computationally constrained devices [45].

In order to increase the performance of the shift convolution, we introduced the Shift Attention Layer (SAL), which can be seen as a selective shift layer. SAL starts with vanilla convolution (i.e., with all the weights kept for each filter) and learns to transform it into a shift layer throughout the training of the network function. The introduced SAL use an attention mechanism [163] that selects the best shift for each feature map of the architecture. Note that this could also be considered a pruning technique as we start with all the weights for each filter and then choose the one to keep. It can significantly outperform the original shift layers from [172] at the cost of requiring more parameters during the training phase. We note that it still ends with fewer parameters during the inference phase.

We depict SAL in Figure 2.24 and provide the code for reproducing our experiments at <https://github.com/eghouti/SAL>. In the next paragraphs, we give more details about their core principle.

2.4.2.1 Methodology

We propose to add a selective tensor \mathbf{A} to standard convolutional layers, in order to identify which weight should be kept for each filter $f \in \mathbb{H}$ where \mathbb{H} is the set of filters of the convolutional layer. As such, we introduce $\mathbf{A} \in \mathbb{R}^{\mathbb{H} \times (f_w \times f_h)}$ a tensor that for each filter in \mathbb{H} contains a matrix of the same size of the filter (f_w by f_h). Each submatrix of \mathbf{A} is then normalized so that each value is between 0 and 1, with the sum of values in the matrix being equal to 1. The values in each submatrix from \mathbf{A} represent how important is the corresponding weight \mathbf{W} in the filter from \mathbb{H} . During the training process, the values of \mathbf{A} are pushed to binarization, until the end of the training process when they are binarized. With the weights in \mathbf{A} binarized, only the corresponding weight for each filter $f \in \mathbb{H}$ are kept.

More precisely, each slice $\mathbf{A}_{\mathbb{H}_i, \cdot, \cdot}$ is normalized with the softmax function from Equation 2.3 with a temperature T . The temperature is decreased smoothly along the training process in order to force the binarization of the softmax outputs. Note that in order to force the mask \mathbf{A} to be selective, we first normalize each slice $\mathbf{A}_{\mathbb{H}_i, \cdot, \cdot}$ so that it has a standard deviation (σ^2) of 1. We summarize the training process of one SAL layer in Algorithm 2.1. At the end of training, the selected weight in each filter \mathbb{H}_i corresponds to the maximum value in $\mathbf{A}_{\mathbb{H}_i, \cdot, \cdot}$.

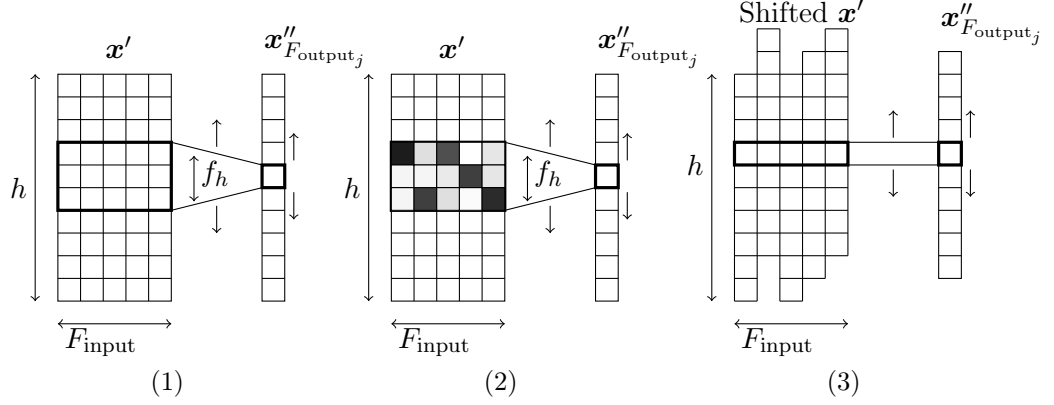


Figure 2.24: Overview of the proposed method: we depict here the computation for a single output feature map F_{output_j} , considering a 1d convolution and its associated shift version. Panel (1) represents a standard convolutional operation: the weight filter $f \in \mathbb{H}$ containing $f_h F_{\text{input}}$ weights is moved along the spatial dimension (h) of the input to produce each output in \mathbf{x}'' . In panel (2), we depict the attention tensor \mathbf{A} on top of the weight filter: the darker the cell, the most important the corresponding weight has been identified to be. At the end of the training process, \mathbf{A} should contain only binary values with a single 1 per slice $\mathbf{A}_{f,\cdot}$. In panel (3), we depict the corresponding obtained shift layer: for each slice along the input feature maps (F_{input}), the cell with the highest attention is kept and the others are discarded. As a consequence, the initial convolution with a feature height f_h has been replaced by a convolution with a feature height 1 on a shifted version of the input \mathbf{x}' . As such, the resulting operation in panel (3) is exactly the same as the shift layer introduced in [172], but here the shifts have been trained instead of being arbitrarily predetermined. Figure and caption adapted from [46]

Note that an advantage of SAL is that the number of filters per direction is not fixed as it was the case in the vanilla shift layers. However, this advantage comes with a drawback: it increases memory usage in order to retain which shift kept for each feature map. We note that this drawback is taken into account in our experiments.

2.4.2.2 Experiments on CIFAR-10/100

Now, we present the benchmarking protocol and a performance comparison of SAL and other shift layers on CIFAR-10/100. We refer the reader to [46] for a more extensive experimental discussion, including tests on the Imagenet dataset and comparison against pruning methods.

Algorithm 2.1 Pseudo-algorithm for one SAL layer, adapted from [46]

Inputs: Input tensor \mathbf{x}' ,

Initial softmax temperature T , Constant $\alpha < 1$.

for each training iteration **do**

$T \leftarrow \alpha T$

$\mathbf{A}' \leftarrow \mathbf{A}$

for $i := 1$ to $|\mathbb{H}|$ **do**

$\mathbf{A}'_{\mathbb{H}_i, \cdot, \cdot} \leftarrow \frac{\mathbf{A}_{\mathbb{H}_i, \cdot, \cdot}}{\sigma^2(\mathbf{A}_{\mathbb{H}_i, \cdot, \cdot})}$

$\mathbf{A}'_{\mathbb{H}_i, \cdot, \cdot} \leftarrow \text{Softmax}(\mathbf{A}'_{\mathbb{H}_i, \cdot, \cdot})$

$\mathbb{H}'_i \leftarrow \mathbb{H}_i \cdot \mathbf{A}'_{\mathbb{H}_i}$

Compute standard convolution as described in Section 2.2.2 using input tensor \mathbf{x}' and the set of filters \mathbb{H}' instead of \mathbb{H} .

Update \mathbb{H} and \mathbf{A} via back-propagation.

In order to promote a fair comparison with the other shift layers [70, 172], we use the same hyperparameters when possible, for example:

- Epochs: 300 epochs;
- Learning rate scheme: the learning rate starts at 0.1 and is divided by 10 after each 100 epochs;
- Batch size: 128;
- Temperature T : T starts at 6.7 and after each parameter update (step) it is updated so that it ends at $T_{\text{final}} = 0.02$.

We present in Table 2.3 a comparison of SAL against the vanilla shift layer in terms of accuracy and number of parameters needed during inference. We observe that our method achieves a better accuracy with fewer parameters than the baseline and other shift-module based methods.

2.5 Robustness

As we previously discussed in Chapter 1, DNNs can provide state-of-the-art performance in many machine learning challenges. This success can be justified based on their universal approximation properties [56], which allow them to approximate any function that associates each training set input to its corresponding class. However, this is also a

Table 2.3: Comparison of accuracy and number of parameters between 3x3 convolution, vanilla shift [172], interpolation shift [70], and ours on CIFAR10 and CIFAR100.

Network	F_{initial}	Convolutional layer	CIFAR10		CIFAR100	
			Accuracy	Params (M)	Accuracy	Params (M)
Resnet20	16	3x3 Convolution	94.66%	1.22	73.7%	1.24
Resnet110	16	Vanilla shift [172]	93.17%	1.2	72.56%	1.23
Resnet20	88	Interpolation shift [70]	94.53%	0.99	76.73%	1.02
Resnet20	83	SAL (ours)	95.52%	0.98	77.39%	1.01

double-edged sword, as the resulting function may not handle well domain shifts (i.e., it does not generalize well to previously unseen inputs). We have previously described this phenomenon in Definition 2.1.5.

Indeed, adversarial attacks, i.e., imperceptible changes to the input explicitly built to fool the network function [40, 153], illustrate the risks of overfitting to \mathcal{D} . More realistic scenarios include isotropic noise [105] and standard corruptions [53] that are also likely to produce similar misclassifications. Robustness to such deviations is, therefore, a key challenge, especially in applications that are very sensitive to errors, such as autonomous vehicles or robotic-assisted surgery.

Note that by using a definition linked to \mathcal{D} , robustness is therefore defined as the resiliency of the network to **corrupted inputs** $\hat{\mathbf{x}}$.

Definition 2.5.1 (corrupted input $\hat{\mathbf{x}}$). We aim to train DNNs to be robust to corrupted inputs $\hat{\mathbf{x}} \notin \mathcal{D}$. These inputs are defined in such a way that there exists a $\mathbf{x} \in \mathcal{D}$ and that $\|\hat{\mathbf{x}} - \mathbf{x}\| \approx \epsilon$ where $\|\cdot\|$ is a measure of distance and ϵ is a small enough threshold. The most common measures of distance in the literature are the L_2 [127] and the L_∞ [104], but it may also refer to more abstract concepts such as the same image but with different levels of contrast/brightness [53].

We note that recent works have also studied the concept of deep neural network robustness in the context of graph neural networks. In this case we have to consider not only corrupted inputs $\hat{\mathbf{x}}$ but also corrupted support graphs. We will not delve into the concept of graph neural network robustness, but we refer the reader to [8] for a more in depth discussion. More details about graph-supported deep neural network methods are available in Chapter 4.

Given our definition of the corrupted inputs $\hat{\mathbf{x}}$, a common approach to increasing the

robustness of DNNs is to concentrate on the Lipschitz constant of the network. Recall that a function f is said to be α -Lipschitz with respect to a norm $\|\cdot\|$ if $\|f(\mathbf{x}_i) - f(\mathbf{x}_j)\| \leq \alpha\|\mathbf{x}_i - \mathbf{x}_j\|, \forall \mathbf{x}_i, \mathbf{x}_j$. Provided α is small, such a function is robust to small deviations around correctly classified inputs, as it holds that: $\|f(\mathbf{x} + \varepsilon) - f(\mathbf{x})\| \leq \alpha\|\varepsilon\|$. One example of such a method that focuses in the Lipschitz constant of f is Parseval Networks [21], where the authors softly enforce the network L_2 and L_∞ Lipschitz constants to be bounded. Another example is [127] where the authors propose to bound only the L_2 norm of the network.

However, imposing a small Lipschitz constraint may be too restrictive of a constraint for the network function f . Indeed, the Lipschitz constant defines the slope of the function everywhere. Nonetheless, given the context of DNNs for classification, it is not unreasonable to expect sharp transitions in the output of f if we are near the class boundaries. In other words, ideally, we would like for the smoothness properties of the network function to be location-dependent (e.g., different behavior close to class boundaries), meaning that global Lipschitz metrics may not be as meaningful.

To illustrate our point, consider a function f that outputs binary label indicator vectors. We can then compute the minimal Lipschitz constraint that allows for outputting these vectors, given the distance between each pair of examples, i.e., the closer a pair of examples of different classes is in the input space, the higher the Lipschitz constant of the network would have to be to allow for outputting label indicator vectors. We are interested in this measure, as the training objective of most classification DNNs is to be able to output this type of vector. We depict in Figure 2.25 the proportion of pairs of training set inputs that are possible for a given Lipschitz constraint. The figure depicts information from the previously introduced datasets, CIFAR-10 and Imagenet32, using the L_∞ norm. We note that for the L_∞ norm a very high Lipschitz constraint is needed in order to correctly output binary label indicator vectors.

In this section we describe our definition of robustness, that was first introduced in our previous work [87]. This definition can be viewed as a *localized* Lipschitz constant of the network function in $\mathcal{D}_{\text{train}}$.

We ensure that any small deviation around a correctly classified training input should not dramatically impact the decision of the network function. Therefore our definition can be seen as a refinement over the previously proposed Lipschitz-based definitions, where we only enforce the Lipschitz constant on the previously defined $\hat{\mathbf{x}}$. We will derive reasonable sufficient conditions to enforce the robustness of a deep learning architecture in the following paragraphs. Using experiments on CIFAR-10 (Section 2.3.1.1) and Imagenet32 (Section 2.3.1.3) we demonstrate that our proposed definition of robustness

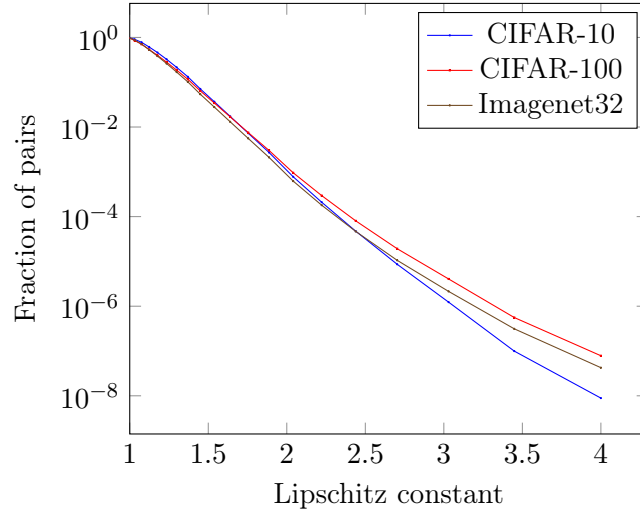


Figure 2.25: Depiction of the proportion of pairs of training examples of distinct classes incompatible with a given Lipschitz constraint on the network function, for various datasets and the L_∞ norm. Figure and caption adapted from [87] ©2019 IEEE.

is correlated to the empirical robustness observed in a series of existing network training methods [21, 90, 104, 127].

Note that the example depicted in Figure 2.25 illustrates that for such a sharp network function, a global Lipschitz constraint is not meaningful: unless the Lipschitz constant is large (e.g., greater than 4) imposing a constraint will prevent the training error from converging to zero. This example also suggests two related principles that can lead to better robustness and motivate our proposed robustness metric: i) robust network functions should be able to yield sharp transitions in boundary regions, and ii) smoothness metrics should be localized.

2.5.1 Defining robustness

We recall that our objective is to train a function f , which maps data from an input space \mathbb{D} into a softmax decision for classification. Therefore f is a function from an input vector space (or tensor space) to \mathbb{R}^c , where c is typically the number of classes. We are interested in the robustness of the network function f . We introduce here a notion of robustness that should account for:

1. A restricted domain R on which it is defined,
2. A locality r around each point in R on which it should be enforced.

More formally, we define robust DNN behavior as follows:

Definition 2.5.2 (α -robustness). We say a network function f is α -**robust** over a domain R and for $r > 0$, and denote $f \in \text{Robust}_\alpha(R, r)$, if:

$$\|f(\mathbf{x} + \varepsilon) - f(\mathbf{x})\| \leq \alpha \|\varepsilon\|, \forall \mathbf{x} \in R, \forall \varepsilon \text{ s.t. } \|\varepsilon\| < r. \quad (2.16)$$

In words, $f \in \text{Robust}_\alpha(R, r)$ if f is locally α -Lipschitz within a radius r of any point in domain R . Note that the following holds: $f \in \text{Robust}_\alpha(\Omega, +\infty)$ if and only if f is α -Lipschitz. Note that as it was previously discussed, we are interested in enforcing robustness for a small radius r around \mathcal{D} that is still inside of \mathbb{D} .

We also define: $\alpha_{\text{lim}}(f, R, r) = \inf\{\alpha : f \in \text{Robust}_\alpha(R, r)\}$, where $\alpha_{\text{lim}}(f, r)$ represents the minimum value α for which a region of radius r is robust. Therefore our robustness definition can leverage a trade-off between the smoothness slope α and a radius r .

Illustrative example: Figure 2.26 (Left) depicts the evolution of $\alpha_{\text{lim}}(\sigma, R, r)$ as function of r for the sigmoid function $\sigma : x \mapsto \frac{1}{1+\exp(-x)}$ and $R = \{-10, 10\}$. Observe that the sigmoid function yields an almost 0-Lipschitz constant around the two points -10 and 10 and for a very small radius r . When the radius starts to increase, the Lipschitz constraint needs to be less restrictive (as the Lipschitz constant increases). The fact that α is almost 0 when r is small is an illustration of robustness around R . The sharp transition occurring for $r \approx 10$ corresponds to a possible boundary between classes.

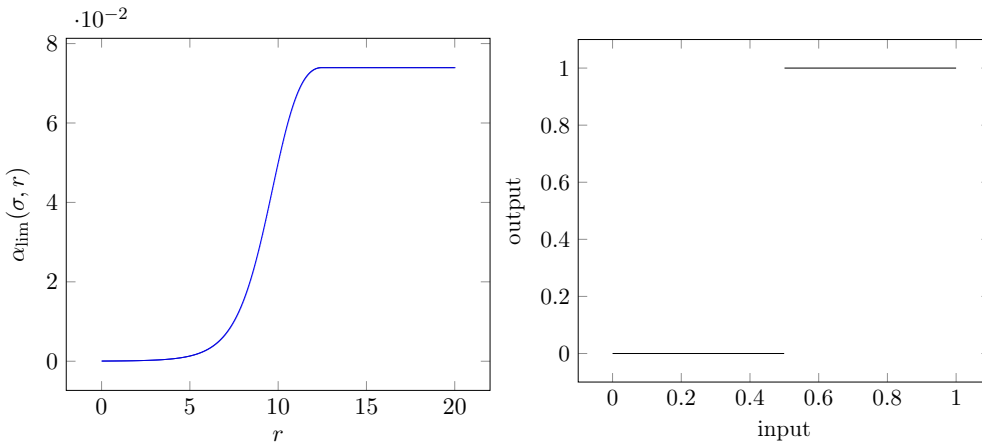


Figure 2.26: **Left:** Evolution of $r \mapsto \alpha_{\text{lim}}(\sigma, R, r)$. **Right:** Representation of the decision of a hyperplane separator between points 0 and 1. Figure and caption extracted from [87] ©2019 IEEE.

2.5.2 Relation with Lipschitz constants

We now describe the relation between our definition of robustness and the Lipschitz constant of f . Note that by particularization, if f is α -Lipschitz then $f \in \text{Robust}_\alpha(R, r), \forall r$. On the other hand, $f \in \text{Robust}_\alpha(R, r)$ for some r does not imply that f is α -Lipschitz. An example of this would be a trivial classification problem where $\mathcal{D}_{\text{train}}$ is composed of two distinct vectors \mathbf{x}_1 and \mathbf{x}_2 of distinct classes. A network function f that uses a hyperplane to separate the space into two halves has no Lipschitz constant because $\alpha \approx \infty$ close to the hyperplane, despite $\alpha_{\text{lim}}(f, R, \|\mathbf{x} - \mathbf{y}\|_2/2) = 0$. See Figure 2.26 (Right) for a 1D depiction of this example.

Note that this relation to the Lipschitz constant is a fundamental result, given that the best Lipschitz constant α of a function f is going to be constrained by \mathcal{D} , i.e., if two training points of different classes are very close to each other then a zero training error classifier will by definition have a large Lipschitz constant near those points (as suggested by Figure 2.25). The proposed robustness described in Definition 2.5.2 is also going to be limited by the construction of \mathcal{D} , but given a small enough r it is easy to imagine that we will be able to reach any small α . Indeed, denote by $\mathbf{c}^{\mathbf{x}}$ the class corresponding to training example \mathbf{x} . Then, if f matches a 1-nn classifier, we obtain that:

$$f \in \text{Robust}_0 \left(\min_{\substack{\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{T} \\ \mathbf{c}^{\mathbf{x}_1} \neq \mathbf{c}^{\mathbf{x}_2}}} \frac{\|\mathbf{x}_1 - \mathbf{x}_2\|}{2} \right), \quad (2.17)$$

and thus any small value for α is achievable within a small radius around examples.

2.5.3 Compositional robustness

Note that directly enforcing a robustness criterion on the entire DNN function f can be hard in practice, because of the numerous intermediate representations that are implied in the process. This is why several works in the literature consider each layer of the architecture separately [21, 90, 127].

Here we will focus on the simple case where f is obtained as the composition of intermediate functions f^ℓ . Note that the results presented here could easily be extended to more generic cases. So we denote by f^ℓ the function corresponding to layer ℓ , by F^ℓ the function $F^\ell = f^\ell \circ \dots \circ f^1$, and we suppose that $F^{\ell_{\text{max}}} = f$ is the DNN function. We define *layer-robustness* as:

Definition 2.5.3 (layer-robustness). We say that an intermediate function $f^{\ell+1}$ is α -robust over $F^\ell(R)$ and for $r > 0$ at depth $\ell + 1$ and we denote $f^{\ell+1} \in \text{Robust}_\alpha(F^\ell(R), r)$ if:

$$\begin{aligned} \|f^{\ell+1}(\mathbf{x}^\ell + \varepsilon) - f^{\ell+1}(\mathbf{x}^\ell)\| &\leq \alpha \|\varepsilon\|, \\ \forall \mathbf{x}^\ell \in F^\ell(R), \forall \varepsilon \text{ s.t. } \|\varepsilon\| &< r. \end{aligned} \quad (2.18)$$

Note that we consider f^0 to be the identity function.

There is a direct relationship between robustness of functions f^ℓ at the various layers of the architecture and that of f , as expressed in the following proposition.

Proposition 1. Suppose that:

$$f^{\ell+1} \in \text{Robust}_{\alpha^{\ell+1}}(F^\ell(R), r \prod_{\lambda \leq \ell} \alpha^\lambda), \forall \ell \text{ s.t. } 0 \leq \ell < \ell_{\max}. \quad (2.19)$$

Denote $\alpha = \prod_{\lambda \leq \ell_{\max}} \alpha^\lambda$, then

$$f = f^{\ell_{\max}} \circ \dots \circ f^1 \in \text{Robust}_\alpha(R, r). \quad (2.20)$$

Proof. Let us fix $\mathbf{x} \in R$. We proceed by induction. Let us show that if:

$$F^\ell \in \text{Robust}_{\prod_{\lambda \leq \ell} \alpha^\lambda}(R, r), \quad (2.21)$$

then

$$F^{\ell+1} \in \text{Robust}_{\prod_{\lambda \leq \ell+1} \alpha^\lambda}(R, r). \quad (2.22)$$

Indeed, let us fix ε s.t. $\|\varepsilon\| < r$, then:

$$\|F^{\ell+1}(\mathbf{x} + \varepsilon) - F^{\ell+1}(\mathbf{x})\| = \|f^{\ell+1}(F^\ell(\mathbf{x} + \varepsilon)) - f^{\ell+1}(F^\ell(\mathbf{x}))\|. \quad (2.23)$$

Note that as $F^\ell \in \text{Robust}_{\prod_{\lambda \leq \ell} \alpha^\lambda}(R, r)$, it holds that:

$$\|F^\ell(\mathbf{x} + \varepsilon) - F^\ell(\mathbf{x})\| \leq \prod_{\lambda \leq \ell} (\alpha^\lambda) \|\varepsilon\|. \quad (2.24)$$

So we can write:

$$F^\ell(\mathbf{x} + \varepsilon) = F^\ell(\mathbf{x}) + \varepsilon', \quad (2.25)$$

where:

$$\|\varepsilon'\| \leq \prod_{\lambda \leq \ell} (\alpha^\lambda) \|\varepsilon\| \leq r \prod_{\lambda \leq \ell} (\alpha^\lambda). \quad (2.26)$$

Finally, we obtain:

$$\begin{aligned} &\|F^{\ell+1}(\mathbf{x} + \varepsilon) - F^{\ell+1}(\mathbf{x})\| \\ &= \|f^{\ell+1}(F^\ell(\mathbf{x}) + \varepsilon') - f^{\ell+1}(F^\ell(\mathbf{x}))\| \\ &\leq \alpha^{\ell+1} \|\varepsilon'\| \leq \prod_{\lambda \leq \ell+1} (\alpha^\lambda) \|\varepsilon\|. \end{aligned} \quad (2.27)$$

□

We note that conditioning the intermediate function $f^{\ell+1}$ is less strict if all the previous layers were already yielding small values of α . In other words, the demanded radius for $f^{\ell+1}$ robustness is smaller. We thus observe there would be multiple possible strategies to enforce the compositional robustness of f in practice: i) forcing all layers to provide similar robustness, and ii) focusing only on a few layers of the architecture. Note that most proposed methods in the literature [21, 90, 127] opt for enforcing the former property as the latter would be too restrictive and would probably prevent the learning procedure from converging. We note that ii) would also be incompatible with the previously described Resnet architecture as it would restrict the ability of residual connections to ignore individual blocks.

2.5.4 Sources of noise

We now present some of the sources of noise that we study in this work and that fulfill Definition 2.5.1. These sources may be deliberate (adversarial attacks), or they could be just circumstantial (Gaussian noise added to a $\mathbf{x} \in \mathcal{D}$).

2.5.4.1 Adversarial attacks

In the literature, several methods have been proposed to measure the robustness of network functions. The first set of approaches [40, 104, 153] proposes to generate perturbations that both maximize the training loss and minimize the distance from the original inputs. This perturbation is generated by backpropagating the gradients through the networks to the inputs. These adversarially generated images are very potent against unprotected DNNs even for very small ε .

Note that in this same vein, various methods were proposed to increase the robustness of DNNs by using adversarial examples as a data augmentation procedure. They do so by adding these adversarial examples to $\mathcal{D}_{\text{train}}$ [40, 104]. Thus, during the training phase, the network function becomes increasingly robust to the corresponding corruptions. However, no guarantee exists that by increasing robustness to a specific type of corruption leads to better performance on other types of corruption, as discussed in [30, 53].

In this work, we consider three gradient-based adversarial attacks:

1. FGSM method [40], that we consider a mean case of adversarial noise, where the adversary can only use one forward and one backward pass to generate the perturbations. This approach is called Fast Gradient Sign Method (FGSM);

2. The DeepFool method [111], that is considered to be a worst case scenario, where the adversary can use multiple forward and backward passes to try to find the smallest perturbation that will fool the network;
3. The PGD method [104], that can be seen a compromise between the mean case and the worst case, where the adversary can do a predefined number of forward and backward passes with a perturbation threshold limit.

A main criticism about adversarial attacks is that they require access to the network function f and its derivative. This is highly improbable for many application cases. As such, some authors prefer to concentrate their studies on natural classifier-agnostic corruptions of data, as described in the next section.

2.5.4.2 Corrupted inputs

Multiple types of corruption may be applied to the inputs during the capture of the data. Images are especially affected by this, as just changing the camera lens or sensor could generate specific artifacts that change the distribution of \mathbb{D} and \mathcal{D} . To deal with this, we introduce 15 common image corruptions, with five levels of severity each. These corruptions were first organized as a benchmark in [53], with releases for the CIFAR-10 and Imagenet datasets. The goal of this benchmark is to compare the performance on the corrupted $\mathcal{D}_{\text{test}}$ and the clean $\mathcal{D}_{\text{test}}$ in order to isolate the original clean set performance from the analysis. We depict the 15 different corruptions in Figure 2.27.

To measure the robustness of the networks to these corruptions, we use the relative mean corrupted error (relative mCE). To compute this measure of performance, one has to follow several steps:

1. Take a trained classifier f and compute the clean $\mathcal{D}_{\text{test}}$ top-1 error rate. We denote this measure clean error rate: E_{clean}^f ;
2. Now test the classifier f on each corruption c and every severity s . We denote this by $E_{s,c}^f$;
3. As different corruptions have different difficulty levels even with the 5 severities s , we now normalize each score using a pre-trained baseline network that we call f_V . This normalized score is called the corruption error or CE and it is defined by the following equation:

$$\text{CE}_c^f = \frac{\sum_{s=1}^5 E_{s,c}^f}{\sum_{s=1}^5 E_{s,c}^{f_V}} \quad (2.28)$$

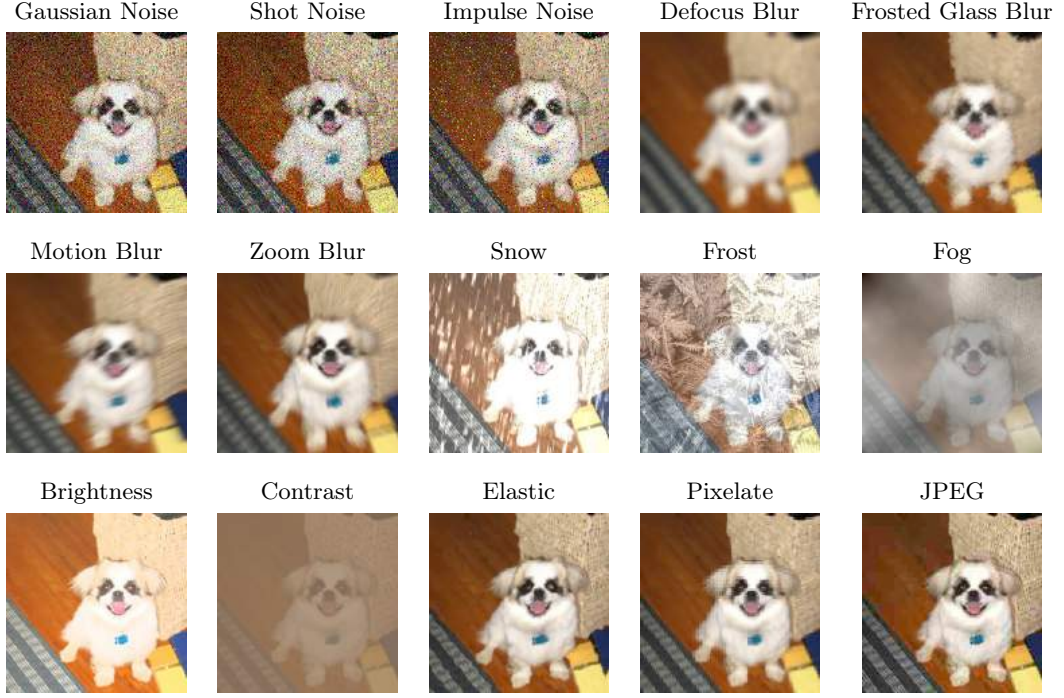


Figure 2.27: The 15 different types of corruptions from [53] applied to a random Imagenet dog image.

4. We can therefore summarize model robustness by averaging over the different CE_c for the same f . This leads to the mean corruption error (mCE);
5. Finally, to ensure that gains in robustness come from the strengths of network f and not as a by product of better performance on the clean $\mathcal{D}_{\text{test}}$ we compute the relative CE

$$\text{Relative CE}_c^f = \frac{\sum_{s=1}^5 E_{s,c}^f - E_{\text{clean}}^f}{\sum_{s=1}^5 E_{s,c}^{f_V} - E_{\text{clean}}^{f_V}}. \quad (2.29)$$

As it was the case CE, we can also average the relative corruption errors to obtain the relative MCE.

Note that this is the standard measure of robustness for this benchmark as introduced in [53].

2.5.5 Existing methods in the lens of our definition of robustness

We now present four of the prior works in the literature from the perspective of our proposed robustness measure, namely: Parseval networks (P) [21], L_2 non-expansive

networks (L2NN) [127], Laplacian networks (L) [90] and Projected Gradient Descent adversarial training (PGD) [104]. See Table 2.4 for a summary.

We recall that in [21], networks are trained to be α -Lipschitz for both L_∞ and L_2 . They do so by adding a regularizer to the training scheme that tries to make each weight matrix of the DNN a Parseval tight frame [80]. Among the four methods we consider, this Parseval was the only one that led to improved performance on the clean test set. However, we note that [21] does not seem to strictly enforce the α -Lipschitz constraint, as it does not consider batch normalization layers. This is why it can achieve excellent memorization (which theoretically, as seen in Figure 2.25, could only be achieved if α is large). This also explains why this method achieves worse results in robustness than L2NN [127]. In terms of our proposed definition of robustness, this is a global method that targets the $\text{Robust}_\alpha(r)$ metric for $r \rightarrow +\infty$, penalizing large slopes in the network function between any two points. We will see that more localized approaches (targeting finite r) achieve improved robustness when compared to Parseval. We denote this method P in the remainder of this work.

L2NN [127], on the other hand, enforces the network to be α -Lipschitz only in terms of the L_2 norm but does it with a stricter criterion. The Lipschitz condition is built into the structure of the network itself. The authors from [127] admit that enforcing a global α -Lipschitz constant is by itself too hard and that the distances between examples should not collapse throughout the network architecture. As such, they also limit the contraction of space. L2NN seems to be the most robust method against L_2 attacks of the four methods we consider. It has also been shown to combine well with PGD training. However, it is also the method that was the worst-performing on the clean test set.

In [90], we applied a regularization at each ReLU activation in the architecture to enforce that the average distance between examples of different classes remains almost constant from layer to layer. We exploit the smoothness of the label indicator vector across the graph generated by intermediate representations at a given layer to enforce this smooth transformation. We detail this contribution in Chapter 5.3. In terms of Definition 2.5.2, this method focuses on pairs of examples of distinct classes and tries to restrict changes in their L_2 distance. Thus, [90] indirectly penalizes changes in local smoothness. If we consider Definition 2.5.2 with $f(\cdot)$ chosen to be the function that assigns to each example its correct label, and we do not allow the average r between opposite class examples to change much, then the corresponding α will change slowly with the training.

Finally, PGD adversarial training [104] is a data augmentation procedure that generates adversarial examples during the training phase, as described in the previous

subsection. Using the PGD data augmentation leads to a min-max game between the network and the examples generation. It works mostly on the domain $\mathcal{D}_{\text{train}}$, as it increases its size and also decreases the difference between $\mathcal{D}_{\text{train}}$ and a noisy test domain. The data augmentation scheme leads to less domain shift to corrupted inputs, but on the other hand, it increases the domain shift to clean images. As a result, the networks perform well against noise (isotropic or adversarial) but have problems with clean examples.

Method	Domain (R)	Slope (α)	Locality (r)	Metric
P	Ω	Yes	No	$L_2 + L_\infty$
L2NN	Ω	Yes	No	L_2
L	T	Approx.	Yes	$L_2 + \cos$
PGD	augmented T	No	Yes	L_∞

Table 2.4: Summary of the methods and the notions of the introduced robustness they consider. Table extracted from [87] ©2019 IEEE.

2.5.6 Empirical evaluation of the proposed robustness metric

We perform experiments to evaluate the relevance of the proposed robustness definition (2.5.2) empirically. Vanilla (V), Parseval (P), and Laplacian (L) refer to the networks that were trained in [90], PGD, and L2NN refer to the networks trained in their original papers [104, 127]. Note that this direct comparison with the baseline is not entirely fair, given that the networks and hyperparameters for different papers are not the same. For example, PGD has more layers and parameters and uses non-adversarial data augmentation during training, while L2NN does not use residual architectures.

We depict in Figure 2.28, as a function of α , the ratio between i) the number of examples within distance d of each other that are not α -robust for L_∞ , and ii) the total number of example pairs. Note that d can be roughly interpreted as a diameter ($2r$) in our definition of robustness. As in the previous Figure 2.25, the output of the network function is a label indicator vector of the corresponding classes. Note that for each choice of d , the curve is initially flat. The flat section means that *all the pairs within d are α -robust*. Note that the number of pairs of examples in distinct classes that are closer than d drops very fast and becomes negligible for $d = 0.3$. In other words, for $d = 0.3$, it is theoretically possible to find a robust f that is compatible with *almost all* pairs of the training set.

In Figure 2.29 we depict the evolution of $\alpha_{\text{lim}}(r)$ as a function of r for the various methods. We use 100 training examples with 1000 Gaussian noise realizations as a proxy

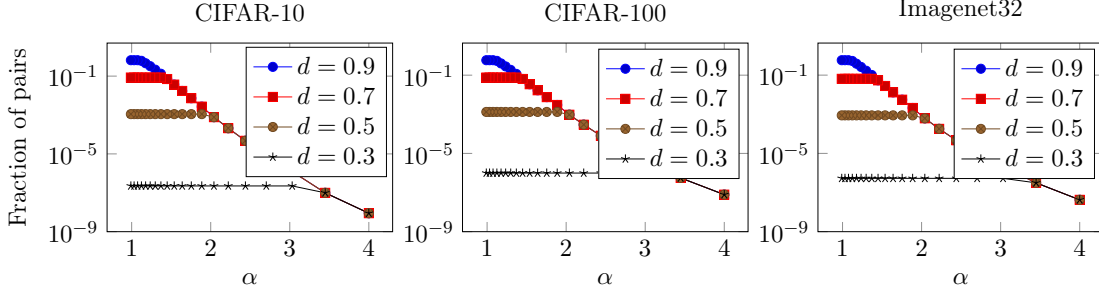


Figure 2.28: Depiction of the proportion of pairs of training examples of distinct classes incompatible with Definition 2.5.2 for the \mathcal{L}_∞ norm, as a function of α and for various values of d .

to estimate $\alpha_{\text{lim}}(\cdot)$ on the CIFAR-10 dataset. Note that for all methods, α increases as a function of r and then achieves its maximum value.

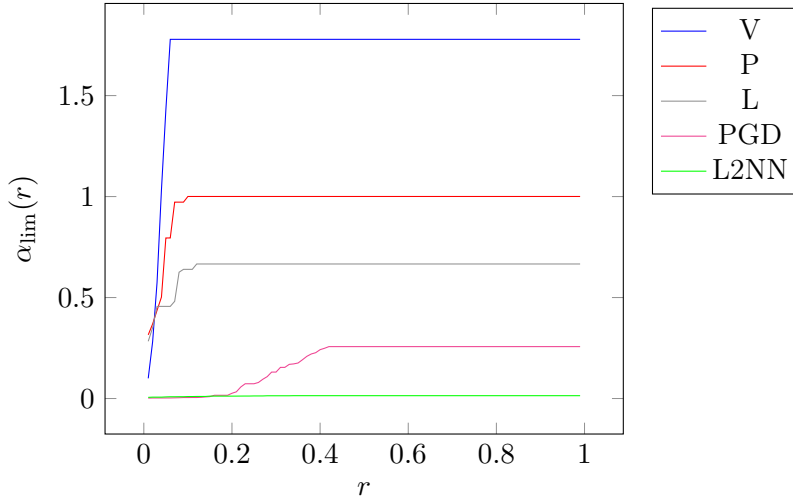


Figure 2.29: Estimations of $\alpha_{\text{lim}}(r)$ obtained for different radius r over training examples with the L_∞ norm. Figure extracted from [87] ©2019 IEEE.

Vanilla (V) is the fastest to saturate and achieves the largest value of α for two reasons i) sharp transitions in the network function over short distances are allowed, and ii) the network function produces outputs closest to the one-hot-bit encoded vector (since V can achieve zero error on the training set). In contrast, the other methods grows at a slower pace and reach smaller maximum values of α_{lim} . This behavior indicates that transitions are not as sharp and that some examples may be misclassified.

Moreover, the fact that, for both P and L, α saturates at larger values of r suggests that the margin between the examples and the boundary is increased compared to Vanilla.

L2NN and PGD saturate at the lowest α values.

We also observe a transition for PGD occurring at around $r = 0.3$ (which is an excellent radius given the aforementioned result that theoretically $d = 0.3$ would be the limit), whereas L2NN remains almost constant. We believe that the L2NN behavior is due to the fact L2NN enforces a strong Lipschitz constraint, on the L_2 norm, everywhere on the function. Therefore the network function described by the L2NN network is almost linear between all training samples, leading to a network that is not as accurate on the clean set. As seen in Figure 2.25, this lower value of α creates strong incompatibilities with the training dataset, which could be the reason why L2NN achieves the worst performance on the clean set (c.f. Table 2.5).

We now compare methods in terms of robustness on the recently proposed benchmark of image corruptions [53]. PGD achieves the best trade-off between accuracy and robustness, as seen in Table 2.5. We note that for PGD, our robustness metric has a relatively small value for α_{\max} , and the slope starts at $0.2 \leq r \leq 0.4$, which corresponds to appropriate values of d as seen in Figure 2.25. The results described in both Table 2.5 and the behavior described in Figure 2.29, seem to suggest that improved robustness is achievable when the network function is smooth locally near the examples, i.e., network functions that favor Definition 2.5.2 were empirically more robust.

Table 2.5: Test set error on the CIFAR-10 dataset under different image conditions. Table and caption extracted from [87] ©2019 IEEE.

Dataset	V	P	L	PGD	L2NN
Clean	11.9%	10.2%	13.2%	12.8%	20.9%
MCE	31.6%	30.5%	31.3%	18.8%	28.5%
relativeMCE	100	103	92	30	39

Finally, we depict in Figure 2.30 the relative robustness performance of each method under the same Gaussian Noise parameters used to generate Figure 2.29².

2.6 Summary of the chapter

In this first chapter, we introduced Deep Neural Networks (DNNs), with a specific focus on Residual architectures. We also introduced the concept of intermediate representations

²We do not report the results for P in the Imagenet32 dataset since we did not find right hyperparameters to obtain a good accuracy on the clean test set. Also, PGD and L2NN results are not reported in the case of CIFAR-100 and Imagenet32 as pre-trained networks were not available.

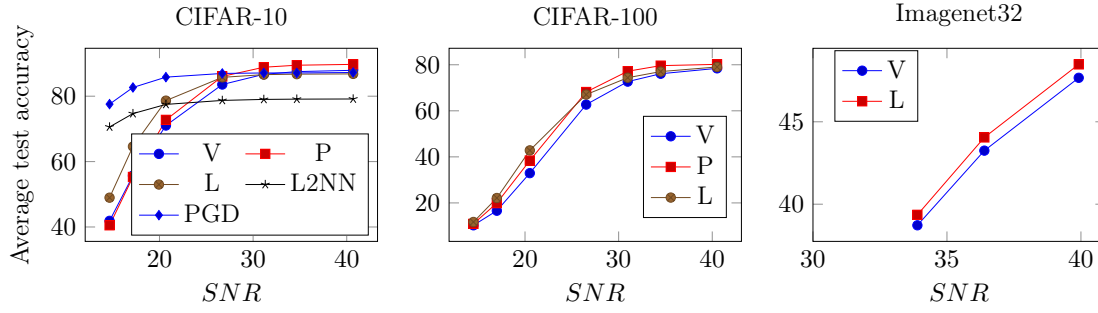


Figure 2.30: Average test set accuracy under Gaussian noise for various datasets and methods. Figure and caption extracted from [87] ©2019 IEEE.

in DNNs, which will be a central part of the following chapters of this thesis. These intermediate representations can be used in order to perform transfer learning as introduced in Section 2.1.2 and also in order to abstract DNNs as feature extractor followed by classifiers.

We presented various problems where DNNs are relevant. These problems are going to be further studied in the next chapters and include the following tasks: i) image retrieval, ii) vision based localization, iii) image classification, iv) neurological task classification, and v) document classification.

We also reviewed some of the literature in neural network compression, including distillation methods and more efficient convolution layers. We introduced SAL, a contribution on the subject of efficient convolution layers, which was subject to the following archival paper:

- Hacene, G. B., Lassance, C., Gripon, V., Courbariaux, M., and Bengio, Y. (2019). Attention based pruning for shift networks. *arXiv preprint arXiv:1905.12300*, to appear in *25th International Conference on Pattern Recognition (ICPR2020)*

Further, we introduced the concept of robustness of a classifier, and demonstrated empirically how it can be connected with an ability to perform well in presence of corrupted inputs. This concept of robustness and its empirical experiments were published in the following conference paper:

- Lassance, C., Gripon, V., Tang, J., and Ortega, A. (2019). Structural robustness for deep learning architectures. In *2019 IEEE Data Science Workshop (DSW)*, pages 125–129

In the following chapter, we introduce the framework of Graph Signal Processing (GSP) which defines a series of concepts and analytical tools. These tools are going to allow us to analyse the intermediate representations of DNNs, and are going to be the cornerstone for our contributions in both Chapter 4 and Chapter 5. Indeed, in Chapter 4 we use the tools of GSP to introduce Graph Neural Networks. In Chapter 5, the framework of GSP will allow us to propose analytical tools to better understand the inner workings of DNNs. We also introduce improvements in the accuracy, robustness and compression of these architectures.

Chapter 3

Graph Signal Processing

3.1	Definitions	65
3.2	Graphs for samples of features	74
3.3	Inferring graph topology from signals	78
3.4	Graph filters	86
3.5	Summary of the chapter	100

We present in this Chapter the concepts of graphs and Graph Signal Processing (**GSP**) that are exploited in this work in order to study the topology of intermediate representations of DNNs. Studying the topology of DNNs will allow us to propose new methods to improve DNNs in Chapter 4 and Chapter 5. This chapter is organized as follows: first we define graphs, graph signals and GSP in Section 3.1. Then we introduce the two types of graphs that we consider in this work in Section 3.2, graph topology inference from data 3.3 and graph filters 3.4. We refer the reader to [148] for a more detailed introduction to GSP.

3.1 Definitions

In general, *graphs* are used as a formalism to represent data and its relationships. More precisely we define a graph as:

Definition 3.1.1 (graph). A graph \mathcal{G} is a tuple of sets $\langle \mathbb{V}, \mathbb{E} \rangle$, such that:

1. The set \mathbb{V} is composed of vertices v_1, v_2, \dots ;
2. The set \mathbb{E} is composed of pairs of vertices of the form (v_i, v_j) called edges.

It is common to represent the set \mathbb{E} using an edge-indicator symmetric adjacency matrix $\mathcal{A} \in \mathbb{R}^{|\mathbb{V}| \times |\mathbb{V}|}$. Note that being symmetric means that in this work we consider only undirected graphs, which allow us to simplify most of our notations. As a quick recall to the reader, in an undirected graph, there is no distinction between edges (v_i, v_j) and (v_j, v_i) .

In some cases, the matrix \mathcal{A} is weighted (it takes values other than 0 or 1) because it not only represents the fact that a pair of vertices $(v_i, v_j) \in \mathbb{E}$ but also the weight associated with that representation, where typically a value closer to 0 corresponds to a vanishing relationship. In other words, each element $\mathcal{A}_{i,j}$ represents the weight of the edge between v_i and v_j .

As is the case in most works in the literature, we consider only graphs with nonnegative weights, and say that two vertices are not connected if their weight is equal to 0.

We can use the \mathcal{A} to define the diagonal **degree matrix** \mathbf{D} of the graph as follows:

$$D_{i,j} = \begin{cases} \sum_{j' \in \mathbb{V}} \mathcal{A}_{i,j'} & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}. \quad (3.1)$$

We also define the r -neighborhood $\mathbb{N}_r(v)$ of a vertex $v \in \mathbb{V}$ as the set of vertices that are at most r -hop away from a vertex $v \in \mathbb{V}$, that is to say that $v' \in \mathbb{N}_r(v)$ if and only if it exists a sequence of at most r vertices $v_{i_1}, v_{i_2}, \dots, v_{i_\rho}$, such that $v_{i_\rho} = v'$, $(v, v_{i_1}) \in \mathbb{E}$ and $(v_{i_j}, v_{i_{j+1}}) \in \mathbb{E}, \forall j$.

A graph typically represents a relation between its vertices. When those vertices are associated with measures (typically scalars), we talk of graph signals. In this thesis we only consider signals supported on the vertices of the graph, but there are also studies that consider signals supported on the edges [141]. In other words, we consider graph signals $\mathbf{s} \in \mathbb{R}^{|\mathbb{V}| \times F}$ where F is the number of realizations of the signal \mathbf{s} for each vertex of the graph. In Figure 3.1 we depict examples of graphs in various machine learning scenarios.

As a natural representation of complex data structures, graphs and graph signals are ubiquitous, in particular in the field of machine learning. In this thesis we focus on two main uses of graphs: i) Graphs that model the inner dependencies of observations, and ii) Graphs that model the relationship between data samples. More details on this subtle distinction are available in Section 3.2.

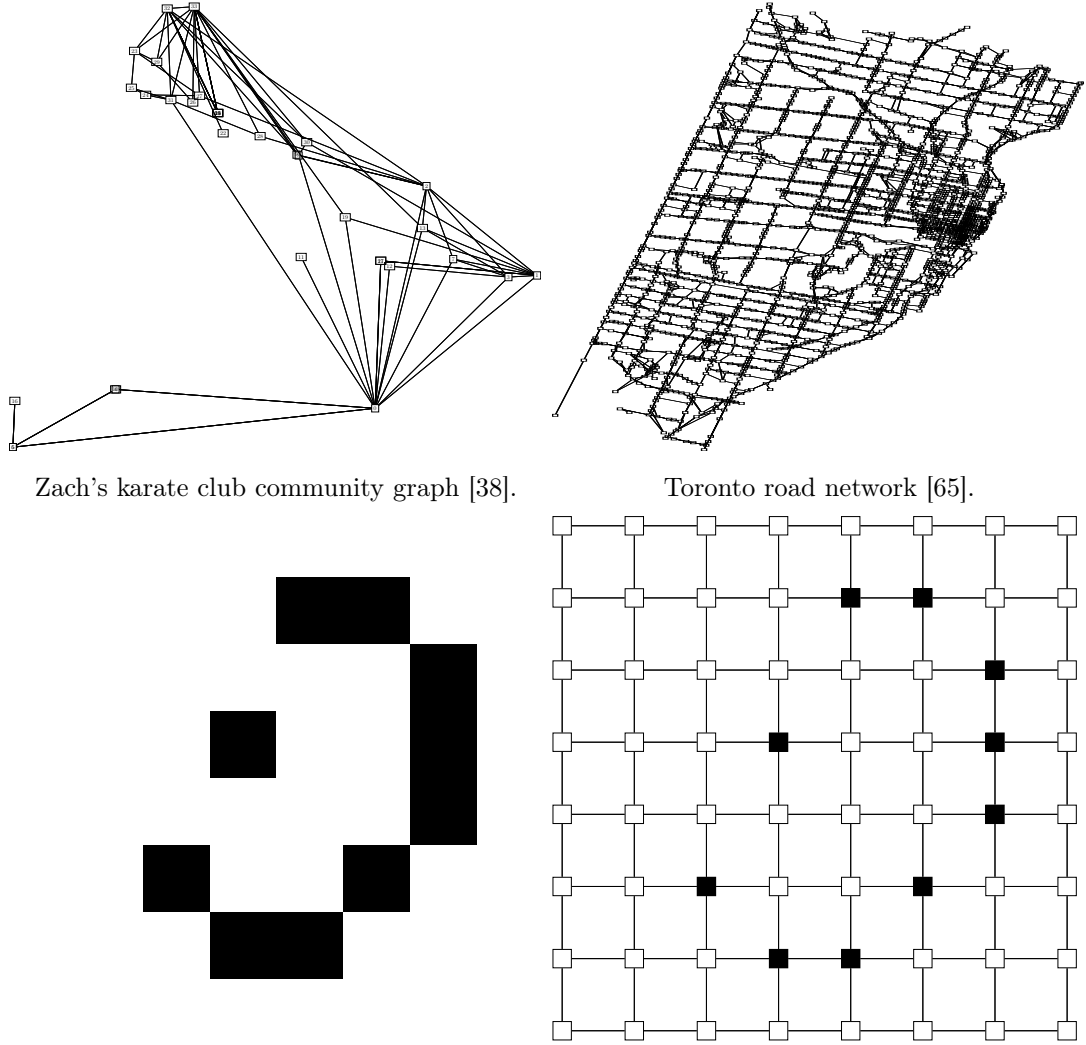


Figure 3.1: Depiction of various examples of graphs commonly used in machine learning problems.

In the following paragraphs we introduce operations on graphs that are going to be useful in the rest of this thesis.

3.1.1 Translations on graphs

In the previous chapter we introduced the ideas of translation-equivariance and translation-invariance and their uses in convolutional networks and computer vision. In the same vein, graphs and graph signals can be subjected to translation. In this subsection we will introduce the concept of graph translations, as was previously defined in the laboratory [118], and define how they are used in this thesis.

First off, we must start by recalling what is the translation operator. Indeed in a discrete euclidean space, translations are quite straightforward to define. For example consider a signal $s(t)$ that evolves over time. This signal $s(t)$ is defined in an 1D euclidean space and can therefore be translated either forward (“advancing in time”) or backwards (“going back in time”). In the same vein, if our signal s is supported on an 2D euclidean space, as it is the case with images, it is quite direct to infer four types of translations, by sending all pixels downwards, upwards, to the right or to the left.

In order to transfer this same translation concept from euclidean spaces to graphs, we choose to use in this work the translation definition and inference methods introduced in [118] that we extend in [91].

Definition 3.1.2 (translation). A *translation* on a graph is a function $\phi : \mathbb{U} \rightarrow \mathbb{V}$, where $\mathbb{U} \subset \mathbb{V}$ and that is:

1. *injective*: $\forall v, v' \in \mathbb{U}, \phi(v) = \phi(v') \Rightarrow v = v'$,
2. *edge-constrained*: $\forall v \in \mathbb{U}, (v, \phi(v)) \in \mathbb{E}$,
3. *strongly neighborhood-preserving*: $\forall v, v' \in \mathbb{U}, (v, v') \in \mathbb{E} \Leftrightarrow (\phi(v), \phi(v')) \in \mathbb{E}$.

We also define the *loss* of a translation as the cardinal $|\mathbb{V} - \mathbb{U}|$ that counts the vertices that are not a part of the translation. We also say that two translations ϕ and ϕ' are *aligned* if $\exists v \in \mathbb{U}, \phi(v) = \phi'(v)$.

Ideally we would also add that translations should be lossless, i.e., $|\mathbb{V}| = |\mathbb{U}|$. Unfortunately it is not possible to guarantee that a given graph \mathcal{G} will be able to admit lossless translations [118]. Therefore in this work we compromise by considering only minimal translations:

Definition 3.1.3 (minimal-translation). A translation is said to be minimal if there is no aligned translation with a strictly smaller loss.

We depict in Figure 3.2 the minimal translations of a grid-graph representing a 2D discrete euclidean space. We note that the inferred translations are exactly the same as previously defined for a 2D euclidean space. We will use this property in Section 4.2 to define convolutions on graphs based on the 2D convolutional layers described in the previous chapter.

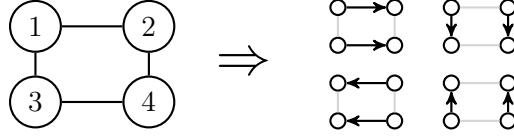


Figure 3.2: Depiction of the minimal translations of a 2x2 grid graph.

3.1.2 Graph Fourier Transform (GFT)

In the domain of signal processing, the **frequency** is one of the most important concepts to analyse a signal and a starting point to introduce many useful tools including filtering. The frequency of a signal can be simplified as its rate of change, or in other words how it varies from one sample to another. Generally speaking, any signal can be decomposed into a continuous sum of sines/cosines by performing the Fourier transform. This decomposition can also be considered as expressing the signal in the frequency domain, as each sine/cosine is characterized by a proper frequency, and the original representation of the signal as its expression on the time domain.

Both the classical and graph Fourier transforms were first introduced to deal with signals defined on time, and then a posteriori extended to deal with other domains. The classical Fourier transform leads itself well to clearly structured domains such as the discrete Euclidean space, while tries to extend the framework developed using the Fourier transform to more loosely defined structures that can be supported on graphs.

We present an example of the Fourier transform and a simple filtering application in a series of figures, from Figure 3.3 to Figure 3.5. First, in the left part of Figure 3.3 we depict an original signal ($\mathbf{s}(t)$), and in the right part of the Figure we depict the same signal with added white noise ($\check{\mathbf{s}}(t)$). We then use the Fourier transform to decompose both signals in the frequency domain (f) that we depict in Figure 3.4. We can observe that, since the white noise has no specific frequency, it has very little impact on the Fourier transform of the signal. As such, it would for example be much easier to perform classification in the frequency domain.

Note how the original signal is mostly defined in the low-frequency side, while the noisy signal has many more high-frequency components. It is therefore straightforward to say that performing a filtering operation to remove the high frequency components in the noisy signal $\check{\mathbf{s}}(t)$ will allow us to retrieve a signal $\hat{\mathbf{s}}(t)$ that is more inline with the original one. The retrieved signal $\hat{\mathbf{s}}(t)$ is also commonly called **denoised signal**. We depict the original signal, the noisy signal and the retrieved signal in Figure 3.5. Note that the retrieved signal is much more inline with the original signal, even if it is not

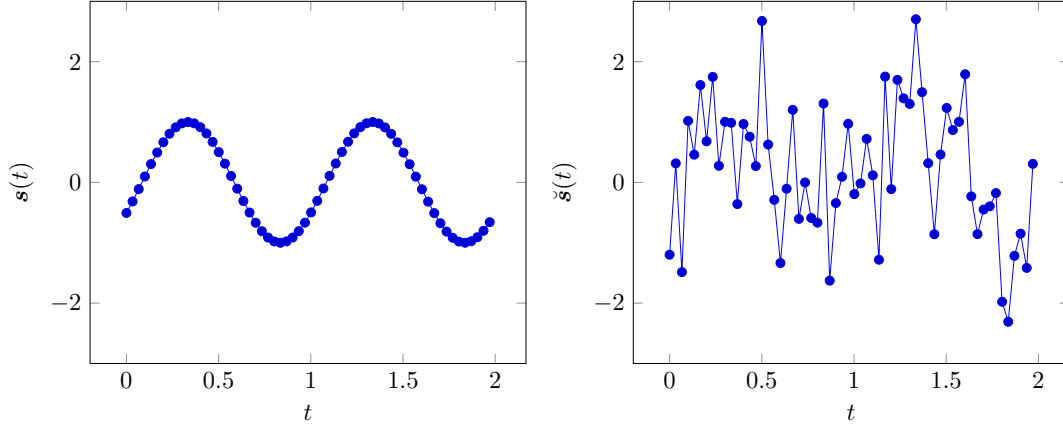


Figure 3.3: Depiction of a signal $s(t)$ and its noisy version $\check{s}(t)$.

exactly the same. Indeed, this denoising operation had the unfortunate effect of lowering the variations of the signal even for small frequencies.

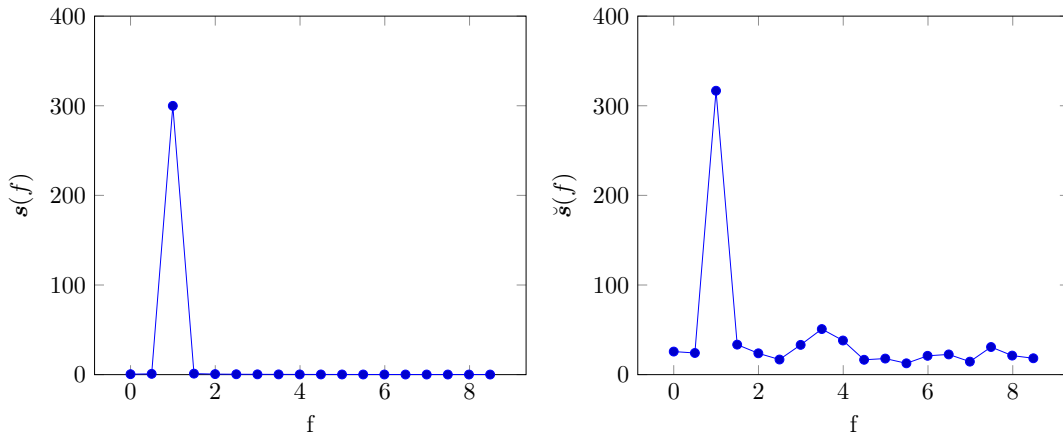


Figure 3.4: Depiction of the Fourier transform of the signal $s(t)$ and its noisy version $\check{s}(t)$.

As we have previously defined, the goal of the **Graph Fourier Transform (GFT)** is to extend the same type of analysis and tools to signals that are described in the domain of the vertices of graphs (i.e. graph signals) instead of the time domain. In the case of graph signals the rate of change will not be evaluated as time evolves¹ but as the vertices evolve (i.e. the relationship between a vertex and its neighbors). In the case of GFT the transform is defined using the **graph Laplacian**.

The graph Laplacian L is defined as:

$$L = D - A, \quad (3.2)$$

¹at least in this thesis, we note that other works consider graph signals that vary on both graph and time domain.

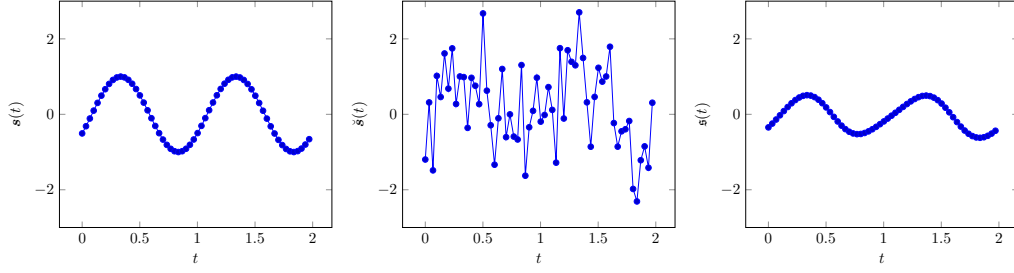


Figure 3.5: Depiction of a signal $\mathbf{s}(t)$, its noisy version $\check{\mathbf{s}}(t)$ and the denoised version $\hat{\mathbf{s}}(t)$.

where \mathbf{D} is the diagonal degree matrix of the graph defined by \mathcal{A} . As the Laplacian matrix is both real and symmetric it can be eigendecomposed into:

$$\mathbf{L} = \mathbf{F}\mathbf{\Lambda}\mathbf{F}^\top, \quad (3.3)$$

where \mathbf{F} are the eigenvectors of \mathbf{L} and $\mathbf{\Lambda}$ are the eigenvalues in crescent order of magnitude. The GFT of a graph signal \mathbf{s} can then be defined as:

$$\tilde{\mathbf{s}} = \mathbf{F}^\top \mathbf{s}, \quad (3.4)$$

where $\tilde{\mathbf{s}} \in \mathbb{R}^{|\mathcal{V}|}$ is the graph signal \mathbf{s} decomposition in the frequency domain, with each dimension corresponds to a specific frequency. Note that the inverse GFT can be similarly defined as:

$$\mathbf{s} = \mathbf{F}\tilde{\mathbf{s}}. \quad (3.5)$$

To illustrate the GFT, let us retake our previous example that used the Fourier transform in the time domain. The time domain could be represented using a simple line graph where each vertex corresponds to a specific time sample and is connected to the consecutive time samples. In that case, the GFT is very similar to the usual Fourier transform in the time domain. This can be seen in Figure 3.6, where we depict the same signal from Figure 3.3 in the left part, we then show a discretization of the signal $\mathbf{s}(t)$ in the center part, where each sample would correspond to a vertex in the line graph, and finally the GFT transform in the right part. As we had done in the previous example, we add white noise to both representations of our signal, and depict the noisy version, the sampled noisy version of the signal and its GFT in Figure 3.7. We can observe very similar behaviors than in the previous “classical” Fourier domain.

We are then able to perform the same filtering operation from before, and depict the results in Figure 3.8. In the left we have the original signal, in the middle the signal that was filtered in the time domain and in the right the signal that was filtered in the graph domain. Note that the results are not exactly the same, as our graph signal is a

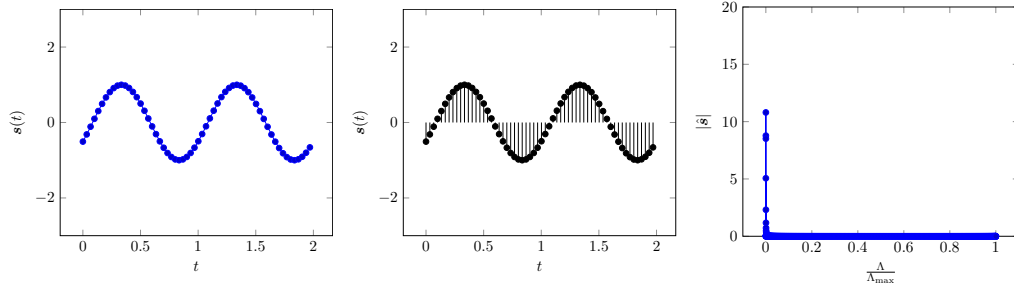


Figure 3.6: Depiction of a signal $s(t)$, its sampled version and the GFT of the graph signal.

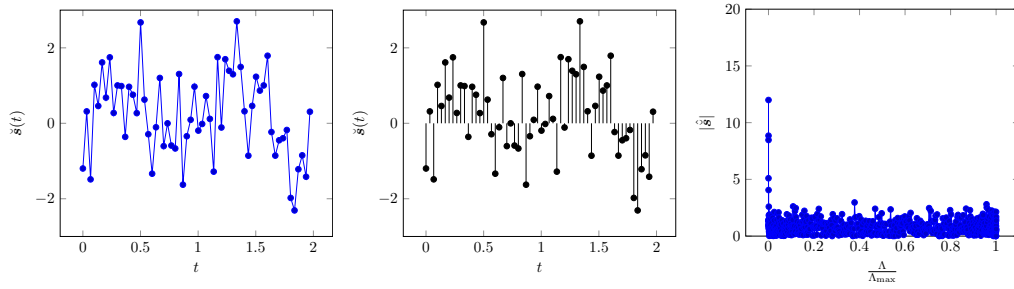


Figure 3.7: Depiction of a noisy signal $\check{s}(t)$, its sampled version and the GFT of the graph signal.

discretization of the real time signal, but the results are very close. We will present in the next sections some of the uses of the GFT and its abstractions.

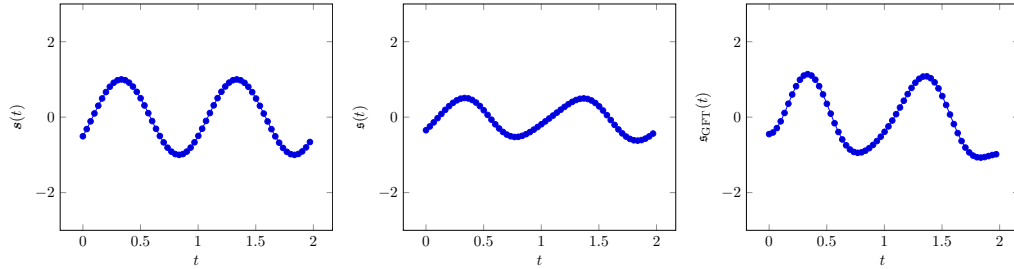


Figure 3.8: Depiction of a signal $s(t)$ and its retrieved time FT ($\check{s}(t)$) and GFT ($\mathfrak{s}_{\text{GFT}}(t)$) filtered signals.

3.1.3 Smoothness of graph signals

Now that we have introduced graphs and the GFT, we can now talk about the **smoothness** of graph signals. This concept is the cornerstone for most of the methods we introduce

in Chapter 5. In this subsection we will first define the concept of smoothness of graph signals and then we present an illustrative example of an application of graph signal smoothness for noise detection/removal.

3.1.3.1 Definition

The smoothness σ of a graph signal \mathbf{s} supported on a graph \mathcal{G} is defined as:

$$\sigma = \mathbf{s} \mathbf{L} \mathbf{s}^\top. \quad (3.6)$$

This is also called the quadratic form of the Laplacian. We note that while we call this measure “smoothness” to be coherent with GSP literature a more adequate name would be “anti-smoothness” as lower values of the smoothness metric are said to be very smooth in respect to the graph and higher values of the smoothness metric are said to be very rough (or unsmooth) in respect to the graph. If we consider the eigendecomposition of the Laplacian matrix (c.f. Equation 3.3), we can also rewrite σ as:

$$\sigma = \mathbf{s} \mathbf{L} \mathbf{s}^\top = \sum_{i=1}^{|\mathbb{V}|} \Lambda_{i,i} \mathbf{s}_i, \quad (3.7)$$

where $\mathbf{\Lambda}$ is a diagonal matrix of the eigenvalues of \mathbf{L} in crescent order. In this way we can see that a smooth signal will be one that is aligned with the first eigenvectors of the Laplacian. Finally, we observe that the smoothness is strongly related to the rate of change of the signal values from one vertex to its neighbors, by rewriting σ as:

$$\sigma = \mathbf{s} \mathbf{L} \mathbf{s}^\top = \mathcal{A}_{i,j} (\mathbf{s}_i - \mathbf{s}_j)^2. \quad (3.8)$$

Considering the smoothness as the rate of change is a very useful abstraction, especially when the signal \mathbf{s} is binary. Indeed if \mathbf{s} is binary, the smoothness can be simplified as the sum of the weights between nodes with different values in \mathbf{s} . If we consider the example where each entry of \mathbf{s} is a binary label indicator vector first defined in Definition 2.1.2 and the graph vertices correspond to samples, the smoothness will be the sum of the weights of edges connecting samples of different classes and the smoothest graph possible would be one that connects only examples of the same class². In other words, smoothness is a measure of discrepancy between a signal and a graph structure.

In the following paragraphs we introduce an application of graph signal smoothness as an illustrative example.

²A graph that has $\mathbb{E} = \emptyset$ would also have $\sigma = 0$, but we do not consider this type of graph in this thesis.

3.1.3.2 Illustrative example

Now that we have defined what is the smoothness of a graph signal, an illustrative example is in order to illustrate its usefulness. Let us reuse the example signal from Figure 3.6 where a signal in the time domain is discretized and converted into a graph signal on a line graph. In Figure 3.9 we depict the original graph signal in the left part, a noisy version of the graph signal in the center and a low-pass filtered version of the noisy graph signal in the right and their smoothness. Note how both the original graph signal and the filtered version are smoother than the noisy version of the signal.

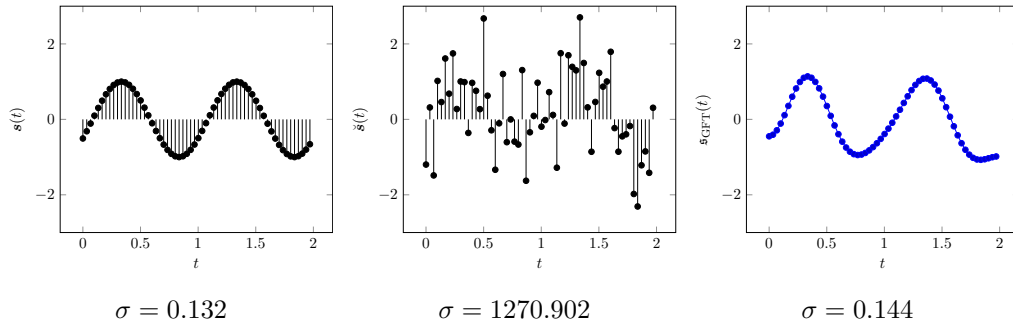


Figure 3.9: Depiction of a sampled signal \mathbf{s} , its noisy version $\tilde{\mathbf{s}}$ and its filtered version \mathbf{s}_{GFT} , with their respective smoothness (σ) values.

In other words the smoothness of a graph signal can also be used to detect if there is noise present in it. In the following of this document we will introduce how this measure can be used to infer a graph from a signal in Section 3.3, how it can be used to improve the robustness of DNNs in Section 5.3 and finally how it can be used to train DNNs for classification in Section 5.2.

3.2 Graphs for samples of features

In this thesis we focus on two main uses of graphs: i) Graphs that model the inner dependencies of observations, and ii) Graphs that model the relationship between data samples.. In the following paragraphs we detail exactly what we understand as each type of graph and give practical examples to illustrate each case.

3.2.1 Modelling inner dependencies of observations with graphs

One of the uses of graphs is modelling the inner dependencies of the observations. In this case each vertex is a coordinate of an observation and the relationship between vertices encode the relationship between the different coordinates. Encoding the relationship between different coordinates can also be seen as exploiting the intrinsic structure of the data samples. This is mostly used for supervised classification of graphs (e.g., protein-protein interaction) and supervised classification of graph signals (e.g., classification of scrambled images as described in Section 4.2).

For example consider that our observations/elements of interest are images. In this case we can create a grid-graph to emulate the intrinsic euclidean 2D structure. The grid representation creates a graph that model the inner dependencies between the pixels (coordinates of an observation) and therefore allow one to extract important information from the intrinsic structure (c.f. Section 4.2). We depict in Figure 3.10 an image and its grid graph representation.

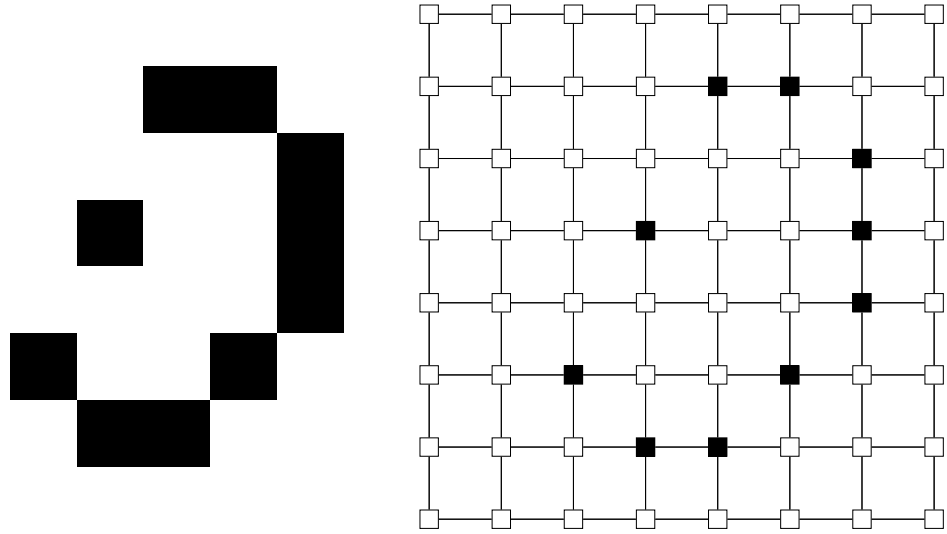


Figure 3.10: Depiction of an image (left) and its grid graph representation (right).

Definition 3.2.1 (2D grid graph). We call (2D) *grid graph* a graph whose vertices are of the form $\{1, \dots, \ell\} \times \{1, \dots, h\}$ where $\ell, h \geq 3$, and edges are added between vertices at Manhattan distance one from each other.

This grid graph image representation has been exploited in multiple GSP works as a basis for explanation and visualization [118, 148]. In the case of an image grid graph, each pixel is a vertex v_i and the RGB values form a graph signal $\mathbf{S} \in \mathbb{R}^{|\mathbb{V}| \times 3}$ where w and h represent the image width and height ($|\mathbb{V}|$) respectively and 3 is the number of the

RGB color channels ($|F|$). In other words, we have $|\mathbb{V}| = |wh|$ coordinates and $|F| = 3$ the number of observations or realizations per coordinate.

Using a low-pass graph filter (that we introduce more formally in Section 3.4) it is possible to exploit the structure of the pixels to remove noise from the image as shown in Figure 3.11. Note how by removing the high frequencies of the graph signal, we obtain an image that is more inline with the original image. Also note that the smoothness value of the recovered graph signal is the same as the original image, while the images are very different. Indeed as the smoothness of a graph signal is a global measure and not a localized one it is easy to see that multiple image configurations are possible for the same value of smoothness.

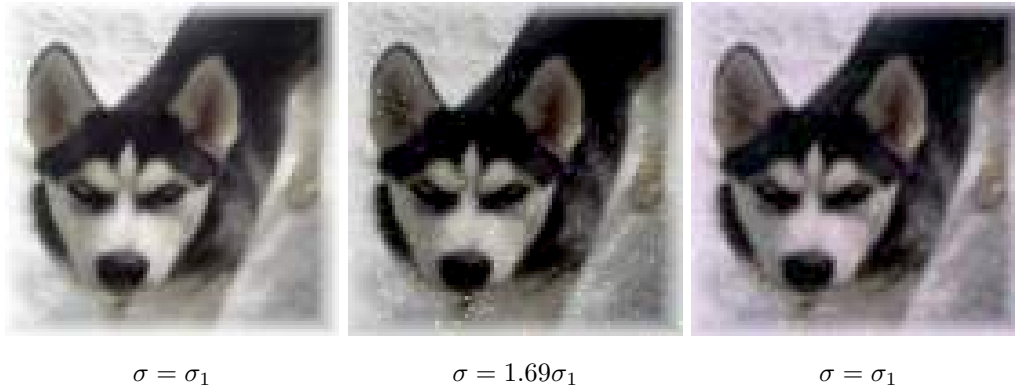


Figure 3.11: Depiction of a dog (left), a noisy realization of the image (center) and the graph filtered image (right).

Finally, it is also possible to infer the graph structure solely from the signal. One naive approach is to compute a similarity metric between the signal coordinates to generate a weighted adjacency matrix and then threshold the k most connected neighbors of each coordinate to generate the graph. Note that the adjacency matrix that is obtained by this approach needs to be then symmetrized. The resulting graph is commonly called a k -nn similarity graph. We delve into more details in graph inference in Section 3.3

3.2.2 Graphs that model the relationship between data samples

The other type of graphs that we consider in this work are graphs that model relationship between different observations. In this case each observation is a vertex and the relationship between the vertices encodes a relationship between them. This relationship depends on the task at hand, for example in the document classification datasets presented in Section 2.3.5, an edge (v_1, v_2) will encode the fact that document v_1 cites v_2 or is cited

by v_2 .

In a graph that models the relationship between data samples, the goal is to exploit the relationship between the different samples rather than the intrinsic structure of each one. This allows us to consider other scenarios such as semi-supervised classification of vertices in the graph. We introduce an example of semi-supervised classification of vertices in the following paragraphs.

Consider that we have a subset of the CIFAR-10 dataset (c.f. Section 2.3.1.1), where we only consider the cat and truck classes, with 10 training examples per class (labeled examples) and 50 test examples per class (unlabeled examples). We can then apply a naive graph inference technique using the cosine similarity between the samples and threshold the k -neighbors in order to generate a graph. We depict the generated graphs in Figure 3.12, where on the left we depict the graph masking the unlabeled examples, on the center the we depict the graph with labels retrieved using the label propagation algorithm (c.f. Section 3.3.4.1 for more details) and finally on the right we depict the ground-truth.

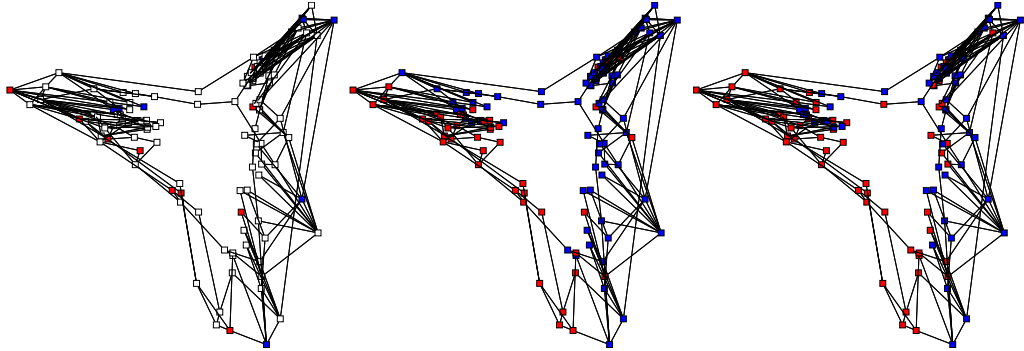


Figure 3.12: Depiction of a graph connecting samples from a subset of the CIFAR-10 dataset, where red represents images of cats, blue represents images of trucks and white represents unlabeled data. On the left we have masked the unlabeled data, while on the center we use a label propagation algorithm to retrieve the labels of the unlabeled data with 70% accuracy and finally on the right we depict the ground truth labels.

Note how just by organizing the data in a graph structure and using Laplacian eigenmaps [6], i.e. taking advantage of the first two eigenvectors of the laplacian³ to position the nodes in a 2D space, we can better understand our data and retrieve a very simple semi-supervised classification baseline. In the following of this thesis we will use this representation in order to solve various tasks.

³we consider the first two eigenvectors to be the eigenvectors associated with the two smallest nonzero eigenvalues.

3.3 Inferring graph topology from signals

In the previous sections we defined graphs, graph signals and gave examples on how they are going to be used in this thesis. However even if a signal has an underlying structure, this support structure is not necessarily explicitly available. If we do not have access to the adjacency matrix that supports a signal, inferring this graph topology is required. We published a large portion of the results in this section in [89].

Many recent works have tackled the problem of inferring the graph topology using graph signals [71, 119, 146]; see also [108] for a recent tutorial treatment. Inferring a graph structure can be performed in a task-agnostic manner, i.e., labeling data is not used. Priors are used to relate observations to the sought task-agnostic graph structure: e.g. smoothness [71] (c.f. Section 3.1.3), stationarity [119], sparsity [146], and probabilistic [29] as well as graph filtering-based [144] generative models, just to name a few.

On the other hand, sometimes it is interesting to consider graphs that are specific to the task at hand. For example in [78] the authors infer graphs for medical search and in [58, 67, 86], the authors aim at improving the accuracy of various classification tasks using inferred graphs: semi-supervised learning, visual based localization and few-shot learning.

As it is more general, task-agnostic graph inference is of particular interest. In this context, the two most common goals are visualization [71] and interpretation [1] using the graph support. On the other hand, it is challenging to compare methods for which there is no ground truth, i.e., that are task agnostic. Unsurprisingly, many works rely on synthetic data to evaluate the ability of their proposed methods in unveiling the topology from the observations. We present a benchmark and delve into more details on the discussion of evaluating task-agnostic methods in Section 3.3.4.

In the following subsections we present a simple framework of graph inference that we call “naive baselines” and we go into more details in the two methods that we are going to use in this work: i) Kalofolias [71], and ii) NNK [146].

3.3.1 Naive baselines for graph topology inference

In this subsection we present a simple framework that allows us to infer graph structures from data with 3 quick steps:

1. Choosing a similarity measure to be applied to the features of each vertex in order to

determine their connectivity. In more details, we consider cosine similarity, sampled covariance or an RBF kernel applied on the L_2 distance between considered items. The result is a square matrix homogeneous to the desired graph structure. Note that some of the similarity measures may require extra hyperparameters such as the γ in an RBF kernel (c.f., Equation 2.5).

2. Choosing a number of neighbors to be kept for each vertex. We simply use a k -nearest neighbor selection. Note that we symmetrize the resulting graph, so that each vertex has at least k neighbors. Note that as it was the case with the similarity measure, in this step we are also adding an extra hyperparameter, which is the number of neighbors to keep.
3. Normalizing the obtained graph adjacency matrix. Note that the decision of normalizing (and which normalization to use) or not the adjacency matrix adds another hyperparameter to the framework.

Note that while this framework is very simple, the amount of decisions and possible combinations is very large. In this subsection as we want to investigate the maximum ability of the methods we test a very extensive number of combinations and always report the best one. In the remaining of this work (excluding the present Section), we will either use a fixed decision on how to infer our graphs from data or perform minimal experiments to show the impacts of this decision.

3.3.2 Graph construction with Non Negative Kernel (NNK) regression

We now consider a recently proposed graph inference method. We choose NNK (Non Negative Kernel regression) [146], due to its simplicity and its demonstrated results on semi-supervised learning tasks. This method can be interpreted as producing representations with orthogonal approximation errors, which in turn favors sparser representations. It has two parameters: k , the maximum degree for each vertex, and σ the minimum value for an edge weight (threshold). In this work we test multiple values of k and fix $\sigma = 10^{-4}$ following [71]. In our experiments we use the authors implementation from https://github.com/STAC-USC/PyNNK_graph_construction.

3.3.3 Graph learning from smooth signals

Graphs can be inferred from graph signals using certain priors. In this subsection we consider that a signal should be smooth (c.f., Section 3.1.3) on its graph support. In this

thesis we rely on a state-of-the-art approach in [71]. It consists in a framework that infers the graph from an underlying set of smooth signals. As it was the case with NNK, it has two parameters: k the desired mean sparsity and σ the minimum value for an edge weight. In this work we test the same values for these two parameters as we did for NNK and keep the best combination. In our experiments we use the implementation from the GSP toolbox [121].

3.3.4 Benchmarking graph topology inference

Now that we have presented the graph topology, we can discuss the problem of comparing these methods. Most methods have to rely on synthetic data in order to evaluate their efficacy and to put forth their pros and cons. While synthetic data are always useful to perform controlled scalability experiments as well as reveal the emerging statistical and computational trade-offs, this validation protocol comes with two shortcomings. First, the models used to generate synthetic data are likely to be biased in favor of the proposed methods. Second, the ability of the proposed method to handle hard real-world problems is often not demonstrated convincingly, e.g., very small datasets or toy data such as black and white digit recognition.

In order to address this problem, standardized benchmarks are required. The main challenge is that benchmarks are necessarily task-specific, and as such they do not encompass the whole potential offered by task-agnostic graph inference state-of-the-art methods. To fill in this gap, in this section we detail a broad collection of benchmarks, that we first introduced in [89] that are specifically designed to compare graph inference algorithms. To this end, we consider three timely problems arising with network data: i) unsupervised clustering of vertices, ii) semi-supervised classification of vertices (with or without vertex features) and, and iii) graph signal denoising. For each problem we introduce an easy-to-use dataset that we release publicly⁴.

We also introduce measures of performance to confront methods and better understand their adequacy to the aforementioned tasks. Furthermore, the released datasets comprise various types of signals, namely natural images, audio, texts, and traffic information. Note that we do not include brain data and protein-protein interactions that are two of the most interesting use-cases of graph inference and classification. Our choice is informed by recent developments in the literature [31, 51], that have found no significant performance gains when graph-based machine learning techniques are brought to bear for some tasks in these areas. We will further discuss this in Chapter 4.

⁴https://github.com/cadurosar/benchmark_graphinference

Our benchmarks are divided into three tasks that encompass two types of machine learning problems. In Tasks 1 and 2, the graph model dependencies between observations, c.f. Section 3.2.2. In the second one (Task 3), the graph models relationships between features as seen in Section 3.2.1. We expect some of the previously presented methods to perform better on the first series of tasks and others to be more adequate to the second one.

3.3.4.1 Benchmarking tasks

3.3.4.1.1 Task 1: Unsupervised Clustering of Vertices (UCV): Consider a dataset composed of $|\mathbb{V}| = N$ observations, each one containing F features. Given a number of classes c , we consider the task of partitioning the N observations into c classes, such that the variability inside classes is smaller than the variability between classes. In practice, variability can be measured using various metrics. For the purpose of obtaining quantified benchmarks, we consider here that the observations belong to c categories (e.g. classes of images or sounds), and that this information is not available when processing the considered methods. So, the performance of a considered method is evaluated by computing the Adjusted Mutual Information score [169] based on the ground truth.

Note that this clustering problem can be treated without a graph structure. Examples are using c -means or DB-Scan algorithms. In the context of this work, we consider using spectral clustering. Spectral clustering consists in creating a graph linking the observations where the edges are inferred from the corresponding features. Then, vertices are projected using the first eigenvectors of the graph Laplacian and clustered using standard non-graph methods. In our work, we use the discretization method first proposed in [178] when features have been projected onto the first c eigenvectors of the graph Laplacian except the very first one. We use the default SciKit-Learn [120] implementation of spectral clustering and of the c -means algorithm in our experiments.

3.3.4.1.2 Task 2: Semi-Supervised Classification of Vertices (SSCV): Consider a dataset composed of $|\mathbb{V}| = N$ observations, each one containing F features. Here, a portion of the N observations are labeled. The task consists in inferring the labels of the other portion of observations. Again, we consider datasets where we have access to the ground truth, and artificially hide the labels of part of the observations when processing the data. The score consists in measuring the accuracy of the classification on initially unlabeled observations.

This problem can be solved without relying on graphs. For example, a common

solution would consist in performing a supervised classification using only the labeled observations. In this work, we consider inferring a graph connecting observations from the features. Then, we use this graph in two settings. In the first setting, we want the graph to fully encompass the information contained in the features, and therefore perform label propagation. Label propagation consists in diffusing the labels from the known observations to the other ones using the inferred graph structure. In a second setting, we use both the graph structure and the features to perform classification. We use the methodology described in [173], called Simplified Graph Convolution (SGC), where the goal is to combine feature diffusion with logistic regression. We delve into more details for this type of model in Section 4.3.

In more details, we use two layers of feature diffusion ($\hat{\mathbf{x}} = \mathcal{A}^2 \mathbf{x}$), followed by a logistic regression. The models are trained for 100 epochs, using Adam optimization with a learning rate of 0.001. We use the average over 100 runs of the accuracy using random splits of 5% training set and 95% test set. We always report the average accuracy and standard deviation. To propagate labels, we simply diffuse the label signal one time using the exponential of the adjacency matrix. We note that SGC models tend to use the “normalized augmented adjacency matrix” $\tilde{\mathcal{A}} = \mathcal{I} + \mathcal{A}$ where \mathcal{I} is the identity matrix. This augmented adjacency matrix is then normalized $\tilde{\mathcal{A}} \leftarrow \mathbf{D}_{\tilde{\mathcal{A}}}^{-1/2} \tilde{\mathcal{A}} \mathbf{D}_{\tilde{\mathcal{A}}}^{-1/2}$. In our work we test both the adjacency matrix and the augmented adjacency matrix and their respective normalizations and we report the best possible combination in terms of mean accuracy.

3.3.4.1.3 Task 3: Denoising of Graph Signals (DGS): Consider a dataset comprising N observations, each one consisting of $|\mathbb{V}| = F$ features. Consider some additive noise generated according to a distribution \mathcal{N} . The task consists in recovering initial observations from their noisy versions. We measure performance by looking at the Signal to Noise Rate.

Here, the graph connects features of observations. The idea is to use the graph structure to easily segregate components of the noise from components of the initial signals. In our work, we use a Simoncelli low-pass filter (c.f., Section 3.4 for more details on graph filtering) on the graph to perform denoising. Note that this filter has a parameter $\tau \in [0, 1]$ that we vary from 0 to 1 in increments of 0.025. We use the noisy signal realization with a SNR (Signal to Noise Ratio) of 7, from [65], and report the best SNR found for each graph construction.

3.3.4.2 Datasets

For the only purpose of benchmarking graph inference methods, we introduce here a few datasets. For Tasks 1 and 2, we use datasets of images, audio and texts (documents). To reduce the difficulty of the tasks in the image and audio domains, we choose to use features extracted from pretrained deep neural networks. Task 3 (DGS) data comes from real life traffic information. Additional details are given in the coming paragraphs.

3.3.4.2.1 Image dataset - flowers102: For the image dataset we use the training set portion of the “102 Category Flower Dataset” (shortened as flowers102) [114]. This split contains $N = 1020$ images of $C = 102$ classes of flowers (10 images per class). The features are extracted from the final pooling layer of the Inceptionv3 architecture [155], which has a size of $F = 2048$ dimensions. Note that Inceptionv3 was trained on the 2012 split of ImageNet challenge, so that the features we obtain are a case of transfer learning. This should be one of the most challenging scenarios we consider, as it provides the highest number of classes and has the highest signal dimension to number of items ratio: 2.

3.3.4.2.2 Audio dataset - ESC-50: For audio data, we use “ESC-50: Dataset for Environmental Sound Classification” [126]. This dataset contains $c = 50$ classes, with 40 audio signals each (2000 in total). It also contains 5 standard splits that are not used here (as we do unsupervised and semi-supervised classification). We use the feature extractor introduced in [83] to generate our dataset, that was trained on AudioSet. Similar to the images data, this can be considered as transfer. At the end we have $N = 2000$ items with $F = 1024$ dimensions each. The signal dimension to number of items ratio is 0.512.

3.3.4.2.3 Text dataset - cora: We use the cora dataset [143] that we have presented in Section 2.3.5.1, which is composed of $N = 2708$ scientific articles of $c = 7$ different domains for document clustering or classification. The features come from a word indicator vector (i.e. bag of words) that indicates if one of the words in the dictionary ($F = 1433$ in total) is present on the title or abstract of the document. The dictionary is built with the most common words in the dataset. The signal dimension to number of items ratio is: 0.53. Note that this dataset is classically used for graph semi-supervised learning as it comes with a citation graph. But in our work we completely disregard this graph. Comparisons between the ground truth graph and inferred ones could be an interesting addition to this work. But since the citation graph is not exactly redundant with the signals, it is expected that inferred graphs and citation ones are quite different.

3.3.5 Toronto traffic data denoising (Toronto)

We use data from the road network of the city of Toronto, from [65]. It describes traffic volume data over a 24 hour period at intersections in the road network of Toronto for a total of $F = 2202$ vertices and $N = 1$ observation. Note that extra information is available, such as the position of each road and intersection, but our baselines only consider the raw signal data. This graph is depicted in Figure 3.1

3.3.5.1 Empirical evaluation of benchmarks

We now present the results for our benchmark evaluation. The tested graph topology methods and the parameters we vary are summarized in Table 3.1. For every test we only display the results obtained by the best combination and we further discuss the effects of parameter choice in Section 3.3.5.2.

Table 3.1: Summary of the tested graph topology inference methods. Table and caption extracted from [89].

Method	Similarity/Distance	k	σ	Adjacency matrices
Naive	Cosine, Covariance, RBF	5, 10, 20, 30, 40, 50,	None	$\mathcal{A}, \mathbf{D}_{\mathcal{A}}^{-1/2} \mathcal{A} \mathbf{D}_{\mathcal{A}}^{-1/2},$ $\tilde{\mathcal{A}}, \mathbf{D}_{\tilde{\mathcal{A}}}^{-1/2} \tilde{\mathcal{A}} \mathbf{D}_{\tilde{\mathcal{A}}}^{-1/2}$
NNK [146]		100, 200, 500, 1000	10^{-4}	
Kalofolias [71]	Square Euclidean distance			

3.3.5.1.1 Task 1: For the UCV task, we display both the results obtained with the inferred graph structures and with a c -means baseline. The results are presented in Table 3.2. We can see that both naive and NNK get the most consistent results, with Kalofolias having difficulties with the cora dataset.

Table 3.2: Results for Task 1. Here we present the best AMI score for each inference method. Table and caption extracted from [89].

Method	Inference/Dataset	ESC-50	cora	flowers102
C -means		0.59	0.10	0.36
Spectral clustering	Naive	0.66	0.34	0.45
	NNK	0.66	0.34	0.44
	Kalofolias	0.65	0.27	0.44

3.3.5.1.2 Task 2: For the SSCV task, the results are presented in Table 3.3. We can see that using a similarity graph as support helps when compared to a simple logistic regression. Note that unfortunately, this is not a 100% fair comparison as the logistic regression is not able to exploit the unsupervised data. In this task we have two methods, Label Propagation and SGC. In the first one, Kalofolias presents the best results for both flowers102 and ESC-50, but still struggles with the cora dataset. In SGC both Kalofolias and NNK seem to not be able to improve that much over the naive baselines.

Table 3.3: Results for Task 2. Here we present the best mean test accuracy and its standard deviation for each inference method. Table and caption extracted from [89].

Method	Inference/Dataset	ESC-50	cora	flowers102
Logistic Regression		52.92% \pm 1.9	46.84% \pm 1.6	33.51% \pm 1.7
Label Propagation	Naive	59.05% \pm 1.8	58.86% \pm 2.9	36.73% \pm 1.6
	NNK	57.44% \pm 2.2	58.66% \pm 2.9	33.57% \pm 1.6
	Kalofolias	59.16% \pm 1.8	58.60% \pm 3.4	37.01% \pm 1.7
SGC	Naive	60.48% \pm 2.0	67.19% \pm 1.5	37.73% \pm 1.5
	NNK	61.38% \pm 2.0	66.58% \pm 1.5	36.81% \pm 1.5
	Kalofolias	59.36% \pm 2.0	66.28% \pm 1.5	37.5% \pm 1.5

3.3.5.1.3 Task 3: For the graph signal denoising task, the results are presented in Table 3.4. In this scenario we are not able to use neither cosine or covariance similarity. We compare our results with the ones we would obtain using the ground truth road map graph. Our RBF baselines were able to reduce the amount of noise, but not at the same level as of the real road graph. The Kalofolias smooth graph was able to achieve a better SNR than the real road graph.

Table 3.4: Results for Task 3. Here we present the best test accuracy for each baseline. Table and caption extracted from [89].

Best SNR	Road graph	Kalofolias	RBF NNK	RBF k -NN
	10.32	10.41	9.99	9.80

3.3.5.2 Discussion

Over all tasks we can extract some lessons on graph inference:

1. **Similarity choice:** If we have multiple non-negative realizations of the signal,

cosine seems the best choice. It has competitive results on all benchmarks and it does not come with a parameter (as does RBF with γ).

2. **Choosing parameter k :** The best amount of sparsity depends not only on the dataset and task, but on the similarity that was chosen. We consider the ESC-50 dataset as an example. In the spectral clustering the best k value for the k -NN graph was 30 for cosine, 5 for RBF and 20 for covariance. We note that in the graph denoising task, the best case was to not perform k -neighbors thresholding.
3. **Normalization:** Note that only our graph denoising task does not expect a normalized graph, therefore most of our better results used normalized graphs. On the graph denoising task, normalized and non-normalized graphs had similar results.
4. **Cora dataset:** The cora dataset is challenging not only because it is not class-balanced, but also because its features are binary (a bag of words, containing 1 if the word is present in the article and 0 if not). This could be a reason for the bad performance of both NNK and Kalofolias in this dataset.
5. **Sparse graphs in semi-supervised problems:** In the semi-supervised tasks, the test accuracy standard deviation over the splits was very high. This could possibly be caused by the fact the sparse graphs we use here have more than one connected component, meaning that sometimes there could be sections of the graph that do not have any labeled vertices. One possible future direction would be to integrate a graph sampling algorithm to the problem in order to select which vertices we should label, instead of doing so randomly.
6. **Naive Baselines vs. optimization approaches:** Over our tests there was no clear winner between simply doing a naive k -NN approach and more advanced graph topology inference techniques. Kalofolias had very good performance on the Label Propagation and Denoising tasks, while NNK was consistent in SGC and Spectral Clustering, but both were not able to consistently beat the naive baseline. On the other hand, there was a clear advantage of both Kalofolias and NNK over the naive baselines when we consider the robustness of both methods to the parameter k selection.

3.4 Graph filters

In the previous sections we touched on graph filters and their applications, without properly defining them. We do so now in this section. Graph filters are modeled using

the same abstraction as traditional signal processing filters and in this work we detail and derive three possible representations of these filters:

1. directly defined in the spectral domain by a diagonal matrix;
2. defined as a function of the filter spectral response;
3. as diffusion operator based on the graph adjacency or Laplacian matrix.

We describe in the following paragraphs the three different types of definitions, the advantages/drawbacks of using each representation and describe some of their applications. We refer the reader to [47, 162] for a more in depth discussion on graph filters.

3.4.1 Defining filters in the spectral domain

The simplest and most general way to define a filter is simply to describe its response to each frequency. In the case of graph signals the frequencies are defined by the diagonal values of the eigenvalue matrix $\mathbf{\Lambda}$. For notation simplicity we consider the frequency vector $\mathbf{\lambda}$ where $\lambda_i = \Lambda_{i,i}$. Recall that we consider that the eigenvalues are ordered in crescent order of magnitude.

We thus define a filter on a graph \mathcal{G} using a diagonal matrix $\mathbf{H}_{\mathcal{G}} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$, where we call each element $H_{i,i}$ the response of the filter to the frequency λ_i . The filter can then be applied to a graph signal \mathbf{s} by first converting the signal to the frequency domain using the GFT:

$$\tilde{\mathbf{s}} = \mathbf{F}^{\top} \mathbf{s} . \quad (3.9)$$

Now that the signal is on the frequency domain, we can apply the filter and obtain the filtered signal $\tilde{\mathbf{s}}$ by a simple matrix multiplication:

$$\tilde{\mathbf{s}} = \mathbf{H}_{\mathcal{G}} \tilde{\mathbf{s}} . \quad (3.10)$$

Finally, we can convert the signal to the vertex domain using the inverse GFT:

$$\mathbf{s} = \mathbf{F} \tilde{\mathbf{s}} . \quad (3.11)$$

Defining filters directly in the spectral domain generates filters that are graph-specific. Indeed, as the frequencies λ_i are discretized, the same filter \mathbf{H} will be applied to very different frequencies for two distinct graphs \mathcal{G} and \mathcal{G}' . Therefore this type of filter tends to be mostly used to remove the lowest or highest frequencies of the graph, without

considering their “true” value. Another drawback is that it is unlikely that this type of filter may be represented with a low order polynomial filter which impacts the complexity (i.e., the possibility of scaling to larger graphs) of applying the filter.

3.4.2 Defining filters using their spectral response

Another more straightforward way to define a filter is by its spectral response, i.e., as a function of the frequency. In this way the filter becomes less graph dependent and more general. One such design is the Simoncelli filter, that we depict in Figure 3.13 and that is defined by the following function:

$$h(\lambda_i) = \begin{cases} 1 & \text{if } \lambda_i \leq \frac{\tau}{2} \\ \cos\left(\frac{\pi}{2} \frac{\log\left(\frac{\lambda_i}{\tau}\right)}{\log(2)}\right) & \text{if } \frac{\tau}{2} < \lambda_i \leq \tau, \\ 0 & \text{if } \lambda_i > \tau \end{cases} \quad (3.12)$$

where $\tau \in [0, 1]$ is a user-defined threshold and λ_i the i -th Laplacian eigenvalue. We consider that the eigenvalues are always normalized by dividing by the largest one, so that $0 \leq \lambda_i \leq 1$. Defining a filter by its spectral response allow for more universal filters (i.e., filters that do not heavily depend on the graph support) and also to more easily represent the filter by a low order polynomial function. Indeed in this thesis we use the PyGSP [25] toolbox to implement this type of graph filters which uses the Chebyshev polynomial approximation in order to apply the filters.

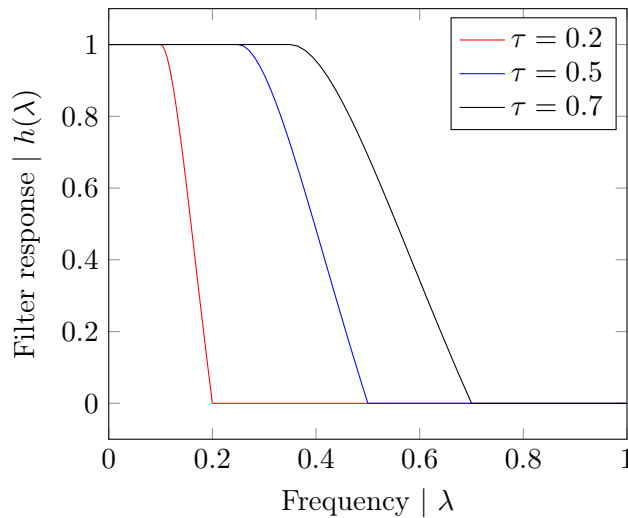


Figure 3.13: Depiction of the Simoncelli filter spectral response for various values of τ .

3.4.3 Defining filters using diffusion operators

It is also possible to define a graph filter using a diffusion operator $\mathbf{S} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$, sometimes also called graph shift operator. This operator is directly applied to a graph signal as follows:

$$\mathbf{s} = \mathbf{S}\mathbf{s} . \quad (3.13)$$

Note that by being applied with just a simple matrix multiplication, this type of filter is easily integrated in deep learning scenarios, where matrix multiplication is king. Indeed, most of the recent developments in graph convolutional layers use this design. We delve into this with more details in Section 4.3.

We have previously introduced the concepts of transfer learning and using DNNs as feature extractors (c.f. Section 2.1.2). In the following subsections we consider the features extracted with the DNNs as graph signals where the graph support is not explicitly provided (i.e., we need to infer the graph structure). We then apply low-pass graph filters to reduce the amount of noise in the extracted features and improve the performance in a downstream task.

3.4.4 Reducing the noise of DNN extracted features for supervised learning and few-shot learning

In this section we analyze the extracted features in two scenarios:

1. *Few-shot classification*: in this task a network has been trained on classes that do not belong to the problem (base) and we want to perform transfer learning to classify new (novel) classes;
2. *Classification*: in this case both base and novel are the same, *i.e* we aim to improve the performance of the original network.

In the case of few-shot classification we evaluate our method three commonly used datasets for this problem:

1. MiniImagenet [131]: A subset of Imagenet, specifically designed for few-shot classification;
2. CIFAR-FS[7]: A reorganization of the CIFAR-100 dataset in order to transform into a few-shot problem (FS = Few-Shot);

3. Caltech-UCSD Birds 200 (CUB) [171]: A dataset of bird classification.

We also add a test called CUB-Cross, where the base classes come from the MiniImageNet dataset, while the novel classes come from the CUB dataset. Finally, we evaluate the classification problem using the previously described CIFAR-10 dataset.

We consider that the features that are extracted from the DNN are signals on graphs, where each node is an example of $\mathcal{D}_{\text{train}}$. The classes are separated in different graphs (1 per class). Our intuition is that by removing the high frequencies from the graph of each class, we can reduce the intra-class noise. With less intra-class noise, we believe that we will be able to increase the classification performance when using these features. In the following subsection we detail our methodology.

3.4.4.1 Methodology

Recall that we can split a deep neural network f into a feature extractor (\mathcal{F}) followed by a classifier (\mathcal{C}), such that $f(\mathbf{x}) = \mathcal{C}(\mathcal{F}(\mathbf{x}))$. We call $\hat{\mathbf{x}}$ the feature tensor extracted using \mathcal{F} . We consider that the input tensor \mathbf{x} contains all elements of the training set $\mathcal{D}_{\text{train}}$. It can be divided in non-intersecting subsets \mathbf{x}_c which contain all elements of class c . Therefore, $\hat{\mathbf{x}}_c = \mathcal{F}(\mathbf{x}_c)$.

Using the cosine similarity we can define the adjacency matrix for the graph of each class c as \mathcal{A}_c . We also call L_c the normalized Laplacian matrix for each graph. This allows us to define low-pass graph class filters. In this section we test two of such filters, one defined directly on the spectral domain and one by its spectral response.

3.4.4.1.1 Filter on the spectral domain: We define our low pass filters on the spectral domain using a diagonal matrix $\mathbf{H}_{\mathcal{G}_c}$ defined as:

$$\mathbf{H}_{\mathcal{G}_c}[i, i] = \begin{cases} 1.0 & \text{if } i \leq F_1 \\ 0.2 & \text{if } F_1 < i \leq F_2 \\ 0 & \text{otherwise} \end{cases} . \quad (3.14)$$

Where (F_1, F_2) are fixed for each scenario. In the few shot scenario we have graphs with 5 nodes and fixed the values to $(1, 3)$. On the other hand, on the classification scenario graphs have 5000 nodes, so we have fixed our values to $(20, 55)$. These filters were designed by hand and a possible extension of this work would be to have an automatic way of choosing these filters or even integrating them as parameters during the learning phase.

3.4.4.1.2 Filter defined by its spectral response: We use the previously defined Simoncelli filter, choosing an $\tau \in [0, 1]$ that is inline with the task at hand. We use $\tau = 0.025$ for the classification scenario (CIFAR-10 dataset) and for the few-shot scenario we use a different value of τ per dataset: i) MiniImagenet, $\tau = 0.3$, ii) CUB, $\tau = 0.3$, iii) CIFAR-FS, $\tau = 0.3$, and iv) CUB-Cross, $\tau = 0.35$.

In the following subsection we use empirical experiments to verify our intuition and show that denoising the representations with low-pass graph filters improves the classification accuracy on both scenarios, allowing us to even beat the state of the art in a competitive image classification benchmark.

3.4.4.2 Empirical experiments

We evaluate the proposed method using two scenarios, Few shot classification (transfer learning) and image classification (improving the original network). The code for all experiments is available at: https://github.com/cadurosar/graph_filter.

3.4.4.2.1 Few shot classification - Transfer learning: In this first case we follow the framework from [106] called few shot with backbone network. In this framework we first pre-train a DNN (backbone) in a self-supervised way on a bigger dataset of base classes. We then use the backbone to perform transfer learning on the novel classes. Recall that the base classes and novel classes do not intersect. In this scenario, in each iteration our training (support) set is composed of 5 examples (shots) for each of 5 classes (ways) that are taken at random from novel classes. The test (query) set is composed of 595 images of each class for the MiniImagenet and CIFAR-FS datasets and 15 images for CUB and CUB-Cross. We perform 100,000 iterations and report the mean accuracy and 95% confidence intervals for each test. We use the pretrained networks that the authors published on https://github.com/nupurkmr9/S2M2_fewshot. We test our method on 3 datasets (MiniImageNet, CUB, CIFAR-FS) for the in-domain Transfer Learning (TL) scenario and on one dataset for a cross domain TL. In the former the pre-trained network base classes come from the same dataset as do the novel classes, while on the latter the base classes come from MiniImagenet and the novel from CUB.

For each iteration we first extract the features $\hat{\mathbf{x}}$ using the pre-trained feature extractor. We then infer 5 graphs of 5 nodes (as we have 5 classes and 5 examples per class), and apply the graph filter. This generates our filtered features \mathbf{s} , that are classified with a simple 1-NN classifier. This is a stress test of the method as it has to be robust to a multitude of different graphs and different features. We compare our method against an

1-NN classifier on $\hat{\mathbf{x}}$, a Nearest Mean Classifier (NMC) on $\hat{\mathbf{x}}$, a Logistic Regression⁵ (LR) trained on $\hat{\mathbf{x}}$ and to the original results from [106]. Note that none of these approaches change the test features. Results are described in Table 3.5. We were able to improve the performance using our filter and an 1-NN classifier in almost all scenarios. We also note that NCM obtained results that are better than the results from the original paper in some scenarios.

Table 3.5: Test error comparison on the few-shot learning task.

Method		In-Domain TL			Cross domain TL
\mathcal{C}	Data	MiniImageNet	CUB	CIFAR-FS	CUB-Cross
Original Paper	$\hat{\mathbf{x}}$	16.82 ± 0.11	9.15 ± 0.44	12.53 ± 0.13	29.56 ± 0.75
1-NN	$\hat{\mathbf{x}}$	21.92 ± 0.04	11.06 ± 0.04	15.86 ± 0.04	36.49 ± 0.06
NCM	$\hat{\mathbf{x}}$	16.73 ± 0.03	8.94 ± 0.03	12.58 ± 0.04	30.00 ± 0.05
LR	$\hat{\mathbf{x}}$	16.49 ± 0.03	8.92 ± 0.03	12.62 ± 0.04	29.17 ± 0.05
1-NN	Spectral Filter \mathfrak{s}	16.53 ± 0.03	8.86 ± 0.03	12.51 ± 0.04	29.13 ± 0.05
1-NN	Simoncelli \mathfrak{s}	16.55 ± 0.03	8.92 ± 0.03	12.53 ± 0.04	29.00 ± 0.05
LR	concatenate($\mathfrak{s}, \hat{\mathbf{x}}$)	16.29 ± 0.03	8.84 ± 0.03	12.5 ± 0.04	28.74 ± 0.05

3.4.4.2.2 Image classification - Improving the original network results: On this second case we use the well known CIFAR-10 dataset and three pre-trained architectures, WideResNet 26-10 [180], ShakeNet [35] and PyramidNet [48]. The first model is trained with traditional data augmentation techniques (namely random crop and horizontal flip) while the latter two⁶ use a stronger learned policy called fast-autoaugment [99] on top of traditional data augmentation. We extract the features $\hat{\mathbf{x}}$, create k -nearest neighbor graphs for each class ($k = 10$), apply the graph filter on each graph and generate our filtered features \mathfrak{s} . We now compare the performance of a 1-NN classifier applied to the filtered features \mathfrak{s} to the same classifiers as in the previous section. The results are described in Table 3.6. The 1-NN classifier on the filtered features was able to improve the performance over both the 1-NN classifier, NCM and even beats the performance of the original network. By using three networks we can show that our filter has to be adapted mostly to the dataset and not exactly to the features that are provided. We also note that we are able to beat the state-of-the art without retraining.

In the following subsection we present a similar method that uses graph filters to improve results in the contexts of Visual-Based Localization (VBL) and Image Retrieval (IR).

⁵using the default parameters of Scikit-Learn

⁶available at: github.com/kakaobrain/fast-autoaugment.

Table 3.6: Test error comparison on the classification task. We note that this task does not have confidence intervals as the objective is to improve the original network that is expensive to train.

Method		DNN architecture		
\mathcal{C}	Data	WideResNet	ShakeNet	PyramidNet
Original Paper	$\hat{\mathbf{x}}$	4.18	2.04	1.44
1-NN	$\hat{\mathbf{x}}$	4.19	2.05	1.46
LR	$\hat{\mathbf{x}}$	4.18	2.02	1.46
NMC	$\hat{\mathbf{x}}$	4.19	2.03	1.48
1-NN	Spectral Filter \mathfrak{s}	4.09	2.03	1.39
1-NN	Simoncelli \mathfrak{s}	4.12	2.02	1.37

3.4.5 Improving VBL and IR using graph filter

We have previously described the tasks of Visual-Based Localization (VBL) and Image Retrieval (IR) in Section 2.3.3 and Section 2.3.2 respectively. In this section we introduce our contribution [86] that aims at combining DNN feature extractors and low-pass graph filters to improve performance on these downstream tasks.

Indeed, using Deep Learning (DL) methods for VBL approaches has recently received a lot of attention. DL can be used to directly map images to poses [13, 72] or to generate latent representation (i.e., use the DNNs as feature extractors) that are resilient to appearance changes in the images [2]. The former comes with major drawbacks. For example, such methods are unable to generalize to previously unseen locations. Furthermore, small differences in query poses can cause significant localization errors. Also, appending new locations to the dataset will require retraining the whole network. On the contrary, representation methods generalize well to new data without the need for this retraining.

Therefore in this section we focus on situations where pose information from a visual query is inferred using $\mathcal{D}_{\text{support}}$, in which images are associated with a pose. This can be seen as an image retrieval problem where the aim is to find images in a set that might have been taken from the same location as that of the query image. Once a match or set of matches is found, the pose for the query image is computed as a function of the poses of the retrieved ones.

The method we propose in this Section takes advantage of the additional information

that might be available for each image in the support set, including GPS coordinates, consecutiveness in the acquisition process or similar latent representations. This is particularly interesting for a robotics setting, where images are almost always acquired sequentially from a camera mounted on a vehicle. This sequential nature of the acquisition process suggests that images closer in time should also have close representations. Additional information such as GPS coordinates, if available, can aid in encoding global relationships between images in the database. We show that by considering such relationships between images, localization accuracy can be increased.

Moreover, enhancements can be achieved using only minor adjustments to the inference process. Specifically, we exploit relationships via a graph filter on top of pre-learned deep representations extracted from deep neural networks [2, 129]. In this graph, each vertex is associated one-to-one with an image (i.e., a graph that models the relationship between data samples). Edges model relations between images and are derived from the additional source of information (e.g. temporal adjacency, GPS, similar latent representations). Interestingly, the proposed method can be seen as a fine-tuning of the representations that does not require additional learning, allowing this operation to be possibly executed on a resource constrained system.

In the following subsections we first introduce how we infer graphs from the extracted features, then we introduce the graph filter that will allow us to improve the performance on the downstream tasks, and finally we derive and discuss experiments.

3.4.5.1 Graph inference

In order to make the graph filter improve the accuracy of VBL, we first need to be sure that the edges of the graph are well chosen to reflect the similarity between two images represented as vertices, as our main goal is to exploit extra information available in the database. In this work, we consider three different sources:

- Metric distance (**dist**): the distance measured by the GPS coordinates between two vertices;
- Sequence (**seq**): the distance in time acquisition between two images (acquired as frames in videos);
- Latent similarity (**latent_sim**): the cosine similarity between latent representations.

The matrix \mathcal{A} can therefore be derived from the three sources as:

$$\mathcal{A} = \mathcal{A}_{\text{dist}} + \mathcal{A}_{\text{seq}} + \mathcal{A}_{\text{latent_sim}}. \quad (3.15)$$

3.4.5.2 Metric distance

In order to transform the metric distance into a similarity, we use an exponential kernel. This is parametrized by a scalar γ that controls the sharpness of the exponential and a threshold parameter $max_{distance}$ that cuts edges between distant vertices:

$$\mathcal{A}_{\text{dist}}[i, j] = \begin{cases} e^{-\gamma dist_{i,j}} & \text{if } dist_{i,j} < max_{distance} \\ 0 & \text{otherwise} \end{cases}. \quad (3.16)$$

3.4.5.3 Sequence

To exploit the information of time acquisition of frames, we use the function $seq(k, \mu, \nu)$ which returns 1 if the frame distance between μ and ν is exactly k and 0 otherwise. We then build a matrix \mathcal{A}_{seq} parametrized by scalars β_k and k_{max} :

$$\mathcal{A}_{\text{seq}}[\mu\nu] = \sum_{k=1}^{k_{max}} \beta_k seq(k, \mu, \nu). \quad (3.17)$$

3.4.5.4 Latent similarity

Finally, we define a matrix $\mathcal{A}_{\text{latent_sim}}$ for the latent representations cosine similarity. This is parametrized by a scalar α that controls the importance of the latent similarity. We only compute this similarity if either the distance similarity or the sequence similarity is nonzero:

$$\mathcal{A}_{\text{latent_sim}}[\mu\nu] = \begin{cases} \alpha sim(\mu, \nu) & \text{if } \mathcal{A}_{\text{dist}}[\mu\nu] > 0 \\ & \text{or } \mathcal{A}_{\text{seq}}[\mu\nu] > 0, \\ 0 & \text{otherwise} \end{cases} \quad (3.18)$$

where sim is the latent similarity function. In this work we use the cosine similarity (sim_{\cos}), but any similarity function could be considered.

3.4.5.5 Graph filter

Given the signal \mathbf{s} and its normalized Laplacian matrix \mathbf{L} , we define our graph low pass filter using a diffusion matrix \mathbf{S} :

$$\mathbf{S} = (\mathcal{I} - a\mathbf{L})^m, \quad (3.19)$$

where $a = 0.1$ and m is an integer that we fix to 20. Note that when $m = 0$ no filtering is performed ($\mathbf{S} = \mathcal{I}$).

3.4.5.6 Experimental results

We first present the results concerning VBL and then we extend our empirical test to the context of IR.

3.4.5.6.1 Visual-based localization: In the context of VBL we can infer our graphs using all three subgraphs (distance, similarity and sequence). We thus have to first search and define the needed parameters. These parameters were obtained using a grid search and keeping the best score on the Adelaide validation query. We then use both the Adelaide test query and the Sidney dataset to ensure that the parameters are not overfitted to the validation query. Note that by using the same parameters for all cities we further validate the fact that our approach does not need to be updated when adding more cities. The parameters we use are $\gamma = 0.1$, $\beta_1 = 0.75$, $\beta_2 = 0.0625$, $\beta_3 = 0.015$, $k_{max} = 3$, $\alpha = 0.66$, $m = 20$.

We test the graph filter in three different cases. First the extra data is available only for the support, second it is available only for the query and finally it is available in both cases. In each case we report two metrics, the median localization error over all the queries and the percentage of localizations that have less than 25m error.

First we perform the tests on the Adelaide dataset and present the results in table 3.7. The graph filter was able to increase performance, even when applied only on the query database, and as expected, adding the graph filter during both query and support gave the best results. Recall that the parameters were defined based on the validation query, under the case where the extra data is available only for the support database.

Second we validate that the operation can be used on other cities and that we do not need to perform an additional grid search for the new data. The results are presented in Table 3.8. As expected the graph filter allowed us to get better performance in both

Table 3.7: Results under different graph filter conditions for the Mapillary Adelaide dataset. GF means Graph Filtering.

Measure	None	GF Support	GF Query	GF S+Q
Validation				
acc < 25m	66.84%	76.23%	69.64%	79.22%
median distance	8.76m	6.90m	13.26m	8.90m
Test				
acc < 25m	44.63%	50.44%	46.25%	52.13%
median distance	110.66m	24.30m	42.03m	22.49m

median distance and accuracy, while using the parameters optimized for the Adelaide dataset. This is inline with our goal that is to have an operation that we do not have to retrain or re-validate parameters for a new dataset. We note that the performance of the hard query set is not inline with a good retrieval system (several kilometers from the correct point), but it is included to show that our method allows us to increase the performance both when the NetVLAD features are already very good for the task and when they are very bad.

Table 3.8: Results under different graph filter conditions for the Mapillary Sydney dataset. GF means Graph Filtering.

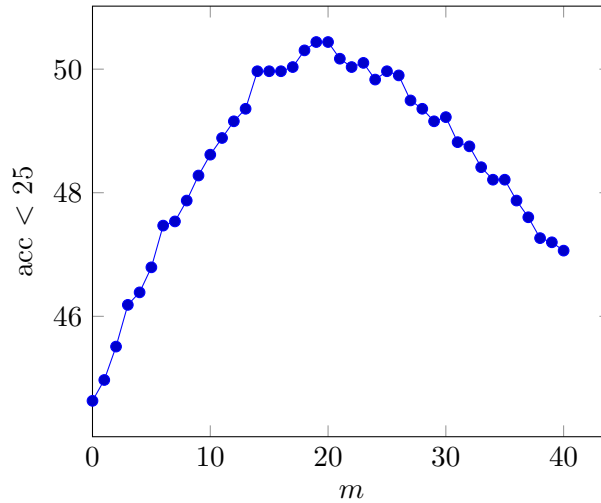
Measure	None	GF Support	GF Query	GF S+Q
Easy				
acc < 25m	49.45%	55.69%	57.21%	64.33%
median distance	28.25m	13.37m	15.89m	11.70m
Hard				
acc < 25m	13.87%	17.29%	17.29%	22.28%
median distance	4000m	3372m	3246m	2226m

Finally we perform ablation tests to ensure that each part of the graph is important, using the Adelaide test set. The results are presented in Table 3.9. The table shows that different sources of information are important, with each one adding to increase in performance. Metric distance and sequence being the most important features and latent similarity being more of a complementary feature (this is expected, as it is being thresholded by the other two features). This is encouraging since in the absence of any other external information (GPS, etc), one can rely on the sequential nature of data collection to get a boost in localization performance. This information is readily available in a robotics setting.

Table 3.9: Ablation study on the Mapillary Adelaide test query. Table extracted from [86]

$\mathcal{A}_{\text{dist}}$	\mathcal{A}_{seq}	$\mathcal{A}_{\text{latent_sim}}$	median distance	acc < 25m
			110.66m	44.63%
X			29.26m	49.42%
	X		39.11m	47.47%
X		X	28.41m	49.56%
X	X		24.35m	50.17%
	X	X	37.34m	47.74%
X	X	X	24.30m	50.44%

We also want to demonstrate the effect of successive filtering. This is achieved by applying the filter m times. Theoretically, this should help increase the performance until it hits a ceiling and then it should start to slowly decrease (as it enforces connected examples of the database to be too similar to each other). The results are presented in Figure. 3.14. As can be seen, there is a clear pattern of increased performance until $m = 20$ after which the performance starts to degrade. It should be noted that even for $m = 40$ the graph filter still performs better than the baseline ($m = 0$).

Figure 3.14: Effect of the parameter m on the retrieval accuracy under 25m for the Adelaide test query. Figure extracted from [86].

3.4.5.6.2 Image retrieval: As a visual localization problem can be seen as an application of Image Retrieval, we test our method in classical Image Retrieval scenarios to verify its genericity. We use the revisited Oxford and revisited Paris datasets [128] with the features from [129]. In this case we do not have the physical distance between the

images to properly create W_{dist} or the image sequence to generate W_{seq} . We therefore use the objects names as classes and our W_{dist} is composed of only 1 (if μ and ν are from the same object) and 0 otherwise. Note that in this way, we differ from traditional methods as they tend to not consider this additional information during training or testing and therefore comparison with other methods is not entirely fair. All the other parameters are the same as in the localization scenario.

In the scenario of Image Retrieval our approach can be categorized as diffusion-based. In the literature there are diffusion methods that are used during the ranking phase with k -NN graphs [66] or that add an additional GCN [74] component that has to be trained in an unsupervised way [101]. In summary, our main contribution is the graph construction (taking advantage of the class data that is available on the support set) and our smoothing/diffusion technique that is based on a low-pass filter.

The results are presented in Table 3.10. Our method was able to increase the mean average precision, with similar results to the approach from [66]. When using in combination with [66] we achieve a similar performance on the Paris dataset to a state of the art approach [101] that requires training an additional GCN network.

Table 3.10: mAP retrieval results comparison, results that do not include our filter are extracted as is from [101]. Table extracted from [86].

Features	Ranking	rOxford		rParis	
		Medium	Hard	Medium	Hard
[129]	sim_{\cos}	64.7	38.5	77.2	56.3
[129]	[66]	69.8	40.5	88.9	78.5
[129] + Our filter [86]	sim_{\cos}	70.58	47.67	87.77	76.04
[129] + Our filter [86]	[66]	71.41	51.27	91.54	81.85
[129] + [101]	[101]	77.8	57.5	92.4	83.5

In the previous paragraphs we showed that using techniques from Graph Signal Processing, the performance of visual based localization and image retrieval can be improved by incorporating additional available information. This additional information acts on the latent representation by making it smoother on a graph designed using all available information, leading to a boost in localization. One encouraging observation is that this additional information can take the form of a simple temporal relationship between surrounding images acquired in a sequence, and still lead to a significant increase in performance.

3.5 Summary of the chapter

In this chapter we have introduced the concepts of graphs and graph signals, alongside with the needed tools from the Graph Signal Processing (GSP) framework. These concepts and tools allow us to perform analysis of deep latent representations and to derive new contributions to the machine learning community that are going to be introduced in the following Chapters.

Some of the tools introduced in this section include the Graph Fourier Transform (GFT) and the analysis of the smoothness of graph signals. We also discuss methods of inferring graphs from data where the graph support structure is not explicitly available, including a novel contribution:

- Lassance, C., Gripon, V., and Mateos, G. (2020b). Graph topology inference benchmarks for machine learning. *arXiv preprint arXiv:2007.08216, to appear in 2020 IEEE 30th International Workshop on Machine Learning for Signal Processing (MLSP)*

Then we derived graph filters, that emulate traditional signal processing filters in the graph domain. These graph filters will be used to link convolutional layers and graph convolutional layers in the next chapter. We also present two applications of graph filters that allow us to reduce the amount of noise of features extracted using DNNs and improve the performance of downstream tasks including i) few-shot learning, ii) image classification, iii) visual-based localization (VBL), and iv) image retrieval (IR). The visual-based localization application was a subject of an archival contribution:

- Lassance, C., Latif, Y., Garg, R., Gripon, V., and Reid, I. (2019b). Improved visual localization via graph smoothing. *arXiv preprint arXiv:1911.02961*

In the following chapter we discuss DNNs that have their latent spaces supported on graphs. To do this we will introduce and discuss the field of Graph Neural Networks (GNNs). First we introduce the needed definitions in Section 4.1 using the concepts introduced in the current chapter and the previous one. These definitions will create a link between convolutional layers and graph convolutional layers that allow us to use an universal framework to represent both types of layers.

We then discuss in Section 4.2 the use of GNNs in the context of supervised classification of graph signals, first doing a sanity check using data that is defined in a regular 2D

Euclidean space (e.g., images) and then using data that is defined in a slightly irregular 3D Euclidean space (e.g., neuroimaging) as a slightly more difficult test. Finally, we will then discuss applications of GNNs in a semi-supervised classification scenario and their pitfalls in Section 4.3.

Chapter 4

Deep Learning for inputs supported on graphs

4.1	Definitions	104
4.2	Supervised classification of graph signals	113
4.3	Semi supervised classification of vertices	121
4.4	Summary of the chapter	128

In the previous chapters, we have introduced and discussed Deep Neural Networks and Graph Signal Processing. In this chapter, we build upon these methods and discuss DNNs that have their inputs supported on graphs. We first introduce the needed definitions in Section 4.1, introducing the graph convolutional layers as graph filters. We also create a link between convolutional layers and graph convolutional layers.

We then discuss in Section 4.2 the use of GNNs in the context of supervised classification of graph signals. We do a quick review of the domain and then present a method to perform the supervised classification of graph signals. The method is evaluated first by a sanity check using data defined in a regular 2D Euclidean space (e.g., images). We then evaluate on a more real-world scenario by using data defined in a slightly irregular 3D Euclidean space (e.g., neuroimaging) as a slightly more difficult test. Finally, we will discuss applications of GNNs in a semi-supervised classification scenario and their pitfalls in Section 4.3.

4.1 Definitions

In this section, we will introduce some of the recent literature in Graph Neural Networks. Presenting all the methods in the literature is very complicated, given the speed of evolution of the field and the amount of already proposed methods [5, 24, 41, 69, 74, 75, 91, 98, 118, 137, 164, 173]. In this thesis, we prefer to present this domain as six methodologies in a logical sequence of developments. Note that even if this can be seen as a logical sequence of developments, where at each time complexity increases, the methods themselves are not in chronological order.

We first present in Section 4.1.1 the use of graph filter as feature extractors followed by a simple classifier. Note that this should not be considered a neural network, but it was introduced as so in [173] that is called “Simplifying Graph Convolutional Networks”. We then present a methodology that use this simple diffusion/filter on a graph inside each layer, leading to *Graph Convolution Layers (GCL)*, and to networks that are called Graph Convolutional Networks (GCN) [74] in Section 4.1.2.

We then present two extensions to the GCL/GCN, first in Section 4.1.3 we showcase methods that improve the GCL by applying multiple graph filters at the same layer in order to increase the degrees of freedom the layer, and second in Section 4.1.4 we show how we can improve the GCNs by modifying the graph during the training, either via learning the weights or by adding additional information.

Finally we present two additional extensions that are orthogonal to the first ones. In section 4.1.5 we show methods that can learn the graph filters directly instead of using predefined filters, and in Section 4.1.6 we showcase methods that try to mimic traditional convolutional layers using the concept of graph translation introduced in Section 3.1.1.

Note that there are multiple ways to represent the above-mentioned methodologies [11, 37, 167, 174, 186, 187]. In this manuscript, we follow a different path to be more inline with the rest of the document. The reader should be informed that similar analysis were proposed in the first months of this year [5, 69].

4.1.1 Using graph filters as feature extractors

As we present the considered methods in ascending order of complexity, we first recall the previously introduced concept of graph filters (c.f. Section 3.4) and then introduce a very simple methodology that consists in using a graph filter as feature extractor. The most relevant part is that the graph filter does not intervene in the training phase of

the classifier and can be seen as a fast pre-processing method. This methodology was popularized by SGC [173] and has been shown to be very efficient not only in the context of inputs supported on graphs but also in contexts where the graph support has to be inferred such as visual based localization (c.f. Section 3.4.5 and few-shot learning [58]).

In the case of SGC, the graph filter that is used is based on the adjacency matrix \mathcal{A} . The first step is to add self-connections to \mathcal{A} , generating the augmented adjacency matrix $\tilde{\mathcal{A}}$:

$$\tilde{\mathcal{A}} = \mathcal{I} + \mathcal{A} . \quad (4.1)$$

then the adjacency matrix is normalized in order to create the diffusion matrix \mathbf{S} as follows:

$$\mathbf{S} = \tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathcal{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}} = \tilde{\mathcal{A}} , \quad (4.2)$$

where $\tilde{\mathbf{D}}$ is the degree matrix of the augmented adjacency matrix. Note that it is common in the literature to represent \mathbf{S} as the $\tilde{\mathcal{A}}$ itself, but here we prefer to use two different symbols. We can then apply the graph filter to the graph signal \mathbf{x} and generate the filtered signal \mathbf{s} as follows:

$$\mathbf{s} = \mathbf{S}^m \mathbf{x} , \quad (4.3)$$

where m represents the amount of “SGC layers” that are applied to the signal. The filtered signal is then used to train a simple logistic regression classifier or even a k -neighbors classifier. In [173] the authors have shown that for various applications the results are as good as multi-layered GCNs, in a fraction of the time (as the graph filter is only applied as a pre-processing). We will delve in more details on this in the following sections.

One advantage of considering this a filtered signal and a graph filter is that we can perform spectral analysis to better understand the underlying functioning of our feature extractor. Indeed we can reformulate equation 4.3 in order to represent our filter using the $h(\lambda)$ function and the graph Fourier transform. First we have to define the symmetric Laplacian of the diffusion matrix as:

$$\tilde{\mathbf{L}} = \mathcal{I} - \mathbf{S} . \quad (4.4)$$

We can then redefine the filtered signal as:

$$\mathbf{s} = \mathbf{S}^m \mathbf{x} \quad (4.5)$$

$$= (\mathcal{I} - \tilde{\mathbf{L}})^m \mathbf{x} \quad (4.6)$$

$$= (\mathcal{I} - \mathbf{F} \tilde{\mathbf{\Lambda}} \mathbf{F}^\top)^m \mathbf{x} \quad (4.7)$$

$$= \mathbf{F} (\mathcal{I} - \tilde{\mathbf{\Lambda}})^m \mathbf{F}^\top \mathbf{x} \quad (4.8)$$

$$= \mathbf{F} \mathbf{H} \mathbf{F}^\top \mathbf{x} , \quad (4.9)$$

where $\mathbf{H} = (\mathbf{I} - \tilde{\mathbf{\Lambda}})^m$ is a diagonal matrix that determines the filter response for each discrete eigenvalue of $\tilde{\mathbf{L}}$. We can then represent this filter by its spectral response with a function $h(\tilde{\lambda})$ as follows:

$$h_{\text{SGC}}(\tilde{\lambda}) = (1 - \tilde{\lambda})^m. \quad (4.10)$$

We represent the spectral response in Figure 4.1 for various values of m . Note that while this should implement a low pass filter, it actually implements a band reject filter, where in the odd values the high frequencies are inverted. This behavior while unexpected is not necessarily bad, as recent theory in deep learning (both applied to computer vision and GSP) seem to converge to robustness is linked to the low frequencies of the signal, while generalization and overfitting are linked to the higher frequencies [34, 63, 170, 177]. We will discuss this difference between the expected low pass and the obtained filter in more details in the application part of this chapter.

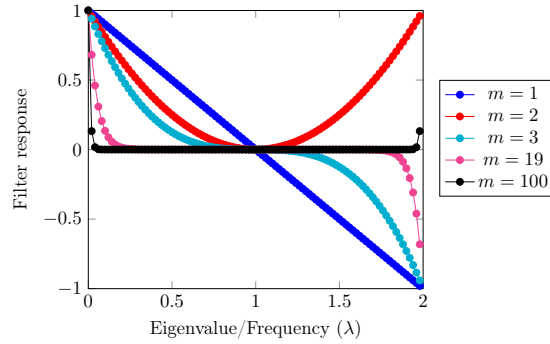


Figure 4.1: Response of the filter used to extract features in SGC [173]. Note that the eigenvalues here are the ones from the Laplacian of the augmented graph

Note that this filter is based on the eigenvalues of the Laplacian of the augmented graph $\tilde{\mathcal{A}}$ and not on the ones of the original graph \mathcal{A} .

Let us now decompose the filter generated with the augmented graph in two filters f_1 and f_2 :

$$\mathbf{s} = \mathbf{S}\mathbf{x} \quad (4.11)$$

$$= \frac{\mathbf{x}}{\mathbf{D} + \mathcal{I}} + (\mathbf{D} + \mathcal{I})^{-\frac{1}{2}} \mathcal{A} (\mathbf{D} + \mathcal{I})^{-\frac{1}{2}} \mathbf{x}; \quad (4.12)$$

$$f_1(\mathbf{x}) = \frac{\mathbf{x}}{\mathbf{D} + \mathcal{I}}; \quad (4.13)$$

$$f_2(\mathbf{x}) = (\mathbf{D} + \mathcal{I})^{-\frac{1}{2}} \mathcal{A} (\mathbf{D} + \mathcal{I})^{-\frac{1}{2}} \mathbf{x}; \quad (4.14)$$

$$\mathbf{s} = f_1(\mathbf{x}) + f_2(\mathbf{x}). \quad (4.15)$$

By decomposing into f_1 and f_2 we can try to analyze them separately, indeed we can see that f_1 attenuates the signal of each sample \mathbf{x} by dividing it by its degree augmented

by one. In the case of f_2 , each node is affected differently depending on its degree and its neighbors. This means that the frequency profile of the filter will be perturbed, unless we have a regular graph (i.e. all nodes have the same degree).

Indeed if the graph represented by matrix \mathcal{A} is regular where each node has degree d , we can rewrite of the augmented graph $\tilde{\mathbf{L}}$ as a function of the Laplacian of the original graph:

$$\mathbf{L} = \mathcal{I} - \mathbf{D}^{-\frac{1}{2}} \mathcal{A} \mathbf{D}^{-\frac{1}{2}} \quad (4.16)$$

$$= \mathcal{I} - \frac{\mathcal{A}}{d} \Rightarrow \mathcal{A} = d\mathcal{I} + d\mathbf{L} ; \quad (4.17)$$

$$\tilde{\mathbf{L}} = \mathcal{I} - (\mathbf{D} + \mathcal{I})^{-\frac{1}{2}} (\mathcal{A} + \mathcal{I}) (\mathbf{D} + \mathcal{I})^{-\frac{1}{2}} \quad (4.18)$$

$$= \mathcal{I} - \frac{\mathcal{A} + \mathcal{I}}{d + 1} \quad (4.19)$$

$$= \frac{d\mathcal{I} - \mathcal{A}}{d + 1} \quad (4.20)$$

$$= \frac{d\mathcal{I} - (d\mathcal{I} - d\mathbf{L})}{d + 1} \quad (4.21)$$

$$= \frac{d\mathbf{L}}{d + 1} , \quad (4.22)$$

$$(4.23)$$

and by writing it as a function of the original graph, we can also write the eigenvalues of the augmented graph as a function of the eigenvalues of the original one:

$$\tilde{\Lambda} = \frac{d}{d + 1} \Lambda , \quad (4.24)$$

finally, we can now write the function $h(\lambda)$ as a function of the eigenvalues of the original graph as follows:

$$h(\lambda) = \left(1 - \frac{d\lambda}{d + 1}\right)^m , \quad (4.25)$$

which still has the same problem of not being exactly a low pass filter. As the maximum eigenvalue of \mathbf{L} is $\lambda_{\max} = 2$, this filter would only be low pass if $2d \leq d + 1$ which is only possible for $d = 0$ and $d = 1$, which would yield graphs that are not fully connected and therefore are not interesting for a GCN application.

4.1.1.1 Other filters proposed in the literature

In the previous paragraphs we have described how one can use a graph filter as a preprocessing step for data, and have developed an analysis using the graph filter defined in [74, 173]. In this section, we will present some other possible choices of filters that were proposed in the literature under a common framework (m is the attenuation factor and

α controls the smoothness of the filter). First, we recall the filter used in Section 3.4.5, which is a proper low pass filter if the parameter α is smaller than 0.5 and that is defined as:

$$\mathbf{S} = (\mathcal{I} - \alpha \mathbf{L})^m \Rightarrow h_{\text{VBL}}(\lambda) = (1 - \alpha\lambda)^m \quad (4.26)$$

Another more frequently used low pass filter is based on the Tikhonov regularization [148] which aims at balancing the amount of change of the filter and the smoothness of the final filtered signal:

$$\underset{h}{\operatorname{argmin}} \|\mathbf{s} - \mathbf{x}\| + \alpha \mathbf{s}^\top \mathbf{L} \mathbf{s} \Rightarrow h_{\text{Tikhonov}}(\lambda) = \frac{1}{1 + \alpha\lambda}, \quad (4.27)$$

In [5] the authors propose to combine four different filters (h_1 to h_4), where h_1 is a low-pass filter, h_2 to h_4 are different band pass filters as follows:

$$h_{\text{bancilar-lowpass}}(\lambda) = \left(\frac{\lambda_{\max} - \lambda}{\lambda_{\max}} \right)^m \quad (4.28)$$

$$h_{\text{bancilar-2}}(\lambda) = \exp(-\alpha(0.25\lambda_{\max} - \lambda)^2) \quad (4.29)$$

$$h_{\text{bancilar-3}}(\lambda) = \exp(-\alpha(0.5\lambda_{\max} - \lambda)^2) \quad (4.30)$$

$$h_{\text{bancilar-4}}(\lambda) = \exp(-\alpha(0.75\lambda_{\max} - \lambda)^2) \quad (4.31)$$

$$(4.32)$$

Finally, there also exists a filter [75] that is based on the personalized PageRank algorithm [115] and the normalized augmented adjacency matrix $\tilde{\mathcal{A}}$ from Equation 4.2:

$$\mathbf{S} = \alpha(\mathcal{I} - (1 - \alpha)\tilde{\mathcal{A}})^{-1} \Rightarrow h_{\text{page}}(\tilde{\lambda}) = \frac{\alpha}{\alpha(1 - \tilde{\lambda}) + \tilde{\lambda}}; \quad (4.33)$$

Also note that the authors originally proposed to apply the filter after the affine transformation of the logistic regression or after the MLP that reduces the number of features to the number of classes. It thus differs substantially from [173]. We will further discuss this in Section 4.3 in order to explicit all the different decisions that are made for each model. We also discuss how to fairly compare them.

Remark 1. While for analysis sake it is better to consider the filters in the spectral form, applying the filters as a diffusion matrix is much more computationally efficient: $O(|\mathbb{V}|^3)$ vs $O(|\mathbb{E}|)$. Therefore it is quite recommended to convert the spectral filter into matrix form as a pre-processing step when possible.

4.1.2 Graph Convolutional Layers (GCL)

Now that we have introduced how to use graph filters to extract relevant features from graph signals, we can extend this to a layer in a neural network. Indeed in the original

GCN paper [74] each layer (that we call GCL in this work) is defined as:

$$f^\ell(\mathbf{x}) = g(\mathbf{S}^m \mathbf{x} \mathbf{W}^\ell + \mathbf{b}^\ell), \quad (4.34)$$

where f^ℓ is the layer function of layer ℓ , \mathbf{W}^ℓ is the weight matrix of the affine transformation of layer ℓ , \mathbf{b}^ℓ is the bias associated with layer ℓ , \mathbf{x} is the input of the layer, g is the nonlinear activation function and \mathbf{S} is the diffusion matrix presented in the previous section based on the augmented adjacency matrix. Note that a GCN model can have more than one layer and that if we had removed the nonlinear activations we would have the same model as the one presented in the previous section as all the weights and bias would be compressed in a single matrix and a single vector that would be responsible for the logistic regression.

In other words, SGC is a specialization of GCNs where there is no concept of network or layers as all the layers are linear. On the other hand, by introducing the nonlinearity we can only approximate the true graph filter, specially in cases where the nonlinear function aggregates over the nodes of the graph as in [137]. In this case we define the filter using \mathbf{S}^m and apply the same filter for each layer.

Note that even if in theory GCNs are able to represent more complex functions than SGC, it is not given which method will perform best. Indeed in [173] the authors show that SGC may present equivalent or even better results than GCN in some tasks. In a recent paper [168], the author perform a more careful evaluation between GCNs and SGC and show that choosing which method to use depends heavily on a simple metric: the higher is the number of samples N per feature F the better GCN performs in comparison with SGC. In other words, when you have $N \gg F$ linear models such as SGC perform very well, on the other hand when $F \gg N$ more complex models such as GCN and its improvements are better.

4.1.3 How to combine multiple filters in each GCL

Given the limitations of the GCN model that implements only one filter per layer and that we are not able to compose it to determine the global filter of the entire network given the nonlinearities, authors have proposed alternative ways to compose multiple filters in each GCL of the GCN. For example, one simple solution introduced in TAGCN [27] is to perform a sum over different filters in each layer as follows:

$$f^\ell \mathbf{x} = \sum_{m \in \mathbb{M}} g(\mathbf{S}^m \mathbf{x} \mathbf{W}^{\ell, m} + \mathbf{b}^\ell), \quad (4.35)$$

where \mathbb{M} is the set of powers of \mathbf{S} that we want to use as filter and there is a different affine transformation for each filter. An extension to this framework is presented in [98]

where the authors use the Lanczos polynomial to accelerate the computation of \mathbf{S}^m in the case where m is very large.

Note that, as before, GCN may be seen as a specialization of this model that only uses one filter per GCL and that, as before, even if this model is more expressive it is not given that it will be better performing as seen in [167, 168].

4.1.4 Modifying the support graph to improve GCNs

Another possibility of improvement to GCNs is to change the graph itself instead of the graph filter. In this case one can even use a different graph for each layer of the architecture. One of the most relevant papers that uses this methodology is GAT [164] where the authors propose to use multiple attention heads [163] in order to learn multiple representations of the graph support. Note that even if the graph support is modified, they do not create new links, they only change the weights of the links that are already present. Their approach may be summarized as follows:

1. For each attention head k , we compute a new diffusion matrix \mathbf{S} as \mathbf{S}_k :
 - (a) each weight of \mathbf{S}_k is equal either to $\alpha_{k,i,j}$ or to 0 if $S_{i,j} = 0$;
 - (b) $\alpha_{k,i,j}$ is computed with the softmax of $e_{i,j}$ so that each node has a similar contribution;
 - (c) $e_{i,j}$ is determined using a shared attention mechanism a that takes the affine transformation \mathbf{W}_k of the features of each node f_i and f_j and outputs an attention coefficient: $e_{i,j} = g(a(\mathbf{W}_k \mathbf{x}_i^\ell || \mathbf{W}_k \mathbf{x}_j^\ell))$, where g is a non-linear activation function.
2. Then, after computing each diffusion matrix the output $\mathbf{x}^{\ell+1}$ is defined as: $\mathbf{x}^{\ell+1} = \text{aggregate}(\mathbf{S}_k \mathbf{W}_k \mathbf{x}^\ell)$ where the aggregation function is either concatenation or average over each of the k heads.

Note that while the way the graph is modified is complex, this methodology allows to use different graphs while still keeping the same overall graph filter that is applied to different eigenvalues. On the other hand, one could argue that keeping the same filter applied to different eigenvalues is not particularly different from applying a different graph filter to the same graph. A similar framework is developed in [69] in order to demonstrate that GAT layers are actually GCN layers with multiple filters per layer that learn the graph as well as the transformations. Note that while this would be the most expressive model of the bunch we presented, it is also the most difficult to train, as both

the number of computations and the sensibility of each parameter increases compared to other methods.

4.1.5 Learning the graph filter directly

Another possibility to improve the GCN/GCL model is to learn the graph filter $h(\lambda)$ directly during the learning of the parameters of the networks. Such methods are mostly constrained by the problem of computing the GFT. Indeed, obtaining the spectral decomposition can quickly become too expensive as the number of vertices in the graph increases. Therefore most of the proposed methods differ in which approximation they use, in ChebNets [24] the authors use the Chebyshev polynomial to approximate any possible filter [47] without needing to explicitly compute the GFT. In Lanczosnet [98] the authors propose to use the Lanczos approximation of the orthonormal decomposition in order to learn the filter during training. A major drawback of these methods is the extra time that it is needed to generate approximations good enough to learn interesting filters. Indeed, Chebnetns are used as the basis of GCN. As a matter of fact, GCN is a faster and more efficient implementation of Chebnetns that uses only a small amount of Chebyshev kernels, and can therefore approximate the filters we desire to extract.

Finally [69] introduces methods to train edge varying filters, including ones based on ARMA graph filters [68]. Note that all approaches cited in this section may be summarized as approximations of learning the matrix \mathbf{H} that represents the filter in equation 4.9.

4.1.6 Using graph translations to generate graph convolutional networks

In this section we present the graph convolutional layers proposed in [118]. The authors propose a method that also learns the graph filter directly, but it vastly differs from the ones presented in the last subsection as the filters here are based on the graph translations instead of on diffusion matrices. In other words the filters described here are way closer to their CNN counterparts than to a spectral definition. Indeed the filters here are obtained as the sum of subfilters applied for each considered translation, which in the case of a 3 by 3 convolutional layer would be i) center (the original node), ii) (node to the) right, iii) left, iv) up, v) down, vi) inferior right diagonal, vii) inferior left diagonal, viii) superior left diagonal, and ix) superior right diagonal. In this case, we generate k translation functions that when applied to \mathcal{A} , generate a set of allocation matrices \mathbb{T} . Each translation function follows Definition 3.1.2.

In this case the layer function is defined as:

$$f^\ell(\mathbf{x}) = \sum_{\mathcal{T}_k \in \mathbb{A}} g(\mathcal{T}_k \mathbf{x} \mathbf{W}^{\ell,k} + \mathbf{b}^\ell), \quad (4.36)$$

where each element \mathcal{T} in \mathbb{T} has a different weight matrix \mathbf{W} . This is very similar to the CNN filters, where each filter has a component for each translation. It is also very simple to extend this concept to multiple feature maps, where each feature map implements its own translation filter. Also note that as it is defined as a sum of translations it is very close to the concept we developed in Shift Attention Layers (SAL, c.f. Section 2.4.2). Indeed in SAL the translations are defined in the 2D grid and the attention kernel chooses which translation should be used for each filter. It would be therefore simple to extend SAL to layers based on graph translations.

4.1.7 Summary of methods presented

In the previous subsections we have presented the evolution of methods from a simple fixed feature extractor to methods that learn the graph filter or even the graph itself during training. We summarize the methods in Table 4.1.

Table 4.1: Summary of the methods presented in this section. The last two columns refer to the fact that the method is used or not in the two following sections.

Methods	DNN	Multiple filters per layer	Learn support graph	Learn Filter	Use translations
SGC [173]					
GCN [74]	X				
TAGCN [27]	X	X			
Lanczosnet [98]	X	X		X	
GAT [164]	X	X	X		
ChebNet [24]	X	X		X	
DSGCN [5]	X	X		X	
Translation [91, 118]	X	X		X	X

While presenting an extensive comparison of all the presented methods would be desirable it would also be impossible to do so in a fair way given the number of methods and hyperparameters that are to optimize. We therefore rather base our discussion on recent contributions. First in Section 4.2 we mainly discuss three papers [31, 51, 91], where the first two show that in some use-cases of supervised classification of graphs and graph signals no significant performance gains were found when graph-based machine learning techniques are brought, and the third one is a contribution of our own aiming at improving the performance of GNNs in the regular 2D-Euclidean space and in an irregular 3D-Euclidean space without using priors about the data.

Then in Section 4.3 we focus on some recent findings from [10, 145, 167]. These papers show that most of the benchmark evaluation that is performed on semi-supervised scenarios is actually not as robust as once thought and that fair comparison is still not available. We build upon this work and propose a framework to verify which one of the graph filters described in Section 4.1.1 is the best performing.

4.2 Supervised classification of graph signals

Let us discuss some of the applications of the previously presented methods. The first one we consider is supervised classification of graph signals.

There is a plethora of applications where one may perform this task, with the two most notable mentions being protein-protein interaction and applications on the medical domain. Unfortunately, recent papers have shown that deep learning techniques fail to surpass simpler methods in these scenarios [31, 51]. These recent findings corroborate with our analysis in the previous section that shows that the graph filter is not well defined in the spectral domain. More so, in [175] the authors have shown that GCNs lack the expressivity to distinguish some types of graphs.

Given these recent developments, we choose to limit our discussion on the supervised classification of graph signals to two cases that serve as a “sanity check” of the expressivity of graph neural networks: i) the classification of images (defined as a 2D square), using oracle (grid) or inferred graphs instead of the underlying 2D structure, and ii) the classification of neuroimaging 3D images (that are not defined on a cube), using inferred graphs. Note that in the first case we would like for networks to be able to get results that are similar to the convolutions defined on the 2D space when we use the oracle (grid-graph) structure and to not lose much performance when we use the inferred graph (i.e. we do not have any prior about the data structure). On the second case, we would also desire a performance on the inferred graph that is close to the one using the 3D structure.

Note that the idea is not to surpass traditional 2D/3D convolutions, but to ensure that the graph convolution is able to represent the same type of functions as the 2D convolution. To do so, we propose to test two avenues: i) use a method to convert the grid/inferred graph structure into a 2D/3D regular structure (a square/cube defined on \mathbb{Z}), and ii) use graph convolutional networks to try to imitate the same functions defined by the CNNs. For the former strategy we will use the embedding methodology we first proposed in [41]. To the later we will use graph convolutional networks using

the convolutions from [24, 91]. Note that we do not include other popular models such as GCN [74] and SGC [173] as it is quite straightforward to see that the models will not be able to express a similar function as they are only able to perform one specific convolutional kernel, while CNNs are normally composed of hundreds of kernels per layer.

4.2.1 Methodology

In this section we describe the considered methodologies.

In Section 4.2.1.1, we present how we infer graphs from the regular 2D/irregular 3D data, then we explain how we can use these graphs in two different approaches from [41, 91] to solve the task of classifying graph signals.

4.2.1.1 Graph Inference

We have two separate methods for inferring graphs, one for the case of images (regular 2D data) and one for the case of fMRI data (irregular 3D data).

4.2.1.1.1 Inferring graphs from images In the case of images, a natural choice for the graph structure would be to use a grid-graph (Definition 3.2.1). We call this an “oracle inference” since we suppose in the following that we ignore the fact we are dealing with images. This setting is only used to provide a best case scenario. Then, we also perform an empirical inference, where we first convert the pixels of all images in the $\mathcal{D}_{\text{train}}$ to a grey-scale. With the pixels now in gray-scale we compute the covariance between all pixels in the image, using each example in $\mathcal{D}_{\text{train}}$ as the different samples. We now have a 2D square matrix of the similarity between the pixels that we use as our adjacency matrix \mathcal{A} . We then threshold the \mathcal{A} so that only the 4 closest neighbors of each vertex (in this case pixel) are connected, symmetrize the resulting matrix and binarize it so that connected elements are connected with weight 1 and unconnected elements have weight 0. Since we are dealing with natural images, we expect that the empirical inference graph is a noisy version of the oracle.

4.2.1.1.2 Inferring graphs from fMRI data For inferring graphs from fMRI data, we use a simple neighborhood graph where nodes are connected if they have are at most at distance d in the 3D-space. Note that this is possible in this case as data was captured from physical sensors that were then masked on the MNI template and resampled on a 16mm cubic grid. The data resampled on the cubic grid is not regular as it does have a

data point per integer coordinate of the cubic grid, but it would be possible to run a 3D CNN on it.

4.2.1.2 Converting a graph to a regular structure defined on \mathbb{Z}

In this subsection we present a first method to classify graph signals using DNNs. This idea was first introduced in [41] where we propose an embedding from the graph structure to \mathbb{Z}^d . The main goal was to define the GFT on \mathcal{G} as a particularization of the classical FT on \mathbb{Z}^d but it may also be used in the context of machine learning to allow the use of 2D CNNs to process data defined on graphs. Indeed, once vertices of the graph have been projected to \mathbb{Z}^d , we can use regular convolutions (2D-ones in the case where $d = 2$).

The method to embed a graph into \mathbb{Z}^d consists in optimizing a cost. Let us first introduce a weighted graph $\mathcal{G} = \langle \mathbb{V}, \mathbb{E} \rangle$.

Definition 4.2.1 (embedding). We call **embedding** a function $\phi : \mathbb{V} \rightarrow \mathbb{Z}^d$, where $d \in \mathbb{N}^*$.

We are specifically interested in embeddings that preserve distances. Specifically, we define the cost $c_\alpha(\phi)$ of an embedding ϕ as the following quantity:

$$c_\alpha(\phi) \triangleq \sum_{v, v' \in \mathbb{V}} | \alpha \| \phi(v) - \phi(v') \|_1 - d_{\mathcal{G}}(v, v') |, \quad (4.37)$$

where $d_{\mathcal{G}}$ is the shortest path distance in \mathcal{G} . In the remaining of this section, we denote $\delta(v, v') = | \alpha \| \phi(v) - \phi(v') \|_1 - d_{\mathcal{G}}(v, v') |$.

Definition 4.2.2 (optimal embedding). Given a fixed value of α , we call *optimal embedding* an embedding with minimum cost.

The choices in this definition are motivated by 5 main reasons:

1. We consider all pairs of vertices and not only edges. Consider for example a ring graph where each vertex has exactly two neighbors. Then there are plenty of embeddings that would minimize the cost if considering only edges, but only a few that minimize the sum over all pairs of vertices.
2. We use a sum and not a maximum. This is because small perturbations of grid graphs would lead to dramatic changes in embeddings minimizing the cost if using a maximum. Consider for instance a 2D grid graph in which an arbitrary edge is removed.

3. We choose embedding in \mathbb{Z}^d instead of in \mathbb{R}^d , as we want to particularize multidimensional *discrete* Fourier transforms.
4. We use the Manhattan distance, as it is more naturally associated with \mathbb{Z}^d than the Euclidean distance. It also ensures there exists natural embeddings for grid graphs with cost 0.
5. Finally, α is a scaling factor.

Note that the question of finding suitable embeddings for graphs is not novel [76]. But to our knowledge enforcing the embedding to be in \mathbb{Z}^d was a novel contribution. Even though Definitions 4.2.1 and 4.2.2 work for any d we are going to focus on the consistency of these definitions by considering the particular case $d = 2$, i.e., a regular 2D Euclidean space.

Note that in the case of grid graphs (Definition 3.2.1) the embedding should be a perfect square, indeed it follows that:

Definition 4.2.3 (natural embedding). We call *natural embedding* of a grid graph (as defined in Definition 3.2.1) the identity function.

Which leads to the following theorem:

Theorem 1. The natural embedding of a grid graph is its only optimal embedding for $\alpha = 1$, up to rotation, translation and symmetry.

Proof. The proof is straightforward, as the cost of the natural embedding is clearly 0. Reciprocally, a cost of 0 forces any group of vertices $\{(x, y), (x, y'), (x', y), (x', y')\}$ to be projected to a translation, rotation and/or symmetry of the corresponding rectangle in \mathbb{Z}^2 . Then any remaining vertex is uniquely defined from these four ones. \square

In more general settings where the graph we are dealing with is not a grid-graph, we have to solve an optimization problem with the expectation of finding a relevant embedding. Once the embedding is found, we simply consider graph signals as images and process them with regular CNNs.

4.2.1.3 Matching CNNs without priors

Another possibility is to build upon translations on graphs we defined in Section 3.1.1, that we used to define convolutions in Equation 4.36. This is the idea we used in [91] to

show that it is possible to approach the performance of CNNs without priors about the fact we are dealing with regular images. An overview of all the steps of the method is presented in Figure 4.2.

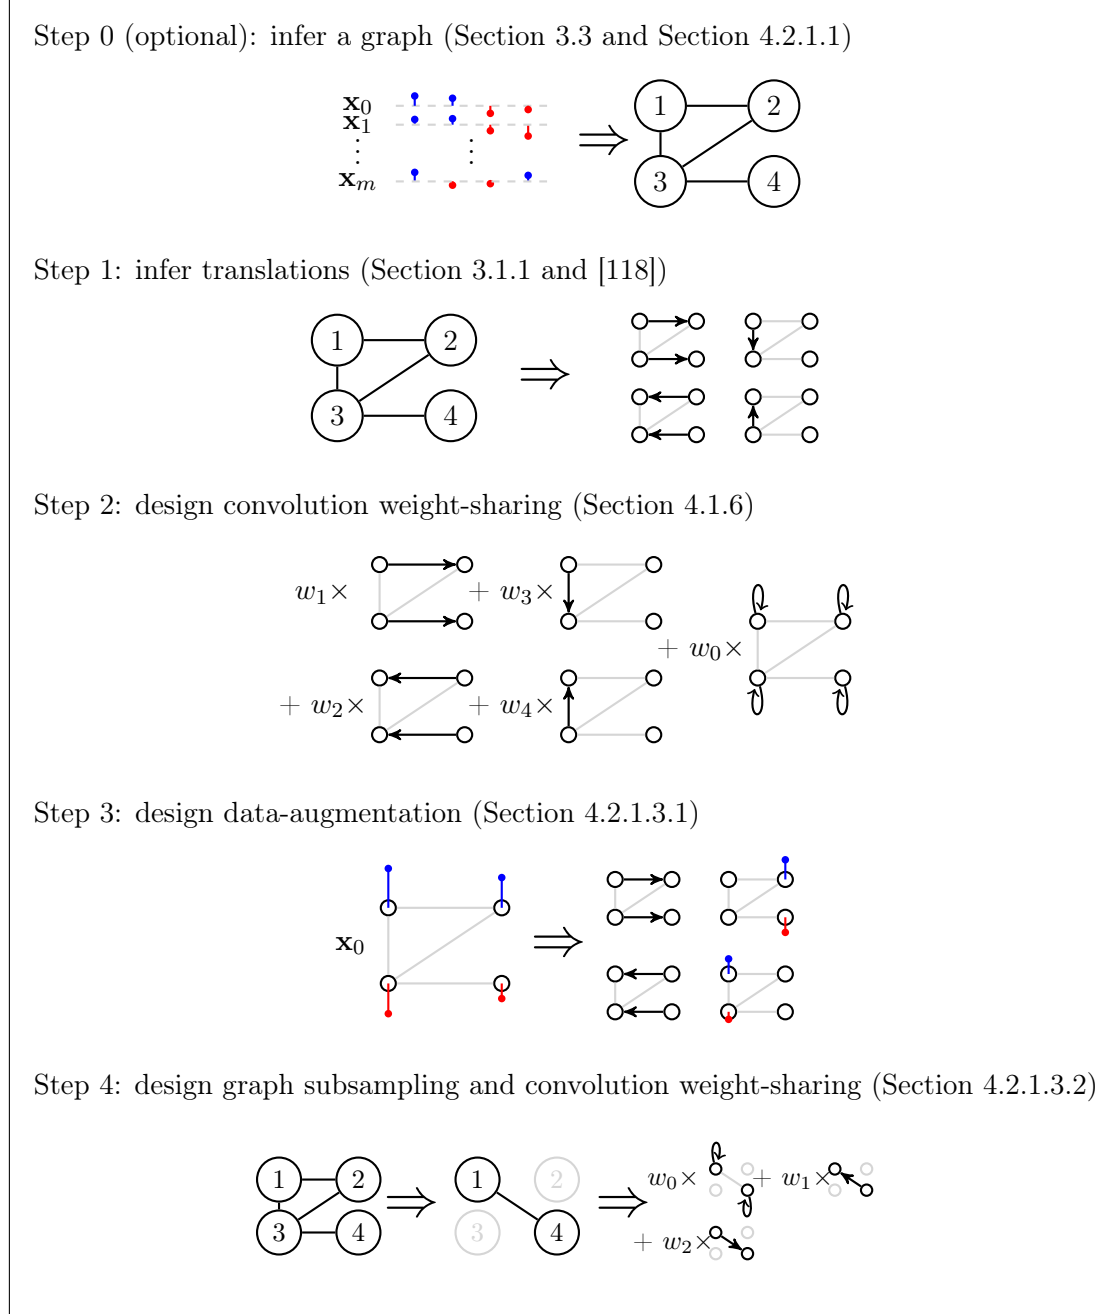


Figure 4.2: Outline of the different steps in designing GNNs based on graph translations. Figure extracted from [91] ©2018 IEEE.

In the next paragraphs, we explain the steps of the method that were not introduced before in the manuscript: data augmentation, strided convolutions and convolutions on the strided graph.

4.2.1.3.1 Defining data augmentation in graph space Once translations are obtained on \mathcal{G} , one can use them to move training vectors, artificially creating new ones. Note that this type of data-augmentation is still not at the same level as the ones for images since no flipping, scaling or rotations are used and that it preceded more recent forms of data augmentation that could be applied to graph data such as [165, 184]. Future work combining such approaches could reduce the gap even further between state of the art CNNs and GNNs.

We depict an example of the translation-based data augmentation in Figure 4.3. Note that in the case of the grid graph the augmented image closely resembles the image augmented using the original 2D support, and that in the inferred case we start losing information.



Figure 4.3: Comparison of image data augmentation and data augmentation using graph translations.

4.2.1.3.2 Strided convolutions and convolutions on strided graphs Downscaling is a tricky part of the process because it supposes one can somehow regularly sample vectors. As a matter of fact, a nonregular sampling is likely to produce a highly irregular downscaled graph, on which looking for translations irremediably leads to poor accuracy, as we noticed in our experiments. We rather define the translations of the strided graph using the previously found translations on \mathcal{G} .

First step: extended convolution with stride r

Given an arbitrary initial vertex $v_0 \in \mathbb{V}$, the set of kept vertices $\mathbb{V}_{\downarrow r}$ is defined inductively as follows:

- $\mathbb{V}_{\downarrow r}^0 = \{v_0\}$,
- $\forall t \in \mathbb{N}, \mathbb{V}_{\downarrow r}^{t+1} = \mathbb{V}_{\downarrow r}^t \cup \{v \in \mathbb{V}, \forall v' \in \mathbb{V}_{\downarrow r}^t, v \notin N_{r-1}(v') \wedge \exists v' \in \mathbb{V}_{\downarrow r}^t, v \in N_r(v')\}$.

This sequence is nondecreasing and bounded by \mathbb{V} , so it eventually becomes stationary and we obtain $\mathbb{V}_{\downarrow r} = \lim_t \mathbb{V}_{\downarrow r}^t$. Figure 4.4 illustrate the first downscaling $\mathbb{V}_{\downarrow 2}$ on a grid

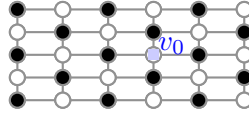


Figure 4.4: Downscaling of the grid graph. Disregarded vertices are filled in. Figure and caption extracted from [91] ©2018 IEEE.

graph.

The output neurons of the extended convolution layer with stride r are $\mathbb{V}_{\downarrow r}$.

Second step: convolutions for the strided graph

Using the proxy-translations on \mathcal{G} , we move a localized r -hop indexing kernel over \mathcal{G} . At each location, we associate the vertices of $\mathbb{V}_{\downarrow r}$ with indices of the kernel, thus obtaining what we define as induced \downarrow_r -translations on the set $\mathbb{V}_{\downarrow r}$. In other words, when the kernel is centered on v_0 , if $v_1 \in \mathbb{V}_{\downarrow r}$ is associated with the index p_0 , we obtain $\phi_{p_0}^{\downarrow r}(v_0) = v_1$. Subsequent convolutions at lower scales are defined using these induced \downarrow_r -translations similarly to Section 4.1.6.

4.2.2 Experiments

In this section we present two types of experiments, first we try to achieve similar performance as 2D convolutions in the CIFAR-10 dataset, we then use the PINES dataset in order to challenge our methods in a less regular domain.

4.2.2.1 CIFAR-10

On the CIFAR-10 dataset, our models are based on the Resnet-18 architecture. We tested different combinations of graph support and data augmentation and present the results in Table 4.2. In particular, it is interesting to note that results obtained without any structure prior (91.07%) are only 2.7% away from the baseline using classical CNNs on images (93.80%). This gap is even smaller (less than 1%) when using the grid prior. Also, without priors the method from Section 4.2.1.3 significantly outperforms the others and the network defined on the translations of the grid graph obtains a result that is slightly better than the original network. Note that this advantage may come from the fact that the networks defined on the translations have slightly more parameters than the original network.

Table 4.2: CIFAR-10 result comparison table. Note that the results using the method from Section 4.2.1.2 and the original CNN are the same in the case of grid graphs (as the embedding of a grid graph is the 2D structure).

Support	MLP[100]	CNN (Resnet-18)	Grid Graph (Oracle)		Covariance Graph (Inferred)	
			Chebnet [24] Section 4.1.5	Section 4.2.1.3	Section 4.2.1.2	Section 4.2.1.3
Random crop + Flip	78.62%	93.80%	85.13%	93.94%	—	92.57%
Random crop	—	92.73%	84.41%	92.94%	—	91.29%
Graph Data Augmentation	—	—	—	92.81%	89.25% ^a	91.07% ^a
None	69.62% ^a	87.78%	—	88.83%	—	85.88% ^a

^a No priors about the structure.

4.2.2.2 PINES

We now test our methods on the previously introduced PINES dataset. We use a shallow network to evaluate our results and compare with a simple MLP, a CNN with 1x1 filters (i.e. that treats each sensor separately before the classification layer) and a 3D CNN with 9x9x9 kernels. In this way we are able to compare with methods that do not take into account the structure (MLP, CNN1x1) and with methods that do, but not optimally (3D CNN applied to irregular inputs). The results are presented in Table 4.3, where we see that both methods based on geometric graphs achieve results that are close to the network that uses the original support. Yet neither the graph supported or the 3D supported methods were able to clearly outperform a method that does not use the structure (CNN 1x1).

Table 4.3: PINES fMRI dataset accuracy comparison table.

Support	None		Irregular 3D structure	Neighborhood Graph		
Method	MLP	CNN1x1	3D CNN	Chebnet [24] Section 4.1.5	Section 4.2.1.3	Section 4.2.1.2
Accuracy	82.62%	84.30%	85.47%	82.80%	85.08%	84.78%

4.2.3 Conclusion

In this section based on recent findings [31, 51, 175] we have shown that graph convolution methods are not able to improve over simple baselines, which is inline with the discussion from the previous section. We then have introduced two methods that try to mitigate these problems by either embedding the graph in a \mathbb{Z}^2 space or infer graph translations that serve as a proxy to translations in the \mathbb{Z}^2 domain. We have demonstrated with experiments that we are able to close the gap between CNNs and GNNs in the case of regular 2D domains, but on the other hand the discussed methods do not seem to generalize well to more irregular graphs (irregular 3D domain).

4.3 Semi supervised classification of vertices

One of the most interesting use of data defined on graph is the semi supervised classification of vertices. Indeed many applications can be resumed into this framework, from social networks (identifying data about the users based on their connections) to citation networks (classifying the article based on its text and citations). It is thus easy to understand why it is the most chosen task to evaluate new methods. Indeed almost all methods discussed in Section 4.1 use this task on their benchmarks in order to empirically evaluate their abilities.

Unfortunately there is a two fold problem in the evaluation of these methods using the standard datasets. Before going deep into the problem, we would like to preface by saying that most of the time these problems do not arise from malice or lack of knowledge, they come simply from a lack of computational power/deadline rush from the authors. The first problem is the choice of the $\mathcal{D}_{\text{train}}, \mathcal{D}_{\text{valid}}, \mathcal{D}_{\text{test}}$ split. While using the same split – as it is done in image applications in order to perform a fair comparison – seems to be a clear method to follow it has been shown in [145] to be problematic as the samples are not independent. Indeed, as all the elements of \mathcal{D} are in the graph \mathcal{G} , just randomly choosing the samples is not enough: the relationships between them will bias the evaluation.

Consider a simple example of a graph with multiple rings, where vertices from one ring are connected with the two successive vertices of the same ring, but with at most one vertex of another ring. In this case the split is very important as methods that give more importance to close nodes would be more effective when vertices from the same ring are present in all sets, while methods that are able to use long connections would be better if only one vertex per ring was chosen for each split. We depict such a ring graph with two different splits in Figure 4.5.

This problem with the \mathcal{D} split has led to the proposal of new benchmarks such as [57], that already implement this multi-split setup by default. Unfortunately, this is only half of the problem with benchmarking GNNs.

There is also the problem of fairly comparing methods using the same underlying architecture/regularization in order to ensure that the improvements come from the proposed method and not from a more efficient regularization. This has already been discussed in several papers such as [75, 145, 167] where the authors have shown that by adding one simple regularization parameter, that we call edge-dropout (as it consists in adding dropout [151] to the edges of S) during training we can have similar performances between models that had been originally very far in performance (GCN [74] and GAT [164]).

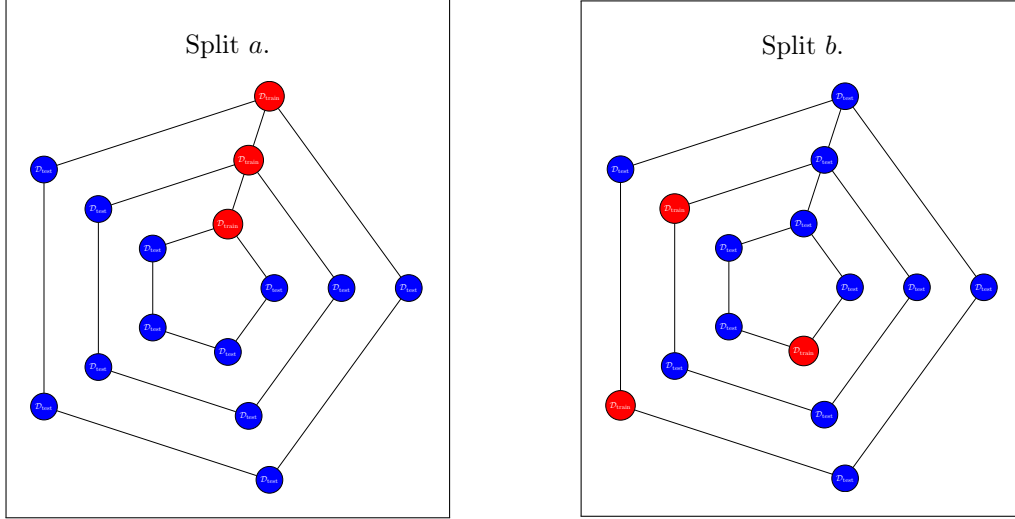


Figure 4.5: Two different train/test splits for the same multi-ring graph. Note how the label will propagate very differently depending on the split. Red vertices are on the training set and blue vertices are on the test set.

More-so, this regularization was present in the GAT model and not in the original GCN.

Indeed this problem still happens, for example consider the recent work in [5] and their low-pass filter method. While we agree with the authors that it is more sound than the GCN filter, we do not agree that they were able to outperform the GCN model in equal conditions. If we take the optimized GCN from [167], where the only difference to the original GCN [74] is the introduction of edge-dropout, their low-pass filter is not able to improve over the optimized GCN, even if the low-pass filter method also applies edge-dropout.

Note that this is only one example of how unfair comparison happens. Making them fair is hard. The easiest way would be to ensure that one only compares to methods that use the same regularization as oneself, but this could generate too big of a barrier, as it would not allow researchers to fully draw the capabilities of their methods. The harder way would be to correctly perform hyperparameter search to ensure that both the proposed method and the one compared with are on the best of their capabilities. This is not only hard in the theoretical sense of what are the hyperparameters one should optimize, but on the practical sense of executing all the computations needed to perform such a research/optimization. Also note that this hyperparameter problem is not only present in this task, but it is amplified by the split problem, which further biases methods for a set of hyperparameters.

In the following paragraphs we delve further in this two-fold problem, by proposing

a framework to benchmark the different graph filters.

4.3.1 A framework for comparing graph filters

An ideal framework for benchmarking the task of semi-supervised classification of vertices in a graph would be one that takes into consideration both problems: i) models should be compared by their performance over multiple cross-validation splits, and ii) models should be compared either to the maximum of their capabilities or by using the exact same architecture and regularization as the method they compare to. Solving the first problem is easy in theory, but hard in practice as it involves running s times more tests where s is the number of splits. Note that s should be a high number in order to avoid biasing the comparison, how high should it be depends on the confidence interval one expects to obtain from the measure.

In the worst case scenario this would lead to an increase of sm to the overall computational time, where m is the number of methods one is comparing too. Note that it could be relaxed to s if the methods are already benchmarked correctly with a common set of s splits. If we consider $s \geq 50$ this would mean that just doing correct split procedure would increase the time needed for running the experiments 50 fold. If we add to this the fact that it is common to rerun the experiments n times to ensure that the results are not obtained thanks to a lucky initialization of parameters, we have even more time of running experiments.

But this is not even the worst part of the computational complexity. Correctly choosing the hyperparameters that one wishes to optimize is a very difficult problem. Here are some questions one should take into account:

- do we add graph filters as pre-processing (SGC [173])?
- do we add graph filters as post-processing (APPNP [75])?
- which neural networks should we use (single-layer, multi-layer, GCN [74])?
- what size we use for each layer? In the case of multi-layer do we use the same?
- do we combine multiple filters in the same layer?
- which graph filter should we use? do we use the same graph filter for each step? what about their hyperparameters?
- which value of dropout to use at the input? do we use the same for the input of each layer?

- which value of edge-dropout should we use? do we use the same for each graph filter?
- do we use L2-regularization? Which value do we choose? The same for all layers?
- which learning rate do we use?

Note that answering all these 12 questions will generate a very large amount of combinations that one should test (2 choices per combination is already 2^{12} possibilities). If we also do it ns times to ensure that it is correctly done in both split and random initialization sides it would be very much intractable.

4.3.2 Relaxed framework for comparing graph filters

Given that it is impossible to tackle all the questions, we thus reduce the number of choices by simplifying the problem with a few simple rules: i) if one element is applied multiple times it is always applied with the same value, ii) only one filter is used to isolate their contribution, iii) only one set of filter hyperparameters is used (SGC with $m = 2$ is considered a different filter from SGC with $m = 3$), iv) we only consider networks with a single hidden layer (hidden size of 64), v) we use the same amount of L2-regularization (0.005) and apply it only to the hidden layer, and vi) we use the same amount of learning rate (0.01) as it is always the same network. Note that by doing so we vastly diminish our ability to ensure that the found solution is the best one. The chosen set parameters for learning rate, hidden layer size and L2-regularization come from [75]. This choices leads us to a simplified framework:

- which graph filters/graph filter hyperparameters do we test? [f choices]
- graph filters as pre-processing or post-processing? [3 choices]
- dropout values for the input? [d_i choices]
- dropout values for the edge-dropouts? [d_e choices]

This leads to a total of $3fd_id_esn$ choices. We can now set these values to amounts that seem reasonable such as 4 for all the simple choices, $s = 50$ and $n = 2$ leaving us with 4,800 tests to execute per filter. Even under this simplified conditions a very quick test that takes 1 second to fully execute (from data load to network training to evaluation and saving the results) would take almost an hour and a half to finish. If we consider a more realistic scenario of 30 seconds per test, we are at almost 2 days of training for a

single filter. Note that this is for a single filter, in a very simplified scenario where there is only the weights of a logistic regression to train and that a great deal of hyperparameters are taken out. It is not surprising that most papers are not able to do this properly if they have to compare their method with others in the literature. Ideally this could be counter-acted by having a common set of hyperparameters one would optimize in order to compare with other methods.

Note that in this experiment scenario, when we use a graph filter for pre-processing we are doing a type of feature denoising. On the other hand, when we use a graph filter for post-processing we are doing a type of label propagation. Finally, if we use both, it is akin to a one layer GCN.

4.3.3 Graph filter comparison

In the previous section we have defined the relaxed framework. In this section we present the results using the cora dataset. We chose a split of 20 examples per class for $\mathcal{D}_{\text{train}}$, 30 examples per class for $\mathcal{D}_{\text{valid}}$ and the rest for the test set as used in [145]. We use a maximum of 10000 epochs with early stopping if the validation accuracy does not improve after 100 epochs (patience threshold). Code for reproducing the experiments is available at <https://github.com/cadurosar> and a summary of the searched hyperparameters is available in Table 4.4.

Table 4.4: Summary of the searched hyperparameters.

Dropout input	0, 0.25, 0.5, 0.75	
Dropout kernel	0, 0.25, 0.5, 0.75	
Graph filter position	Pre-processing, post-processing, both	
Filters	h_{SGC}	$\alpha = 1$ and $m = 2$
	h_{Tikhonov}	$\alpha = \{10, 50\}$
	h_{VBL}	$\alpha = 0.1$ and $m = 20$
	$h_{\text{baleilar-lowpass}}$	$\alpha = 1$ and $m = \{5, 10\}$
	h_{Page}	$\alpha = 0.1$

We first present the mean validation and test set accuracies, alongside the 95% confidence interval in Table 4.5. Note that we only show the results for the hyperparameters that achieved the best validation accuracy for each considered graph filter. In this scenario, the best performing filter of our hyperparameter search was the PageRank filter introduced in [75], but contrary from the original paper, it performed best when used only for feature

denoising instead of label propagation. Also note that unfortunately the results are not 100% conclusive, as seen by the confidence intervals. We believe that adding more initializations per split could allow us to retrieve conclusive results (more than the 95% confidence interval).

Table 4.5: Performance comparison between different filters using 50 different splits and 2 initializations per split. Results presented as mean accuracy \pm 95% confidence interval.

Graph filter	Graph filter position	Dropout input	Dropout kernel	Validation Set	Test set
$h_{\text{bilateral-lowpass}}(\alpha = 1, m = 10)$	Both	0.5	0.25	83.73 ± 0.51	79.28 ± 0.31
$h_{\text{VBL}}(\alpha = 0.1, m = 20)$	Both	0.5	0.5	84.46 ± 0.53	80.14 ± 0.34
$h_{\text{Tikhonov}}(\alpha = 10)$	Pre-processing	0.5	0.5	84.78 ± 0.47	80.19 ± 0.36
$h_{\text{SGC}}(\alpha = 1, m = 2)$	Both	0.75	0	84.96 ± 0.51	80.52 ± 0.33
$h_{\text{Page}}(\alpha = 0.1)$	Pre-processing	0.25	0.5	85.24 ± 0.47	80.92 ± 0.3

4.3.4 Results on the planetoid split

Now that we have selected the hyperparameters on the cora dataset, we can check the performance of our found hyperparameters on the split defined by [176]. This split is used as the de-facto comparison in most papers in the literature. Note that we do this only to verify the performance when compared to other papers as this is not the ideal comparison method. We test 100 different initializations and report the mean test set accuracy alongside the 95% confidence interval in Table 4.6. We note that the relative order of the methods stays the same as per our hyperparameter search and that the performance of the best method (Pagerank) rivals with the best performances found using GCNs and GATs [5, 74, 145, 164, 167] while only applying the graph filtering operation as a pre-processing. Finally, the smaller values for the 95% confidence interval when compared to the difference previous test could be linked to the difference in amount of initializations and splits, which allowed us to have more significant results.

Table 4.6: Performance comparison between different filters on the split from [176]. Results presented as mean accuracy \pm 95% confidence interval.

Graph filter	Graph filter position	Dropout input	Dropout kernel	Test set accuracy
$h_{\text{bilateral-lowpass}}(\alpha = 1, m = 10)$	Both	0.5	0.25	81.27 ± 0.20
$h_{\text{VBL}}(\alpha = 0.1, m = 20)$	Both	0.5	0.5	82.42 ± 0.22
$h_{\text{Tikhonov}}(\alpha = 10)$	Pre-processing	0.5	0.5	82.52 ± 0.19
$h_{\text{SGC}}(\alpha = 1, m = 2)$	Both	0.75	0	82.62 ± 0.19
$h_{\text{Page}}(\alpha = 0.1)$	Pre-processing	0.25	0.5	83.53 ± 0.16

4.3.5 Ablation results

We now present ablation results for two of our variables, the dropout possibilities and the position of the graph filter. First concerning the different dropouts possibilities and then for the three different filter behaviours (feature denoising, label propagation and GCN-like).

4.3.5.1 Dropout ablation

We investigate the effect of using dropout on the inputs, the graph kernel or both. Note that the former may be seen as a type of data augmentation of the inputs, while the latter as a form of data augmentation of the graphs (increasing the amount of graphs that we can use on the train set). We display the mean test set accuracy for the best performing hyperparameters of each condition in Table 4.7. The presence of dropout seems to improve the performance on the tested cases, with the dropout on the inputs being more important than on the kernel, but the use of both improves the performance for all graph filters save for SGC.

Table 4.7: Performance comparison between different filters using 50 different splits and 2 initializations per split for different dropout conditions. Results presented as mean accuracy \pm 95% confidence interval.

Graph filter	No dropout	Dropout input only	Dropout kernel only	Dropout both
$h_{\text{baleilar-lowpass}}$	77.80 ± 0.35	78.90 ± 0.38	78.42 ± 0.34	79.28 ± 0.31
h_{VBL}	78.75 ± 0.37	79.90 ± 0.34	79.52 ± 0.33	80.14 ± 0.34
h_{Tikhonov}	79.43 ± 0.32	80.14 ± 0.35	80.11 ± 0.31	80.19 ± 0.36
h_{SGC}	79.94 ± 0.35	80.52 ± 0.33	80.00 ± 0.36	80.41 ± 0.34
h_{Page}	79.91 ± 0.32	80.76 ± 0.33	80.51 ± 0.51	80.92 ± 0.3

4.3.5.2 Graph filter position

Now we investigate the effect of using the graph filter as a pre-processing (directly at the input as a sort of feature denoising), as a post-processing step (directly at the output as a sort of label propagation) and in both, which in this case is close to a one layer GCN. We display the mean test set accuracy for the best performing hyperparameters of each condition in Table 4.8. We can see that for most filters, being applied as a post-processing works slightly better than pre-processing, but more tests would be needed to have any

significant conclusion. We also note that surprisingly the PageRank filter had a worse performance when it is used as both pre and post-processing. Further analysis on this effect is needed.

Table 4.8: Performance comparison between different filters using 50 different splits and 2 initializations per split for different dropout conditions. Results presented as mean accuracy \pm 95% confidence interval. Note that differently from the validation set, in the test set the page rank filter obtains a better result at the post-processing instead of the pre-processing, however for ensuring the correct rigor we use the pre-processing version for all other results as it was the best results on the *validation set*.

Graph filter	Pre-processing only	Post-processing only	Pre and post-processing
$h_{\text{bancilar-lowpass}}$	77.44 ± 0.33	77.29 ± 0.34	79.28 ± 0.31
h_{VBL}	78.45 ± 0.37	78.69 ± 0.33	80.15 ± 0.34
h_{Tikhonov}	80.18 ± 0.36	80.36 ± 0.39	80.11 ± 0.36
h_{SGC}	79.67 ± 0.33	79.73 ± 0.33	80.52 ± 0.34
h_{Page}	80.92 ± 0.3	81.09 ± 0.32	80.05 ± 0.32

In the previous paragraphs we have introduced what we believe would be the correct framework to evaluate the different graph filters, but had to use a downgraded version in order to keep it under reasonable timing constraints. We were able to see that the PageRank filter introduced in [75] had the best results between the analysed filters, but that further testing would be needed to understand all the effects that we saw in the results.

4.4 Summary of the chapter

In this chapter we have delved into the domain of deep neural networks defined on graphs. We have built upon the concepts from the previous chapters in order to define recent methods in a single graph filter framework that we have presented in increasing order of complexity in Section 4.1. While this framework is not exactly novel, we have extended it to more methods and have introduced a discussion on the drawbacks of these methods.

We then discussed applications of DNNs defined on graphs in the context of the supervised classification of graph signals in Section 4.2. We have discussed recent contributions that show the drawbacks of the current approaches in this domain and then introduced two of our contributions. Their aim is to close the gap between graph convolutions and classic 2D/3D convolutions. Our two introduced contributions were

published in conferences as follows:

- Grelier, N., Lassance, C. E. R. K., Dupraz, E., and Gripon, V. (2018). Graph-projected signal processing. In *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 763–767. IEEE
- Lassance, C. E. R. K., Vialatte, J.-C., and Gripon, V. (2018b). Matching convolutional neural networks without priors about data. In *2018 IEEE Data Science Workshop (DSW)*, pages 234–238. IEEE

Finally, we have discussed the applications on the context of semi-supervised classification of vertices. We have first discussed the problem of fairly evaluating the different GNN methods on this task. While it is not a novel problem with the domain, recent works still commit the two most common pitfalls: i) using only one train/valid/test split that has been already shown to bias results in [145], and ii) not comparing the methods fairly, e.g., method A performs better than method B, but it is mostly due to adding dropout than due to the method itself. Note that these problems are not necessarily due to malpractice or to malice, but mostly from the sheer computational problem that would be necessary to perform everything correctly. Indeed, we propose a framework in order to solve both i) and ii), but show that we would never be able to execute the optimal version in reasonable delay. We added a relaxed framework and present our results on the cora dataset.

Chapter 5

Deep Neural Networks latent spaces supported on graphs

5.1	Characterizing DNN behavior via smoothness of intermediate representations	132
5.2	Smoothness as an objective function for DNNs	138
5.3	Controlling DNN smoothness to improve robustness	146
5.4	Using intermediate representation graphs to compress DNNs	163
5.5	Summary of the chapter	172

In the first two chapters we introduced the concepts of DNNs and GSP. We then combined these concepts in the previous chapter to introduce neural networks on inputs defined on graphs. In this chapter we present applications where graphs can be used to represent the topology of intermediate DNN activations, even if the inputs are not defined on a graph support.

We mainly use the smoothness of intermediate representation graphs to characterize DNN behavior in this chapter. We consider multiple goals. The first one is to determine whether a DNN is overfitted, underfitted or achieves a good fit in Section 5.1. We then focus on applying this concept during the training of DNNs, showing that we can obtain performing feature extractors in Section 5.2 or that we can improve robustness to attacks and deviations in Section 5.3. Finally, we show that using intermediate representation graphs to define the topology of intermediate spaces we are able to improve knowledge

distillation performance in Section 5.4.

5.1 Characterizing DNN behavior via smoothness of intermediate representations

We have previously discussed the problem of overfitting in Chapter 2 and have defined our view in Definition 2.1.5. One of the problems that arises from overfitting is that most of the ways to address it revolve around two main concepts: i) better understanding of the domain in order to generate artificial data points, c.f. Section 2.1.3, and ii) removing data points from the trainable ($\mathcal{D}_{\text{train}}$) subset of \mathcal{D} and allocating them to $\mathcal{D}_{\text{valid}}$ and $\mathcal{D}_{\text{test}}$. While the former strategy has been shown to help prevent overfitting and improve generalization performance, the latter has been demonstrated to reduce overall generalization performance, leading to the new paradigm of only separating data into $\mathcal{D}_{\text{train}}$ and $\mathcal{D}_{\text{test}}$ which has lead to improved generalization performance, but may lead to overfitting to the $\mathcal{D}_{\text{test}}$ (i.e. the network may lose performance if we resample $\mathcal{D}_{\text{test}}$ from \mathbb{D}).

In order to address this problem, one possible solution is to develop analytical tools that try to infer overfitting using only data from $\mathcal{D}_{\text{train}}$. While it is hard to develop such tools with a very strong scientific background, we develop in this Section two empirical frameworks in order to analyse the state of DNNs a posteriori: i) use the evolution of the smoothness of intermediate representations to characterize the state of a DNN into underfitted, ok, overfitted and strongly overfitted., and ii) use the lessons learned from (i) to find a metric that is correlated to the performance on the $\mathcal{D}_{\text{test}}$ using only data on the $\mathcal{D}_{\text{train}}$.

5.1.1 Predicting DNN behavior using GSP

In this section we present the idea from [43] that was the cornerstone for our interest in using graphs and GSP to analyze the intermediate representations of DNNs. The authors propose to analyse the evolution of the smoothness of the binary label indicator signal (c.f. Definition 2.1.2) on graphs that are inferred at the output of each layer of a DNN at each training epoch. We recall that the smoothness of a signal on a graph represents how the signal and the graph are aligned. Also, using a binary label indicator vector as our signal means that the smoothness metric can be considered as the sum of the edges of nodes in different classes. Finally, as our graphs are inferred using the similarity between nodes and are thresholded to form a k -NN graph, a perfectly smooth graph would be one where either the similarity between examples of different classes is zero or there are no

connections between examples of different classes.

The goal is to try to find if there are differences in behavior when analyzing the smoothness metric for each case. Indeed, if there is a clear difference in behavior, this metric could potentially be used to identify, using only $\mathcal{D}_{\text{train}}$, which condition the network seems to favor and take the correct steps in order to transform an overfitted network to a “good fit” network, without needing to use $\mathcal{D}_{\text{valid}}$ or $\mathcal{D}_{\text{test}}$. In the following paragraphs, we present and discuss their results.

5.1.1.1 Experiments

The goal of the experiments is to analyze the evolution of the smoothness based metric over the training epochs. To do so, a Resnet-18 with preact blocks is used as the architecture, data augmentation is used as well in order to try to avoid overfitting to $\mathcal{D}_{\text{train}}$ and the experiments are performed on the CIFAR-10 dataset. There are four main different conditions that are analyzed:

1. “Good fit”/Reference: The reference training procedure;
2. Underfitting: An underfitted architecture where we divide by 10 the number of feature maps in each convolutional layer.
3. Extreme overfit, In this case we randomize the labels on $\mathcal{D}_{\text{train}}$. The accuracy on the randomized $\mathcal{D}_{\text{train}}$ is therefore 100% but on $\mathcal{D}_{\text{test}}$ it is as good as a random guess). This is inline with the findings from [183];
4. Overfitting: A slightly overfitted condition in which data augmentation and regularization are removed.

In Figure 5.1 we depict the evolution of train and test accuracy during the training of the networks. Note that as expected, the best test accuracy is detected on the “reference” case and that we are able to train the extreme overfit very quickly when compared with the other networks.

We then present the smoothness evolution in Figure 5.2, note how there is a very distinct behavior over the 4 conditions. Each representation is the output of a Resnet block (c.f. Section 2.1.1). That difference in behavior could indicate that it is indeed possible to evaluate the status of the network, using only information from $\mathcal{D}_{\text{train}}$. Note that there are important changes of dimension occurring multiple times throughout the process, which seems to be inline with the concept of group of blocks of the Resnet, in

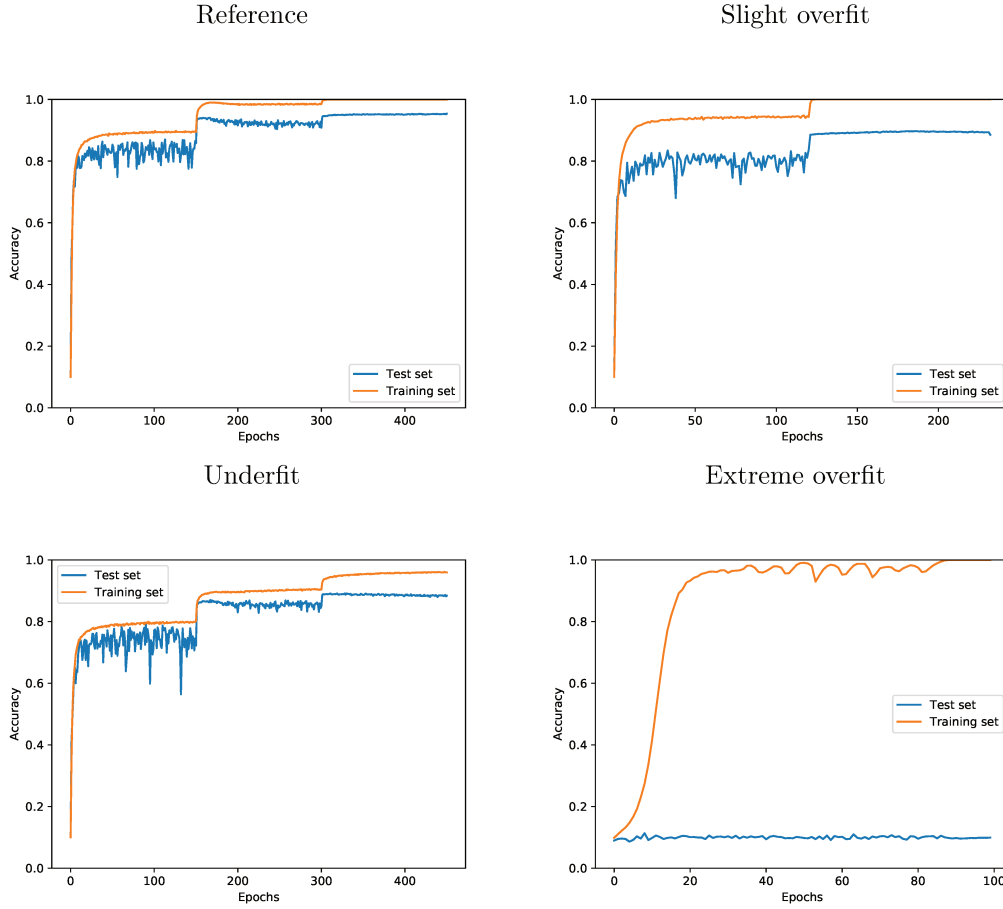


Figure 5.1: Train and test accuracy evolution on a Resnet18 under different conditions. Figure adapted from [43] ©2018 IEEE.

representations 3, 5 and 7. As a result, it is better to consider the representations as groups: i) 1, 2, and 3, ii) 4 and 5, iii) 6 and 7, and iv) 8, 9, and 10.

Interestingly, we are able to observe significant changes between the different conditions, in particular in the last group. In the reference case the three layers have reasonably similar label smoothness, whereas in other conditions we see important gaps, in particular between representations 8 and 9. Indeed, the lack of a gap between label smoothness between the penultimate representation and the output one seems to be a good indicator of overfitting.

On the other hand, it is also fair to say that the reference condition is slightly overfitting, considering the very high score on the training set. Recall that DNNs are normally trained with a learning rate that decays very quickly, and it seems that these sudden changes in the learning rate specializes the last group in clustering properly the examples whereas it effects the contrary for previous representations. Another interesting

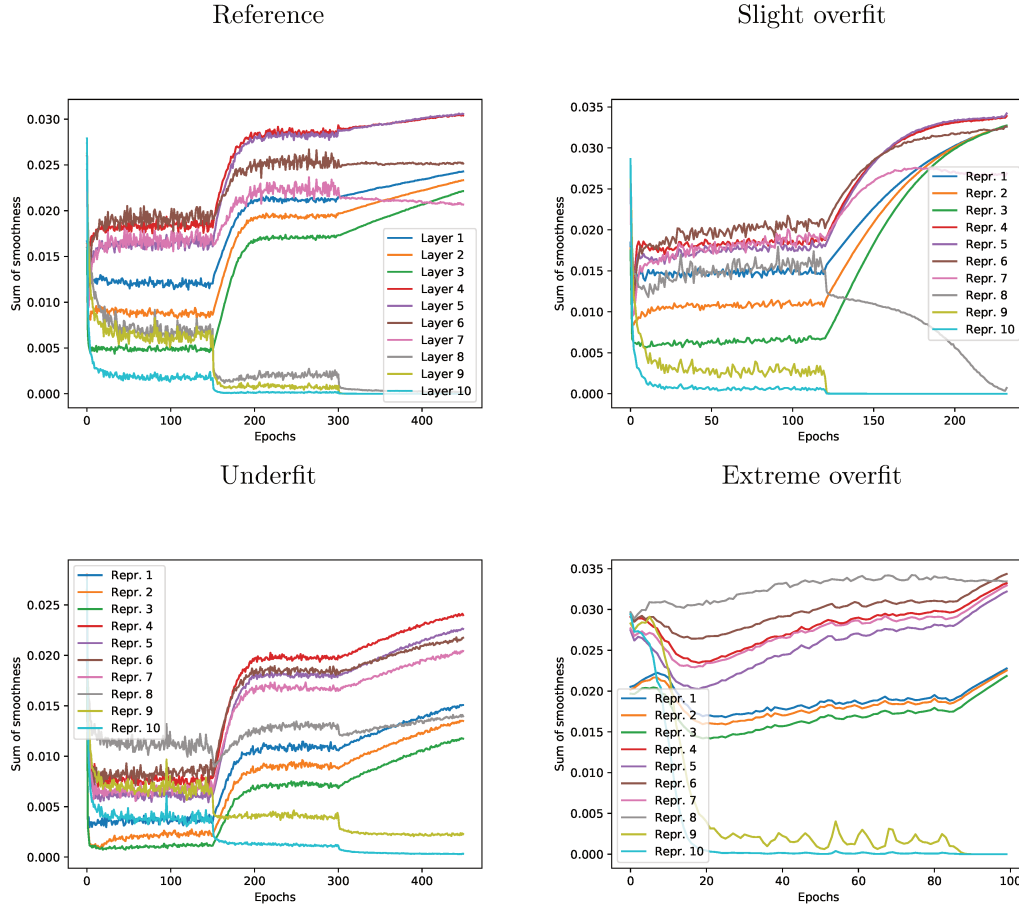


Figure 5.2: Smoothness evolution over the different representations of a Resnet-18 under different conditions. Figure adapted from [43] ©2018 IEEE.

observation is the fact that the label smoothness continues to evolve even if training and test accuracies remain constant. Indeed, this motivates the idea that even if the training accuracy is optimal, there are some representations that are still changing, and therefore, the label smoothness could also be used as a secondary measure to see if the network has converged. In the following section we present a continuation of this work, where we use the same smoothness metric in order to try to predict the test accuracy of the network.

5.1.2 Predicting under/overfitting using graph signal smoothness

To avoid underfitting and overfitting often boils down to performing crossvalidation, where one needs to split \mathcal{D} into $\mathcal{D}_{\text{train}}$ and $\mathcal{D}_{\text{valid}}$ in order to assess the generalization performance of a DNN trained on $\mathcal{D}_{\text{train}}$. However, this requires to reduce the size of the training set, thus leading to globally poorer accuracy, and it does not guarantee the architecture will be the best for a set distinct from the validation one. In this subsection

we use the GSP framework to develop a measure that we call “Smoothness Gap” in order to analyse the state of a DNN into underfitted, ok or overfitted. We do so using the knowledge acquired from the work detailed in the previous subsection [43] which shows that a neural network that generalizes well seems to have a very smooth k -neighbours graph using the features from its last layers. This is in line with previous works in the regularization of DNNs [95], which use the classification results based on the features of the intermediate layers of a network as a regularizer. The ideas of this section were accepted to a workshop without published proceedings [92].

5.1.2.1 Smoothness Gap

First we formally define what we call smoothness gap. Let us consider M example inputs for each of the C classes which we then use to generate intermediate representations across a given trained DNN. We then use the same strategy from Section 3.4.4 to generate k -nearest neighbor graphs for each intermediate representation of the DNN and consider the binary label indicator vector from Definition 2.1.2 as our graph signal. Note that by using k -NN graphs, each \mathcal{A} contains less than $2MCk$ nonzero elements and that choosing the correct value of k is a well known problem as we have previously discussed in the previous chapters.

With a graph for each intermediate representation and a graph signal we can then compute the smoothness (σ) of the graph signal for each intermediate representation. We recall that the smoothness of a label signal is a direct measure of how well the examples of this class are separated from the other classes, and that a global label smoothness of 0 indicates pairs of examples belonging to distinct classes are not connected in the graph or are completely orthogonal.

We define the smoothness gap as the difference between the smoothness of the representations on the last layer of the network (i.e. the classification layer) and the representations of the penultimate layer (i.e. the representations after the global average pooling). Note that this is influenced by our architecture, in our case we use Resnet-18 as defined in Chapter 2. Finally, in order to compare smoothness of a given signal on various graphs with possibly very different weightings, we choose to normalize smoothness by its maximum possible value. In our case, we rather use an upper-bound which is $2MCk$.

5.1.2.2 Experiments

We train our DNNs on a portion of the CIFAR-10 dataset [81]. To estimate label smoothness, we sample 50 examples from each class to generate our graphs ($M = 50$ and $C = 10$). We repeat this sampling 10 times. We evaluate our measure using graphs and we also compute the R^2 coefficient obtained by a linear regression over our measures to further stress the correlation. The results reported are the mean label smoothness over the 10 graphs. In order to controllably generate our underfitting and overfitting conditions, we proceed as follows:

1. **Overfitting:** we use only a portion of the training set ranging from 21% to 99% by 2% increments;
2. **Underfitting:** convolutional layers come with a hyperparameter which is the number of feature maps. In order to easily vary the number of trainable parameters without changing the global architecture, we thus vary the number of feature maps. In the chosen architecture, the number of feature maps on the first convolutional layer determines all the others. We thus vary it from 5 to 64, its default value.

We considered various values of k (10, 20, $M = 50$, $MC = 500$). Most consistent results were obtained with a value of 20. Using 10 would incur on a lot of points being concentrated with approximately zero smoothness. This is not surprising as it tends to select only the very nearest neighbors. Using M or MC leads to a lot of noise in the measures as there are many more pairwise distances to take into consideration.

In Figure 5.3 we show that by varying the size of the dataset we can generate highly overfitted DNNs. Moreover, there is a correlation between generalization abilities reported by the test accuracy score and the smoothness gap δ_s . We stressed this fact by computing a linear regression and obtained a R^2 coefficient of 68%.

Now we study the case of underfitted/properly fitted DNNs. First we test the case where we vary the amount of parameters on the network following the traditional scaling of Resnet-18 (c.f. Section 2.1.1) and we depict the experiments in Figure 5.4, where we show that we can also obtain a strong correlation between the test accuracy and the smoothness gap δ_s , as seen by the $R^2 = 84\%$ coefficient of its linear regression. These results show a very high predictability of the test error given the smoothness gap δ_s . It is very interesting to see this high predictability as the computation of the smoothness gap does not require any knowledge about the test set.

However, we note that for the underfitted condition, the performance of the network

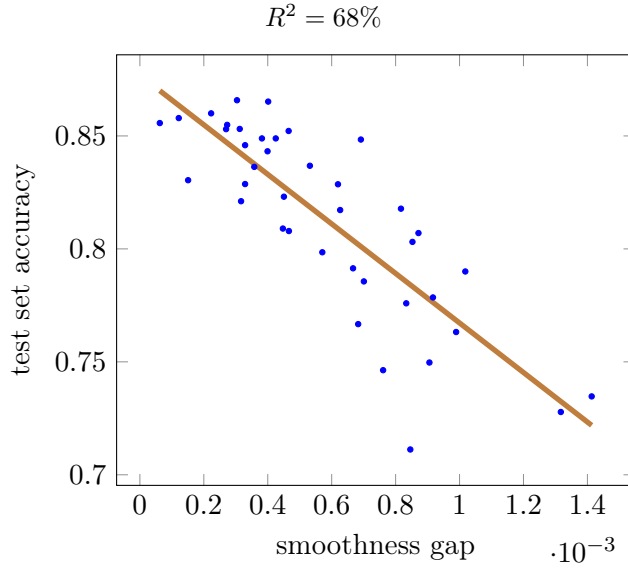


Figure 5.3: Results generated by varying the size of the dataset. Figure extracted from [92].

is also correlated to the number of parameters ($R^2 = 50\%$), even if this relation is less strict. In order to be sure that our measure was not only an indirect measure of the size of the network, we performed additional experiments where the number of feature maps at each layer is changed independently, resulting in almost no correlation between the global number of parameters and test accuracy ($R^2 = 14\%$), while still maintaining a good correlation between the smoothness gap and test accuracy ($R^2 = 67\%$). We depict the relationships between size, smoothness and test set accuracy in Figure 5.5.

In this subsection we have proposed to measure the smoothness gap and have shown via experiments that there exists a strong correlation between this measure and the generalization of DNNs. While this seems very promising, it is of utmost importance to be careful and not overpromise as further study is still needed to see if this is an useful correlation or if it is a subproduct of a possibly poor experimental design. Future work includes developing an understanding of why the training set smoothness is correlated with the test set accuracy, using this measure explicitly when performing hyperparameter search, and studying how to use this measure during the training phase.

5.2 Smoothness as an objective function for DNNs

In the previous section we described how the smoothness of graphs generated by the intermediate features of DNNs may be linked with their generalization abilities. In

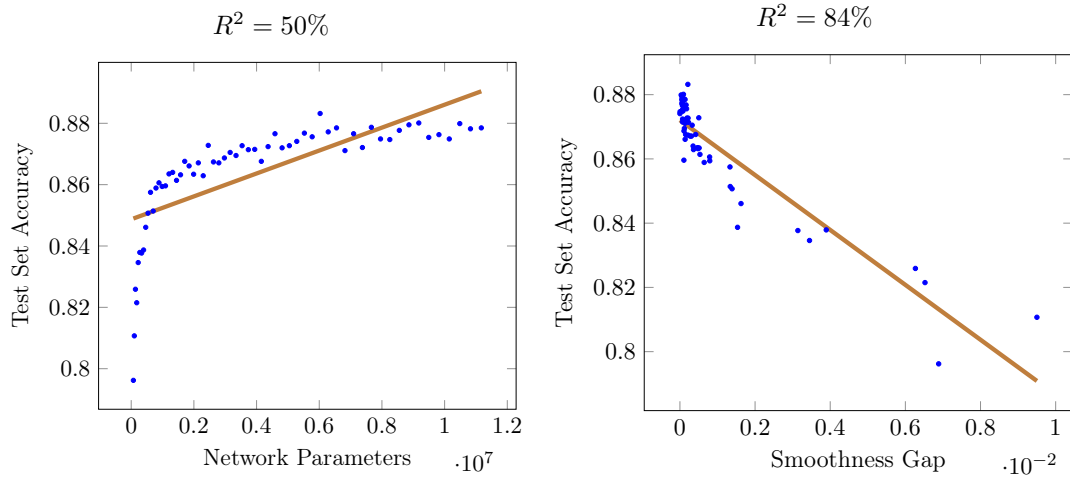


Figure 5.4: Results generated by varying the size of DNN following the traditional Resnet scaling pattern of the DNNs. In the left we depict the correlation between network size and test set performance, while on the right we show the correlation between the smoothness gap and test set performance. Figure extracted from [92].

this section we inverse this concept, by training the network to directly minimize the smoothness generated by a DNN in order to train a feature extractor, which was the central theme of our contribution [10].

In machine learning, classification is one of the most studied problems and cross-entropy is the most popular loss function for computer vision tasks. Cross-entropy is often preferred over mean squared error because it converges faster and tends to reach better accuracy. However, cross-entropy requires the outputs of the network to be label indicator vectors of the classes. We believe this decision comes with noticeable drawbacks:

- The dimension of the output vectors has to be equal to the number of classes, preventing an easy adaptation to the introduction of new classes. In scenarios where the number of classes is large, this also causes the last layer of the network to contain a lot of parameters.
- Inputs of the same class are forced to be mapped to the same output, even if they belong to distinct clusters in the input space. This might cause severe distortions in the topological space that are likely to create vulnerabilities to small deviations of the inputs.
- The arbitrary choice of the one-hot-bit encoding is independent of the distribution of the input and of the initialization of the network parameters, which can slow and harden the training process.

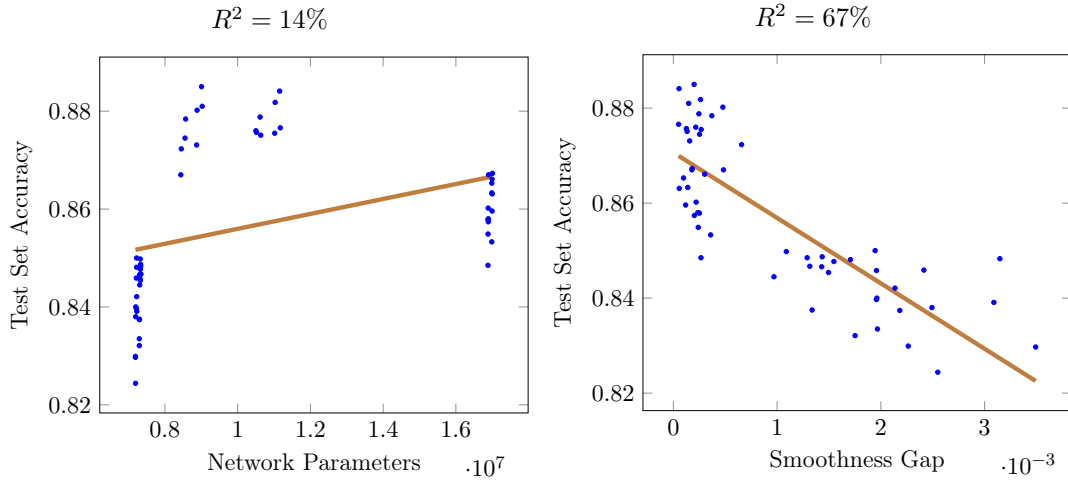


Figure 5.5: Results generated by varying the size of DNN without following the traditional Resnet scaling pattern of the DNNs. In the left we depict the correlation between network size and test set performance, while on the right we show the correlation between the smoothness gap and test set performance. Figure extracted from [92].

To overcome these drawbacks, authors have proposed several solutions, some of which we present here:

1. Some works propose to train DNNs solely as feature extractors, one example is [54], where the authors use triplets, where the first element is the example to train, the second belongs to the same class and the last to another class. They then enforce that the first is closer to the second than to the last.
2. Other authors try to smooth the outputs of the DNN, either using smoother representations [155] or by using a teacher network to define the output vector [55] or just by using smoother vectors.
3. Finally, works such as [26, 160] propose to use error correcting codes to generate outputs of the network.

In this section, we tackle the problem of training deep learning architectures to generate features that are easy to classify without relying on arbitrary choices for the representation of the output. We introduce a loss function that aims at maximizing the distances between outputs of different classes. It is expressed using the smoothness of a label signal on similarity graphs built at the output of the network. The proposed criterion does not force the output dimension to match the number of classes, can result in distinct clusters in the output domain for a same class, and builds upon the distribution of the inputs and the initialization of the network parameters. We demonstrate the ability

of the proposed loss function to train networks with state-of-the-art accuracy on common computer vision benchmarks and its ability to yield increased robustness to deviations of the inputs.

5.2.1 Methodology

Consider the problem of training a classification function f using $\mathcal{D}_{\text{train}}$, where $\mathcal{D}_{\text{train}}$ is made of n elements. In this case we can say that \mathbf{x}_μ refers to an input tensor and \mathbf{y}_μ to the corresponding output vector. We denote C the number of classes. In the context of deep learning, \mathbf{y}_μ is typically a binary label indicator vector of its class ($\mathbf{y}_\mu \in \mathbb{R}^C$) and the network function f is trained to minimize the *cross-entropy loss* as previously defined in Equation 2.4.

We will consider graphs \mathcal{G} , defined by their weighted adjacency matrix \mathcal{A} , where $\mathcal{A}_{\mu,\nu}$ is the weight of the edge between vertices μ and ν , or 0 if no such edge exists. We also define the Laplacian \mathbf{L} as: $\mathbf{L} = \mathbf{D} - \mathcal{A}$ where \mathbf{D} is the degree matrix of the graph.

Given a graph \mathcal{G} and a graph signal vector $\mathbf{s} \in \mathbb{R}^V$, we can compute the *graph signal smoothness*, c.f. Section 3.1.3. Finally, we call *label signal* associated with the class c the binary indicator vector \mathbf{s}_c of elements of class c . Hence, $\mathbf{s}_{c_\mu} = 1$ if and only if \mathbf{x}_μ is in class c (c.f. Definition 2.1.2).

5.2.2 Proposed graph smoothness loss

We propose to replace the cross-entropy loss with a graph smoothness loss. Consider a fixed metric $\|\cdot\|$. We can compute the distances/similarities between the representations $f(\mathbf{x}_\mu), \forall \mu \in \mathcal{D}_{\text{train}}$. Using this information, we build a k -nearest neighbor graph.

To generate this graph, denoted \mathcal{G} , we compute the similarity between each representation using an RBF-kernel parameterized by α , and then threshold to the closest k -neighbors. This leads us to the following \mathcal{A} :

$$\mathcal{A}_{\mu,\nu} \neq 0 \Rightarrow \mathcal{A}_{\mu,\nu} = \exp(-\alpha \|f(\mathbf{x}_\mu) - f(\mathbf{x}_\nu)\|), \forall \mu, \forall \nu. \quad (5.1)$$

We can then define our graph smoothness loss as follows.

Definition 5.2.1 (graph smoothness loss). We call *graph smoothness loss* of f the quantity:

$$\mathcal{L}_{\mathcal{G}} = \mathbf{s}^\top \mathbf{L} \mathbf{s}$$

$$= \sum_{\substack{\mathbf{x}_\mu, \mathbf{x}_\nu, \mathbf{W}_k [\mu\nu] \neq 0 \\ \mathbf{s}_c [\mu] \mathbf{s}_c [\nu] = 0, \forall c}} \exp(-\alpha \|f(\mathbf{x}_\mu) - f(\mathbf{x}_\nu)\|) .$$

sum over inputs of distinct classes

In the following subsection, we motivate the use of this loss.

5.2.3 Properties of the graph smoothness loss

The cross-entropy loss introduced in Equation 2.4 aims at mapping inputs of the network to arbitrarily chosen one-hot-bit encoded vectors representing the corresponding classes. Our proposed loss function differs from the cross-entropy loss in three main aspects:

- The cross-entropy loss forces a mapping from the input to a single point for each class. This might force the network to considerably distort space, for example in the case where a class is made of several disjoint clusters. The use of k -nearest neighbors gives more flexibility to the proposed loss: using a small value of k , it is possible to minimize the graph smoothness loss with multiple clusters of points for each class;
- The cross-entropy loss requires to arbitrarily choose the outputs of the network, disregarding the dataset and the initialization of the network. In contrast, the proposed loss is only interested in relative positioning of outputs with regards to one another, and can therefore build upon the initial distribution yielded by the network;
- To use the cross-entropy loss we are obliged to use an output vector whose dimension is the number of classes of the problem at hand. It is thus required to modify the network to accommodate for new classes (e.g. in an incremental scenario). The dimension of the network output d is less tightly tied to the number of classes with the proposed loss.

It is important to note that the capacity and the dimension of the output space need to be bounded by the problem at hand. If the output dimension is too small, it is likely that the network will not be able to converge (i.e., underfit): consider a toy example, where we try to separate n samples so that they are all at the same distance in the output space. It is only possible to suffice this condition if the output dimension is at least $n - 1$. On the other hand, if the capacity is too large, we can simply scatter each point in the output space, so that the distance between the image of any two inputs is large, but not significant as it has the same behavior for any input. This relation between the dimension

of the output space and the ability of the network to classify is further discussed in the experiments.

5.2.4 Experiments

We evaluate the performance of the proposed loss using three common datasets of image classification: i) CIFAR-10, ii) CIFAR-100, and iii) SVHN. For each dataset, we follow the same experimental process: i) We define the architecture we are going to use, following networks that are known to provide a good result when using the cross-entropy loss, ii) We train two networks of the same architecture (number of layers, number of features per layer) and hyperparameters (number of epochs, learning rate, gradient descent algorithm, mini-batch size, weight decay, weight normalization), but one is trained with the cross-entropy loss and the other with the proposed graph smoothness loss, and iii) We then tune the additional hyperparameters of the proposed loss (k, α, d). When performing classification, we train a simple classifier on top of the network to measure its accuracy. Note that all input images are normalized before being processed. It is important to keep in mind that by choosing this methodology, we bias the experiments in favor of using the cross-entropy loss, since the chosen architectures have been designed for its use.

The network architecture we use is Resnet-18 [50], as previously defined in Chapter 2. The network is trained for 200 epochs using 100 examples per mini-batch. SVHN and CIFAR-10 networks are trained with SGD, using a learning rate that starts at 0.1 and is divided by 10 at epochs 100 and 150, with a weight decay factor of 10^{-4} and a Nesterov momentum of 0.9. On the other hand, CIFAR-100 is trained with the Adam optimizer [73], using a learning rate that starts at 0.001 and is divided by 10 at epochs 100 and 150. Note that due to computational constraints we built a graph for each mini-batch (i.e., graph smoothness is calculated on a graph of 100 vertices that changes at each mini-batch).

In the original version of the chosen architecture, the linear function of the last layer outputs a C dimensional vector on which a softmax function is applied. When using the proposed loss, the linear function outputs a d dimensional vector, where d is an hyperparameter, normalized with respect to the L_2 norm. We use this normalization to constrain the outputs to remain in a compact subset of the output space. As previously discussed in Section 5.2.1, if we did not normalize the output, and since we use the L_2 metric to build the graphs in our experiments, the network would likely converge to a trivial solution that would scatter the outputs far away from each other in the output domain, regardless of their class.

5.2.4.1 Visualization

We first compare the embedding obtained using the proposed loss and $d = 2$ with the one obtained when putting a bottleneck layer of the same dimension $d = 2$ using the cross-entropy loss. Results on CIFAR-10 are depicted in Figure 5.6. Exceptionally, for this experiment we do normalize the output of the last layer of the network using batch norm instead of L_2 norm. This is because using $d = 2$ with a L_2 normalization would reduce the output space dimension to 1, which would likely be too small to allow the training loss to descend to 0. We observe that in the third column of Figure 5.6, our method creates clusters whereas the baseline method creates lines. This reflects the choice of the distance metric: our method uses the L_2 distance, whereas the baseline seems to use the cosine distance instead. Figure 5.6 shows that training examples are better clustered at the end of the training process when using the proposed loss than with the cross-entropy loss.

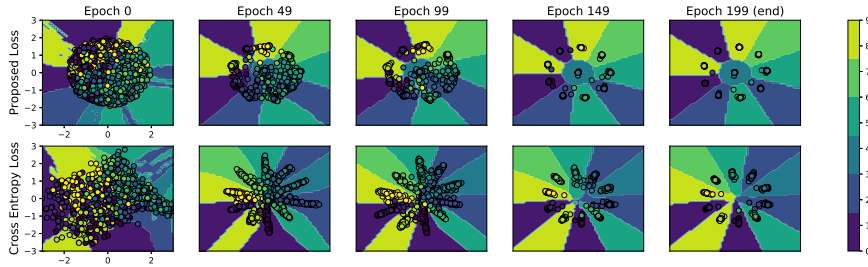


Figure 5.6: Embeddings of CIFAR-10 training set learned using the proposed graph smoothness loss with $d = 2$ (top row) compared with the ones obtained using a bottleneck layer and cross-entropy with the same architecture (bottom row). Figure and caption extracted from [10]. ©2019, IEEE.

5.2.4.2 Classification

We evaluate the influence on classification performance of the three hyperparameters of the proposed loss: the number of neighbors k to consider in the similarity graph \mathcal{G} , the number of dimensions d coming out of the network and the scaling parameter α used to define the weights of the graph. When varying k , we fix d to be the number of classes and $\alpha = 2$; when varying d , we fix k to the maximum value and $\alpha = 2$ and when varying α , we fix d to be the number of classes and k to the maximum value. The results are summarized in Figure 5.7. Note that a 10-NN classifier was used to obtain the accuracy. We observe that the higher k is, the higher the test accuracy is, even if

the sensitivity to k is lower when k is larger than the number of classes. As soon as d becomes large enough to accommodate for the number of classes, we observe that the test accuracy starts dropping slowly. Therefore, because using a larger value of d does not seem particularly harmful, applications where the number of classes is unknown (such as in incremental learning) should use a high d . Similarly, there is almost no dependence to α as long as its value is small enough. Indeed, when α is large, the loss tends to be close to 0 even if the corresponding distances are still relatively small.

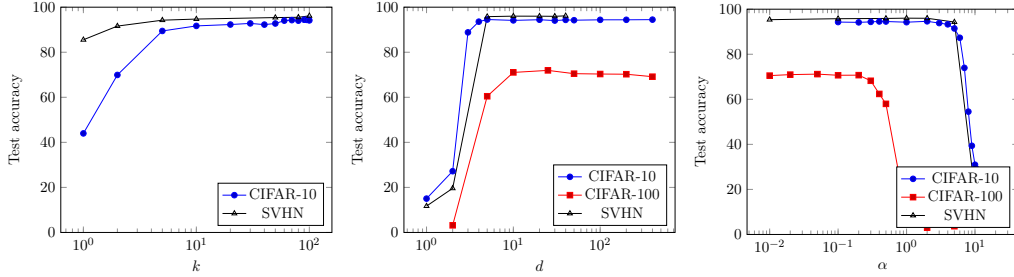


Figure 5.7: Test set accuracy as a function of the different parameters k , d and α . Plots extracted from [10] ©2019, IEEE.

We next evaluate the performance of the graph smoothness loss for classification. To this end, we compare its accuracy to that achieved with optimized network architectures using a cross-entropy loss (CE). We use various classifiers on top of the graph smoothness loss-trained architectures: a 1-nearest neighbors classifier (1-NN), a 10-nearest neighbors classifier (10-NN) and a support vector classifier (SVC) using radial basis functions. The results are summarized in Table 5.1. We observe that the test error obtained with the proposed loss is close to the CE test error, suggesting that the proposed loss is able to compete in terms of accuracy with the cross-entropy. Interestingly, we do not observe a significant difference in accuracy between the classifiers. Besides, both losses require the same training time.

5.2.4.3 Robustness

We now evaluate the robustness of the trained architectures using the robustness benchmark defined in Section 2.5.4.2. We report the results in Table 5.2. We first report the error rate on the clean test set for which we observe a small drop in performance when using the proposed loss. However, this drop is compensated by a better accommodation to deviations of the inputs, as reported by the Mean Corruption Error (MCE) scores (see [53]). Such a trade-off between accuracy and robustness has been discussed in [32]. For this experiment, we fixed k to its maximum value, $d = 200$, $\alpha = 2$ and we used 10-NN as a

Table 5.1: Test errors on CIFAR-10, CIFAR-100 and SVHN datasets. The top contains the test error of the optimized network architectures for a cross-entropy loss (CE). The bottom contains the test error of the same network architectures for our proposed graph smoothness loss, associated with three different classifiers. Table and caption extracted from [10] ©2019, IEEE.

Loss - Classifier	CIFAR-10	CIFAR-100	SVHN
CE - Argmax	5.06%	27.92%	3.69%
Proposed - 1-NN	5.63%	29.17%	3.84%
Proposed - 10-NN	5.48%	28.82%	3.34%
Proposed - RBF SVC	5.50%	30.55%	3.40%

classifier when using the graph smoothness loss.

Table 5.2: Robustness comparison on the 15 corruptions benchmarks from [53] on the CIFAR-10 dataset. Table and caption extracted from [10] ©2019, IEEE.

Method	Clean test error	MCE	relative MCE
Cross-entropy	5.06%	100	100
Proposed	5.60%	95.28	90.33

In the previous paragraphs, we have introduced a loss function that consists in minimizing the graph smoothness of label signals on similarity graphs built at the output of a deep learning architecture. We discussed several interesting properties of this loss when compared to using the classical cross-entropy. We have shown empirically that the proposed loss can reach similar performance as cross-entropy, while providing more degrees of freedom and increased robustness to deviations of the inputs.

5.3 Controlling DNN smoothness to improve robustness

In the previous section we concentrated in the effects of changing the whole training objective of a DNN to the minimization of graph signal signal smoothness. Now, we go back to the ideas presented in Section 5.1 and study the robustness effect of controlling the evolution of DNN smoothness. The contents described in this section were made available in our archival contribution [90], and were the starting stone for defining the robustness metric in [87].

As we have previously discussed, the ability of DNNs to achieve good generalization

is closely related to the amount of data available. This strong dependency on data may lead to selection of biased features of the training dataset, resulting in a lack of robustness in classification performance. In this work robustness has been defined to be the ability of a classifier to infer correctly even when the inputs (or the parameters of the classifier) are subject to perturbations. These perturbations can be due to general factors –such as noise, quantization of inputs or parameters, and adversarial attacks– as well as application specific ones –such as the use of a different camera lens, brightness exposure, or weather, in an imaging task, c.f. Section 2.5 for a more in-depth discussion.

In this section we propose to introduce a regularizer that penalizes large deformations of the class boundaries throughout the network architecture, independently of the types of perturbations that we expect to face when the system is deployed. It also enforces a large margin r (i.e., mid-distance between examples of distinct classes) at each layer of the architecture. Note that we have already discussed some of the properties and results of this regularizer in Section 2.5, but here we detail its methodology and provide experiments to support our claims.

To understand the intuition behind our proposed regularizer, first recall that networks are typically trained with the objective of yielding zero error for the training set. If error on the training set is (approximately) zero then any two examples with different labels can be separated by the network, even if these examples are close to each other in the original domain. This means that the network function can create significant deformations of the space (i.e., small distances in the original domain map to larger distances in the final layers) and explains how an adversarial attack with small changes to the input can lead to class label changes. Our proposed regularizer penalizes big changes at the boundaries between classes. By forcing boundary deformations to evolve smoothly across the architecture, and at the same time by maintaining a large margin, the proposed regularizer therefore favors smooth variations. We argue that favoring smooth variations leads to better robustness, as per Definition 2.5.2. We will empirically demonstrate this claim on classical vision datasets.

The proposed regularizer is based on a series of graphs, one for each layer of the DL architecture, where each graph captures the similarity between training examples given their intermediate representation at that layer. Our regularizer favors small changes, from one layer to the next, in the distances between pairs of examples in different classes. Note that the distance between any two examples at a certain layer depends on their positions in the original domain and the network function applied up to that layer. Thus, constraints on the distances lead to constraints on the parameters of the network function. It achieves so by penalizing large changes in the smoothness (computed using the Laplacian quadratic

form) of the class indicator vectors (viewed as “graph signals”). As a result, the margin is kept almost constant across layers, and the deformations of space are controlled at the boundary regions, as illustrated in Figure 5.8. This regularizer draws heavily from the analysis derived in Section 5.1, and uses the robustness definition that was previously introduced in Section 2.5.

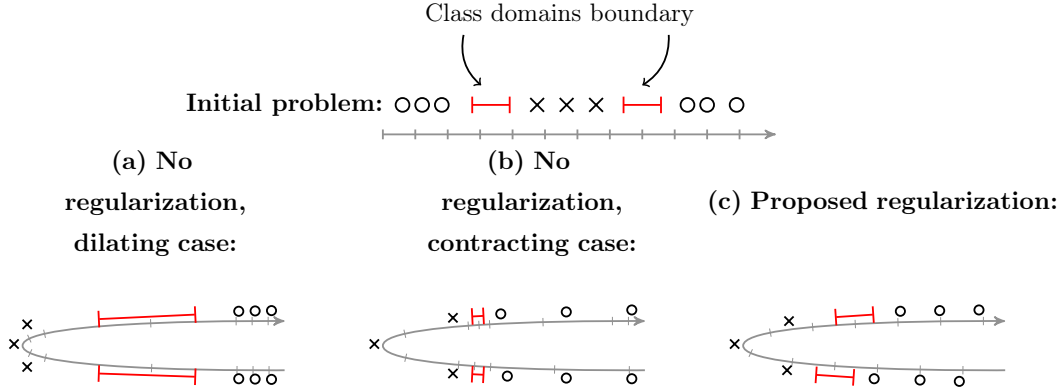


Figure 5.8: Illustration of the effect of our proposed regularizer. In this example, the goal is to classify circles and crosses (top). Without use of regularizers (bottom left), the resulting embedding may considerably stretch the boundary regions. Consequently the risk is to obtain sharp transitions in the network function (that would correspond to a large value of α in Equation (2.16)). Another possible issue would be to push inputs closer to the boundary (bottom center), thus reducing the margin (that would correspond to a small value of r in Equation (2.16)). Forcing small variations of smoothness of label signals (bottom right), we ensure the topology is not dramatically changed in the boundary regions. Figure and caption extracted from [90].

Another example of prior work that is related to our regularizer is [152] where the authors exploit graph convolutional layers. This leads to smoothing the latent representations of the inferred images using similar images from the training set, in order to increase the robustness of the network. Note that this could be described as a denoising of the inference (test) image, using the training ones. However this differs from the proposed regularizer as our work focuses on generating a smooth network function and their work focuses on combining inputs in order to generate a smooth network function.

In the following of this section we first present a quick recall of our robustness definition and then introduce our regularizer that enforce this property using similarity graphs. We then demonstrate, using readily-available image classification datasets, the robustness of the proposed regularizer to the following common perturbations: i) noise [53, 105], for which we show reductions in relative error increase, ii) adversarial attacks [40, 154], for which the median defense radius [111] is increased by 50% in comparison with the

baseline and by 12% in comparison with another method in the literature [21], and iii) implementation defects, which result in only approximately correct computations [60], for which we increase the median accuracy by 48% relative to the baseline and 26% relative to another method in the literature [21].

5.3.1 Methodology

In this subsection, we first recall our robustness definition and then present the proposed regularizer.

5.3.1.1 Robustness definition

Recall that a deep neural network architecture can be entirely described by its associated “network function”. In most cases, the network function f receives an input \mathbf{x} and outputs a class-wise classification score $f(\mathbf{x})$. Typically this output is a vector with as many coordinates as the number of classes in the problem, where the highest valued coordinate is the decision of the network (i.e. $\arg \max$ classifier). This function is constructed via the composition of multiple intermediate functions f^ℓ :

$$f = f^L \circ f^{L-1} \circ \dots \circ f^1, \quad (5.2)$$

where each function f^ℓ is highly constrained, typically as the concatenation of a parameter-free nonlinear function with a parameterized linear function.

The function f is typically obtained based on a very large number of parameters, which are tuned during the learning phase. During this phase, a loss function is minimized over a set of training examples using a variant of the stochastic gradient descent algorithm. At the end of the training process, each training example is associated through f with a vector whose largest value is the actual class of that example, leading to an accuracy close to 100% on the training set. Importantly, the loss function usually targets a specific margin in the output domain. For example, when using the classical cross-entropy loss, the loss function is minimized when the output of the training examples are the one-hot-bit vectors of their corresponding class [39], which corresponds to a margin in the output domain of about $\sqrt{2}/2$ for the L_2 norm.

We use the α -robust concept introduced in Definition 2.5.2. Recall that we can say that a network is α -robust if f is locally α -Lipschitz within a radius r of any point in domain R . Obviously, we would like to obtain a function f that is α -robust for any valid

input. But since we only have access to training samples, we only enforce the property over the training set.

This definition captures a compromise between margin (represented by r) and slope (represented by α) of the network function. This is in contrast to other works [21] where robustness is directly linked to the Lipschitz constant of the network function. The main motivation for introducing this weaker definition of robustness is that we do not want network functions to be contractive *everywhere*. Indeed, if all mappings are contractive everywhere we cannot hope to separate some samples in different classes. A more in-depth discussion of this is available in Section 2.5.

In what follows, we introduce regularizers that enforce this property using similarity graphs.

5.3.1.2 Intermediate representation graphs

First let us recall the concept of intermediate representation graphs and its notations. Consider a deep learning network architecture. Such a network is obtained by assembling layers of various types. A layer can be represented by a function $f^\ell : \mathbf{x}^\ell \mapsto \mathbf{x}^{\ell+1}$ where \mathbf{x}^ℓ is the intermediate representation of the input at layer ℓ . Assembling can be achieved in various ways: composition, concatenation, sums, etc so that we obtain a global function f that associates an input tensor \mathbf{x} to an output tensor $\mathbf{y} = f(\mathbf{x})$. In practice a batch of b inputs $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_b\}$ is processed concurrently.

Given a (meaningful) similarity measure sim on tensors, we can define the similarity matrix of the intermediate representations at layer ℓ as:

$$\mathcal{A}_{i,j}^\ell = \text{sim}(\mathbf{x}_i^{\ell+1}, \mathbf{x}_j^{\ell+1}), \forall 1 \leq i, j \leq b, \quad (5.3)$$

where $\mathcal{A}^\ell[i, j]$ denotes the element at line i and column j in \mathcal{A}^ℓ . In our experiments we mostly focus on the use of cosine similarity, which is widely used in computer vision. It is often the case that the output $\mathbf{x}^{\ell+1}$ is obtained right after using a ReLU function, that forces all its values to be nonnegative, so that all values in \mathcal{A}^ℓ are also nonnegative. We then use \mathcal{A}^ℓ to define a weighted graph $\mathcal{G}^\ell = \langle \mathbb{V}, \mathbb{E}^\ell \rangle$, where $\mathbb{V} = \{1, \dots, b\}$ is the set of vertices and \mathbb{E} the set of edges defined with \mathcal{A} .

5.3.1.3 Smoothness of label signals

Given a weighted graph: \mathcal{G}^ℓ , the Laplacian of \mathcal{G}^ℓ is the matrix:

$$\mathbf{L}^\ell = \mathbf{D}^\ell - \mathbf{A}^\ell. \quad (5.4)$$

Consider a graph signal $\mathbf{s} \in \mathbb{R}^b$, we define $\hat{\mathbf{s}}$ the Graph Fourier Transform (GFT) of \mathbf{s} on \mathcal{G}^ℓ as [148]:

$$\hat{\mathbf{s}} = \mathbf{F}^\top \mathbf{s}. \quad (5.5)$$

Assume the order of the eigenvectors is chosen so that the corresponding eigenvalues are in ascending order. If only the first few entries of $\hat{\mathbf{s}}$ are nonzero then \mathbf{s} is said to be low frequency (i.e., smooth) on the graph. In the extreme case where only the first entry of $\hat{\mathbf{s}}$ is nonzero we have that \mathbf{s} is constant (maximum smoothness). Recall that the smoothness $\sigma^\ell(\mathbf{s})$ of a signal \mathbf{s} can be measured using the Laplacian quadratic form:

$$\sigma^\ell(\mathbf{s}) = \mathbf{s}^\top \mathbf{L}^\ell \mathbf{s} = \sum_{i,j=1}^b \mathcal{A}_{i,j}^\ell (\mathbf{s}_i - \mathbf{s}_j)^2. \quad (5.6)$$

In this section, we are particularly interested in smoothness of the label signals. Label signals are also called binary label indicator vector, as we have previously defined in Definition 2.1.2. Recall that when we are dealing with binary signals, the smoothness of the signal is given by the sum of similarities between examples in distinct classes (since $\mathbf{s}_i - \mathbf{s}_j$ is zero when i and j have the same label). Thus, a total smoothness of 0 means that all examples in distinct classes have 0 similarity.

Next we introduce a regularizer that limits how much σ^ℓ can vary from one layer to the next, thus leading to a network that is more inline with Definition 2.5.2. This will be shown later to improve robustness in Section 5.3.2.

5.3.1.4 Proposed regularizer

5.3.1.4.1 Definition: We propose to measure the deformation induced by a given layer ℓ by computing the difference between label signal smoothness before and after the layer for all labels:

$$\delta_\sigma^\ell = \sum_c \left| \sigma^\ell(\mathbf{s}_c) - \sigma^{\ell-1}(\mathbf{s}_c) \right|. \quad (5.7)$$

These quantities are used to regularize modifications made by each of the layers during the learning process. The pseudo-code of Algorithm 5.1 describes how we use the proposed regularizer to compute the loss.

5.3.1.4.2 Illustrative example: In Figure 5.8 we depicted a toy illustrative example to motivate the proposed regularizer. We consider here a one-dimensional two-class problem. To linearly separate circles and crosses, it is necessary to group all circles. Without regularization the resulting embedding is likely to either considerably increase the distance between examples in different classes (case (a)), thus producing sharp transitions in the network function, or to reduce the margin (case (b)). In contrast, by penalizing large variations of the smoothness of label signals (case (c)), the average distance between examples in different classes must be preserved in the embedding domain, resulting in a more precise control of distances within the boundary region.

Remark 2. Since we only consider label signals, we solely depend on the similarities between examples of distinct classes. As such, the regularizer only focuses on the boundary, and does not vary if the distance between examples of the same label grows or shrinks.

Remark 3. Compared with [21], there are key differences that characterize the proposed regularizer:

1. Only pairwise distances between examples are taken into account. This has the effect of controlling space deformations only in the directions of training examples;
2. The network is forced to maintain a minimum margin by keeping the smoothness small at each layer of the architecture, thus controlling both contraction and dilatation of space at the boundary. This is illustrated in Figure 5.8, where [21] is represented by b) and our method by c);
3. The proposed criterion is an average (sum) over all distances, rather than a stricter criterion (e.g. maintaining a small Lipschitz constant), which would force each pair of vectors $(\mathbf{x}_i, \mathbf{x}_j)$ to obey the constraint.

In summary, by enforcing small variations of smoothness across the layers of the network, the proposed regularizer maintains a large enough r so that Equation (2.16) can hold, while also controlling dilatation. Combining it with Parseval [21] would allow for a better control of the α parameter in the other directions of the input space.

5.3.2 Experiments

In the following subsections we evaluate the proposed method using various tests. We use the well known CIFAR-10 dataset [81] as a first benchmark and we demonstrate that our proposed regularizer can improve robustness as defined in Section 2.5.

Algorithm 5.1 Loss function of the regularized network

Inputs:
 \mathbf{x} : list of all the representations of the network.

 ReLUs, the list containing the positions of all the ReLU activations on f .

 \mathbf{y} , the output of the network

 \mathbf{s} , the label signal of the batch, i.e., the ground truth labels of the examples of the batch

 m , the power of the Laplacian for which we wish to compute the smoothness;

 γ , the scaling coefficient of the regularizer loss.

procedure LOSS(\mathbf{x} , ReLUs, \mathbf{y} , \mathbf{s} , m , γ)

for $\ell \in \text{ReLUs}$ **do**
 $\sigma^\ell \leftarrow \text{Smoothness}(\mathbf{x}^\ell, \mathbf{s}, m)$
 $\Delta \leftarrow \frac{\sum_{\ell \in \text{ReLUs}} |\sigma^\ell - \sigma^{\ell-1}|}{|\text{ReLUs}| - 1}$
return CategoricalCrossEntropy(\mathbf{s} , \mathbf{y}) + $\gamma\Delta$
procedure SMOOTHNESS(\mathbf{x}^ℓ , \mathbf{s} , m)

 $\mathcal{A}^\ell \leftarrow$ Pairwise similarity of \mathbf{x}^ℓ (we use cosine similarity in our work)

 $\mathbf{D}^\ell \leftarrow$ Diagonal degree matrix of \mathcal{A}^ℓ
 $\mathbf{L}^\ell \leftarrow \mathbf{D}^\ell - \mathbf{M}^\ell$
 $\sigma^\ell \leftarrow \text{Trace}(\mathbf{s}^\top (\mathbf{L}^\ell)^m \mathbf{s})$
return σ^ℓ

In summary, in Section 5.3.2.1 we first verify that the proposed regularizer favors Definition 2.5.2. We then show in Section 5.3.2.2 that by using the proposed regularizer we are able to increase robustness for random perturbations and weak adversarial attacks. In Section 5.3.2.3, we challenge our method on more competitive benchmarks. Finally, in Section 5.3.2.4 we extend the analysis to CIFAR-100 [81] and Imagenet32x32 [19] to validate the generality of the method. These experiments demonstrate that DNNs trained with the proposed regularizer lead to improved robustness.

To measure accuracy, we average over 10 runs each time, unless mentioned otherwise. In all reports, P stands for Parseval [21] trained networks, R for networks trained with the proposed regularizer and V for vanilla (i.e. baseline) networks. The corresponding code is available at https://github.com/cadurosar/laplacian_networks.

5.3.2.1 Robustness of trained architectures

First we verify that the proposed regularizer improves robustness as defined in Section 2.5. For various values of r , we estimate $\alpha_{\min}(r) = \arg \min_{\alpha} \{f \in \text{Robust}_{\alpha}(r)\}$. We use 1000 training examples and generate 100 uniform noises to estimate $\alpha_{\min}(\cdot)$. Results are shown in Figure 5.9. We observe that networks trained with the proposed regularizer allow for smaller α values when the radius r increases. The Parseval method achieves better (smaller) Lipschitz constant than Vanilla, as suggested by the large values of r . However, we observe that α_{\min} grows fast when using Parseval, suggesting that sharp transitions are allowed in the vicinity of trained examples.

5.3.2.2 Experiments on perturbations and adversarial attacks

In this subsection we verify the ability of the proposed regularizer to increase robustness, while retaining acceptable accuracy on the clean test set, on the CIFAR-10 dataset without any type of data augmentation.

5.3.2.2.1 Clean test set Before checking the robustness of the network, we first test the performance on clean examples. In the second column of Table 5.3, we show the baseline accuracy of the models on the clean CIFAR-10 test set (no perturbation is added at this point). These experiments agree with the claim from [21] where the authors show that they are able to increase the performance of the network on the clean test set. We observe that the proposed method leads to a minor decrease of performance on this test. However, we see in the following experiments that this is compensated by an

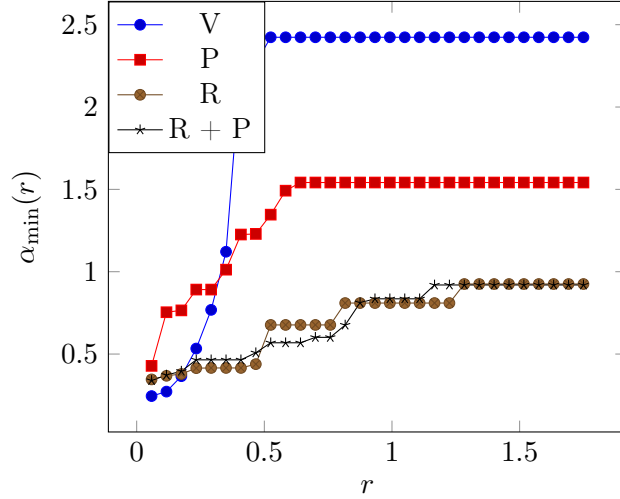


Figure 5.9: Estimations of $\alpha_{\min}(r)$ obtained for different radius r over training examples. The proposed regularizer allows for smaller α values when r increases. Figure and caption extracted from [90].

increased robustness to perturbations. Such a trade-off between robustness and accuracy has already been discussed in the literature [32].

Table 5.3: Network mean Relative Error Inflation (mREI) under different types of perturbation. Bottom line represents the corresponding median Cosine Distance (mCD) (at the highest perturbation severity) between corrupted and clean images. Table and caption extracted from [90].

Network	Clean set $mREI$		Noise			Blur				Weather				Digital			
			Gauss.	Shot	Impulse	Defocus	Glass	Motion	Zoom	Snow	Frost	Fog	Bright	Contrast	Elastic	Pixel	JPEG
Vanilla (V)	11.9%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Parseval (P)	10.3%	0.29	0.71	0.48	0.57	0.10	1.01	0.15	0.11	0.16	0.17	-0.02	0.04	0.13	0.14	0.25	0.28
Regularizer (R)	13.2%	-0.29	-1.12	-0.86	0.10	-0.03	-0.65	-0.09	-0.19	-0.30	-0.61	0.17	0.01	0.40	-0.15	-0.50	-0.50
P and R	12.8%	-0.35	-1.33	-1.00	0.05	-0.09	-0.75	-0.18	-0.31	-0.41	-0.67	0.13	0.04	0.48	-0.17	-0.47	-0.53
mCD 10^{-3}			18	16	37	5	24	15	17	15	20	51	14	57	14	6	3

5.3.2.2.2 Perturbation robustness In order to assess the effectiveness of the various methods when subject to perturbations, we use the benchmark proposed in [53], and previously described in Section 2.5.4.2. The benchmark consists of 15 different perturbations, with 5 levels of severity each (note that they are referred to as “corruptions” in [53]). Perturbations test the robustness of the network to noise when compared to its clean test set performance.

In more details, we are interested in the mean Relative Error Inflation (mREI).

To define it, consider $E_{\text{net}}^{\text{per},\text{sev}}$ the error rate of a network **net** (V,P,R or P+R), under perturbation type **per** and severity **sev**. Denote E_{net} the error rate of the network **net** on the clean set. We first define Error Inflation (EI) as:

$$EI_{\text{net}}^{\text{per},\text{sev}} = \frac{E_{\text{net}}^{\text{per},\text{sev}}}{E_{\text{net}}}.$$

Then the Relative Error Inflation REI is defined as:

$$REI_{\text{net}}^{\text{per},\text{sev}} = EI_{\text{net}}^{\text{per},\text{sev}} - EI_V^{\text{per},\text{sev}}.$$

Finally, mREI is obtained by averaging over all severities. Note that this is different from the traditional MCE metric, but we believe that this metric is more inline with our objective here.

The results are described in 5.3 for the CIFAR-10 dataset. The raw error rates under each type of perturbations can be found in the original paper. We observe that Parseval alone is not able to help with the mREI, despite reducing the clean set error. On the other hand, the proposed regularizer and its combination with Parseval training decreases the clean set accuracy but increases the relative performance under perturbations by a significant amount.

This experiment supports the fact that the proposed regularizer can significantly improve robustness to most types of perturbations introduced in [53]. It is worth pointing out that this finding does not hold for Impulse Noise, Fog, and Contrast. Looking more into details, we observe that Impulse noise shifts some values on the image to either its maximum possible value or the minimum possible value, while Fog and Contrast perform a re-normalization of the image. In those cases perturbations have the effect of creating noisy inputs that are far away (in terms of the cosine distance) from the original images, as supported by the last line of the table. This is in contrast to the other types of perturbations in the experiment. Because they can be far away, these perturbations do not fulfill Definition 2.5.2, where there is a maximum radius r for which robustness is enforced around the examples. In other words, our robustness definition is focusing on small deviations/distances as those are more likely to characterize noise (i.e., we focus on distances that are too small to change the class of the image).

5.3.2.2.3 Adversarial Robustness We next evaluate robustness to adversarial inputs, which are specifically built to fool the network function. Such adversarial inputs can be generated and evaluated in multiple ways. Here we implement three approaches: i) a mean case of adversarial noise, where the adversary can only use one forward and one backward pass to generate the perturbations, ii) a worst case scenario, where the adversary

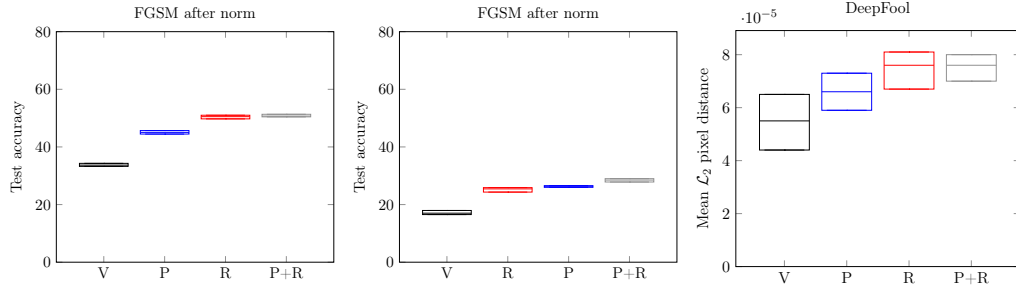


Figure 5.10: Robustness against an adversary measured by the test set accuracy under FGSM attack in the left and center plots and by the mean \mathcal{L}_2 pixel distance needed to fool the network using DeepFool on the right plot. Figure and caption extracted from [90].

can use multiple forward and backward passes to try to find the smallest perturbation that will fool the network, and iii) a compromise between the mean case and the worst case, where the adversary can do a predefined number of forward and backward passes with a perturbation threshold limit.

For the first approach, we add the scaled gradient sign (FGSM attack) to the input [84], so that we obtain a target SNR of 33. This is inline with previous works [21]. Obtained results are introduced in the left and center plots of Figure 5.10. In the left plot the noise is added after normalizing the input, whereas on the middle plot it is added before normalizing it. As with the perturbation tests, a combination of the Parseval method and our proposed approach yields the most robust architecture.

In regards to the second approach, where a worst case scenario is considered, we use the Foolbox [130] implementation of DeepFool [111]. Due to time constraints we sample only $\frac{1}{10}$ of the test set images for this test. The conclusions we can draw are similar (right plot of Figure 5.10) to those obtained for the first adversarial attack approach. Finally, for the third approach we use the PGD (Projected Gradient Descent) attack introduced in [104]. PGD is an iterative version of FGSM, which loops for a maximum number of it iterations. For each iteration it moves by a distance of $step$ in the direction of the gradient, provided it does not move away from the original image by a distance greater than ϵ . Our experiments, described in Table 5.4, show that the proposed regularizer increases robustness against a PGD attack, for an epsilon corresponding to an SNR of about 33 ($it = 20, step = 0.002, \epsilon = 0.01$).

A common pitfall in evaluating robustness to adversarial attacks comes from the fact the gradient of the architecture can be masked due to the introduced method. As a consequence, generated attacks become weaker compared to those on the vanilla architecture. So, to further verify that the obtained results are not only due to gradient

Table 5.4: Median test set accuracy on the CIFAR-10 dataset against the PGD attack. Table and caption extracted from [90].

Model	PGD Accuracy
V	1.18%
P	1.72%
R	5.2%
P+R	5.6%

masking, we perform tests with black box FGSM, where the target attacked network is not the same as the source of the adversarial noise. This way, all networks are tested against the same attacks.

For this test we continue to use an SNR of about 33 with the FGSM method. We choose the network with the best performance for each of the tested methods. The results are depicted in Table 5.5. In our experiments, we found that the combination of our method with Parseval is the most robust to noise coming from other sources. This demonstrates that the improvements are not caused by gradient masking, but are caused by the increased robustness of the proposed method and Parseval’s. Interestingly, the noise created by both Parseval and our method did not challenge the other methods as well as the one created by Vanilla, justifying a posteriori the interest of this experiment.

Table 5.5: Comparison of CIFAR-10 test set accuracy under the black box FGSM attack. The most robust target for a given source is bolded, while the strongest source for a target is in italic. Table and caption extracted from [90].

Target	Source			
	V	P	R	P+R
V	X	<i>60.74</i>	61.49	72.51
P	<i>57.82</i>	X	68.21	73.87
R	<i>69.72</i>	74.96	X	73.56
P+R	75.35	76.11	<i>70.22</i>	X

5.3.2.2.4 Robustness to parameter and activation noises In a third series of experiments we aim at evaluating the robustness of the architecture to noise on parameters and activations. We consider two types of noises: i) erasures of the memory (dropout), and ii) quantization of the weights [60].

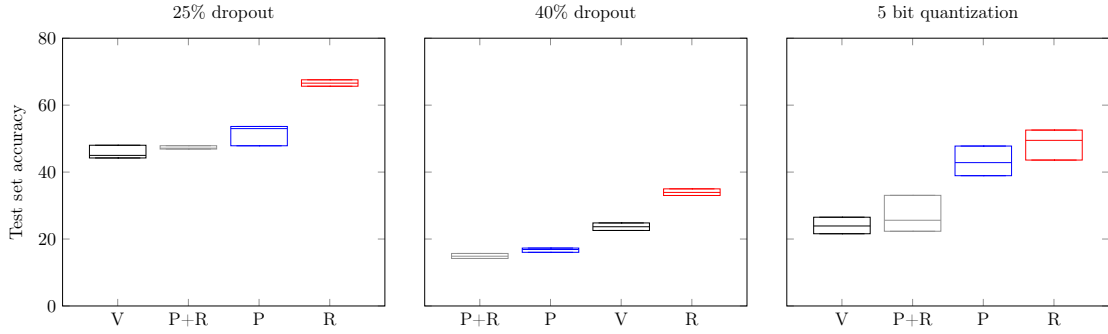


Figure 5.11: CIFAR-10 test set accuracy under different types of implementation related noise. Figure and caption extracted from [90].

In the dropout case, we compute the test set accuracy when the network has a probability of either 25% or 40% of dropping an intermediate representation value after each block of computation in the architecture. We average over a run of 40 experiments. Results are depicted in the left and center plots of Figure 5.11. It is interesting to note that the Parseval trained functions seem to collapse as soon as we reach 40% probability of dropout, providing an average accuracy smaller than the vanilla networks. In contrast, the proposed method is the most robust to these perturbations.

For the quantization of the weights, we aim at compressing the network size in memory by a factor of 6. We therefore quantize the weights using 5 bits (instead of 32) and re-evaluate the test set accuracy. The right plot of Figure 5.11 shows that the proposed method is providing a better robustness to this perturbation than the tested counterparts.

Overall, these experiments confirm previous ones in the conclusion that the proposed regularizer obtains the best robustness compared to Parseval and Vanilla architectures.

5.3.2.3 Experiments on challenging benchmarks

In this subsection we verify the ability of the proposed regularizer to increase robustness on the CIFAR-10 dataset while being combined with recent techniques of adversarial data augmentation. This is important as those methods are seen as the state of the art for adversarial robustness. We recall that adversarial data augmentation consists in augmenting the training set during the training stage by using the same kind of attacks as those described in the last subsection. We refer to techniques using adversarial data augmentation using the letter A.

5.3.2.3.1 Tests with FGSM adversarial data augmentation We first perform experiments with adversarial data augmentation as suggested in [84]. To be more precise we use the method they advise which is called “step1.1” using $\epsilon = \frac{8}{255}$.

A first test consists in measuring the accuracy of these methods when the test set inputs are modified with additive Gaussian noise with various SNRs. As expected, we observe in Figure 5.12 that training with adversarial examples helps in this case, as it adds more variation to the training set. Yet it reduces the accuracy on the clean set (left plot). Note that combining our method with adversarial training results in the best median accuracy.

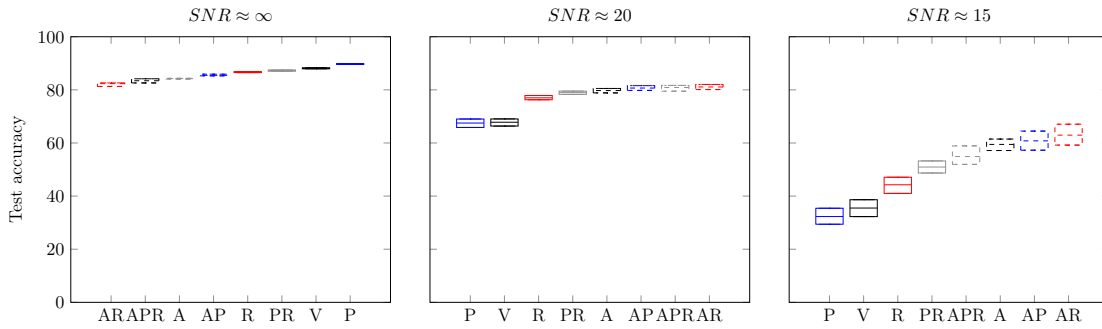


Figure 5.12: Test set accuracy under Gaussian noise with varying Signal-to-Noise Ratio (SNR). Figure and caption extracted from [90].

About robustness to adversarial attacks, the obtained results are depicted in Figure 5.13. We observe that adding FGSM adversarial training does not generalize well to other types of attack (which is readily seen in the literature [104]). Overall, the models using the proposed regularizer are the most robust again.

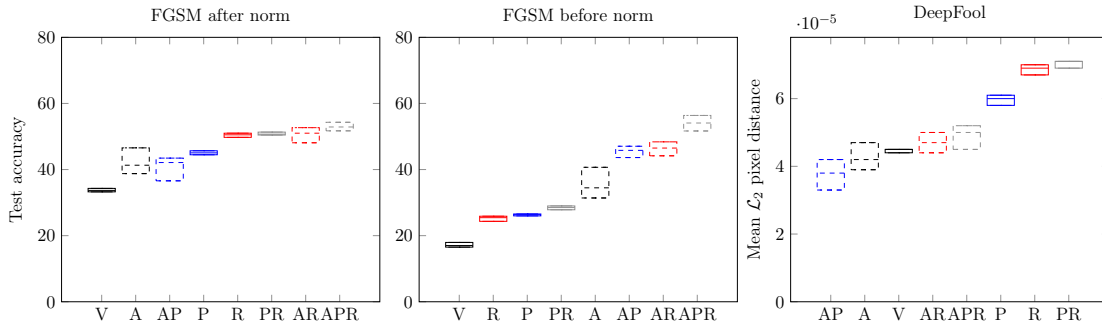


Figure 5.13: Robustness against an adversary measured by the test set accuracy under FGSM attack in the left and center plots and by the mean \mathcal{L}_2 pixel distance needed to fool the network using DeepFool on the right plot. Figure and caption extracted from [90].

Finally, when considering implementation related perturbations, the results depicted

in Figure 5.14 are consistent with the ones from the previous section, in which is shown that the proposed regularizer helps improving robustness to this type of noise.

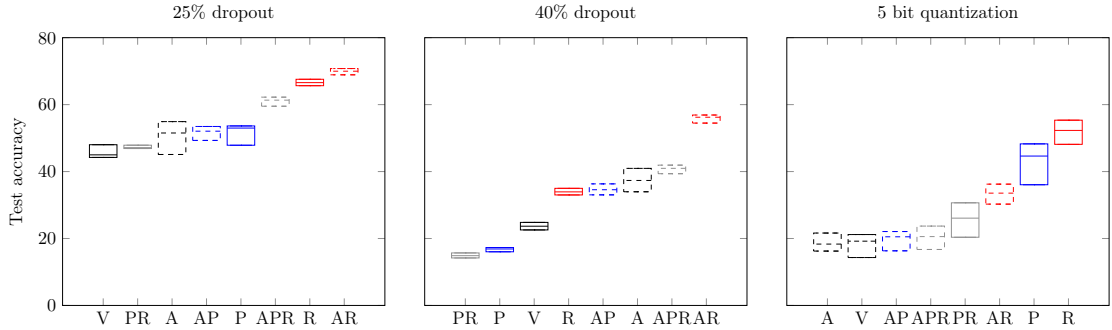


Figure 5.14: Test set accuracy under different types of implementation related noise. Figure and caption extracted from [90].

In summary, even when adding adversarial training, the proposed regularizer is either the most robust in median, or capable of improving the robustness when combined with the other methods.

5.3.2.3.2 Tests with PGD adversarial data augmentation Most of our adversarial tests are performed with FGSM because of its simplicity and speed, even though it has already been shown (e.g: [104]) that FGSM is weak as an attack and as a defense mechanism. Despite the fact we do not only target adversarial defense, we further stress the ability of the proposed regularizer to improve it and to combine with other methods. To this end we perform experiments against the PGD (Projected Gradient Descent) attack.

As the proposed regularizer can be combined with FGSM defense, it is natural to also test it alongside PGD training. We use the parameters advised in [104]: 7 iterations with $step = 2/255$, and $\epsilon = 8/255$. The results depicted in Table 5.6 show that using our regularizer increases robustness of networks trained with PGD. Note that Dropout and Gaussian Noise were applied ten times to each of the networks and the results are displayed as the mean test set accuracy under these perturbations. A rate of 40% was used for dropout. The PGD attack uses the following parameters: $it = 20$, $step = \frac{2}{255}$, $\epsilon = \frac{8}{255}$.

5.3.2.4 Experiments with other datasets

In this final subsection, we test the generality of the method using the CIFAR-100 and ImageNet32x32 datasets, with a subset of the perturbations used for CIFAR-10. Gaussian

Table 5.6: Test set accuracy results on the CIFAR-10 dataset with PGD training. Table and caption extracted from [90].

	Clean	Gaussian	PGD	Dropout
A	76.39%	71.25%	32.78%	35.20%
A + R	76.36%	72.26%	33.72%	55.63%

Noise is applied ten times to each of the networks for a total of 30 different runs. A SNR of 33 is used for FGSM and 15 for Gaussian Noise. Images are normalized in the same way as the experiments with CIFAR-10. Standard data augmentation is used for CIFAR-100.

Results on CIFAR-100 are shown in Table 5.7 as the mean over three different initializations. We observe that as it was the case on CIFAR-10, the proposed method and the combination of the methods is the most robust on these test cases.

Table 5.7: Test set accuracy results on the CIFAR-100 dataset. Table and caption extracted from [90].

Model	Clean Set	Gaussian Noise	FGSM
Vanilla (V)	78.7%	12.6%	20.5%
Parseval (P)	80.1%	14.8%	22.0%
Regularizer (R)	79.4%	15.9%	23.0%
P+R	79.5%	19.1%	24.4%

We then use Imagenet32x32, a downscaled version of Imagenet [19] which can be used as an alternative to CIFAR-10 while maintaining a similar computational budget [19]. We use the same network and training hyperparameters of the original paper. Gaussian Noise and Dropout are applied 40 times to each of the networks. Gaussian noise is applied with SNR=33 whereas Dropout is applied with 15%.

Results are shown in Table 5.8. We observe that as it was the case on CIFAR-10 and CIFAR-100, the proposed method provides more robustness in all of these test cases. Note that we had trouble fine-tuning the β parameter for the Parseval criterion, explaining the poor performance of Parseval and its combination with our proposed regularizer.

In this Section we have introduced a definition of robustness alongside an associated regularizer. The former takes into account both small variations around the training set examples and the margin. The latter enforces small variations of the smoothness of label

Table 5.8: Test set accuracy results on the Imagenet32x32 dataset. Table and caption extracted from [90].

Model	Clean	Gaussian Noise	Dropout
Vanilla (V)	52.1%	36.8%	2.3%
Parseval (P)	48.1%	34.10%	3.71%
Regularizer (R)	52.4%	37.4%	7.0%
P+R	43.80%	29.87%	5.0%

signals on similarity graphs obtained at intermediate layers of a deep learning network architecture. We have empirically shown with our tests that the proposed regularizer can lead to improved robustness in various conditions compared to existing counterparts. We also demonstrated that combining the proposed regularizer with existing methods can result in even better robustness for some conditions. Future work includes a more systematic study of the effectiveness of the method with regards to other datasets, models and perturbations. Recent works shown adversarial noise is partially transferable between models and dataset and therefore we are confident about the generality of the method in terms of models and datasets.

5.4 Using intermediate representation graphs to compress DNNs

In the previous sections we have shown the interest of using the concepts of GSP in order to analyze and improve DNNs. In this section we present an introductory work that specializes a knowledge distillation framework, called Relational Knowledge Distillation (RKD), to the graph domain. We call this new framework Graph Knowledge Distillation (GKD). In other words, we present a technique that allows us to compress DNNs by using graphs to represent the intermediate spaces of neural networks. We presented this work in a recent contribution [88] and note two works from the same period that proposed similar ideas [96, 103].

As we have previously discussed in Section 2.4, the success of DNNs is heavily linked to the availability of large amounts of data and special purpose hardware, e.g., graphics processing units (GPUs) allowing significant levels of parallelism. However, this need for a significant amount of computation is a limitation in the context of embedded systems, where energy and memory are constrained. As a result, numerous recent works have focused on compressing deep learning architectures, some of them using the distillation

technique.

In a quick recall from Section 2.4.1, one approach to distillation is performing Individual Knowledge Distillation (IKD) [4, 55, 133]. Initial IKD techniques [55] focused on using the output representations of the teacher as a target for the smaller architecture, while more recent works have reached better accuracy by performing this process layer-wise, or block-wise for complex architectures [79, 133]. However, IKD can be directly performed layer-wise only if the student and the teacher have inner data representations with the same dimension [79], or if transformations are added [133].

In an effort to allow distillation to be performed layer-wise on architectures with varying dimensions, recent works [117] have introduced distillation in a dimension-agnostic manner. To do so, these methods focus on the relative distances of the intermediate representations of training examples, rather than on the exact positions of each example in their corresponding domains. These methods are referred to as relational knowledge distillation (RKD) in the literature.

In this section, we present our work in which we extend this notion of RKD by introducing *graph knowledge distillation* (GKD). As in the previous sections of this chapter, we construct graphs where vertices represent training examples, and the edge weight between two vertices is a function of the similarity between the representations of the corresponding examples at a given layer of the network architecture. The main motivation for this choice is that even though representations generally have different dimensions in each architecture, the size of the corresponding graphs is always the same (since the number of nodes is equal to the number of training examples). Thus, information from graphs generated from the teacher architecture can be used to train the student architecture by introducing a discrepancy loss between their respective adjacency matrices during training.

In other words, we introduce a layer-wise distillation process using graphs, extending the RKD framework, and we demonstrate that this method can improve the accuracy of students trained in the context of distillation, using standard vision benchmarks. The reported gains are about twice as important as those obtained by using standard RKD instead of no distillation.

5.4.1 Methodology

In this section we first introduce RKD and recall some of the notations from Section 2.4.1, then we introduce the methodology used to define GKD.

5.4.1.1 Relational Knowledge Distillation (RKD)

Recall that T and S denote teacher and student architectures, respectively. The goal of distillation is to transfer knowledge from T to S , where S typically contains fewer parameters than T . For presentation simplicity, we assume that both architectures generate the same number of inner representations. In the context of distillation, we consider that the teacher has already been trained, and that we want to use both the training set and the inner representations of the teacher in order to train the student. This is an alternative to directly training the student using only the training data (which we refer to as “baseline” in our experiments). Also recall, that we use the following loss to train the student:

$$\mathcal{L} = \mathcal{L}_{\text{task}} + \lambda_{\text{KD}} \cdot \mathcal{L}_{\text{KD}} . \quad (5.8)$$

We denote $\mathbf{X} \in \mathcal{D}_{\text{train}}$ a batch of input examples and \mathbb{X}' the set of intermediate representations generated using \mathbf{X} that are used for inferring knowledge. RKD approaches consider relative metrics between the respective inner representations of the networks to be compared. In the specific case of RKD-D [117], the mathematical formulation is:

$$\mathcal{L}_{\text{RKD-D}} = \sum_{X' \in \mathbb{X}'} \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \dot{X}'} \mathcal{L}_d \left(\frac{\|\mathbf{x}_i^S - \mathbf{x}_j^S\|_2}{\Delta'^S}, \frac{\|\mathbf{x}_i^T - \mathbf{x}_j^T\|_2}{\Delta'^T} \right), \quad (5.9)$$

where \dot{X}' is the set of all possible pairs from X' , Δ'^A is the average distance between all couples $(\mathbf{x}_i^A, \mathbf{x}_j^A)$ for each $X' \in \mathbb{X}'$ for the given architecture, and \mathcal{L}_d is the Huber loss [62]. The main advantage of using RKD is that it allows to distillate knowledge from an inner representation of the teacher to one of the student, even if their respective dimensions are different.

5.4.1.2 Proposed Approach: Graph Knowledge Distillation (GKD)

We now introduce our proposed approach. Instead of directly trying to make the distances between data points in the student match those of the teacher, we consider the problem from a graph perspective. Given an architecture A , a batch of inputs X , we compute the corresponding inner representations $X'^A = f'^A([\mathbf{x}, \mathbf{x} \in X])$. We can then choose a set of layers that we want to consider and create a set \mathbb{X}' containing these intermediate representations. These representations are then used to define a similarity graph $\mathcal{G}^A(X')$, for each $X' \in \mathbb{X}'$. The graph contains a node for each input in the batch, and the edge weight $\mathcal{A}^A(X')_{i,j}$ represents the similarity between the i -th and the j -th elements of X' from architecture A . In this work, we use the cosine similarity. Finally, in order to control the importance of outliers, we also normalize the adjacency matrix.

While training the student, we input our training batch into both the student architecture and the (now fixed) previously trained teacher architecture. This provides a similarity graph for each representation X' from the set of representations to consider \mathbb{X}' . The loss we aim to minimize combines the task loss, as expressed in Equation 5.8, with the following graph knowledge distillation (GKD) loss:

$$\mathcal{L}_{\text{GKD}} = \sum_{X' \in \mathbb{X}'} \mathcal{L}_d(\mathcal{G}^S(X'), \mathcal{G}^T(X')) . \quad (5.10)$$

In our work, we mainly consider the case where \mathcal{L}_d is the Frobenius norm between the adjacency matrices. The GKD loss measures the discrepancy between the adjacency matrices of teacher and student graphs. In this way the geometry of the latent representations of the student will be forced to converge to that of the teacher. Our intuition is that since the teacher network is expected to generalize well to the test, mimicking its latent representation geometry should allow for better generalization of the student network as well. An equivalent definition of our proposed loss is:

$$\mathcal{L}_{\text{GKD}} = \sum_{X' \in \mathbb{X}'} \|\mathcal{A}^S(X') - \mathcal{A}^T(X')\|_2^2 . \quad (5.11)$$

A first obvious advantage of GKD with respect to RKD-D is the fact it has a more natural normalization over the batch of inputs, yielding to a more robust process. This is discussed in Section 5.4.2.3. Amongst other degrees of freedom that become available when using graphs, we focus on three possible variations of the method:

1. Task specific: considering only examples of the same (resp. distinct) classes when creating the edges of the graph, thus focusing on the clustering (resp. margin) of classes,
2. Localized: weighting differently the closest and furthest neighbors of each node in the graph, in an effort to focus on locality, or to the contrary on remoteness,
3. Smoothed: taking powers p of the normalized adjacency matrix of considered graphs before computing the loss. By considering higher powers of \mathcal{A} , we consider smoothed relations between inner representations of inputs.

5.4.2 Experiments

In this section, we perform two types of experiments. First we compare the accuracy of RKD-D and GKD using the CIFAR-10 and CIFAR-100 datasets [81], analyze the impact of the normalization of the similarities, compare the consistency with the teacher and

perform spectral analysis of the different graphs and graph signals. We then look at proposed variations of GKD: task specific, localized and smoothed.

5.4.2.1 Hyperparameters

We train our CIFAR-10/100 networks for 200 epochs, using standard Stochastic Gradient Descent (SGD) with batches of size 128 ($|X| = 128$) and an initial learning rate of 0.1 that is decayed by a factor of 0.2 at epochs 60, 120 and 160. We also add a momentum of 0.9 and follow standard data augmentation procedure. We use a ResNet26-1 architecture for our teacher network, while the student network uses a Resnet26-0.5. In terms of scale, ResNet26-0.5 has approximately 27% of the operations and parameters of ResNet26-1. All these architectures are particularly small compared to the ones achieving state-of-the-art performance. We use a network of same size of the students but trained without a teacher as a baseline that we call Vanilla. Our RKD-D [117] students are trained with the parameters from [117], $\lambda_{\text{RKD-D}} = 25$ and applied to the output of each block. We applied the same values for GKD. Note that all these choices were made to remain as consistent as possible with existing literature. For each student network we run either 10 (CIFAR-10) or 3 (CIFAR-100) tests and report the median value. The code for reproducing the experiments and boxplots for each experiment is available at https://github.com/cadurosar/graph_kd.

5.4.2.2 Direct comparison between GKD and RKD-D

In a first experiment we simply evaluate the test set error rate when performing distillation. Results are summarized in Table 5.9. We compare student sized networks trained without distillation, that we call baseline, with GKD and RKD-D [117] trained networks. We also report the performance of the teacher. We note that RKD-D [117] by itself provides a small gain in error rate with respect to the Baseline approach, while GKD outperforms RKD-D by almost the same gain.

5.4.2.3 Effect of the normalization

To better understand why GKD performed better than RKD-D we analyze the contribution of each example in a batch in both the GKD loss and the RKD-D one. If our premise from Section 5.4.1.2 is correct, by using a degree normalized adjacency matrix instead of the distance pairs directly, most examples will be able to contribute to the optimization. To do so, we compute the respective loss, for each block, using 50 batches of 1000 training

Table 5.9: Error rate comparison of GKD and RKD-D. Table and caption extracted from [88] ©2020 IEEE.

Method	CIFAR-10	CIFAR-100	Relative size
Teacher	7.27% (± 0.26)	31.26%	100%
Baseline	10.34% (± 0.27)	38.50%	27%
RKD-D [117]	10.05% (± 0.28)	38.26%	27%
GKD	9.71% (± 0.27)	38.17%	27%

set examples and analyze the median amount of examples that are responsible for 90% of the loss at each block. In Table 5.10, we present the results. As we suspected for GKD, it shows a significant advantage on the number of examples responsible for 90% of the loss.

Table 5.10: Comparison of the effect of the normalization on the amount of examples that it takes to achieve 90% of the total loss value. Table and caption extracted from [88] ©2020 IEEE.

Block position in the architecture	RKD-D	GKD
Middle	83.70%	86.50%
Final	82.05%	83.60%

5.4.2.4 Classification consistency

We now take our trained students and compare their outputs to the trained teacher’s outputs. For the output of each WideResNet block we compute the classification of a simple Logistic Regression, while the network’s final output is already a classifier. The ideal scenario would be one where the student is 100% consistent with the teacher’s decision on the test set, as this would greatly improve the classification performance when compared to the baseline. The results are depicted in Figure 5.15. As expected the GKD was able to be more consistent with the teacher than the RKD-D.

5.4.2.5 Spectral analysis

Given that we have introduced intermediate representation graphs, it is quite natural to analyze performance from a GSP perspective [148]. We propose to do so by considering specific graph signals \mathbf{s} and computing their respective smoothness on each of the two graphs. We create graphs with 1000 examples chosen at random from the training set. The signals that we consider are i) the label binary indicator signal, and ii) the Fiedler

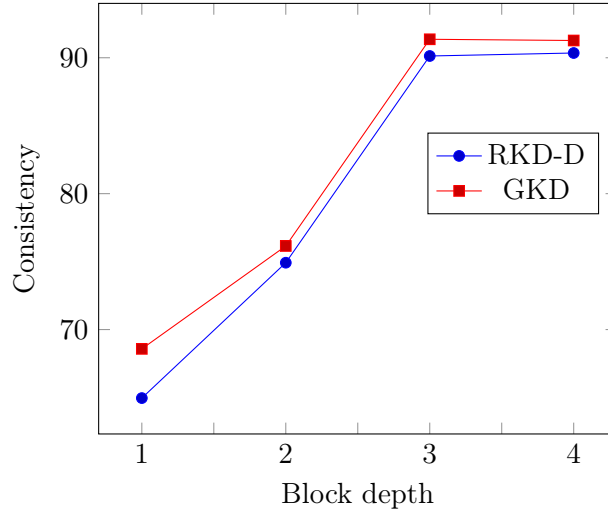


Figure 5.15: Analysis of the consistency of classification compared to the teacher, across blocks of RKD-D and GKD students. We consider the classification layer of the network as the “fourth block”. Figure and caption adapted from [88] ©2020 IEEE.

eigenvectors from each intermediate representation in the teacher, which allow us to compare the clustering of both networks and how they evolve over successive blocks. The results are depicted in Figure 5.16. We can see that both signals have more smoothness in the graphs generated by GKD. This means that the geometry of the latent spaces from GKD are more aligned to those of the teacher.

5.4.2.6 Task specific graph signals

We now consider variations of the proposed GKD method. The first one are the effects of considering only intra or inter-class distances. If we consider only inter-class distances we can focus mostly on having a similar margin in both teacher and student. On the other hand, considering only intra-class distances would force both networks to perform the same type of clustering on the classes. The results are presented in Figure 5.17. In this case, focusing on the margin helped decrease both median test error rate and its standard deviation, while concentrating on the clustering was not effective. This result is similar to what we found in our prior work (Section 5.3), which shows that the margin is a better tool to interpret the network results than the class clustering.

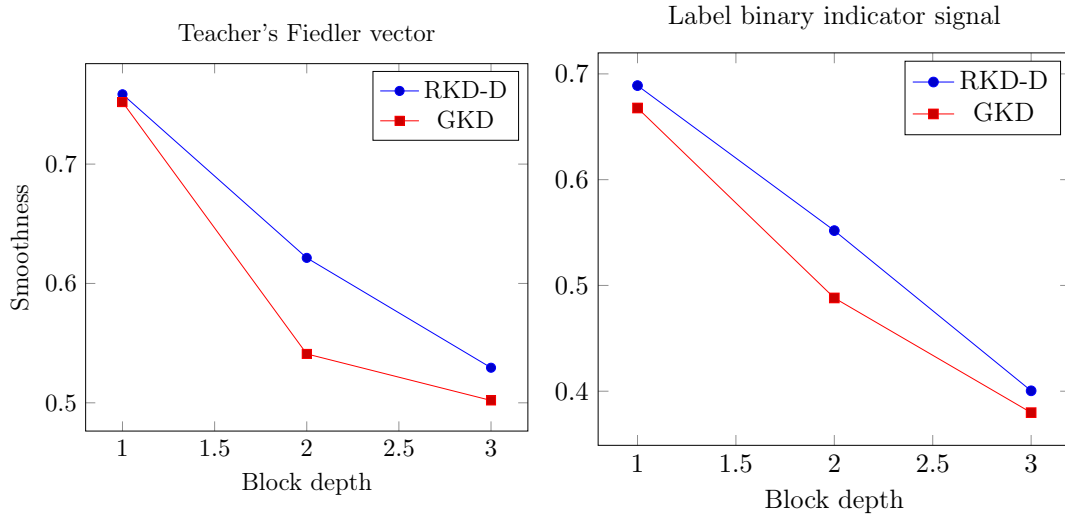


Figure 5.16: Analysis of the smoothness evolution across blocks of the students. In the right we have the label binary indicator signal and in the left we use the Teacher's Fiedler vector as a signal. Figure and caption extracted from [88] ©2020 IEEE.

5.4.2.7 Effect of locality

To study the effect of locality, we partition the graph edges in two parts: 1) the ones corresponding to the k nearest neighbors of each vertex $\mathcal{A}^A(k)$ and 2) the other ones $\overline{\mathcal{A}^A(k)}$. Consequently, we can write $\mathcal{A}^A = \mathcal{A}^A(k) + \overline{\mathcal{A}^A(k)}$. We then introduce the new adjacency matrix $\mathcal{A}^A(k, \alpha) = \alpha \mathcal{A}^A(k) + (1 - \alpha) \overline{\mathcal{A}^A(k)}$, where α scales the importance of 1) with respect to 2). So choosing $\alpha = 0$ means to disregard nearest neighbors while $\alpha = 1$ corresponds to focusing only on them. Results are summarized in Figure 5.18. We observe that for small value of k , small values of α lead to the best performance, whereas for $k = |X|/2$ larger values of α are better. This is similar to results such as [54], where the authors show that one should not concentrate on the hardest/easiest cases, but on the intermediate cases.

5.4.2.8 Smoothed representations

Finally, we study the effect of varying the power of adjacency matrices p . This allows us to consider smoothed relations between inner representation of inputs when compared to fixing p to 1. The results are presented in Table 5.11. Smoothed relations do not seem to help the transfer of knowledge. One possible reason is that larger powers have the effect of drowning out the information.

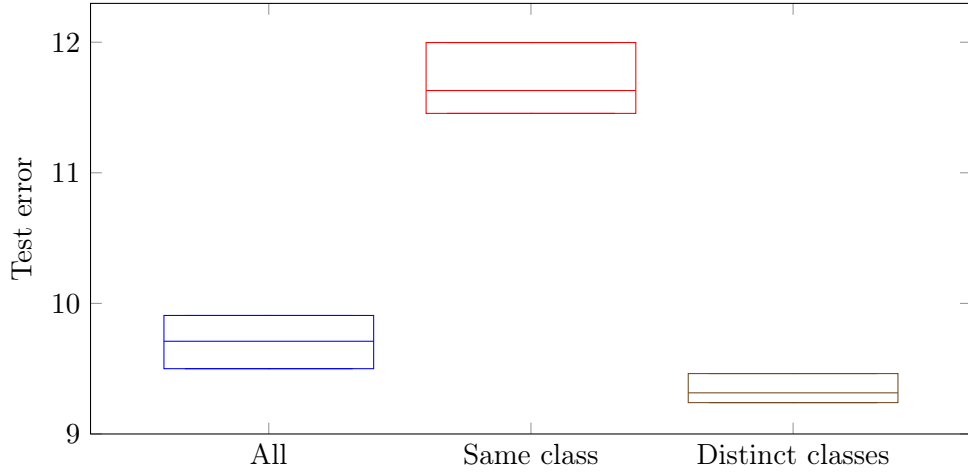


Figure 5.17: Analysis of the effect of task specific graph signals. Figure and caption extracted from [88] @2020 IEEE.

Table 5.11: Analysis of the effect of varying p on the error rate. Table and caption extracted from [88] @2020 IEEE.

p	1	2	3
Error Rate	9.71%	10.20%	10.07%

In the previous paragraphs we have introduced graph knowledge distillation (GKD), a method using graphs to transfer knowledge from a teacher architecture to a student one. By using graphs, the method opens the way to numerous variations that can significantly benefit the accuracy of the student, as demonstrated by our experiments. We note that we are not the first to propose such an extension, but that nonetheless this is an interesting research direction. In future work we consider: i) using more appropriate graph distances, such as in [16, 142], ii) doing a more in-depth exploration of how to properly scale the student network, e.g. following [157], and iii) combining with approaches such as [10, 54] to train a teacher network in a layer-wise fashion.

5.5 Summary of the chapter

Differently from the previous ones, in this chapter we have mainly presented our contributions in the domain of “Deep Neural Networks latent spaces supported on graphs”. While this domain is not very developed, we hope that our contributions may shine a light and allow for more development on it, as we believe there are a lot of interesting

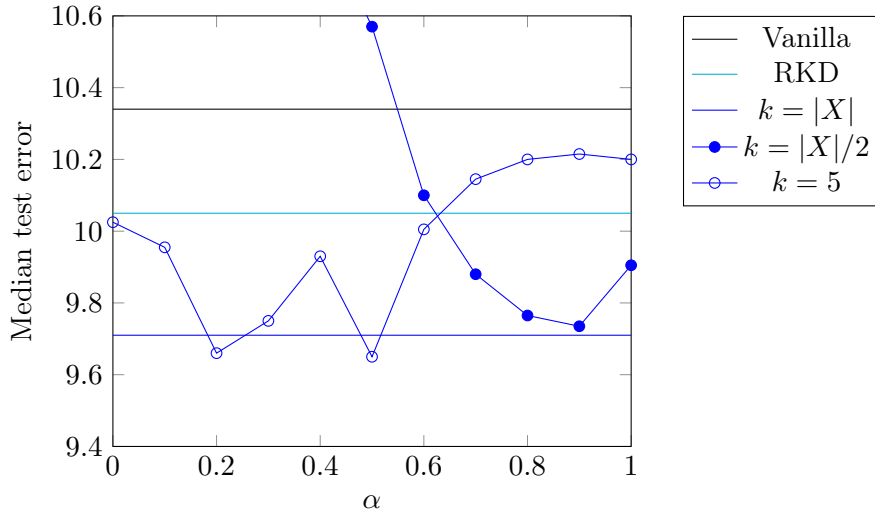


Figure 5.18: Median test error for different values of k and α . Figure and caption extracted from [88] ©2020 IEEE.

contributions to pursue.

We have first introduced the work that we believe was the cornerstone for our interest in the domain [43], in which the authors have shown that it was possible to characterize different DNN behaviors by analyzing the evolution of the graph signal smoothness over their representations. We then built upon this work to propose a measure that is empirically correlated (but that we are not able to ensure causation) with the generalization performance of DNNs. This was subject of a contribution to a non-archival conference:

- Lassance, C. E. R. K., Gripon, V., and Ortega, A. (2018c). Predicting under and overfitting in deep neural networks using graph smoothness. *2018 Graph Signal Processing Workshop (Non-archival)* available at <https://cadurosar.github.io/papers/GSP2018.pdf>

Then we concentrated on possible uses of graph signal smoothness during the training of neural networks. First we showed that we are able to train good feature extractors by training the network to minimize the smoothness of the label indicator signals on the graphs generated by their outputs. This new objective function has three important features that are not present in the traditional cross entropy loss and we demonstrate using experiments that we are able to obtain networks that are more robust, without losing too much generalization performance. Second, we propose to use a regularizer in order to control how the smoothness of the label indicator signals evolve over graphs that are

generated by the intermediate representations of DNNs. We show that these regularizers are not only theoretically inline with our definition of robustness (Definition 2.5.2), but also that we can demonstrate empirically their efficacy when compared (or added) to other methods in the literature. These two uses of graph signal smoothness were subjects of contributions, one to a conference and the other is under the review process of a journal:

1. Bontonou*, M., Lassance*, C., Hacene, G. B., Gripon, V., Tang, J., and Ortega, A. (2019). Introducing graph smoothness loss for training deep learning architectures. In *2019 IEEE Data Science Workshop (DSW)*, pages 160–164, * authors contributed equally
2. Lassance, C. E. R. K., Gripon, V., and Ortega, A. (2018a). Laplacian networks: Bounding indicator function smoothness for neural network robustness. *arXiv preprint arXiv:1805.10133*, under journal review since 01/2020

Finally we have presented a method that does not build from the GSP framework, but that allows us to use the GSP framework on previously defined techniques. In other words, we have specialized the RKD framework as GKD, which we have shown empirically and analytically to improve the performance of the compressed networks. This introductory work was published at a conference as:

- Lassance, C., Bontonou, M., Hacene, G. B., Gripon, V., Tang, J., and Ortega, A. (2020a). Deep geometric knowledge distillation with graphs. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8484–8488. IEEE

In summary, the works in this chapter aim at proposing and expanding the domain of “Deep Neural Networks latent spaces supported on graphs”, while showing the possible improvements this domain can bring to the overall Deep Learning community. In the next chapter we provide a summary of the overall thesis, in order to conclude the work and present the research directions that we have opened but were not able to explore yet.

Chapter 6

Conclusion

6.1	Summary of contributions	176
-----	------------------------------------	-----

6.2	Perspectives and future work	178
-----	--	-----

6.3	Discussion and considerations about the field	181
-----	---	-----

In this section, we first present a quick summary of the thesis. We then present a summary of the contributions presented in this document and the perspectives/research directions that we consider now open in the context of this work. We close this document with a discussion and considerations on the overall field of deep learning. The main idea that we pursued during the last three years was to tackle some shortcomings of deep learning architectures by looking at their intermediate representations.

To perform our analyses, we used the framework of Graph Signal Processing, in which graphs are used to represent the topology of a complex domain (here: latent spaces of deep learning architectures). We have considered deep learning applications within three machine learning domains: i) representation/transfer learning, ii) compression of deep learning architectures, and iii) study of overfitting (generalization and robustness). In the following paragraphs, we present a summary of the contributions introduced in the document, grouped according to the above-mentioned domains.

6.1 Summary of contributions

6.1.1 Representation and transfer learning

In Section 3.3 we introduced a benchmark to allow the comparison of graph inference methods. Being able to measure and determine the most effective technique for graph inference is of utmost importance, given that most of the analysis we performed during the PhD depends on such inferred graphs. Our findings are that the “naive” baseline (k -nn symmetric similarity graphs) can achieve good enough performance when well-tuned compared to more principled approaches. This is even more important given the fact these naive baselines can typically be computed very efficiently, allowing for deployment during the learning of deep neural networks. We consider that the experiments and findings presented in this document could be the starting point to interesting extensions, including taking into account the specific task that graphs are inferred for when they are created. We shall discuss this point later in this chapter.

Then in Section 4.2, we introduced two techniques that allow us to perform the supervised classification of graph signals. First, we presented an embedding technique which aim at representing a graph in a 2D Euclidean space. As most conventional deep learning architectures are already well adapted to inputs on the 2D Euclidean space, it is quite straight-forward then to use these representations in a classical deep convolutional neural network. We then introduced a set of methods that attack the question of whether one can match the performance of CNNs without using priors about the data structure. Three ideas were discussed in this context: i) using a graph convolution scheme based on translations, which were first defined in Section 4.1.6, ii) introducing a method for performing downsampling on those graphs, and iii) introducing a data augmentation scheme based on graph translations. These three improvements allowed us to reduce the gap from convolutions in a 2D space to graph convolutions to a drop in accuracy of only about 2.7% in a competitive image classification benchmark.

We also considered the task of learning representations that are suited for the classification task. Namely, we introduced a graph smoothness loss for deep learning architectures in Section 5.2. Using this loss to train a deep learning architecture enforces that the output should be able to generate graphs that are smooth with regards to the label signal. This graph smoothness loss has three properties that are interesting in the context of representation learning: i) no contraction, as we only force elements of different classes to be distant, but not all elements of the same class to be close, ii) no arbitrary decision of where the points of each class should be, and iii) no restriction on the amount

of dimensions of the output. We demonstrate via experiments that networks trained with the proposed loss are more robust while achieving similar accuracy than networks trained with the standard cross-entropy loss.

Finally, we considered transfer learning. In the case of transfer learning, the goal is to use representations learned in a first task where data is abundant in a second problem where data tends to be scarce. In Section 3.4 we presented graph filters and how they can be used to improve the pre-trained representations by either combining additional information (Section 3.4.5) or by denoising the learned features in the new class domain (Section 3.4.4). We showed through our experiments that graph filters are well suited to the transfer learning task, allowing us to improve the performance in various tasks ranging from visual-based localization to few-shot learning.

6.1.2 Neural network compression

In the context of neural network compression, we presented two types of contributions: i) efficient layers, and ii) distillation. First, in Section 2.4.2, we presented Shift Attention Layers (SAL). This efficient layer scheme reduces the amount of weights per convolutional filter using the concept of attention so that only the most important weights are kept per filter at the end of the training. The goal of SAL is to start with a vast optimization space to optimize our network (more parameters). As the network evolves, we thin out the optimization space, pruning out the less important parameters. We compared with similar shift layers and showed that SAL can improve the performance of the compressed networks.

We then focused on distillation. In Section 5.4, we presented the Graph Knowledge Distillation (GKD) framework that uses graphs to represent latent spaces. GKD has the advantage of disregarding the dimensionality difference between teacher and student's latent spaces while distilling the learned structure from the teacher to the student. Using graphs also allowed us to propose newer additions to the distillation framework. For example, we proposed to consider only the edges between elements of distinct classes, which we showed to improve the performance of GKD.

6.1.3 Generalization and robustness

In this manuscript, we also considered the concept of overfitting, as stated in Definition 2.1.5. Per our definition, overfitting is closely linked to both robustness (the network is considered to be overfitted to the \mathcal{D}) and generalization (we say a network generalizes

well if it is not overfitted to $\mathcal{D}_{\text{train}}$ or $\mathcal{D}_{\text{valid}}$).

We first formally defined what we call robustness in Section 2.5. We also introduced the concept of α -robustness. The principle is to not only focus on controlling the maximum perturbation a noise ϵ applied to the input may cause to the output but also on bounding the radius r where this control should be applied. The goal is to enforce a small α around the examples (small r), while allowing more significant transitions on the parts of the space that are not represented in \mathcal{D} . We analyzed the representations of four recently proposed methods to increase the robustness of DNNs and found out that the methods that follow our definition (i.e., are more α -robust) are the ones that are more empirically robust.

We then presented in Section 5.3 a regularizer that, when applied to networks, creates Laplacian networks. In a Laplacian network, the smoothness of the label signal should evolve slowly across sequential intermediate representations. Therefore, a small perturbation applied to the input or one of the representations should not be able to impact the overall classification significantly. We analyzed how this is linked to the previously presented robustness definition and showed through experiments how it empirically improves the performance under various perturbations.

Finally, we also studied the problem of analyzing the generalization of DNNs, when no extra labeled data (or a $\mathcal{D}_{\text{valid}}$) is available, in Section 5.1. We used the smoothness of the label signal as our metric. We first analyzed qualitatively (i.e., via graphs) if there is a difference in behavior from a reference network to controlled scenarios where we know the network will be underfitted or overfitted. We verified that the smoothness gap between the last layers of the network is a good indicator of generalization. We then performed a more quantitative analysis. We trained networks varying multiple hyperparameters and verified that there is indeed a correlation between the smoothness gap and the generalization of the network.

6.2 Perspectives and future work

In this thesis, we have introduced many contributions, each opening their research directions and perspectives. We have already described these perspectives during the introduction of each work. Therefore, we now present more high-level point of view. There are four main perspectives that we discuss: i) graph inference, ii) extending our framework to other tasks, iii) graphs and data acquisition, and iv) communication of results.

6.2.1 Graph inference

First, recall that our goal was to study the intermediate representations of DNNs. In most cases, we have introduced graphs that are inferred from data via a similarity metric. We have discussed such construction in Section 3.3, where we have shown that, when correctly tuned, these graph representations can rival more principled techniques. On the other hand, it is fair to say that improvements in graph construction should lead to improvements in most methods presented in this thesis. One such improvement would be to take into account the task at hand when creating the graph, e.g., considering the label information during graph inference allowed us to improve our graph distillation in Section 5.4.

We already started working on this problem, and thinking of possible solutions to infer graphs given two objectives: matching the representations and helping in the considered downstream task. As a matter of fact, very often authors in the domain of graph inference introduce priors to solve what is in-fine an ill-posed problem: there are infinitely many graphs that would correspond to a dataset. For example, in [119], the authors showed that when the prior is that signals are stationary on the graph, there is a polytope of possible graph structures that would fit the provided data. Choosing a point in the polytope boils down to favorizing a specific key property of the sought graph structure. In their work, the authors consider sparsity or simplicity for example. We believe that using the task as a prior could lead to an interesting tradeoff between matching the signals and helping in solving the considered task.

6.2.2 Extension to other tasks

Note that in this manuscript, we mainly focused on semi-supervised and supervised classification. Extending our framework, which uses graphs that represent latent spaces, to other tasks, would be interesting future work. For example, recent literature in self-supervised learning uses a technique close to the presented graph smoothness loss, where augmented examples from the same original image should be closer in latent space than examples from different images [18, 42].

More generally, the consideration of hyperbolic spaces in the design and the optimization of deep learning architectures has become increasingly popular [102, 166]. Graphs could be considered as a natural way to reformulate or improve these methods.

Consider manifold mixup for instance, that is being used in the training of many modern state-of-the-art classifiers [106]. The principle is to interpolate inputs, outputs

and intermediate representations to augment the training set. Instead of using a naive linear interpolation, using graphs instead could lead to a more accurate generation of augmented inputs.

6.2.3 Graphs and data acquisition

In the previous subsection, we described how our framework may be adapted to tasks that we did not consider in this document. An orthogonal problem would be to use the GSP framework during the data acquisition and or labeling phase.

Consider the few-shot learning task. In this task, the goal is to learn from a few labeled samples. There are already many techniques to tackle this problem, including the ones presented in this thesis. The labeled samples are either chosen at random or predefined by the benchmark. This procedure leads to two drawbacks:

1. The algorithm’s results are very dependent on exactly which are the labeled samples, requiring a large number of random initializations (i.e., drawing the few labeled samples) to have a good enough confidence interval for comparing two methods;
2. The procedure is not inline with real-world scenarios. A more realistic scenario is to acquire a collection of unlabelled samples. One can then either: i) choose in which order it should label the samples (active learning), and ii) receive a small subset of labeled examples and be able to exploit both labeled and unlabelled data (semi-supervised learning).

Note that it is common to call “few labels” the scenario of “few-shot semi-supervised learning”. Even if we do not explicitly treat the few-labels task in this thesis, we note that using similarity graphs (akin to those we use in this thesis) improves accuracy on the few-labels scenario [58]. Moreover, we believe that extending this framework to the active learning scenario¹ should lead to improvements when compared to the traditional semi-supervised setting.

Indeed, imagine that the similarity graph we generate is not well-behaved (i.e., either disconnected or with a high variance of degrees). In this case, correctly choosing which node to label is of significant importance because some label information may be lost in the case of an unlucky random sample. To mitigate this problem, we believe that using a graph sampling algorithm [158] should allow us to accurately select the correct nodes to

¹more precisely the “pool setting”, where the learner is given the set of unlabelled samples and can then iteratively choose which points that it wants to label.

label, reducing the number of labels needed for adequate performance and reducing the variance caused by the random sampling.

6.2.4 Communication of results

Finally, we have to analyze the diffusion of the ideas presented here. Parallel to the writing of this thesis, we are also preparing a book chapter that presents the domain of “Graphs for deep learning latent representations” in a more concise way, focusing less on our contributions and more on the domain itself. The goal is to create a more straightforward introduction of the concepts described here to diffuse our contributions and inspire more interest in the presented domain. Further diffusing our findings with introductory courses to deep learning and GSP would be advisable as well.

6.3 Discussion and considerations about the field

Deep Learning has attracted a lot of attention in the past few years, and the trend is increasing. Consider for instance the number of papers submitted to NeurIPS, which is an iconic conference of the domain, in the past few years, showed in Figure 6.1. We can clearly see how popular the domain has become.

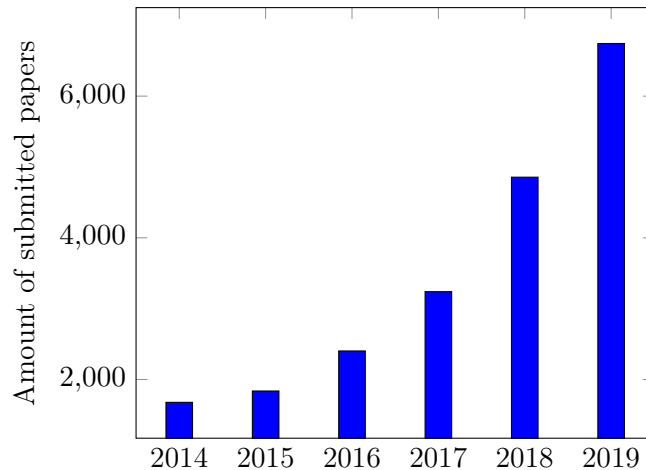


Figure 6.1: NeurIPS paper submission evolution from 2014 to 2019. Data extracted from <https://www.openresearch.org/wiki/NIPS> and <https://medium.com/@NeurIPSCConf/what-we-learned-from-neurips-2019-data-111ab996462c>

Such a sudden popularity is not without drawbacks. For example, we noticed that despite being mostly driven by experiments, publications to the top-tier conferences in machine learning tend to provide weak statistical guarantees about the improvements

they claim to obtain in their work. Too often we found papers in which the confidence interval is not given, the code made available (if made available) does not reproduce the presented results, and the mathematical formulation is disconnected from the presented experimental results.

This is not surprising given the fact that acceptance of papers in these venues is becoming increasingly important for applying to some companies or even academic positions. The problem is that the more papers are submitted, the more reviewers are needed, and of course when the numbers explode that fast, it is not possible to ensure a fair and balanced process.

In this context, most of the references that we cited in this document are very recent (> 2017) and they might contain contradictory results and claims. This is highly problematic in the context of a PhD in science, where we should be more focused on reproducibility and generalization of the results introduced in our contributions.

One such example happened at the start of this PhD thesis. One of our initial goals was to study graph neural networks. Very quickly we started to realize that there was a problem in the way that papers compared with each other. This greatly impacted our vision of the domain and was further confirmed by studies such as [31, 145, 167]. We have previously discussed this in Section 4.3, but I believe that this required a more in-depth discussion in this conclusion. The fact that the problems come not only from the benchmarks (that is tackled by recent contributions such as [57]) but from the experimental design is a clear signal that the domain could be evolving too fast. Note that this is not exclusive to GNNs. The same problems have been found in deep metric learning, where in [135] the authors show that the gap between older methods and more recent contributions is smaller than it is advertised in the recent papers. This is due mostly to improvements that are not linked with the more recent methods themselves, but in data acquisition/pre-processing.

We tackled many important problems in this thesis. These problems are very relevant for a safe and trustworthy deployment of deep learning solutions in the society. Of course we would not pretend that the work that we did was not without failures or better than the rest of the literature. Indeed, some of our critics, also apply to some of our work. Yet depending on the opportunities that are going to appear in the continuity of my career, I hope I will be able to continue in this direction of research always striving to perform my work with scientific rigor.

Bibliography

- [1] Anirudh, R., Bremer, P., Sridhar, R., and Thiagarajan, J. (2017). Influential sample selection: A graph signal processing approach. Technical report, Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States).
- [2] Arandjelovic, R., Gronat, P., Torii, A., Pajdla, T., and Sivic, J. (2016). Netvlad: Cnn architecture for weakly supervised place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5297–5307.
- [3] Ardakani, A., Condo, C., and Gross, W. J. (2017). Sparsely-connected neural networks: towards efficient vlsi implementation of deep neural networks. *International Conference on Learning Representations (ICLR)*.
- [4] Ba, J. and Caruana, R. (2014). Do deep nets really need to be deep? In *Advances in neural information processing systems*, pages 2654–2662.
- [5] Balcilar, M., Renton, G., Héroux, P., Gaüzère, B., Adam, S., and Honeine, P. (2020). Bridging the gap between spectral and spatial domains in graph neural networks. *arXiv preprint arXiv:2003.11702*.
- [6] Belkin, M. and Niyogi, P. (2003). Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, **15**(6), 1373–1396.
- [7] Bertinetto, L., Henriques, J. F., Torr, P., and Vedaldi, A. (2019). Meta-learning with differentiable closed-form solvers. In *International Conference on Learning Representations*.
- [8] Bojchevski, A. and Günnemann, S. (2019). Certifiable robustness to graph perturbations. In *Advances in Neural Information Processing Systems 32*, pages 8319–8330. Curran Associates, Inc.
- [9] Bontonou, M., Lassance, C., Gripon, V., and Farrugia, N. (2019). Comparing linear structure-based and data-driven latent spatial representations for sequence prediction. In *Wavelets and Sparsity XVIII*, volume 11138, page 111380Z. International Society for Optics and Photonics.
- [10] Bontonou*, M., Lassance*, C., Hacene, G. B., Gripon, V., Tang, J., and Ortega, A. (2019). Introducing graph smoothness loss for training deep learning architectures. In *2019 IEEE Data Science Workshop (DSW)*, pages 160–164, * authors contributed equally.

- [11] Bontonou, M., Lassance, C., Vialatte, J.-C., and Gripon, V. (2019). A unified deep learning formalism for processing graph signals. *arXiv preprint arXiv:1905.00496*.
- [12] Bottou, L. (2010). Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pages 177–186. Springer.
- [13] Brahmabhatt, S., Gu, J., Kim, K., Hays, J., and Kautz, J. (2018). Geometry-aware learning of maps for camera localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2616–2625.
- [14] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., *et al.* (2020). Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- [15] Bulat, A. and Tzimiropoulos, G. (2019). Xnor-net++: Improved binary neural networks. *arXiv preprint arXiv:1909.13863*.
- [16] Bunke, H. and Shearer, K. (1998). A graph distance metric based on the maximal common subgraph. *Pattern recognition letters*, **19**(3-4), 255–259.
- [17] Chang, L. J., Gianaros, P. J., Manuck, S. B., Krishnan, A., and Wager, T. D. (2015). A sensitive and specific neural signature for picture-induced negative affect. *PLoS biology*, **13**(6), e1002180.
- [18] Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. (2020). A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*.
- [19] Chrabaszcz, P., Loshchilov, I., and Hutter, F. (2017). A downsampled variant of imagenet as an alternative to the cifar datasets. *arXiv preprint arXiv:1707.08819*.
- [20] Cireşan, D. C., Giusti, A., Gambardella, L. M., and Schmidhuber, J. (2013). Mitosis detection in breast cancer histology images with deep neural networks. In *International conference on medical image computing and computer-assisted intervention*, pages 411–418. Springer.
- [21] Cisse, M., Bojanowski, P., Grave, E., Dauphin, Y., and Usunier, N. (2017). Parseval networks: Improving robustness to adversarial examples. In *International Conference on Machine Learning*, pages 854–863.
- [22] Cubuk, E. D., Zoph, B., Mane, D., Vasudevan, V., and Le, Q. V. (2019). Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 113–123.
- [23] De, S. and Smith, S. L. (2020). Batch normalization biases deep residual networks towards shallow paths. *arXiv preprint arXiv:2002.10444*.
- [24] Defferrard, M., Bresson, X., and Vandergheynst, P. (2016). Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, pages 3837–3845.

- [25] Defferrard, M., Martin, L., Pena, R., and Perraudin, N. (2017). Pygsp: Graph signal processing in python. *Zenodo*.
- [26] Dietterich, T. G. and Bakiri, G. (1994). Solving multiclass learning problems via error-correcting output codes. *Journal of artificial intelligence research*, **2**, 263–286.
- [27] Du, J., Zhang, S., Wu, G., Moura, J. M., and Kar, S. (2017). Topology adaptive graph convolutional networks. *arXiv preprint arXiv:1710.10370*.
- [28] Dyson, F. (2004). A meeting with enrico fermi.
- [29] Egilmez, H. E., Pavez, E., and Ortega, A. (2017). Graph learning from data under Laplacian and structural constraints. *IEEE Journal on Selected Topics in Signal Processing*.
- [30] Engstrom, L., Ilyas, A., and Athalye, A. (2018). Evaluating and understanding the robustness of adversarial logit pairing. *arXiv preprint arXiv:1807.10272*.
- [31] Errica, F., Podda, M., Bacciu, D., and Micheli, A. (2020). A fair comparison of graph neural networks for graph classification. In *International Conference on Learning Representations*.
- [32] Fawzi, A., Fawzi, O., and Frossard, P. (2018). Analysis of classifiers’ robustness to adversarial perturbations. *Machine Learning*, **107**(3), 481–508.
- [33] Friedman, J., Hastie, T., and Tibshirani, R. (2001). *The elements of statistical learning*, volume 1. Springer series in statistics New York.
- [34] Gama, F., Bruna, J., and Ribeiro, A. (2019). Stability properties of graph neural networks. *arXiv preprint arXiv:1905.04497*.
- [35] Gastaldi, X. (2017). Shake-shake regularization. *arXiv preprint arXiv:1705.07485*.
- [36] Giles, C. L., Bollacker, K. D., and Lawrence, S. (1998). Citeseer: An automatic citation indexing system. In *Proceedings of the third ACM conference on Digital libraries*, pages 89–98.
- [37] Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. (2017). Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1263–1272.
- [38] Girvan, M. and Newman, M. E. (2002). Community structure in social and biological networks. *Proceedings of the national academy of sciences*, **99**(12), 7821–7826.
- [39] Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.
- [40] Goodfellow, I. J., Shlens, J., and Szegedy, C. (2014). Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.
- [41] Grelier, N., Lassance, C. E. R. K., Dupraz, E., and Gripon, V. (2018). Graph-projected signal processing. In *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 763–767. IEEE.

- [42] Grill, J.-B., Strub, F., Alth  , F., Tallec, C., Richemond, P. H., Buchatskaya, E., Doersch, C., Pires, B. A., Guo, Z. D., Azar, M. G., *et al.* (2020). Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*.
- [43] Gripon, V., Ortega, A., and Girault, B. (2018). An inside look at deep neural networks using graph signal processing. In *Proceedings of ITA*.
- [44] Hacene, G. B. (2019). *Processing and learning deep neural networks on chip*. Ph.D. thesis, Ecole nationale sup  rieure Mines-T  l  com Atlantique.
- [45] Hacene, G. B., Gripon, V., Arzel, M., Farrugia, N., and Bengio, Y. (2018). Quantized guided pruning for efficient hardware implementations of convolutional neural networks. *arXiv preprint arXiv:1812.11337*.
- [46] Hacene, G. B., Lassance, C., Gripon, V., Courbariaux, M., and Bengio, Y. (2019). Attention based pruning for shift networks. *arXiv preprint arXiv:1905.12300*, to appear in *25th International Conference on Pattern Recognition (ICPR2020)*.
- [47] Hammond, D. K., Vandergheynst, P., and Gribonval, R. (2011). Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis*, **30**(2), 129–150.
- [48] Han, D., Kim, J., and Kim, J. (2017). Deep pyramidal residual networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [49] Hansen, K. A., Kay, K. N., and Gallant, J. L. (2007). Topographic organization in and near human visual area v4. *Journal of Neuroscience*, **27**(44), 11896–11911.
- [50] He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- [51] He, T., Kong, R., Holmes, A. J., Nguyen, M., Sabuncu, M. R., Eickhoff, S. B., Bzdok, D., Feng, J., and Yeo, B. T. (2020). Deep neural networks and kernel regression achieve comparable accuracies for functional connectivity prediction of behavior and demographics. *NeuroImage*.
- [52] Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., and Meger, D. (2018). Deep reinforcement learning that matters. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [53] Hendrycks, D. and Dietterich, T. (2019). Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*.
- [54] Hermans, A., Beyer, L., and Leibe, B. (2017). In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*.
- [55] Hinton, G., Vinyals, O., and Dean, J. (2014). Distilling the knowledge in a neural network. *Neural Information Processing Systems 2014 Deep Learning Workshop*.

- [56] Hornik, K., Stinchcombe, M., and White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural networks*, **2**(5), 359–366.
- [57] Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. (2020a). Open graph benchmark: Datasets for machine learning on graphs. *arXiv preprint arXiv:2005.00687*.
- [58] Hu, Y., Gripon, V., and Pateux, S. (2020b). Exploiting unsupervised inputs for accurate few-shot classification. *arXiv preprint arXiv:2001.09849*.
- [59] Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., and Bengio, Y. (2016). Binarized neural networks. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 4107–4115. Curran Associates, Inc.
- [60] Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., and Bengio, Y. (2017). Quantized neural networks: Training neural networks with low precision weights and activations. *The Journal of Machine Learning Research*, **18**(1), 6869–6898.
- [61] Hubel, D. H. and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of physiology*, **160**(1), 106–154.
- [62] Huber, P. J. (1992). Robust estimation of a location parameter. In *Breakthroughs in statistics*, pages 492–518. Springer.
- [63] Ilyas, A., Santurkar, S., Tsipras, D., Engstrom, L., Tran, B., and Madry, A. (2019). Adversarial examples are not bugs, they are features. In *Advances in Neural Information Processing Systems*, pages 125–136.
- [64] Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, pages 448–456.
- [65] Irion, J. and Saito, N. (2016). Efficient approximation and denoising of graph signals using the multiscale basis dictionaries. *IEEE Transactions on Signal and Information Processing over Networks*.
- [66] Iscen, A., Tolias, G., Avrithis, Y., Furon, T., and Chum, O. (2017). Efficient diffusion on region manifolds: Recovering small objects with compact cnn representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2077–2086.
- [67] Iscen, A., Tolias, G., Avrithis, Y., and Chum, O. (2019). Label propagation for deep semi-supervised learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- [68] Isufi, E., Loukas, A., Simonetto, A., and Leus, G. (2016). Autoregressive moving average graph filtering. *IEEE Transactions on Signal Processing*, **65**(2), 274–288.

- [69] Isufi, E., Gama, F., and Ribeiro, A. (2020). Edgenets: Edge varying graph neural networks. *arXiv preprint arXiv:2001.07620*.
- [70] Jeon, Y. and Kim, J. (2018). Constructing fast network through deconstruction of convolution. In *Advances in Neural Information Processing Systems*, pages 5951–5961.
- [71] Kalofolias, V. and Perraudin, N. (2019). Large scale graph learning from smooth signals. In *International Conference on Learning Representations*.
- [72] Kendall, A., Grimes, M., and Cipolla, R. (2015). Posenet: A convolutional network for real-time 6-dof camera relocalization. In *Proceedings of the IEEE international conference on computer vision*, pages 2938–2946.
- [73] Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- [74] Kipf, T. N. and Welling, M. (2016). Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- [75] Klicpera, J., Bojchevski, A., and Günnemann, S. (2019). Predict then propagate: Graph neural networks meet personalized pagerank. In *International Conference on Learning Representations*.
- [76] Kobourov, S. G. (2013). Force-directed drawing algorithms. In *Handbook of Graph Drawing and Visualization*, pages 383–408. CRC Press.
- [77] Kolesnikov, A., Beyer, L., Zhai, X., Puigcerver, J., Yung, J., Gelly, S., and Houlsby, N. (2019). Large scale learning of general visual representations for transfer.
- [78] Koopman, B., Zuccon, G., Bruza, P., Sitbon, L., and Lawley, M. (2016). Information retrieval as semantic inference: a graph inference model applied to medical search. *Information Retrieval Journal*.
- [79] Koratana, A., Kang, D., Bailis, P., and Zaharia, M. (2019). Lit: Learned intermediate representation training for model compression. In *International Conference on Machine Learning*, pages 3509–3518.
- [80] Kovačević, J. and Chebira, A. (2008). An introduction to frames. *Foundations and Trends in Signal Processing*, 2(1), 1–94.
- [81] Krizhevsky, A. and Hinton, G. (2009). Learning multiple layers of features from tiny images. <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>.
- [82] Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105.
- [83] Kumar, A., Khadkevich, M., and Fügen, C. (2018). Knowledge transfer from weakly labeled audio using convolutional neural network for sound events and scenes. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.

- [84] Kurakin, A., Goodfellow, I. J., and Bengio, S. (2017). Adversarial machine learning at scale. *International Conference on Learning Representations*.
- [85] Lassance, C., Bontonou, M., Hacene, G. B., Gripon, V., Tang, J., and Ortega, A. (2019a). Deep geometric knowledge distillation with graphs. *arXiv preprint arXiv:1911.03080*.
- [86] Lassance, C., Latif, Y., Garg, R., Gripon, V., and Reid, I. (2019b). Improved visual localization via graph smoothing. *arXiv preprint arXiv:1911.02961*.
- [87] Lassance, C., Gripon, V., Tang, J., and Ortega, A. (2019). Structural robustness for deep learning architectures. In *2019 IEEE Data Science Workshop (DSW)*, pages 125–129.
- [88] Lassance, C., Bontonou, M., Hacene, G. B., Gripon, V., Tang, J., and Ortega, A. (2020a). Deep geometric knowledge distillation with graphs. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8484–8488. IEEE.
- [89] Lassance, C., Gripon, V., and Mateos, G. (2020b). Graph topology inference benchmarks for machine learning. *arXiv preprint arXiv:2007.08216*, to appear in *2020 IEEE 30th International Workshop on Machine Learning for Signal Processing (MLSP)*.
- [90] Lassance, C. E. R. K., Gripon, V., and Ortega, A. (2018a). Laplacian networks: Bounding indicator function smoothness for neural network robustness. *arXiv preprint arXiv:1805.10133*, under journal review since 01/2020.
- [91] Lassance, C. E. R. K., Vialatte, J.-C., and Gripon, V. (2018b). Matching convolutional neural networks without priors about data. In *2018 IEEE Data Science Workshop (DSW)*, pages 234–238. IEEE.
- [92] Lassance, C. E. R. K., Gripon, V., and Ortega, A. (2018c). Predicting under and overfitting in deep neural networks using graph smoothness. *2018 Graph Signal Processing Workshop (Non-archival)* available at <https://cadurosar.github.io/papers/GSP2018.pdf>.
- [93] LeCun, Y., Bengio, Y., *et al.* (1995). Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, **3361**(10), 1995.
- [94] LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, **86**(11), 2278–2324.
- [95] Lee, C.-Y., Xie, S., Gallagher, P., Zhang, Z., and Tu, Z. (2015). Deeply-supervised nets. In *Artificial Intelligence and Statistics*, pages 562–570.
- [96] Lee, S. and Song, B. (2019). Graph-based knowledge distillation by multi-head attention network. *arXiv preprint arXiv:1907.02226*.
- [97] Li, H., Kadav, A., Durdanovic, I., Samet, H., and Graf, H. P. (2017). Pruning filters for efficient convnets. *International Conference on Learning Representations (ICLR)*.
- [98] Liao, R., Zhao, Z., Urtasun, R., and Zemel, R. (2019). Lanczosnet: Multi-scale deep graph convolutional networks. In *International Conference on Learning Representations*.

- [99] Lim, S., Kim, I., Kim, T., Kim, C., and Kim, S. (2019). Fast autoaugment. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [100] Lin, Z., Memisevic, R., and Konda, K. (2015). How far can we go without convolution: Improving fully-connected networks. *arXiv preprint arXiv:1511.02580*.
- [101] Liu, C., Yu, G., Chang, C., Rai, H., Ma, J., Gorti, S., and Volkovs, M. (2019a). Guided similarity separation for image retrieval. In *NeurIPS*.
- [102] Liu, Q., Nickel, M., and Kiela, D. (2019b). Hyperbolic graph neural networks. In *Advances in Neural Information Processing Systems*, pages 8230–8241.
- [103] Liu, Y., Cao, J., Li, B., Yuan, C., Hu, W., Li, Y., and Duan, Y. (2019c). Knowledge distillation via instance relationship graph. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7096–7104.
- [104] Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. (2018). Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*.
- [105] Mallat, S. (2016). Understanding deep convolutional networks. *Phil. Trans. R. Soc. A*, **374**(2065), 20150203.
- [106] Mangla, P., Singh, M., Sinha, A., Kumari, N., Balasubramanian, V. N., and Krishnamurthy, B. (2019). Charting the right manifold: Manifold mixup for few-shot learning. *CoRR*, **abs/1907.12087**.
- [107] Manning, C. D., Raghavan, P., and Schütze, H. (2008). *Introduction to information retrieval*. Cambridge university press.
- [108] Mateos, G., Segarra, S., Marques, A. G., and Ribeiro, A. (2019). Connecting the dots: Identifying network structure via graph signal processing. *IEEE Signal Processing Magazine*.
- [109] McCallum, A. K., Nigam, K., Rennie, J., and Seymore, K. (2000). Automating the construction of internet portals with machine learning. *Information Retrieval*, **3**(2), 127–163.
- [110] Miotto, R., Wang, F., Wang, S., Jiang, X., and Dudley, J. T. (2018). Deep learning for healthcare: review, opportunities and challenges. *Briefings in bioinformatics*, **19**(6), 1236–1246.
- [111] Moosavi DeZfooli, S. M., Fawzi, A., and Frossard, P. (2016). Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [112] Namata, G., London, B., Getoor, L., Huang, B., and EDU, U. (2012). Query-driven active surveying for collective classification. In *ICML Workshop on MLG*.
- [113] Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. (2011). Reading digits in natural images with unsupervised feature learning. *Neural Information Processing Systems Workshop on Deep Learning and Unsupervised Feature Learning*.

- [114] Nilsback, M.-E. and Zisserman, A. (2008). Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics and Image Processing*.
- [115] Page, L., Brin, S., Motwani, R., and Winograd, T. (1999). The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab.
- [116] Papernot, N. and McDaniel, P. (2018). Deep k-nearest neighbors: Towards confident, interpretable and robust deep learning. *arXiv preprint arXiv:1803.04765*.
- [117] Park, W., Kim, D., Lu, Y., and Cho, M. (2019). Relational knowledge distillation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3967–3976.
- [118] Pasdeloup, B. (2017). *Extending convolutional neural networks to irregular domains through graph inference*. Ph.D. thesis, Ecole nationale supérieure Mines-Télécom Atlantique Bretagne Pays de la Loire.
- [119] Pasdeloup, B., Gripon, V., Mercier, G., Pastor, D., and Rabbat, M. G. (2017). Characterization and inference of graph diffusion processes from observations of stationary signals. *IEEE Transactions on Signal and Information Processing over Networks*.
- [120] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*.
- [121] Perraudin, N., Paratte, J., Shuman, D., Martin, L., Kalofolias, V., Vandergheynst, P., and Hammond, D. K. (2014). GSPBOX: A toolbox for signal processing on graphs. *ArXiv e-prints*.
- [122] Pham, H., Guan, M. Y., Zoph, B., Le, Q. V., and Dean, J. (2018). Efficient neural architecture search via parameter sharing. *arXiv preprint arXiv:1802.03268*.
- [123] Philbin, J., Chum, O., Isard, M., Sivic, J., and Zisserman, A. (2007). Object retrieval with large vocabularies and fast spatial matching. In *2007 IEEE conference on computer vision and pattern recognition*, pages 1–8. IEEE.
- [124] Philbin, J., Chum, O., Isard, M., Sivic, J., and Zisserman, A. (2008). Lost in quantization: Improving particular object retrieval in large scale image databases. In *2008 IEEE conference on computer vision and pattern recognition*, pages 1–8. IEEE.
- [125] Piasco, N., Sidibé, D., Demonceaux, C., and Gouet-Brunet, V. (2018). A survey on visual-based localization: On the benefit of heterogeneous data. *Pattern Recognition*, **74**, 90 – 109.
- [126] Piczak, K. J. (2015). ESC: Dataset for Environmental Sound Classification. In *Proceedings of the 23rd Annual ACM Conference on Multimedia*. ACM Press.
- [127] Qian, H. and Wegman, M. N. (2019). L2-nonexpansive neural networks. In *International Conference on Learning Representations*.

- [128] Radenović, F., Iscen, A., Tolias, G., Avrithis, Y., and Chum, O. (2018). Revisiting oxford and paris: Large-scale image retrieval benchmarking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5706–5715.
- [129] Radenović, F., Tolias, G., and Chum, O. (2019). Fine-tuning cnn image retrieval with no human annotation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **41**(7), 1655–1668.
- [130] Rauber, J., Brendel, W., and Bethge, M. (2017). Foolbox: A python toolbox to benchmark the robustness of machine learning models. In *Reliable Machine Learning in the Wild Workshop, 34th International Conference on Machine Learning*.
- [131] Ravi, S. and Larochelle, H. (2017). Optimization as a model for few-shot learning. *International Conference on Learning Representations*.
- [132] Recht, B., Roelofs, R., Schmidt, L., and Shankar, V. (2019). Do imagenet classifiers generalize to imagenet? *arXiv preprint arXiv:1902.10811*.
- [133] Romero, A., Ballas, N., Kahou, S. E., Chassang, A., Gatta, C., and Bengio, Y. (2015). Fitnets: Hints for thin deep nets. In *International Conference on Learning Representations*.
- [134] Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, **65**(6), 386.
- [135] Roth, K., Milbich, T., Sinha, S., Gupta, P., Ommer, B., and Cohen, J. P. (2020). Revisiting training strategies and generalization performance in deep metric learning. In *Proceedings of Machine Learning and Systems 2020*, pages 3006–3016. PMLR.
- [136] Roux, L., Racoceanu, D., Loménie, N., Kulikova, M., Irshad, H., Klossa, J., Capron, F., Genestie, C., Le Naour, G., and Gurcan, M. N. (2013). Mitosis detection in breast cancer histological images an icpr 2012 contest. *Journal of pathology informatics*, **4**.
- [137] Ruiz, L., Gama, F., Marques, A. G., and Ribeiro, A. (2019). Invariance-preserving localized activation functions for graph neural networks. *IEEE Transactions on Signal Processing*, **68**, 127–141.
- [138] Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986). Learning representations by back-propagating errors. *nature*, **323**(6088), 533–536.
- [139] Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., *et al.* (2015). Imagenet large scale visual recognition challenge. *International journal of computer vision*, **115**(3), 211–252.
- [140] Santurkar, S., Tsipras, D., Ilyas, A., and Madry, A. (2018). How does batch normalization help optimization? In *Advances in Neural Information Processing Systems*, pages 2483–2493.
- [141] Schaub, M. T. and Segarra, S. (2018). Flow smoothing and denoising: graph signal processing in the edge-space. In *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 735–739. IEEE.

- [142] Segarra, S., Huang, W., and Ribeiro, A. (2015). Diffusion and superposition distances for signals supported on networks. *IEEE Transactions on Signal and Information Processing over Networks*, **1**(1), 20–32.
- [143] Sen, P., Namata, G., Bilgic, M., Getoor, L., Galligher, B., and Eliassi-Rad, T. (2008). Collective classification in network data. *AI magazine*, **29**(3), 93–93.
- [144] Shafipour, R., Segarra, S., Marques, A., and Mateos, G. (2018). Identifying the topology of undirected networks from diffused non-stationary graph signals. *IEEE Transactions on Signal Processing*.
- [145] Shchur, O., Mumme, M., Bojchevski, A., and Günnemann, S. (2018). Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868*.
- [146] Shekkizhar, S. and Ortega, A. (2019). Graph construction from data using non negative kernel regression (NNK graphs). *arXiv preprint arXiv:1910.09383*.
- [147] Shimodaira, H. (2000). Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of statistical planning and inference*, **90**(2), 227–244.
- [148] Shuman, D. I., Narang, S. K., Frossard, P., Ortega, A., and Vandergheynst, P. (2013). The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE signal processing magazine*, **30**(3), 83–98.
- [149] Simonyan, K. and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- [150] Sivic, J., Russell, B. C., Efros, A. A., Zisserman, A., and Freeman, W. T. (2005). Discovering objects and their location in images. In *Tenth IEEE International Conference on Computer Vision (ICCV’05) Volume 1*, volume 1, pages 370–377. IEEE.
- [151] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, **15**(1), 1929–1958.
- [152] Svoboda, J., Masci, J., Monti, F., Bronstein, M., and Guibas, L. (2019). Peernets: Exploiting peer wisdom against adversarial attacks. In *International Conference on Learning Representations*.
- [153] Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. (2013). Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*.
- [154] Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. (2014). Intriguing properties of neural networks. In *International Conference on Learning Representations*.
- [155] Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. (2016). Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.

- [156] Tan, C., Sun, F., Kong, T., Zhang, W., Yang, C., and Liu, C. (2018). A survey on deep transfer learning. In *International conference on artificial neural networks*, pages 270–279. Springer.
- [157] Tan, M. and Le, Q. V. (2019). Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*.
- [158] Tanaka, Y., Eldar, Y. C., Ortega, A., and Cheung, G. (2020). Sampling on graphs: From theory to applications. *arXiv preprint arXiv:2003.03957*.
- [159] Tang, Y. (2013). Deep learning using linear support vector machines. *arXiv preprint arXiv:1306.0239*.
- [160] Tigréat, P., Lassance, C. R. K., Jiang, X., Gripon, V., and Berrou, C. (2016). Assembly output codes for learning neural networks. In *2016 9th International Symposium on Turbo Codes and Iterative Information Processing (ISTC)*, pages 285–289. IEEE.
- [161] Torralba, A., Fergus, R., and Freeman, W. T. (2008). 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **30**(11), 1958–1970.
- [162] Tremblay, N., Gonçalves, P., and Borgnat, P. (2018). Design of graph filters and filterbanks. In *Cooperative and Graph Signal Processing*, pages 299–324. Elsevier.
- [163] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- [164] Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. (2018). Graph attention networks. In *International Conference on Learning Representations*.
- [165] Verma, V., Qu, M., Lamb, A., Bengio, Y., Kannala, J., and Tang, J. (2019a). Graphmix: Regularized training of graph neural networks for semi-supervised learning. *arXiv preprint arXiv:1909.11715*.
- [166] Verma, V., Lamb, A., Beckham, C., Najafi, A., Mitliagkas, I., Lopez-Paz, D., and Bengio, Y. (2019b). Manifold mixup: Better representations by interpolating hidden states. In *International Conference on Machine Learning*, pages 6438–6447.
- [167] Vialatte, J.-C. (2018). *On convolution of graph signals and deep learning on graph domains*. Ph.D. thesis, IMT Atlantique.
- [168] Vignac, C., Ortiz-Jiménez, G., and Frossard, P. (2020). On the choice of graph neural network architectures. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8489–8493.
- [169] Vinh, N. X., Epps, J., and Bailey, J. (2009). Information theoretic measures for clusterings comparison: is a correction for chance necessary? In *Proceedings of the 26th annual international conference on machine learning*.

- [170] Wang, H., Wu, X., Huang, Z., and Xing, E. P. (2020). High-frequency component helps explain the generalization of convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8684–8694.
- [171] Welinder, P., Branson, S., Mita, T., Wah, C., Schroff, F., Belongie, S., and Perona, P. (2010). Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology.
- [172] Wu, B., Wan, A., Yue, X., Jin, P., Zhao, S., Golmant, N., Gholaminejad, A., Gonzalez, J., and Keutzer, K. (2018). Shift: A zero flop, zero parameter alternative to spatial convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9127–9135.
- [173] Wu, F., Souza, A., Zhang, T., Fifty, C., Yu, T., and Weinberger, K. (2019). Simplifying graph convolutional networks. In *International Conference on Machine Learning*.
- [174] Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., and Philip, S. Y. (2020). A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*.
- [175] Xu, K., Hu, W., Leskovec, J., and Jegelka, S. (2019). How powerful are graph neural networks? In *International Conference on Learning Representations*.
- [176] Yang, Z., Cohen, W. W., and Salakhutdinov, R. (2016). Revisiting semi-supervised learning with graph embeddings. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning-Volume 48*, pages 40–48.
- [177] Yin, D., Lopes, R. G., Shlens, J., Cubuk, E. D., and Gilmer, J. (2019). A fourier perspective on model robustness in computer vision. In *Advances in Neural Information Processing Systems*, pages 13276–13286.
- [178] Yu and Shi (2003). Multiclass spectral clustering. In *Proceedings Ninth IEEE International Conference on Computer Vision*.
- [179] Yun, S., Han, D., Oh, S. J., Chun, S., Choe, J., and Yoo, Y. (2019). Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6023–6032.
- [180] Zagoruyko, S. and Komodakis, N. (2016). Wide residual networks. *arXiv preprint arXiv:1605.07146*.
- [181] Zeiler, M. D. and Fergus, R. (2014). Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer.
- [182] Zhai, X., Puigcerver, J., Kolesnikov, A., Ruysen, P., Riquelme, C., Lucic, M., Djolonga, J., Pinto, A. S., Neumann, M., Dosovitskiy, A., Beyer, L., Bachem, O., Tschannen, M., Michalski, M., Bousquet, O., Gelly, S., and Houlsby, N. (2019). A large-scale study of representation learning with the visual task adaptation benchmark.

- [183] Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2017). Understanding deep learning requires rethinking generalization. *International Conference on Learning Representations*.
- [184] Zhang, H., Cisse, M., Dauphin, Y. N., and Lopez-Paz, D. (2018). mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*.
- [185] Zhang, H., Dauphin, Y. N., and Ma, T. (2019a). Residual learning without normalization via better initialization. In *International Conference on Learning Representations*.
- [186] Zhang, S., Tong, H., Xu, J., and Maciejewski, R. (2019b). Graph convolutional networks: a comprehensive review. *Computational Social Networks*, **6**(1), 11.
- [187] Zhang, Z., Cui, P., and Zhu, W. (2020). Deep learning on graphs: A survey. *IEEE Transactions on Knowledge and Data Engineering*.

Titre : Graphes pour représenter les espaces latents des réseaux neuronaux profonds

Mots clés : Apprentissage statistique ; Réseaux de neurones profonds ; Traitement de signal sur graphe

Résumé : Ces dernières années, les méthodes d'apprentissage profond ont atteint l'état de l'art dans une vaste gamme de tâches d'apprentissage automatique, y compris la classification d'images et la traduction automatique. Ces architectures sont assemblées pour résoudre des tâches d'apprentissage automatique de bout en bout. Afin d'atteindre des performances de haut niveau, ces architectures nécessitent souvent d'un très grand nombre de paramètres. Les conséquences indésirables sont multiples, et pour y remédier, il est souhaitable de pouvoir comprendre ce qui se passe à l'intérieur des architectures d'apprentissage profond. Il est difficile de le faire en raison de: i) la dimension élevée des représentations ; et ii) la stochasticité du processus de formation. Dans cette thèse, nous étudions ces architectures en introduisant un formalisme à base de graphes, s'appuyant notamment sur les récents progrès du traitement de signaux sur graphe (TSG). À savoir, nous utilisons des graphes pour représenter les espaces latents des réseaux neuronaux profonds. Nous montrons que ce formalisme des graphes nous permet de répondre à diverses questions, notamment: i) mesurer des capacités de généralisation ; ii) réduire la quantité de choix arbitraires dans la conception du processus d'apprentissage ; iii) améliorer la robustesse aux petites perturbations ajoutées sur les entrées ; et iv) réduire la complexité des calculs.

Title : Graphs for deep learning representations

Keywords : Machine Learning ; Deep neural networks; Graph signal processing

Abstract : In recent years, Deep Learning methods have achieved state of the art performance in a vast range of machine learning tasks, including image classification and multilingual automatic text translation. These architectures are trained to solve machine learning tasks in an end-to-end fashion. In order to reach top-tier performance, these architectures often require a very large number of trainable parameters. There are multiple undesirable consequences, and in order to tackle these issues, it is desired to be able to open the black boxes of deep learning architectures. Problematically, doing so is difficult due to the high dimensionality of representations and the stochasticity of the training process. In this thesis, we investigate these architectures by introducing a graph formalism based on the recent advances in Graph Signal Processing (GSP). Namely, we use graphs to represent the latent spaces of deep neural networks. We showcase that this graph formalism allows us to answer various questions including: ensuring generalization abilities, reducing the amount of arbitrary choices in the design of the learning process, improving robustness to small perturbations added to the inputs, and reducing computational complexity.