

A Spline Theory of Deep Learning

Randall Balestrieri and Richard G. Baraniuk
Rice University

October 30, 2017

Abstract

This report presents a thoroughly analysis of deep neural networks via the introduction and use of spline operators. Rewriting deep neural networks via spline operators will allow us to derive numerous results s.a. Lipschitz constant of deep networks, conditions for convexity and template matching inference or understanding generalization in this context. Hence, the derivations enable us not only to understand deep networks as nonlinear best template matching machines but also explain many observed behaviors s.a. adversarial examples, the impact of the pooling function and the learning phase. For each section we will provide experiments on MNIST and CIFAR10 as well as details for implementations.

Contents

1	Introduction	5
2	Background: Deep Neural Networks for Function Approximation	8
2.1	Layers Description	9
2.2	Deep Convolutional Network	13
2.3	Learning	14
3	Understand Deep Neural Networks Internal Mechanisms	16
3.1	Spline Operators[FINI]	17
3.2	Linear Spline Operator: Generalized Neural Network Layers	20
3.2.1	Nonlinearity layers	20
3.2.2	Sub-Sampling layers	22
3.2.3	Linear layers:FC and convolutional	23
3.3	Deriving Analytical DNNs Mappings to Explicit their Faculty to Perform Template Matching	24
3.3.1	Composition of Splines for Explicit DNN Template Matching	24
3.3.2	Deep Convolutional Networks	26
3.3.3	Deep Residual Networks	27
3.3.4	Deep Recurrent Networks	28
3.4	Template Matching with DNN: How and Why	29
3.4.1	Template Matching in the Context of Splines	30
3.4.2	DNNs Are Composition of Adaptive Partitioning Splines	31
3.4.3	Input Encoding and Template Visualization	32
4	Theoretical Results: Generalization, Optimal Learning, Memorization	35

4.1	What is Generalization for DNN and why Regularization is Key	35
4.2	Learning Optimal Templates	37
4.2.1	Unregularized Learning Solution: Unstable Training	38
4.2.2	Regularized Learning: Global Optimum, Robust, implies Dataset Memorization	38
4.3	Adversarial Examples Are Natural, being Fooled is Not	40
4.3.1	Unregularized Optimal Templates Imply Adv. Noise	41
4.3.2	Ensuring Adversarial Noise Robustness via Lipschitz Constant minimization: Contractive DNNs	41
5	Extension	44
5.1	DNN Inversion: Input Reconstruction is Necessary for Generalization	44
5.1.1	Optimal DNN Leads to Input Reconstruction	46
5.1.2	Boundary Inversion Method, Implementation and Loss Function	47
5.1.3	Semi-Sup Experiments for State-of-the-art Performances across Topologies	49
5.1.4	Extensions	49
6	Conclusion	53
A	Extra Material	60
A.1	Spline Operator	60
A.1.1	Spline Functions[FINI]	60
A.1.2	Multivariate spline functions[FINI]	62
A.2	Dataset Memorization Proof	65
A.3	Conditions for local to global inference	68
A.4	Space Contraction and Adversarial Examples	70
A.5	Function Approximation, Orbits, Class Separation, Generalization and Activation Graph	72

A.5.1	Function Approximation and Orbits for Invariant Learning	72
A.5.2	Activation Graph, Paths	75
A.5.3	Experiments	76
B	Dataset and Model Description	78
B.1	CIFAR10	78
B.2	Networks Description and Training Details	78

1 Introduction

Deep Neural Networks (DNNs) are universal function approximators providing state-of-the-art solutions on wide range of applications. Common perceptual tasks such as speech recognition, image classification, and object tracking are now commonly tackled via DNNs. Some fundamental problems remain: (1) the lack of a mathematical framework providing an explicit and interpretable input-output formula for any topology, (2) quantification of DNNs stability regarding adversarial examples (i.e. modified inputs fooling DNN predictions whilst undetectable to humans), (3) absence of generalization guarantees and controllable behaviors for ambiguous patterns, (4) leverage unlabeled data to apply DNNs to domains where expert labeling is scarce as in the medical field. Answering those points would provide theoretical perspectives for further developments based on a common ground. Furthermore, DNNs are now deployed in tremendous societal applications, pushing the need to fill this theoretical gap to ensure control, reliability, and interpretability.

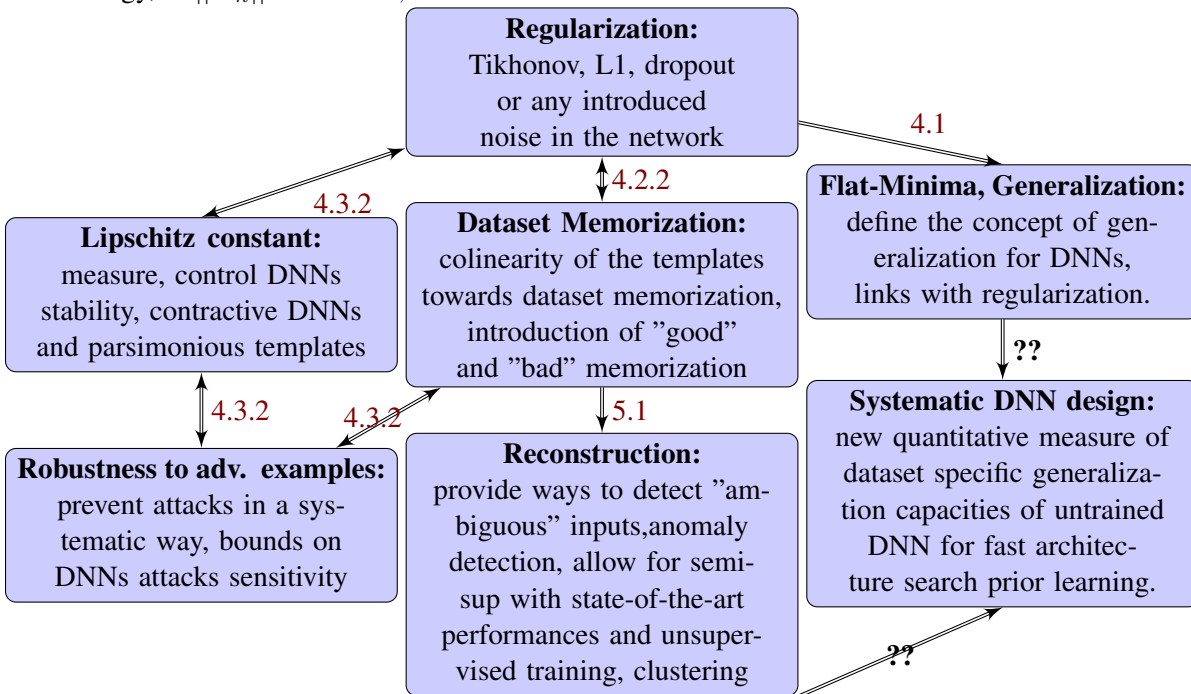
DNNs are models involving compositions of nonlinear and linear transforms. (1) We will provide a straightforward methodology to express the nonlinearities as affine spline functions. The linear part being a degenerated case of spline function, we can rewrite any given DNN topology as succession of such functionals making the network itself a piecewise linear spline. This formulation provides a universal piecewise linear expression of the input-output mapping of DNNs, clarifying the role of its internal components. (2) In functional analysis, the regularity of a mapping is defined via its Lipschitz constant. Our formulation eases the analytical derivation of this stability variable measuring the adversarial examples sensitivity. For any given architecture, we provide a measure of risk to adversarial attacks. (3) Recently, the deep learning community has focused on the reminiscent theory of flat and sharp minima to provide generalization guarantees. Flat minima are regions in the parameter space associated with great generalization capacities. We will first, prove the equivalence between flat minima and spline smoothness. After bridging those theories, we will motivate a novel regularization technique pushing the learning of DNNs towards flat minima, maximizing generalization performances. (4) From (1) we will reinterpret DNNs as template matching algorithms. When coupled with insights derived from (2), we will integrate unlabeled data information into the network during learning. To do so, we will propose to guide DNNs templates towards their input via a scheme assimilated as a reconstruction formula for DNNs. This inversion can be computed efficiently by back-propagation leading to no computational overhead. From this, any semi-supervised technique can be used out-of-the-box with current DNNs where we provide state-of-the-art results. Unsupervised tasks would also become reachable to DNNs, a task considered as the keystone of learning for the neuro-science community. To date, those problematics have been studied independently leading to over-specialized solutions generally topology specific and cumbersome to incorporate into a pre-existing pipeline. On the other hand, all the proposed solutions necessitate negligible software updates, suited for efficient large-scale deployment.

Non-exhaustive list of the main contributions

1) We first develop spline operators (SOs) A.1.1, a natural generalization of multivariate spline functions as well as their linear case (LSOs). LSOs are shown to "span" DNNs layers, being restricted cases of LSOs 3.2. From this, composition of those operators lead to the explicit analytical input-output formula of DNNs, for any architecture 3.3. We then dive into some analysis:

- Interpret DNNs as template matching machines, provide ways to visualize and analyze the inner representation a DNN has of its input w.r.t each classes and understand the prediction 3.4.
- Understand the impact of design choices such as skip-connections and provide conditions for "good" weight initialization schemes 3.3.
- Derive a simple methodology to compute the Lipschitz constant of any DNN, quantifying their stability and derive strategies for adversarial example robustness 4.3.2.
- Study the impact of depth and width for generalization and class separation, orbit learning A.5.

2) Secondly, we prove the following implications for any DNN with the only assumption that all inputs have same energy, as $\|X_n\|^2 = K > 0, \forall n$.



3) We then leverage our understanding of DNNs to develop layers with inhibitor connections, leading to deep inhibitor neural networks (DINN) ?? via development of a new type of layer called PSD layer ?? . While being biologically plausible and outperforming all DNNs topologies ??, DINN provide powerful properties:

- Guarantees of sparse activation during forward pass, extension for practical application with sparse network implementations. Non-sparse backpropagation, synonym of absence of vanishing gradient, dying neuron problems and faster learning ??; analysis of winner-share-all (DINN) strategies versus winner-takes-all (DNN) ??
- Linearization and energy diffusion properties through the layers ??.
- Robustness to "detrimental" overfitting and wrong hyper-parameter selection such as learning rate; reduced number of degrees of freedom than DNNs with same topology by removal of the bias units

Symbols

x	”Dummy” variable representing an input/observation
$\hat{y}(x)$	”Dummy” variable representing an output/prediction associated to input x
X_n	Observation n of shape (K, I, J) .
Y_n	Target variable associated to X_n , for classification $Y_n \in \{1, \dots, C\}$, $C > 1$, for regression $Y_n \in \mathbb{R}^C$, $C \geq 1$.
\mathcal{D} (resp. \mathcal{D}_s)	Labeled training set with N (resp. N_s) samples $\mathcal{D} = \{(X_n, Y_n)_{n=1}^N\}$.
\mathcal{D}_u	Unlabeled training set with N_u samples $\mathcal{D}_u = \{(X_n)_{n=1}^{N_u}\}$.
$f_{\theta^{(\ell)}}$	Layer at level ℓ with internal parameters $\theta^{(\ell)}$, $\ell = 1, \dots, L$.
Θ	Collection of all parameters $\Theta = \{\theta^{(\ell)}, \ell = 1, \dots, L\}$.
f_{Θ}	Deep Neural Network mapping with $f_{\theta} : \mathbb{R}^D \rightarrow \mathbb{R}^C$
$(C^{(\ell)}, I^{(\ell)}, J^{(\ell)})$	Shape of the representation at layer ℓ with $(C^{(0)}, I^{(0)}, J^{(0)}) = (K, I, J)$ and $(C^{(L)}, I^{(L)}, J^{(L)}) = (C, 1, 1)$.
$D^{(\ell)}$	Dimension of the flattened representation at layer ℓ with $D^{(\ell)} = C^{(\ell)} I^{(\ell)} J^{(\ell)}$, $D^{(0)} = D$ and $D^{(L)} = C$.
$z^{(\ell)}(x)$	Representation of x at layer ℓ in an unflattened format of shape $(C^{(\ell)}, I^{(\ell)}, J^{(\ell)})$, with $z^{(0)}(x) = x$
$z_{c,i,j}^{(\ell)}(x)$	Value at channel c and spatial position (i, j) .
$\mathbf{z}^{(\ell)}(x)$	Representation of x at layer ℓ in a flattened format of dimension $D^{(\ell)}$
$\mathbf{z}_d^{(\ell)}(x)$	Value at dimension d

2 Background: Deep Neural Networks for Function Approximation

Most of applied mathematics interests take the form of function approximation. Two main cases arise, one where the target function f to approximate is known and one where only a set of samples $(X_n, f(X_n))_{n=1}^N$ are observed, providing a limited information on the domain-codomain structure of f . The latter case is the one of supervised learning. Given this *training set* $\mathcal{D} := \{(X_n, Y_n)_{n=1}^N\}$ with $Y_n := f(X_n)$, the unknown functional f is estimated through the approximator \hat{f} . Finding an approximant \hat{f} with correct behaviors on \mathcal{D} is usually an ill-posed problem with many possible solutions. Yet, each one might behave differently for new observations, leading to different generalization performances. Generalization is the ability to replicate the behavior of f on new inputs not present in \mathcal{D} thus not used to obtain \hat{f} . Hence, one seeks for an approximator \hat{f} having the best generalization performance. In some applications, the unobserved f is known to fulfill some properties such as boundary and regularity conditions for PDE approximation. In machine learning however, the lack of physic based principles does not provide any property constraining the search for a good approximator \hat{f} except the performance measure based on the training set \mathcal{D} and an estimate of generalization performance based on a *test set*. To tackle this search, one commonly resorts to a parametric functional \hat{f}_Θ where Θ contains all the free parameters controlling the behavior of \hat{f}_Θ . The task thus "reduces" to finding the optimal set of parameters Θ^* minimizing the empirical error on the training set and maximizing empirical generalization performance on the test set. We now refer to this estimation problem as a regression problem if Y_n is continuous and a classification problem if Y_n is categorical or discrete. We also restrict ourselves to \hat{f}_Θ being a Deep Neural Network (DNN) and denote $f_\Theta := \hat{f}_\Theta$. Also, x is used for a generic input as opposed to the n^{th} given sample X_n .

DNNs are a powerful and increasingly applied machine learning framework for complex prediction tasks like object and speech recognition. In fact, they are proven to be universal function approximators [Cybenko, 1989, Hornik et al., 1989], fitting perfectly the context of function approximation of supervised learning described above. There are many flavors of DNNs, including convolutional, residual, recurrent, probabilistic, and beyond. Regardless of the actual network topology, we represent the mapping from the input signal $x \in \mathbb{R}^D$ to the output prediction $\hat{y} \in \mathbb{R}^C$ as $f_\Theta : \mathbb{R}^D \rightarrow \mathbb{R}^C$. By its parametric nature, the behavior of f_Θ is governed by its underlying parameters Θ . All current deep neural networks boil down to a composition of L "layer mappings" denoted by

$$f_\Theta(x) = (f_{\theta^{(L)}}^{(L)} \circ \dots \circ f_{\theta^{(1)}}^{(1)})(x), \quad \Theta = \{\theta^{(1)}, \dots, \theta^{(L)}\}. \quad (1)$$

In all the following cases, a neural network layer at level (ℓ) is an operator $f_{\theta^{(\ell)}}^{(\ell)}$ that takes as input a vector-valued signal $z^{(\ell-1)}(x) \in \mathbb{R}^{D^{(\ell-1)}}$ which at $\ell = 0$ is the input signal $z^{(0)}(x) := x$ and produces a vector-valued output $z^{(\ell)}(x) \in \mathbb{R}^{D^{(\ell)}}$. This succession of mappings is in general non-commutative, making the analysis of the complete sequence of generated signals crucial, denoted by

$$z^{(\ell)}(x) = (f_{\theta^{(\ell)}}^{(\ell)} \circ \dots \circ f_{\theta^{(1)}}^{(1)})(x), \ell \in \{1, \dots, L\}. \quad (2)$$

For concreteness, we will focus on processing K -channel inputs x , such as RGB images, stereo signals, as well as multi-channel representations $z^{(\ell)}, \ell = 1, \dots, L$ which we refer to as a "signal". This signal is indexed $z_{c,i,j}^{(\ell)}, c = 1, \dots, C^{(\ell)}, i = 1, \dots, I^{(\ell)}, j = 1, \dots, J^{(\ell)}, \ell = 1, \dots, L$, where i, j are usually spatial coordinates, and c is the channel. Any signal with greater index-dimensions fall under the following analysis by adaptation of the notations and operators. Hence, the volume $z^{(\ell)}$ is of shape $(C^{(\ell)}, I^{(\ell)}, J^{(\ell)})$ with $(C^{(0)}, I^{(0)}, J^{(0)}) = (K, I, J)$ and $(C^{(L)}, I^{(L)}, J^{(L)}) = (C, 1, 1)$. For consistency with the introduced




Figure 1: Reshaping of the multi-channel signal $z^{(\ell)}$ of shape $(3, 3, 3)$ to form the vector $z^{(\ell)}$ of dimension 27.

layer mappings, we will use $z^{(\ell)}$, the flattened version of $z^{(\ell)}$ as depicted in Fig. 1. The dimension of $z^{(\ell)}$ is thus $D^{(\ell)} = C^{(\ell)}I^{(\ell)}J^{(\ell)}$. In this section, we introduce the basic concepts and notations of the main used layers enabling to create state-of-the-art DNNs as well as standard training techniques to update the parameters Θ .

2.1 Layers Description

In this section we describe the common layers one can use to create the mapping f_{Θ} . The notations we introduce will be used throughout the report. We now describe the following: Fully-connected; Convolutional; Nonlinearity; Sub-Sampling; Skip-Connection; Recurrent layers.

Fully-Connected Layer A Fully-Connected (FC) layer is at the origin of DNNs known as Multi-Layer Perceptrons (MLPs) composed exclusively of FC-layers and nonlinearities. This layer performs a linear transformation of its input as

$$f_W^{(\ell)}(z^{(\ell-1)}(x)) = W^{(\ell)}z^{(\ell-1)}(x) + b^{(\ell)}. \quad (3)$$

The internal parameters $\theta^{(\ell)} = \{W^{(\ell)}, b^{(\ell)}\}$ are defined as $W^{(\ell)} \in \mathbb{R}^{D^{(\ell)} \times D^{(\ell-1)}}$ and $b \in \mathbb{R}^{D^{(\ell)}}$. This linear mapping produces an output vector $z^{(\ell)}$ of length $D^{(\ell)}$. In current topologies, FC layers are used at the end of the mapping, as layer L and $L - 1$, for their capacity to perform nonlinear dimensionality reduction in order to output C final values. However, due to their high number of degrees of freedom ($D^{(\ell)} \times D^{(\ell-1)} + D^{(\ell)}$) coupled with the unconstrained internal structure of $W^{(\ell)}$, MLPs inherit poor generalization performances for common perception tasks as demonstrated on computer vision tasks in [Zhang et al., 2016].

Convolutional Layer The greatest accuracy improvements in DNNs occurred after the introduction of the *convolutional layer*. Through convolutions, it leverages one of the most natural operation used for decades in signal processing and template matching. In fact, as opposed to the FC-layer, the convolutional layer is the corestone of DNNs dealing with perceptual tasks by leveraging the ability of convolutions to perform local feature extractions from their input. It is defined as

$$\boxed{f_C^{(\ell)}(\mathbf{z}^{(\ell-1)}(x)) = \mathbf{C}^{(\ell)} \mathbf{z}^{(\ell-1)}(x) + b^{(\ell)}}. \quad (4)$$

where a special structure is defined on $\mathbf{C}^{(\ell)}$ so that it performs multi-channel convolutions on the vector $\mathbf{z}^{(\ell-1)}$. To highlight this fact, we first remind the multi-channel convolution operation performed on the unflattened input $z^{(\ell-1)}(x)$ of shape $(C^{(\ell-1)}, I^{(\ell-1)}, J^{(\ell-1)})$ given a filter bank $W^{(\ell)}$ composed of $C^{(\ell)}$ filters, each being a 3D tensor of shape $(C^{(\ell-1)}, M^{(\ell)}, N^{(\ell)})$ with $M^{(\ell)} \leq I^{(\ell-1)}, N^{(\ell)} \leq J^{(\ell-1)}$. Hence $W \in \mathbb{R}^{C^{(\ell)} \times C^{(\ell-1)} \times M^{(\ell)} \times N^{(\ell)}}$ with $C^{(\ell-1)}$ representing the filters depth, equal to the input number of channels, and $(M^{(\ell)}, N^{(\ell)})$ the spatial size. The application of the linear filters on the signal form another multi-channel signal as

$$\begin{aligned} (W^{(\ell)} \star z^{(\ell-1)}(x))_{c,i,j} &= \sum_{k=1}^{C^{(\ell-1)}} (W_{c,k}^{(\ell)} \star z_k^{(\ell-1)}(x))_{i,j} \\ &= \sum_{k=1}^{C^{(\ell-1)}} \sum_{m=1}^{M^{(\ell)}} \sum_{n=1}^{N^{(\ell)}} W_{c,k,m,n}^{(\ell)} z_{k,i-m,j-n}^{(\ell-1)}(x), \end{aligned} \quad (5)$$

where the output of this convolution contains $C^{(\ell)}$ channels, the number of filters in $W^{(\ell)}$. Then a bias term is added for each output channel, shared across spatial positions. We denote this bias term as $\xi \in \mathbb{R}^{C^{(\ell)}}$. As a result, to create channel c of the output, we perform a 2D convolution of each channel $k = 1, \dots, C^{(\ell-1)}$ of the input with the impulse response $W_{c,k}^{(\ell)}$ and then sum those outputs element-wise over k to finally add the bias leading to $z^{(\ell)}(x)$ as

$$z_c^{(\ell)}(x) = \sum_{k=1}^{C^{(\ell-1)}} (W_{c,k}^{(\ell)} \star z_k^{(\ell-1)}(x)) + \xi_c. \quad (6)$$

In general, the input is first transformed in order to apply some boundary conditions such as zero-padding, symmetric or mirror. Those are standard padding techniques in signal processing [Mallat, 1999]. We now describe how to obtain the matrix $\mathbf{C}^{(\ell)}$ and vector $b^{(\ell)}$ corresponding to Eq. 6 but applied on the flattened input $\mathbf{z}^{(\ell-1)}(x)$ and producing $\mathbf{z}^{(\ell)}$. The matrix $\mathbf{C}^{(\ell)}$ is obtained by replicating the filter weights $W_{c,k}^{(\ell)}$ into the circulant-block-circulant matrices $\mathbf{W}_{c,k}^{(\ell)}, c = 1, \dots, C^{(\ell)}, k = 1, \dots, C^{(\ell-1)}$ [Jayaraman et al., 2009] and stacking them into the super-matrix $\mathbf{C}^{(\ell)}$

$$\mathbf{C}^{(\ell)} = \begin{bmatrix} \mathbf{W}_{1,1}^{(\ell)} & \mathbf{W}_{1,2}^{(\ell)} & \cdots & \mathbf{W}_{1,C^{(\ell-1)}}^{(\ell)} \\ \mathbf{W}_{2,1}^{(\ell)} & \mathbf{W}_{2,2}^{(\ell)} & \cdots & \mathbf{W}_{2,C^{(\ell-1)}}^{(\ell)} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{W}_{C^{(\ell)},1}^{(\ell)} & \mathbf{W}_{C^{(\ell)},2}^{(\ell)} & \cdots & \mathbf{W}_{C^{(\ell)},C^{(\ell-1)}}^{(\ell)} \end{bmatrix}. \quad (7)$$

We provide an example in Fig. 2 for $\mathbf{W}_{c,k}^{(\ell)}$ and $\mathbf{C}^{(\ell)}$. By the sharing of the bias across spatial positions,

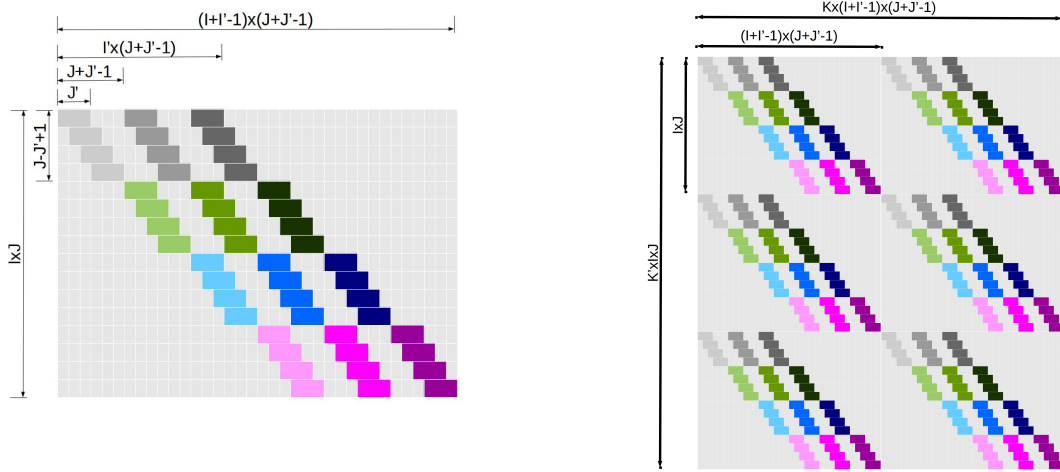


Figure 2: Left: depiction of one convolution matrix $W_{k,l}$. Right: depiction of the super convolution matrix C .

the bias term $b^{(\ell)}$ inherits a specific structure. It is defined by replicating $\xi_c^{(\ell)}$ on all spatial position of each output channel c :

$$b_d^{(\ell)} = \xi_c^{(\ell)}, d = I^{(\ell)}J^{(\ell)}(c-1) + 1, \dots, I^{(\ell)}J^{(\ell)}(c-1), c = 1, \dots, C^{(\ell)}. \quad (8)$$

The internal parameters of a convolutional layer are $\theta^{(\ell)} = \{W^{(\ell)}, \xi^{(\ell)}\}$. The number of degrees of freedom for this layer is much less than for a FC-layer, it is of $C^{(\ell)}C^{(\ell-1)}M^{(\ell)}N^{(\ell)} + C^{(\ell)}$. If the convolution is circular, the spatial size of the output is preserved leading to $(I^{(\ell)}, J^{(\ell)}) = (I^{(\ell-1)}, J^{(\ell-1)})$ and thus the output dimension only changes in the number of channels. Taking into account the special topology of the input by constraining the $C^{(\ell)}$ matrix to perform convolutions coupled with the low number of degrees of freedom while allowing a high-dimensional output leads to very efficient training and generalization performances in many perceptual tasks which we will discuss in details in 4.1. While there are still difficulties to understand what is encoded in the filters $W^{(\ell)}$, it has been empirically shown that for images, the first filter-bank $W^{(1)}$ converges toward an over-complete Gabor filter-bank, considered as a natural basis for images [Meyer, 1993, Olshausen et al., 1996]. Hence, many signal processing tools and results await to be applied for analysis.

Element-wise Nonlinearity Layer A *scalar/element-wise nonlinearity* layer applies a nonlinearity σ to each entry of a vector and thus preserve the input vector dimension $D^{(\ell)} = D^{(\ell-1)}$. As a result, this layer produces its output via application of $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ across all positions as

$$f_{\sigma}^{(\ell)}(z^{(\ell-1)}(x))_d = \sigma(z_d^{(\ell-1)}(x)), d = 1, \dots, D^{(\ell)}. \quad (9)$$

The choice of nonlinearity greatly impacts the learning and performances of the DNN as for example sigmoids and tanh are known to have vanishing gradient problems for high amplitude inputs, while ReLU based activation lead to unbounded activation and dying neuron problems. Typical choices include

- Sigmoid: $\sigma_{sig}(u) = \frac{1}{1+e^{-u}}$,
- tanh: $\sigma_{tanh}(u) := 2\sigma_{sig}(2u) - 1$,
- ReLU: $\sigma_{relu}(u) = \max(u, 0)$,
- Leaky ReLU: $\sigma_{lrelu}(u) = \max(u, 0) + \min(\eta u, 0)$, $\eta > 0$,
- Absolute Value: $\sigma_{abs}(u) = \max(u, 0) + \max(-u, 0)$.

The presence of nonlinearities in DNNs is crucial as otherwise the composition of linear layers would produce another linear layer, with factorized parameters.

Pooling Layer A *pooling* layer operates a sub-sampling operation on its input according to a sub-sampling policy ρ and a collection of regions on which ρ is applied. We denote each region to be sub-sampled by $R_d, d = 1, \dots, D^{(\ell)}$ with $D^{(\ell)}$ being the total number of pooling regions. Each region contains the set of indices on which the pooling policy is applied leading to

$$\boxed{f_{\rho}^{(\ell)}(\mathbf{z}^{(\ell-1)}(x))_d = \rho(\mathbf{z}_{R_d}^{(\ell-1)}(x)), d = 1, \dots, D^{(\ell)}} \quad (10)$$

where ρ is the pooling operator and $\mathbf{z}_{R_d}^{(\ell-1)}(x) = \{\mathbf{z}_i^{(\ell-1)}(x), i \in R_d\}$. Usually one uses mean or max pooling defined as

- Max-Pooling: $\rho_{max}(\mathbf{z}_{R_d}^{(\ell-1)}(x)) = \max_{i \in R_d} \mathbf{z}_i^{(\ell-1)}(x)$,
- Mean-Pooling: $\rho_{mean}(\mathbf{z}_{R_d}^{(\ell-1)}(x)) = \frac{1}{Card(R_d)} \sum_{i \in R_d} \mathbf{z}_i^{(\ell-1)}(x)$.

The regions R_d can be of different cardinality $\exists d_1, d_2 | Card(R_{d_1}) \neq Card(R_{d_2})$ and can be overlapping $\exists d_1, d_2 | Card(R_{d_1}) \cap Card(R_{d_2}) \neq \emptyset$. However, in order to treat all input dimension, it is natural to require that each input dimension belongs to at least one region: $\forall k \in \{1, \dots, D^{(\ell-1)}\}, \exists d \in \{1, \dots, D^{(\ell)}\} | k \in R_d$. The benefits of a pooling layer are three-fold. Firstly, by reducing the output dimension it allows for faster computation and less memory requirement. Secondly, it allows to greatly reduce the redundancy of information present in the input $\mathbf{z}^{(\ell-1)}$. In fact, sub-sampling, even though linear, is common in signal processing after filter convolutions. Finally, in case of max-pooling, it allows to only backpropagate gradients through the pooled coefficient enforcing specialization of the neurons. The latter is the corestone of the winner-take-all strategy stating that each neuron specializes into what it performs best.

Skip-Connection A skip-connection layer can be considered as a bypass connection added between the input of a layer and its output. Hence, it allows for the input of a layer such as a convolutional layer or FC-layer to be linearly combined with its own output. The added connections lead to better training stability and overall performances as there always exists a direct linear link from the input to all inner layers. Simply written, given a layer $f_{\theta^{(\ell)}}^{(\ell)}$ and its input $\mathbf{z}^{(\ell-1)}(x)$, the skip-connection layer is defined as

$$\boxed{f_s^{(\ell)}(\mathbf{z}^{(\ell-1)}(x); f_{\theta^{(\ell)}}^{(\ell)}) = \mathbf{z}^{(\ell-1)}(x) + f_{\theta^{(\ell)}}^{(\ell)}(\mathbf{z}^{(\ell-1)}(x))} \quad (11)$$




Figure 3: Depiction of a simple RNN with 3 layers. Connections highlight input-output dependencies.

In case of shape mis-match between $z^{(\ell-1)}(x)$ and $f_{\theta^{(\ell)}}^{(\ell)}(z^{(\ell-1)}(x))$, a ”reshape” operator is applied to $z^{(\ell-1)}(x)$ before the element-wise addition. Usually this is done via a spatial down-sampling and/or through a convolutional layer with filters of spatial size (1, 1).

Recurrent Finally, another type of layer is the recurrent layer which aims to act on time-series. It is defined as a recursive application along time $t = 1, \dots, T$ by transforming the input as well as using its previous output. The most simple form of this layer is a fully recurrent layer defined as

$$z^{(1,t)}(x) = \sigma \left(W^{(in,h_1)} x^t + W^{(h_1,h_1)} z^{(1,t-1)}(x) + b^{(1)} \right) \quad (12)$$

$$z^{(\ell,t)}(x) = \sigma \left(W^{(in,h_\ell)} x^t + W^{(h_{\ell-1},h_\ell)} z^{(\ell-1),t}(x) + W^{(h_\ell,h_\ell)} z^{(\ell,t-1)}(x) + b^{(\ell)} \right) \quad (13)$$

while some applications use recurrent layers on images by considering the serie of ordered local patches as a time serie, the main application resides in sequence generation and analysis especially with more complex topologies such as LSTM [Graves and Schmidhuber, 2005] and GRU [Chung et al., 2014] networks. We depict the topology example in Fig. 3.

2.2 Deep Convolutional Network

The combination of the possible layers and their order matter greatly in final performances, and while many newly developed stochastic optimization techniques allow for faster learning, a sub-optimal layer chain is almost never recoverable. We now describe a ”typical” network topology, the deep convolutional

network (DCN), to highlight the way the previously described layers can be combined to provide powerful predictors. Its main development goes back to [LeCun et al., 1995] for digit classification. A DCN is defined as a succession of blocks made of 3 layers : Convolution \rightarrow Element-wise Nonlinearity \rightarrow Pooling layer. In a DCN, several of such blocks are cascaded end-to-end to create a sequence of activation maps followed usually by one or two FC-layers. Using the above notations, a single block can be rewritten as $(f_\rho \circ f_\sigma \circ f_C)$. Hence a basic model with 2 blocks and 2 FC-layers is defined as

$$\mathbf{z}^{(4)}(x) = \underbrace{(f_W^{(4)} \circ f_W^{(3)})}_{\text{MLP}} \circ \underbrace{(f_\rho^{(2)} \circ f_\sigma^{(2)} \circ f_C^{(2)})}_{\text{Block 2}} \circ \underbrace{(f_\rho^{(1)} \circ f_\sigma^{(1)} \circ f_C^{(1)})}_{\text{Block 1}}(x), \quad (14)$$

where for convenience we treat the linear transformation, nonlinearity and pooling as a one layer mapping and the subscript represents the role of each one. The astonishing results that a DCN can achieve comes intuitively from the ability of the blocks to convolve the learned filter-banks with their input, "separating" the underlying features present relative to the task at hand. This is followed by a nonlinearity and a spatial sub-sampling to compress and reduce the redundant representation. Finally, the MLP part simply acts as a nonlinear classifier, the final key for prediction. The duality in the representation/high-dimensional mappings followed by dimensionality reduction/classification is a core concept in machine learning referred as: pre-processing-classification.

2.3 Learning

In order to optimize all the weights Θ leading to the predicted output $\hat{y}(x)$, one disposes of (1) a labeled dataset $\mathcal{D} = \{(X_n, Y_n), n = 1, \dots, N\}$, (2) a loss function $\mathcal{L} : \mathbb{R}^C \times \mathbb{R}^C \rightarrow \mathbb{R}$, (3) a learning policy to update the parameters Θ . In the context of classification, the target variable Y_n associated to an input X_n is a categorical target $Y_n \in \{1, \dots, C\}$. In order to predict such target, the output of the last layer of a network $\mathbf{z}^{(L)}(X_n)$ is transformed via a softmax nonlinearity [de Brébisson and Vincent, 2015]. It is used to transform $\mathbf{z}^{(L)}(X_n)$ into a probability distribution defined as

$$\hat{y}_c(X_n) = \frac{e^{\mathbf{z}_c^{(L)}(X_n)}}{\sum_c e^{\mathbf{z}_c^{(L)}(X_n)}}, \in (0, 1), \forall c \quad (15)$$

thus leading to $\hat{y}_c(X_n)$ representing $\mathbb{P}(\text{class of } X_n \text{ is } c | X_n)$, with $\sum_{c=1}^C \hat{y}_c(X_n) = 1$. The used loss function quantifying the distance between $\hat{y}(X_n)$ and Y_n is the cross-entropy (CE) defined as

$$\mathcal{L}_{CE}(Y_n, \hat{y}(X_n)) = -\log(\hat{y}_{Y_n}(X_n)) \quad (16)$$

$$= -\mathbf{z}_{Y_n}^{(L)}(X_n) + \log\left(\sum_{c=1}^C e^{\mathbf{z}_c^{(L)}(X_n)}\right). \quad (17)$$

For regression problems, the target Y_n is continuous and thus the final DNN output is taken as the prediction $\hat{y}(X_n) = \mathbf{z}^{(L)}(X_n)$. The loss function \mathcal{L} is usually the ordinary squared error (SE) defined as

$$\mathcal{L}_{SE}(Y_n, \hat{y}(X_n)) = \sum_{c=1}^C (Y_{n,c} - \hat{y}(X_n)_c)^2. \quad (18)$$

Since all of the operations introduced above in standard DNNs are differentiable almost everywhere with

respect to their parameters and inputs, given a training set and a loss function, one defines an update strategy for the weights Θ . This takes the form of an iterative scheme based on a first order iterative optimization procedure. Updates for the weights are computed on each input and usually averaged over *mini-batches* containing B exemplars with $B \ll N$. This produces an estimate of the "correct" update for Θ and is applied after each mini-batch. Once all the training instances of \mathcal{D} have been seen, this terminates an *epoch*. The dataset is then shuffled and this procedure is performed again. Usually a network needs hundreds of epochs to converge. For any given iterative procedure, the updates are computed for all the network parameters by *backpropagation* [Hecht-Nielsen et al., 1988], which follows from applying the chain rule of calculus. Common policies are Gradient Descent (GD) [Rumelhart et al., 1988] being the simplest application of backpropagation, Nesterov Momentum [Bengio et al., 2013] that uses the last performed updates in order to "accelerate" convergence and finally more complex adaptive methods which can have internal parameters updated based on the weights/updates statistics such as Adam [Kingma and Ba, 2014], Adadelta [Zeiler, 2012], Adagrad [Duchi et al., 2011], RMSprop [Tieleman and Hinton, 2012], ... For the remaining of the report and experiments we will adopt the Adam method, details are provided in the appendix. Finally, to measure the actual performance of a trained network in the context of classification, one uses the accuracy loss defined as

$$\mathcal{L}_{AC}(Y_n, \hat{y}(X_n)) = -1_{\{Y_n = \operatorname{argmax}_c \hat{y}_c(X_n)\}}, \quad (19)$$

defined such that smaller value is better,

3 Understand Deep Neural Networks Internal Mechanisms

In this section we develop spline operators, a generalization of spline function which are also generalized DNN layers. By doing so, we will open DNNs to explicit analysis and especially understand their behavior and potential through the spline region inference. DNNs will be shown to leverage template matching, a standard technique in signal processing to tackle perception tasks.

Let first motivate the need to adopt the theory of spline functions for deep learning and machine learning in general. As described in Sec. 2, the task at hand is to use parametric functionals f_Θ to be able to understand [Cheney, 1980], predict, interpolate the world around us [Reinsch, 1967]. For example, partial differential equations allow to approximate real world physics [Bloor and Wilson, 1990, Smith, 1985] based on grounded principles. For this case, one knows the underlying laws that must be fulfilled by f_Θ . In machine learning however, one only disposes of N observed inputs $X_n, n = 1, \dots, N$ or input-output pairs $(X_n, Y_n)_{n=1}^N$. To tackle this approximation problem, splines offer great advantages. From a computational regard, polynomials are very efficient to evaluate via for example the Horner scheme [Peña, 2000]. Yet, polynomials have "chaotic" behaviors especially as their degree grows, leading to the Runge's phenomenon [Boyd and Xu, 2009], synonym of extremely poor interpolation and extrapolation capacities. On the other hand, low degree polynomials are not flexible enough for modeling arbitrary functionals. Splines, however, are defined as a collection of polynomials each one acting on a specific region of the input space. The collection of possible regions $\Omega = \{\omega_1, \dots, \omega_R\}$ forms a partition of the input space. On each of those regions ω_r , the associated and usually low degree polynomial ϕ_r is used to transform the input x . Through the per region activation, splines allow to model highly nonlinear functional, yet, the low-degree polynomials avoid the Runge phenomenon. Hence, splines are the tools of choice for functional approximation if one seeks robust interpolation/extrapolation performances without sacrificing the modeling of potentially very irregular underlying mappings f .

In fact, as we will now describe in details, current state-of-the-art DNNs are linear spline functions and we now proceed to develop the notations and formulation accordingly. Let first remind briefly the case of multivariate linear splines.

Definition 1. Given a partition $\Omega = \{\omega_1, \dots, \omega_R\}$ of \mathbb{R}^D , we denote **multivariate spline** with local mappings $\Phi = \{\phi_1, \dots, \phi_R\}$, with $\phi_{\omega_r} : \mathbb{R}^D \rightarrow \mathbb{R}$ the mapping

$$s[\Phi, \Omega](x) = \sum_{r=1}^R \phi_r(x) 1_{\{x \in \omega_r\}} \quad (20)$$

$$= \phix, \quad (21)$$

where the input dependent selection is abbreviated via

$$\phi[x] := \phi_r \text{ s.t. } x \in \omega_r. \quad (22)$$

If the local mappings ϕ_r are linear we have $\phi_r(x) = \langle a_r, x \rangle + b_r$ with $a_r \in \mathbb{R}^D$ and $b_r \in \mathbb{R}$. We denote this

functional as a **multivariate linear spline**:

$$s[\mathbf{a}, \mathbf{b}, \Omega](x) = \sum_{r=1}^R (\langle a_r, x \rangle + b_r) 1_{\{x \in \omega_r\}} \quad (23)$$

$$= a[x]^T x + b[x]. \quad (24)$$

where we explicit the polynomial parameters by $\mathbf{a} = \{a_1, \dots, a_R\}$ and $\mathbf{b} = \{b_1, \dots, b_R\}$.

In the next sections, we study the capacity of linear spline operators to span standard DNN layers. All the development of the spline operator as well as a detailed review of multivariate spline functions is contained in Appendix A.1.1. Afterwards, the composition of the developed linear operators will lead to the explicit analytical input-output mapping of DNNs allowing to derive all the theoretical results in the remaining of the report. In the following sections, we omit the cases of regularity constraints on the presented functional thus leading to the most general cases.

3.1 Spline Operators[FINI]

A natural extension of spline functions is the **spline operator** (SO) we denote $S : \mathbb{R}^D \rightarrow \mathbb{R}^K, K > 1$. We present here a general definition and propose in the next section an intuitive way to construct spline operators via a collection of multivariate splines, the special case of current DNNs.

Definition 2. A spline operator is a mapping $S : \mathbb{R}^D \rightarrow \mathbb{R}^K$ defined by a collection of local mappings $\Phi^S = \{\phi_r^S : \mathbb{R}^D \rightarrow \mathbb{R}^K, r = 1, \dots, R\}$ associated with a partition of \mathbb{R}^D denoted as $\Omega^S = \{\omega_r^S, r = 1, \dots, R\}$ s.t.

$$\begin{aligned} S[\Phi^S, \Omega^S](x) &= \sum_{r=1}^R \phi_r^S(x) 1_{\{x \in \omega_r^S\}} \\ &= \phi^Sx, \end{aligned}$$

where we denoted the region specific mapping associated to the input x by $\phi^S[x]$.

A special case occurs when the mappings ϕ_r^S are linear. We thus define in this case the **linear spline operator** (LSO) which will play an important role for DNN analysis. In this case, $\phi_r^S(x) = A_r x + b_r$, with $A_r \in \mathbb{R}^{K \times D}, b_r \in \mathbb{R}^K, \forall r$. As a result, a LSO can be rewritten as

$$\begin{aligned} S[\mathbf{A}, \mathbf{b}, \Omega^S](x) &= \sum_{r=1}^R (A_r x + b_r) 1_{\{x \in \omega_r^S\}} \\ &= A[x]x + b[x] \end{aligned}$$

where we denoted the collection of intercept and biases as $\mathbf{A} = \{A_r, r = 1, \dots, R\}$, $\mathbf{b} = \{b_r, r = 1, \dots, R\}$ and finally the input specific activation as $A[x]$ and $b[x]$.

Such operators can also be defined via a collection of multivariate polynomial (resp. linear) splines. Given K multivariate spline functions $s[\Phi_k, \Omega_k] : \mathbb{R}^D \rightarrow \mathbb{R}, k = 1, \dots, K$, their respective output is

”stacked” to produce an output vector of dimension K . The internal parameters of each multivariate spline are Ω_k , a partition of \mathbb{R}^D with $Card(\Omega_k) = R_k$ and $\Phi_k = \{\phi_1, \dots, \phi_{R_k}\}$. Stacking their respective output to form an output vector leads to the induced spline operator $S \left[(s[\Phi_k, \Omega_k])_{k=1}^K \right]$.

Definition 3. The spline operator $S : \mathbb{R}^D \rightarrow \mathbb{R}^K$ defined with K multivariate splines $(s[\Phi_k, \Omega_k])_{k=1}^K$ with $s[\Phi_k, \Omega_k] : \mathbb{R}^D \rightarrow \mathbb{R}$ is defined as

$$S \left[(s[\Phi_k, \Omega_k])_{k=1}^K \right] (x) = \begin{bmatrix} s[\Phi_1, \Omega_1](x) \\ \vdots \\ s[\Phi_K, \Omega_K](x) \end{bmatrix}. \quad (25)$$

with $\Omega_k = \{\omega_{k,1}, \dots, \omega_{k,R_k}\}$, $\Phi_k = \{\phi_{k,r} \in \mathbb{R}^D, r = 1, \dots, R_k\}$, $k = 1, \dots, K$.

The use of K multivariate splines to construct a SO does not provide directly the explicit collection of mappings and regions Φ^S, Ω^S . Yet, it is clear that the SO is jointly governed by all the individual multivariate splines. Let first present some intuitions on this fact. The spline operator output is computed with each of the K splines having ”activated” a region specific functional depending on their own input space partitioning. In particular, each of the region ω_r^S of the input space leading to a specific joint configuration ϕ_r^S is the one of interest, leading to Ω^S and Φ^S . We can thus write explicitly the new regions of the spline operator based on the ensemble of partition of all the involved multivariate splines as

$$\Omega^S = \left(\bigcup_{(\omega_1, \dots, \omega_K) \in \Omega_1 \times \dots \times \Omega_K} \left\{ \bigcap_{k \in \{1, \dots, K\}} \omega_k \right\} \right) \setminus \{\emptyset\}. \quad (26)$$

We also denote the number of region associated to this SO as $R^S = Card(\Omega^S)$. From this, the local mappings of the SO ϕ_r^S correspond to the joint mappings of the splines being activated on ω_r^S we denote

$$\phi^S[\omega_r^S](x) = \begin{bmatrix} \phi_1[\omega_r^S](x) \\ \vdots \\ \phi_K[\omega_r^S](x) \end{bmatrix}, \quad (27)$$

with $\phi_k[\omega_r^S] = \phi_{k,q} \in \Phi_k$ s.t. $\omega_r^S \subset \omega_{k,q}$. In fact, for each region ω_r^S of the SO there is a unique region ω_{k,q_k} for each of the splines $k = 1, \dots, K$, such that it is a subset as $\omega_r^S \subset \omega_{k,q_k}$ and it is disjoint to all others $\omega_r^S \cap \omega_{k,l} = \emptyset, \forall l \neq q_k, k = 1, \dots, K$. In other word we have the following property:

$$\forall (r, k) \in \{1, \dots, R^S\} \times \{1, \dots, K\}, \exists! q_k \in \{1, \dots, R_k\} \text{ s.t. } \omega_r^S \cap \omega_{k,l} = \begin{cases} \omega_r^S, & l = q_k \\ \emptyset, & l \neq q_k \end{cases}, \quad (28)$$

as we remind $\omega_r^S \cap \omega_{k,l} = \omega_r^S \iff \omega_r^S \subset \omega_{k,q_k}$.




Figure 4: Illustrative examples of the new partition Ω^S given two partitions Ω_1, Ω_2 with $I_1 = 2, I_2 = 3$.

This leads to the following SO formulation

$$\begin{aligned}
S \left[(s[\Phi_k, \Omega_k])_{k=1}^K \right] (x) &= \sum_{r=1}^{R^S} \begin{bmatrix} s[\Phi_1, \Omega_1](x) \\ \vdots \\ s[\Phi_K, \Omega_K](x) \end{bmatrix} 1_{\{x \in \omega_r^S\}} \\
&= \sum_{r=1}^{R^S} \begin{bmatrix} \phi_1[\omega_r^S](x) \\ \vdots \\ \phi_K[\omega_r^S](x) \end{bmatrix} 1_{\{x \in \omega_r^S\}} \\
&= \sum_{r=1}^{R^S} \phi^S[\omega_r^S](x) 1_{\{x \in \omega_r^S\}} \\
&= \phi^Sx, \tag{29}
\end{aligned}$$

We can now study the case of linear splines leading to LSOs. If a SO is constructed via aggregation of linear multivariate splines,. The linear property allows notation simplifications. It is defined as

$$\begin{aligned}
S \left[(s[\mathbf{a}_k, \mathbf{b}_k, \Omega_k])_{k=1}^K \right] (x) &= \sum_{r=1}^{R^S} \begin{bmatrix} s[\mathbf{a}_1, \mathbf{b}_1, \Omega_1](x) \\ \vdots \\ s[\mathbf{a}_K, \mathbf{b}_K, \Omega_K](x) \end{bmatrix} 1_{\{x \in \omega_r^S\}} \\
&= \sum_{r=1}^{R^S} \begin{bmatrix} \langle a_1[\omega_r^S], x \rangle + b_1[\omega_r^S] \\ \vdots \\ \langle a_K[\omega_r^S], x \rangle + b_K[\omega_r^S] \end{bmatrix} 1_{\{x \in \omega_r^S\}} \\
&= \sum_{r=1}^{R^S} \left(\begin{bmatrix} a_1[\omega_r^S]^T \\ \vdots \\ a_K[\omega_r^S]^T \end{bmatrix} x + \begin{bmatrix} b_1[\omega_r^S] \\ \vdots \\ b_K[\omega_r^S] \end{bmatrix} \right) 1_{\{x \in \omega_r^S\}} \\
&= \sum_{r=1}^{R^S} (A[\omega_r^S]x + b[\omega_r^S]) 1_{\{x \in \omega_r^S\}} \\
&= A[x]x + b[x], \tag{30}
\end{aligned}$$

with $\Omega_k = \{\omega_{k,1}, \dots, \omega_{k,R_k}\}$, $\mathbf{a}_k = \{a_{k,r} \in \mathbb{R}^D, r = 1, \dots, R_k\}$, $\mathbf{b}_k = \{b_{k,r} \in \mathbb{R}, r = 1, \dots, R_k\}$.

As it is clear, the collection of matrices and biases and the partitions completely define a LSO. Hence, we denote the set of all possible matrices and biases as $\mathbf{A} = \{A[\omega], \omega \in \Omega^S\}$, $\mathbf{b} = \{b[\omega], \omega \in \Omega^S\}$. Any LSO is thus written as $S[\mathbf{A}, \mathbf{b}, \Omega^S]$.

ReLU [Glorot et al., 2011]	LReLU [Xu et al., 2015]	Abs.Value
$\Omega_d = \{\omega_{d,1}, \omega_{d,2}\},$ $\omega_{d,1} = \{x \in \mathbb{R}^d : x_d > 0\},$ $\omega_{d,2} = \{x \in \mathbb{R}^d : x_d \leq 0\}$ $a_{d,1} = \mathbf{e}_d, a_{d,2} = 0,$ $b_{k,1} = b_{k,2} = 0,$	$\Omega_d = \{\omega_{d,1}, \omega_{d,2}\},$ $\omega_{d,1} = \{x \in \mathbb{R}^d : x_d > 0\},$ $\omega_{d,2} = \{x \in \mathbb{R}^d : x_d \leq 0\}$ $a_{d,1} = \mathbf{e}_d, a_{d,2} = \eta \mathbf{e}_d, \eta > 0,$ $b_{k,1} = b_{k,2} = 0,$	$\Omega_d = \{\omega_{d,1}, \omega_{d,2}\},$ $\omega_{d,1} = \{x \in \mathbb{R}^d : x_d > 0\},$ $\omega_{d,2} = \{x \in \mathbb{R}^d : x_d \leq 0\}$ $a_{d,1} = \mathbf{e}_d, a_{d,2} = -\mathbf{e}_d,$ $b_{k,1} = b_{k,2} = 0,$

Table 1: Example of multivariate splines associated with standard DNN nonlinearities, note that they are exact by definition as opposed to a linear spline approximating a sigmoid function for example, for which approximation can be made arbitrarily close.

3.2 Linear Spline Operator: Generalized Neural Network Layers

In this section we demonstrate how current DNNs are expressed as composition of LSOs. We first proceed to describe layer specific notations and analytical formula to finally perform composition of LSOs providing analytical DNN mappings in the next section.

3.2.1 Nonlinearity layers

We first analyze the elementwise nonlinearity layer. Our analysis deals with any given nonlinearity. If this nonlinearity is by definition a spline s.a. with ReLU, leaky-ReLU, absolute value, they fall directly into this analysis. If not, arbitrary functions such as tanh, sigmoid are approximated via linear splines. We remind that a nonlinearity layer $f_\sigma^{(\ell)}$ is defined by applying a nonlinearity σ on each input dimension of its input $\mathbf{z}^{(\ell-1)} \in \mathbb{R}^{D^{(\ell-1)}}$ and produces a new output vector $\mathbf{z}^{(\ell)} \in \mathbb{R}^{D^{(\ell)}}$. While in general the used nonlinearity is the same applied on each dimension we present here a more general case where one has a specific σ_d per dimension. In addition, we present the case where the nonlinearity might not act on only one input dimension but any part of it. We thus define by $\sigma^{(\ell)}[\mathbf{a}_d, \mathbf{b}_d, \Omega_d] : \mathbb{R}^{D^{(\ell-1)}} \rightarrow \mathbb{R}$ the nonlinearity acting on the d^{th} input dimension, $\mathbf{z}_d^{(\ell-1)}$. We provide illustration of famous nonlinearities in Table 1 being cases where the output dimension at position d only depends on the input dimension of d .

Given a collection of $\mathcal{D}^{(\ell)}$ such nonlinearities $(\sigma[\mathbf{a}_d, \mathbf{b}_d, \Omega_d])_{d=1}^{D^{(\ell)}}$ of such linear splines, we define the spline operator which defined an actual nonlinear layer as $S_\sigma^{(\ell)}[\mathbf{A}_\sigma, \mathbf{b}_\sigma, \Omega_\sigma^S]$ with the induced matrix and vector parameters as defined in the previous section. Hence we have

$$\begin{aligned}
f_\sigma^{(\ell)}(\mathbf{z}^{(\ell-1)}) &= S_\sigma^{(\ell)}[\mathbf{A}_\sigma, \mathbf{b}_\sigma, \Omega_\sigma^S](\mathbf{z}^{(\ell-1)}) \\
&= \mathbf{A}_\sigma^{(\ell)}[\mathbf{z}^{(\ell-1)}] + \mathbf{b}[\mathbf{z}^{(\ell-1)}].
\end{aligned} \tag{31}$$

We also have for the provided examples

$$\begin{aligned}\mathbf{b}_{relu} &= \mathbf{b}_{lrelu} = \mathbf{b}_{abs} = \{\mathbf{0}\}, \\ \mathbf{A}_{relu} &= \{\text{diag}(v) | v \in \{0, 1\}^K\}, \\ \mathbf{A}_{lrelu} &= \{\text{diag}(v) | v \in \{\eta, 1\}^K\}, \\ \mathbf{A}_{abs} &= \{\text{diag}(v) | v \in \{-1, 1\}^K\}.\end{aligned}$$

In order to better demonstrate the underlying spline operator mappings induced by typical DNN nonlinearities we provide a detailed example for the ReLU case below with $D^{(\ell-1)} = 3, D^{(\ell)} = 3$, we now omit the layer notation for the example. We have according to the presented definition that the slopes of the splines are

$$\mathbf{A}_{relu} = \left\{ \begin{aligned} A_{(2,2,2)} &= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, A_{(2,2,1)} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}, A_{(2,1,2)} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}, A_{(2,1,1)} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \\ A_{(1,2,2)} &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, A_{(1,2,1)} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}, A_{(1,1,2)} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}, A_{(1,1,1)} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \end{aligned} \right\},$$

$$\mathbf{b}_{relu} = \{\mathbf{0}\},$$

$$\begin{aligned}\Omega_{relu}^S &= \{\omega_{(2,2,2)}^S, \omega_{(2,2,1)}^S, \omega_{(2,1,2)}^S, \omega_{(2,1,1)}^S, \omega_{(1,2,2)}^S, \omega_{(1,2,1)}^S, \omega_{(1,1,2)}^S, \omega_{(1,1,1)}^S\}, \\ \omega_{(2,2,2)}^S &= \{x \in \mathbb{R}^3 : x_1 < 0, x_2 < 0, x_3 < 0\}, \\ \omega_{(2,2,1)}^S &= \{x \in \mathbb{R}^3 : x_1 < 0, x_2 < 0, x_3 \geq 0\}, \\ \omega_{(2,1,2)}^S &= \{x \in \mathbb{R}^3 : x_1 < 0, x_2 \geq 0, x_3 < 0\}, \\ \omega_{(2,1,1)}^S &= \{x \in \mathbb{R}^3 : x_1 < 0, x_2 \geq 0, x_3 \geq 0\}, \\ \omega_{(1,2,2)}^S &= \{x \in \mathbb{R}^3 : x_1 \geq 0, x_2 < 0, x_3 < 0\}, \\ \omega_{(1,2,1)}^S &= \{x \in \mathbb{R}^3 : x_1 \geq 0, x_2 < 0, x_3 \geq 0\}, \\ \omega_{(1,1,2)}^S &= \{x \in \mathbb{R}^3 : x_1 \geq 0, x_2 \geq 0, x_3 < 0\}, \\ \omega_{(1,1,1)}^S &= \{x \in \mathbb{R}^3 : x_1 \geq 0, x_2 \geq 0, x_3 \geq 0\},\end{aligned}$$

Max-Pooling	Mean-Pooling
$\Omega_d = \{\omega_{d,1}, \dots, \omega_{d,Card(R_d)}\},$ $\omega_{d,r} = \{x \in \mathbb{R}^d r = \operatorname{argmax} x_{R_d}\}$ $[a_{d,r}]_l = \mathbf{e}_{R_d(r)}, b_{d,r} = 0,$	$\Omega_k = \{\mathbb{R}^d\},$ $a_{d,1} = \frac{1}{Card(R_d)} \left(\sum_{i \in R_d} \mathbf{e}_i \right), b_{k,1} = 0, \forall k$

Table 2: Example of multivariate splines associated with standard DNN pooling.

and we thus have when given some inputs that

$$\begin{aligned}
S_{relu}[\mathbf{A}_{relu}, \mathbf{b}_{relu}, \Omega_{relu}^S]([1, 2, 3]^T) &= A_{(1,1,1)}[1, 2, 3]^T + b_{(1,1,1)} \\
&= [1, 2, 3]^T, \\
S_{relu}[\mathbf{A}_{relu}, \mathbf{b}_{relu}, \Omega_{relu}^S]([1, -2, 3]^T) &= A_{(1,2,1)}[1, -2, 3]^T + b_{(1,2,1)} \\
&= [1, 0, 3]^T, \\
S_{relu}[\mathbf{A}_{relu}, \mathbf{b}_{relu}, \Omega_{relu}^S]([-1, -2, -3]^T) &= A_{(2,2,2)}[-1, -2, -3]^T + b_{(2,2,2)} \\
&= [0, 0, 0]^T.
\end{aligned}$$

3.2.2 Sub-Sampling layers

We now study the case of sub-sampling layers. We first remind briefly that a pooling layer $f_\rho^{(\ell)}$ is defined by a pooling policy ρ and a collection of regions $R_d, d = 1, \dots, d^{(\ell)}$. Each of these regions contain indices on which ρ will be applied to form the output vector $\mathbf{z}^{(\ell)} \in \mathbb{R}^{D^{(\ell)}}$. As for the nonlinear layer, it is common to use the same pooling policy across all regions, yet we now formula the spline functional for the more general case of a pooling per region ρ_d by $\rho^{(\ell)}[\mathbf{a}_d, \mathbf{b}_d, \Omega_d] : \mathbb{R}^{D^{(\ell-1)}} \rightarrow \mathbb{R}$. We provide illustration of standard cases in Table 2. Similarly to the nonlinearity case we can now define the spline operator. Again, given a collection $\left(\rho[\mathbf{a}_d, \mathbf{b}_d, \Omega_d] \right)_{d=1}^{D^{(\ell)}}$ of such affine splines we can create the spline operator denoted by $S_\rho[\mathbf{A}_\rho, \mathbf{b}_\rho, \Omega_\rho^S]$. For the max-pooling policy we have

$$\mathbf{A}_{max} = \left\{ \left[\begin{array}{c} e_1^T \\ \vdots \\ e_{D^{(\ell)}}^T \end{array} \right], e_d \in \{\mathbf{e}_d, d \in R_d\}, d = 1, \dots, D^{(\ell)} \right\}, \mathbf{b}_{max} = \{\mathbf{0}\}, \quad (32)$$

We now present an illustrative example of the max-pooling LSO with $D^{(\ell-1)} = 4, D^{(\ell)} = 2$ and $R_1 = \{1, 2\}, R_2 = \{3, 4\}$ which corresponds to pooling over non-overlapping regions of size 2. We thus have

$$\mathbf{A}_{max} = \left\{ A_{(1,1)} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}, A_{(1,2)} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \right. \\ \left. A_{(2,1)} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}, A_{(2,2)} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \right\},$$

$$\mathbf{b}_{max} = \{\mathbf{0}\},$$

$$\mathbf{\Omega}_{max} = \{\omega_{(1,1)}^S, \omega_{(1,2)}^S, \omega_{(2,1)}^S, \omega_{(2,2)}^S\}, \\ \omega_{(1,1)}^S = \{x \in \mathbb{R}^4 : 1 = \operatorname{argmax}_{d \in \{1,2\}} x_d, 3 = \operatorname{argmax}_{d \in \{3,4\}} x_d\}, \\ \omega_{(1,2)}^S = \{x \in \mathbb{R}^4 : 1 = \operatorname{argmax}_{d \in \{1,2\}} x_d, 4 = \operatorname{argmax}_{d \in \{3,4\}} x_d\}, \\ \omega_{(2,1)}^S = \{x \in \mathbb{R}^4 : 2 = \operatorname{argmax}_{d \in \{1,2\}} x_d, 3 = \operatorname{argmax}_{d \in \{3,4\}} x_d\}, \\ \omega_{(2,2)}^S = \{x \in \mathbb{R}^4 : 2 = \operatorname{argmax}_{d \in \{1,2\}} x_d, 4 = \operatorname{argmax}_{d \in \{3,4\}} x_d\},$$

and we thus have when given some inputs that

$$S_{max}[\mathbf{A}_{max}, \mathbf{b}_{max}, \mathbf{\Omega}_{max}]([1, 2, 3, 4]^T) = A_{(2,2)}[1, 2, 3, 4]^T + b_{(2,2)} \\ = [2, 4]^T, \\ S_{max}[\mathbf{A}_{max}, \mathbf{b}_{max}, \mathbf{\Omega}_{max}]([1, -2, 3, 1]^T) = A_{(1,1)}[1, -2, 3, 1]^T + b_{(1,1)} \\ = [1, 3]^T, \\ S_{max}[\mathbf{A}_{max}, \mathbf{b}_{max}, \mathbf{\Omega}_{max}]([-1, -2, -3, 0]^T) = A_{(1,2)}[-1, -2, -3, 0]^T + b_{(1,2)} \\ = [-1, 0]^T.$$

3.2.3 Linear layers:FC and convolutional

Finally, in order to provide a complete spline interpretation of DNN layers we present the case of linear layers as convolutional and FC layers. By definition of being linear mappings, they are equivalent to a spline operator with one region corresponding to the input space. We thus define this operator as

$$S_W^{(\ell)}[\{W^{(\ell)}\}, \{b^{(\ell)}\}, \{\mathbb{R}^{D^{(\ell-1)}}\}](z^{(\ell-1)}) = W^{(\ell)}z^{(\ell-1)} + b^{(\ell)}, \quad \text{FC-layer} \quad (33)$$

$$S_C^{(\ell)}[\{C^{(\ell)}\}, \{b^{(\ell)}\}, \{\mathbb{R}^{D^{(\ell-1)}}\}](z^{(\ell-1)}) = C^{(\ell)}z^{(\ell-1)} + b^{(\ell)}, \quad \text{convolutional layer} \quad (34)$$

where we shall omit the trivial parameters and denote $S_W^{(\ell)} := S_W^{(\ell)}[\{W^{(\ell)}\}, \{b^{(\ell)}\}, \{\mathbb{R}^{D^{(\ell-1)}}\}]$ and $S_C^{(\ell)} := S_C^{(\ell)}[\{C^{(\ell)}\}, \{b^{(\ell)}\}, \{\mathbb{R}^{D^{(\ell-1)}}\}]$. We derived all the notations and gave examples on how to define standard DNNs layers via linear spline operators. We can now move to the composition of such operators defining

the complete DNN mappings.

3.3 Deriving Analytical DNNs Mappings to Explicit their Faculty to Perform Template Matching

To perform perceptual tasks such as object recognition, a standard technique is template matching. It aims at detecting the presence in the input of a class specific template even if the template in the input has suffered some perturbation. Template matching is well studied when the template perturbation belongs to the standard groups of natural deformations s.a. translation, rotation for example and this process is usually referred as elastic matching. There are also been extension to perform a hierarchical elastic matching in [Zhang et al., 1997, Bajcsy and Kovačič, 1989, Burr, 1981] by marginalizing out layer after layer all the possible local perturbation. Many extensions have also been studied to model more complex diffeomorphisms as in [Korman et al., 2013, Kim and De Araújo, 2007]. A detailed review of elastic matching is proposed in [Uchida and Sakoe, 2005]. This task can also be formulated as a problem of best basis selection where the optimal atom is the correct template with the input adapted perturbation. Concept of input dependent basis has been well studied for example in [Coifman and Wickerhauser, 1992, Tropp, 2004, Mallat, 2008, Berger et al., 1994]. Yet, the need for exact mathematical modeling of the template transformations limit the ability to produce algorithms flexible enough to learn classes of diffeomorphisms in a complete data driven, parametric learning approach. As we will see, this is performed by state-of-the-art DNNs.

3.3.1 Composition of Splines for Explicit DNN Template Matching

As demonstrated in the previous sections, neural network layers are special cases of LSOs. From the derived notation and LSOs of the previous section, we can now proceed to rewrite the complete DNN mapping as composition of such operators. Firstly, we define $S^{(\ell)}[\mathbf{A}^{(\ell)}, \mathbf{b}^{(\ell)}, \Omega^{S^{(\ell)}}] := S_{\theta^{(\ell)}}^{(\ell)}$. We thus have for any DNN

$$f_{\Theta}(x) \approx (S_{\theta^{(L)}}^{(L)} \circ \dots \circ S_{\theta^{(1)}}^{(1)})(x), \quad \Theta = \{\theta^{(1)}, \dots, \theta^{(L)}\}, \quad (35)$$

where the approximation becomes an equality if the used layers are splines s.a. with ReLU, max-pooling. If not, arbitrary close approximation schemes can be found. This provides a very intuitive result from this composition of linear mappings.

Theorem 1. *Any deep network f_{Θ} made of LSOs s.a. max-pooling, ReLU, leaky ReLU, ... is itself a LSO of the form*

$$f_{\Theta}(x) = A[x]x + b[x], \forall x. \quad (36)$$

For the case where the layers are not natural LSOs s.a. with tanh, sigmoid nonlinearities, then, it can always be approximated arbitrarily closely by a affine spline operator and thus

$$f_{\Theta}(x) \approx A[x]x + b[x], \forall x. \quad (37)$$

In fact, it has been shown that linear splines can approximate any functions arbitrarily closely [Nishikawa, 1998]. In addition, using linear spline approximations is computationally efficient. For example, using ultra fast sigmoid (a linear spline version) instead of the standard sigmoid results in almost $2 \times$ speedup for 100M float64 elements on a Core2 Duo @ 3.16 GHz [Bastien et al., 2012].

As opposed to previous work studying DNNs as composition of linear mappings in the special case of ReLU coupled with mean or max pooling [Rister and Rubin, 2017], we extend the results to arbitrary DNNs by allowing linear spline approximation of non spline functionals as well as general piecewise linear splines.

We thus propose to bridge the concept of input adaptive representations with DNNs. To do so we leverage the fact shown above that any DNN can be rewritten as a linear mapping as $f_{\Theta}(x) = A[x]x + b[x]$, with input dependent intercept $A[x]$ and biases $b[x]$. We thus propose the following definition.

Definition 4. *As any DNN can be rewritten $f_{\Theta}(x) = A[x]x + b[x]$, we denote $A[x]$ as the template of the DNN mapping, and specifically $A[x]_c$, the template associated with class c . By nature of the underlying LSOs, the input adaptive template matching is induced by the per region coefficients making DNNs effective hierarchical template matching algorithms.*

We now proceed to write the analytical mapping defined by this composition of LSOs. For clarity, we only set one MLP layer at the end of the DNNs, extensions to any number of MLP layers is straightforward by adding a simple product term over those layers. Finally, while providing formula for 3 standard topologies, we also aim at presenting the methodology in order for one to generalize the presented results to any used topology, as only replacement of some operators will lead to any possible DNN topology. With this layer notation we can naturally derive the formula for the output $z^{(L)}(x)$ of any DNN with L being the number of layers in the mapping. In fact, as $S_{\theta^{(\ell)}}^{(\ell)}$ is a LSO at each layer ℓ we have the output expression given by

$$z^{(L)}(x) = W^{(L)} \left(\underbrace{\left(\prod_{\ell=L-1}^1 A_{\theta^{(\ell)}}^{(\ell)} \right) x + \sum_{\ell=1}^{L-1} \left(\prod_{j=L-1}^{\ell+1} A_{\theta^{(j)}}^{(j)} \right) b_{\theta^{(\ell)}}^{(\ell)}}_{\text{Convolutional Layers}} \right) + b^{(L)} \quad (38)$$

$$= \underbrace{W^{(L)} \left(\prod_{\ell=L-1}^1 A_{\theta^{(\ell)}}^{(\ell)} \right)}_{\text{Template Matching}} x + \underbrace{W^{(L)} \sum_{\ell=1}^{L-1} \left(\prod_{j=L-1}^{\ell+1} A_{\theta^{(j)}}^{(j)} \right) b_{\theta^{(\ell)}}^{(\ell)}}_{\text{Bias}} + b^{(L)}. \quad (39)$$

Given this interpretation we now proceed to derive explicitly what are the templates and biases for some standard topologies below as well as emphasizing the methodology for one to adapt the result to specific cases.

3.3.2 Deep Convolutional Networks

We first study the case of standard DCNs as described in 2.2. A DCN is composed of B blocks of 3 layers defined as

$$\begin{aligned}
\mathbf{z}^{(\ell)}(x) &= S_{\theta^{(\ell)}}^{(\ell)}(\mathbf{z}^{(\ell-1)}(x)) \\
&= (S_{\rho}^{(\ell)} \circ S_{\sigma}^{(\ell)} \circ S_{\mathbf{C}}^{(\ell)})(\mathbf{z}^{(\ell-1)}(x)) \\
&= A_{\rho}^{(\ell)}[\mathbf{z}^{(\ell-1)}(x)] \left(A_{\sigma}^{(\ell)}[\mathbf{z}^{(\ell)}(x)] \left(\mathbf{C}^{(\ell)} \mathbf{z}^{(\ell-1)}(x) + b^{(\ell)} \right) + b_{\sigma}^{(\ell)}[\mathbf{z}^{(\ell)}(x)] \right) + b_{\rho}^{(\ell)}[\mathbf{z}^{(\ell)}(x)] \\
&= \underbrace{A_{\rho}^{(\ell)} A_{\sigma}^{(\ell)} \mathbf{C}^{(\ell)}}_{\text{Template Matching}} \mathbf{z}^{(\ell-1)}(x) + \underbrace{A_{\rho}^{(\ell)} A_{\sigma}^{(\ell)} b^{(\ell)} + A_{\rho}^{(\ell)} b_{\sigma}^{(\ell)} + b_{\rho}^{(\ell)}}_{\text{Bias}} \\
&:= A_{\theta^{(\ell)}}^{(\ell)} \mathbf{z}^{(\ell-1)}(x) + b_{\theta^{(\ell)}}^{(\ell)}, \tag{40}
\end{aligned}$$

hence for a convolutional block, we have

$$A_{\theta^{(\ell)}}^{(\ell)} = A_{\rho}^{(\ell)} A_{\sigma}^{(\ell)} \mathbf{C}^{(\ell)}, \tag{41}$$

$$b_{\theta^{(\ell)}}^{(\ell)} = A_{\rho}^{(\ell)} A_{\sigma}^{(\ell)} b^{(\ell)} + A_{\rho}^{(\ell)} b_{\sigma}^{(\ell)} + b_{\rho}^{(\ell)}. \tag{42}$$

The topology implies the input conditioning of the spline tp depend on the previous layer output hence $A^{(\ell)}[\mathbf{z}^{(\ell-1)}] := A^{(\ell)}$. Using Eq. 38, we can write the overall DCN mapping as

$$\begin{aligned}
\mathbf{z}_{CNN}^{(L)}(x) &= W^{(L)} \underbrace{\prod_{\ell=L-1}^1 A_{\rho}^{(\ell)} A_{\sigma}^{(\ell)} \mathbf{C}^{(\ell)}}_{\text{Template Matching}} x \\
&\quad + W^{(L)} \underbrace{\sum_{\ell=1}^{L-1} \left(\prod_{j=L-1}^{\ell+1} A_{\rho}^{(j)} A_{\sigma}^{(j)} \mathbf{C}^{(j)} \right) \left(A_{\rho}^{(\ell)} A_{\sigma}^{(\ell)} b^{(\ell)} + A_{\rho}^{(\ell)} b_{\sigma}^{(\ell)} + b_{\rho}^{(\ell)} \right) + b^{(L)}}_{\text{Bias}}.
\end{aligned}$$

For cases of unbiased nonlinearities and pooling s.a. ReLU and max-pooling, this formula simplifies to

$$\mathbf{z}_{CNN}^{(L)}(x) = W^{(L)} \underbrace{\prod_{\ell=L-1}^1 A_{\rho}^{(\ell)} A_{\sigma}^{(\ell)} \mathbf{C}^{(\ell)}}_{\text{Template Matching}} x + W^{(L)} \underbrace{\sum_{\ell=1}^{L-1} \left(\prod_{j=L-1}^{\ell+1} A_{\rho}^{(j)} A_{\sigma}^{(j)} \mathbf{C}^{(j)} \right) \left(A_{\rho}^{(\ell)} A_{\sigma}^{(\ell)} b^{(\ell)} \right) + b^{(L)}}_{\text{Bias}}.$$

(43)

3.3.3 Deep Residual Networks

We now present the case of Residual Networks. A generic residual layer [He et al., 2016] is defined as

$$\begin{aligned}
z^{(\ell)}(x) &= S_{\theta^{(\ell)}}^{(\ell)}(z^{(\ell-1)}(x)) \\
&= A_{\rho}^{(\ell)} \left(A_{\sigma, in}^{(\ell)} \left(C_{in}^{(\ell)} z^{(\ell-1)}(x) + b_{in}^{(\ell)} \right) + C_{out}^{(\ell)} z^{(\ell-1)}(x) + b_{out}^{(\ell)} + b_{\sigma, in}^{(\ell)} \right) + b_{\rho}^{(\ell)} \\
&= A_{\rho}^{(\ell)} \underbrace{\left(A_{\sigma, in}^{(\ell)} C_{in}^{(\ell)} + C_{out}^{(\ell)} \right)}_{\text{Template Matching}} z^{(\ell-1)}(x) + \underbrace{A_{\rho, in}^{(\ell)} A_{\sigma, in}^{(\ell)} b_{in}^{(\ell)} + A_{\rho, in}^{(\ell)} b_{out}^{(\ell)} + A_{\rho, in}^{(\ell)} b_{\sigma, in}^{(\ell)} + b_{\rho}^{(\ell)}}_{\text{Bias}} \\
&:= A_{\theta^{(\ell)}}^{(\ell)} z^{(\ell-1)}(x) + b_{\theta^{(\ell)}}^{(\ell)}, \tag{44}
\end{aligned}$$

with the nonlinearity conditioned on $C_{in}^{(\ell)} z^{(\ell-1)}(x) + b_{in}^{(\ell)}$ and the pooling on $A_{\sigma, in}^{(\ell)} \left(C_{in}^{(\ell)} z^{(\ell-1)}(x) + b_{in}^{(\ell)} \right) + C_{out}^{(\ell)} z^{(\ell-1)}(x) + b_{out}^{(\ell)} + b_{\sigma, in}^{(\ell)}$. Hence for a residual block, we have

$$A_{\theta^{(\ell)}}^{(\ell)} = A_{\rho}^{(\ell)} A_{\sigma, in}^{(\ell)} C_{in}^{(\ell)} + C_{out}^{(\ell)}, \tag{45}$$

$$b_{\theta^{(\ell)}}^{(\ell)} = A_{\rho}^{(\ell)} A_{\sigma, in}^{(\ell)} b_{in}^{(\ell)} + A_{\rho}^{(\ell)} b_{out}^{(\ell)} + A_{\rho}^{(\ell)} b_{\sigma, in}^{(\ell)} + b_{\rho}^{(\ell)}. \tag{46}$$

Using Eq. 38, we can write the overall Resnet mapping as

$$\begin{aligned}
z_{RES}^{(L)}(x) &= W^{(L)} \underbrace{\left[\prod_{\ell=L-1}^1 A_{\rho}^{(\ell)} \left(A_{\sigma, in}^{(\ell)} C_{in}^{(\ell)} + C_{out}^{(\ell)} \right) \right]}_{\text{Template Matching}} x \\
&\quad + \underbrace{\sum_{\ell=L-1}^1 \left(\prod_{j=L-1}^{\ell+1} A_{\rho}^{(j)} \left(A_{\sigma, in}^{(j)} C_{in}^{(j)} + C_{out}^{(j)} \right) \right) \left(A_{\rho}^{(\ell)} A_{\sigma, in}^{(\ell)} b_{in}^{(\ell)} + A_{\rho}^{(\ell)} b_{out}^{(\ell)} + A_{\rho}^{(\ell)} b_{\sigma, in}^{(\ell)} + b_{\rho}^{(\ell)} \right)}_{\text{Bias}} + b^{(L)}.
\end{aligned}$$

It is common in Resnet to not have a pooling operation $A_{\rho}^{(\ell)} = I, b_{\rho}^{(\ell)} = 0$ but instead to apply a linear sub-sampling via the stride parameter of the convolution. Standard convolutions have a stride of $(1, 1)$ corresponding to no sub-sampling. Stride of (k, k) naturally correspond to (k, k) linear sub-sampling. Also, if no bias nonlinearity is used, then the Resnet recursion simplifies to

$$\boxed{z_{RES}^{(L)}(x) = W^{(L)} \underbrace{\left[\prod_{\ell=L-1}^1 \left(A_{\sigma, in}^{(\ell)} C_{in}^{(\ell)} + C_{out}^{(\ell)} \right) \right]}_{\text{Template Matching}} x + \underbrace{\sum_{\ell=L-1}^1 \left(\prod_{i=L-1}^{\ell+1} \left(A_{\sigma, in}^{(i)} C_{in}^{(i)} + C_{out}^{(i)} \right) \right) \left(A_{\sigma, in}^{(\ell)} b_{in}^{(\ell)} + b_{out}^{(\ell)} \right)}_{\text{Bias}} + b^{(L)}. \tag{47}}$$

Interestingly one can rewrite the Resnet formulation in the case $C_{out}^{(\ell)} = I$ as

$$\begin{aligned}
z_{RES}^{(L)}(x) &= x + \sum_{\ell=1}^{L-1} \left(\prod_{k=\ell}^1 A_{\sigma,in}^{(k)} W_{in}^{(k)} \right) x + \sum_{\ell=L-1}^1 \left(\prod_{i=L-1}^{\ell+1} (C_{out}^{(i)} + A_{\sigma,in}^{(i)} W_{in}^{(i)}) \right) (b_{out}^{(\ell)} + A_{\sigma,in}^{(\ell)} b_{in}^{(\ell)}) \\
&= \sum_{\ell=L-1}^1 \left(\prod_{i=L-1}^{\ell+1} (C_{out}^{(i)} + A_{\sigma,in}^{(i)} W_{in}^{(i)}) \right) (b_{out}^{(\ell)} + A_{\sigma,in}^{(\ell)} b_{in}^{(\ell)}) \\
&\quad \left. \begin{aligned}
&+ x \\
&+ A_{\sigma,in}^{(1)} W_{in}^{(1)} x \\
&+ A_{\sigma,in}^{(2)} W_{in}^{(2)} A_{\sigma,in}^{(1)} W_{in}^{(1)} x \\
&+ A_{\sigma,in}^{(3)} W_{in}^{(3)} A_{\sigma,in}^{(2)} W_{in}^{(2)} A_{\sigma,in}^{(1)} W_{in}^{(1)} x \\
&\vdots \\
&+ \left(\prod_{k=L-1}^1 A_{\sigma,in}^{(k)} W_{in}^{(k)} \right) x
\end{aligned} \right\} \text{Ensemble of Models}
\end{aligned}$$

In fact, it has been shown in [Veit et al., 2016] that deep residual networks behave like ensemble of relatively shallow models.

3.3.4 Deep Recurrent Networks

Similarly, we can derive the one step of a standard fully recurrent neural network [Graves, 2013] as

$$\begin{aligned}
z_{RNN}^{(1,t)} &= A_{\sigma}^{(1,t)} [W_{in}^{(1)} x^t + W_{rec} z_{RNN}^{(1,t-1)} + b^{(1)}] (W_{in}^{(1)} x^t + W_{rec} z_{RNN}^{(1,t-1)} + b^{(1)}) + b_{\sigma}^{(1,t)}, \\
z_{RNN}^{(\ell,t)} &= A_{\sigma}^{(\ell,t)} [W_{in}^{(\ell)} x^t + W_{rec} z_{RNN}^{(\ell,t-1)} + W_{up} z_{RNN}^{(\ell-1,t)} + b^{(\ell)}] (W_{in}^{(\ell)} x^t + W_{rec} z_{RNN}^{(\ell,t-1)} \\
&\quad + W_{up} z_{RNN}^{(\ell-1,t)} + b^{(\ell)}) + b_{\sigma}^{(\ell,t)}, \ell > 1.
\end{aligned}$$

By the double recursion of the formula (in time and in depth) we first proceed by writing the time unrolled RNN mapping as

$$\begin{aligned}
z_{RNN}^{(1,T)}(x) &= \sum_{t=T}^1 \left(\prod_{k=T}^{t+1} A_{\sigma}^{(1,k)} W_{rec}^{(1)} \right) A_{\sigma}^{(1,t)} (W_{in}^{(1)} x^t + b_{\sigma}^{(1,t)} + A_{\sigma}^{(1,t)} b^{(1)}) \\
&= \sum_{t=T}^1 \left(\prod_{k=T}^{t+1} A_{\sigma}^{(1,k)} W_{rec}^{(1)} \right) A_{\sigma}^{(1,t)} W_{in}^{(1)} x_t + \sum_{t=T}^1 \left(\prod_{k=T}^{t+1} A_{\sigma}^{(1,k)} W_{rec}^{(1)} \right) [b_{\sigma}^{(1,t)} + A_{\sigma}^{(1,t)} b^{(1)}] \\
z_{RNN}^{(\ell,T)}(x) &= \sum_{t=T}^1 \left(\prod_{k=T}^{t+1} A_{\sigma}^{(\ell,k)} W_{rec}^{(\ell)} \right) A_{\sigma}^{(\ell,t)} W_{in}^{(\ell)} x^t \\
&\quad + \sum_{t=T}^1 \left(\prod_{k=T}^{t+1} A_{\sigma}^{(\ell,k)} W_{rec}^{(\ell)} \right) [b_{\sigma}^{(\ell,t)} + A_{\sigma}^{(\ell,t)} b^{(\ell)} + A_{\sigma}^{(\ell,t)} W_{up} z_{RNN}^{(\ell-1,t)}(x)], \ell > 1.
\end{aligned}$$

The presented formula unrolled in time are still recursive in depth. While the exact unrolled version would be cumbersome for any layer ℓ we propose a simple way to find the analytical formula based on the possible paths an input can take till the final time representation of layer ℓ . To do so, one can look in Fig. 5.6. Hence we can thus decompose all those paths by blocks of forward interleave with upward paths. With this, we can




Figure 5: At any given layer, there always exists a direct input to representation path and recursion.




Figure 6: In addition, all the possible input path going through the hidden layers. Combinatorial number of paths yet fully determined by the succession of forward in time or upward in layer successions.

see that the possible paths are all the path from the input to the final nodes, they can not got back in time nor down in layers. Hence they are all the possible combinations for forward in time or upward in depth. We can thus find the exact output formula below for RNN as

$$z_{RNN}^{(2,T)}(x) = \sum_{t=1}^T \sum_{k_1=t}^T \left(\prod_{q=T}^{k_1+1} A_{\sigma}^{(2,q)} W_{rec}^{(2)} \right) A_{\sigma}^{(2,k_1)} W_{up}^{(2)} \left(\prod_{q=t+1}^{k_1-1} A_{\sigma}^{(1,q)} W_{rec}^{(1)} \right) A_{\sigma}^{(1,t)} W_{in}^{(1)} x^t \quad (48)$$

$$z_{RNN}^{(3,T)}(x) = \sum_{t=1}^T \sum_{k_1=t}^T \sum_{k_2 \geq k_1}^T \left(\prod_{q=T}^{k_2+1} A_{\sigma}^{(3,q)} W_{rec}^{(3)} \right) A_{\sigma}^{(3,k_2)} W_{up}^{(3)} \left(\prod_{q=k_1}^{k_2-1} A_{\sigma}^{(2,q)} W_{rec}^{(2)} \right) A_{\sigma}^{(2,k_1)} W_{up}^{(2)} \left(\prod_{q=t+1}^{k_1-1} A_{\sigma}^{(1,q)} W_{rec}^{(1)} \right) A_{\sigma}^{(1,t)} W_{in}^{(1)} x^t \quad (49)$$

$$z_{RNN}^{(4,T)}(x) = \sum_{t=1}^T \sum_{k_1=t}^T \sum_{k_2 \geq k_1}^T \sum_{k_3 \geq k_2}^T \left(\prod_{q=T}^{k_3+1} A_{\sigma}^{(4,q)} W_{rec}^{(4)} \right) A_{\sigma}^{(4,k_3)} W_{up}^{(4)} \left(\prod_{q=T}^{k_2+1} A_{\sigma}^{(3,q)} W_{rec}^{(3)} \right) A_{\sigma}^{(3,k_2)} W_{up}^{(3)} \left(\prod_{q=k_1}^{k_2-1} A_{\sigma}^{(2,q)} W_{rec}^{(2)} \right) A_{\sigma}^{(2,k_1)} W_{up}^{(2)} \left(\prod_{q=t+1}^{k_1-1} A_{\sigma}^{(1,q)} W_{rec}^{(1)} \right) A_{\sigma}^{(1,t)} W_{in}^{(1)} x^t \quad (50)$$

...

3.4 Template Matching with DNN: How and Why

We have seen in the last section the template matching formulation of DNNs via LSOs simply as being the slope of the linear transform. By definition of template matching, there exists an internal "matching" procedure performed by the DNN. We propose to study this inference problem in this section in standard DNN and why can we label DNNs as template matching machines. As we will see, a greedy, per layer, maximization problem is governing the spline selection and thus template inference. We then study the impact of choosing different LSOs, such as ReLU or absolute value and their impact in the inference problem

each one performs. Deriving such results will allow two main applications. Firstly, with the convexity property, one can derive arbitrary splines with regions that can be implicitly changed and learned "online", as the selection will become intrinsically partition agnostic, known as adaptive partitioning [Hannah and Dunson, 2013]. Secondly, the inference problem will be of great interest when dealing with deep neural networks analysis in further sections. We first briefly describe some theoretical results to link inference-LSOs-template matching.

3.4.1 Template Matching in the Context of Splines

We study in this context under what condition spline functions can be considered to perform template matching. Let first define what do we refer to as template inference.

Definition 5. For a spline functional (univariate;multivariate;SO), given a partition of the input space denoted by $\Omega = \{\omega_1, \dots, \omega_R\}$ and local mappings ϕ_1, \dots, ϕ_R , the inference problem refers to, given an input x , finding the region in which it belongs:

$$\text{Given } x: \text{ Find } \omega \in \Omega \text{ s.t. } x \in \omega.$$

This region is then used to perform the actual mapping via $\phi[x]$.

As we now describe, this problem can represent very interesting behaviors linked with template selection in some cases, especially when the functional is convex. We study in this section the convexity criteria for spline operators and the associated spline inference problem. Note that we focus now on linear functionals, provided results can easily be extended.

Theorem 2. Given a linear multivariate spline $s[\mathbf{a}, \mathbf{b}, \Omega]$ we have

$$s[\mathbf{a}, \mathbf{b}, \Omega](x) = a[x]^T x + b[x] \tag{51}$$

$$= \max_{r=1, \dots, R} a_r^T x + b_r, \forall x, \tag{52}$$

if and only if $s[\mathbf{a}, \mathbf{b}, \Omega]$ is a convex function [Hannah and Dunson, 2013].

This theorem states that given a convex spline function, finding the region to which an input x belongs to is equivalent to finding the region in which the mapping leads to the highest output. We provide an illustrative example in Fig. 7. This result provides ways to create adaptive partitioning convex splines simply by learning the collection of hyperplanes with the mappings defined as the maximum of the hyperplane projections. We can now extend the result to LSOs made of a collection of K multivariate splines.

Theorem 3. Given an LSO defined $S \left[(s[\mathbf{a}_k, \mathbf{b}_k, \Omega_k])_{k=1}^K \right]$, with all internal linear multivariate splines $s[\mathbf{a}_k, \mathbf{b}_k, \Omega_k]$ being convex, we have

$$S[\mathbf{A}, \mathbf{b}, \Omega^S](x) = \Phi^*(x), \tag{53}$$

with $\phi^{S*} = \operatorname{argmax}_{\phi^S \in \Phi^S} \langle \phi^S(x), 1 \rangle$, with $\Phi^S = \{\phi_r^S, r = 1, \dots, R\}$ and with $\phi_r^S(x) = A_r x + b_r$.




Figure 7: Illustrative examples of the theorem where one can see that the sub-region/sub-function associated with the query point is the one returning the maximum value among all the possible sub-functions.

Proof.

$$\begin{aligned}
\phi^{S^*} = \operatorname{argmax}_{\phi^S \in \Phi^S} \langle \phi^S(x), 1 \rangle &= \operatorname{argmax}_{[\phi_1, \dots, \phi_K] \in \Phi_1 \times \dots \times \Phi_K} \langle \Phi(x), 1 \rangle \\
&= \begin{bmatrix} \operatorname{argmax}_{\phi_1 \in \Phi_1} \sum_{k=1}^K \phi_k(x) \\ \vdots \\ \operatorname{argmax}_{\phi_K \in \Phi_K} \sum_{k=1}^K \phi_k(x) \end{bmatrix} \\
&= \begin{bmatrix} \operatorname{argmax}_{\phi_1 \in \Phi_1} \phi_1(x) \\ \vdots \\ \operatorname{argmax}_{\phi_K \in \Phi_K} \phi_K(x) \end{bmatrix} \\
&= \begin{bmatrix} \phi_1[x] \\ \vdots \\ \phi_K[x] \end{bmatrix}, \quad s_k \text{ convex } \forall k \\
&= \Phi[x]
\end{aligned} \tag{54}$$

□

This last theorem leverages the independence between the multivariate splines making up the spline operator. As a result, the per multivariate spline region selection solved via the max operator in case the are convex can be done for all multivariate spline simultaneously via the $\langle \cdot, 1 \rangle$ operator, leading to the sum of the output dimensions.

3.4.2 DNNs Are Composition of Adaptive Partitioning Splines

As demonstrated in the last section, convex splines defined through a max over hyperplanes projections is defined as adaptive partitioning as changing the hyperplanes induces changes in the input space partitioning. Hence for regression problems for example, optimal partitions can be found in this manner simply by tweaking the hyperplanes parameters [Hannah and Dunson, 2013, Magnani and Boyd, 2009] and has been shown to be very performant in the context of Nonlinear Least Square Regression. In fact, hyperplanes combination to solve function approximation problems go back to [Breiman, 1993] reinforcing the fact that we can now see current state-of-the-art (sota) DNN as efficient composition of such approaches. We now

describe this last statement in details. Current sota DNNs leverage the Relu or LReLU nonlinearity, both convex, as well as max and/or mean-pooling, being also convex mappings. Based on the results drawn from the last section we can thus see that all the succession of linear layers such as FC-layer of convolutional layer followed by nonlinearities and possibly sub-sampling correspond to adaptive partitioning multivariate spline function. In fact, one has in those cases

Theorem 4. *DNNs with convex activation functions s.a. Relu or LReLU, and/or convex sub-sampling s.a. mean or max pooling applied on linear layers are composition of partition adaptive splines [Hannah and Dunson, 2013, Magnani and Boyd, 2009],*

$$S_{\sigma}(Wx + b) = A_{\sigma, r^*}x + b_{\sigma, r^*}, \quad \text{with } r^* = \underset{r=1, \dots, R_{\sigma}^S}{\operatorname{argmax}} \langle A_r(Wx + b) + b_r, 1 \rangle. \quad (55)$$


Given the last theorem, one might wonder if the local per layer partition optimization can be extended to a global adaptive partitioning. This question is answered in Appendix A.3 where we provide sufficient condition to obtain a globally convex DNN, hence making the last theorem not only applicable on a per layer basis but overall the mapping. In fact, composition of such layers are in general not globally convex with unconstrained weights. We now have linked DNN to known powerful frameworks for function approximation and can now provide ways to visualize the final inferred templates with standard DNNs. We propose to do so in the net section in order to highlight the extrem adaptivity DNNs have with this regard.

3.4.3 Input Encoding and Template Visualization

In this section we present experiments on MNIST and CIFAR10 to provide visualization of the adapted templates given few samples. We also provide a simple methodology to compute the templates. Since the final DNN can be expressed as $f_{\Theta}(x) = A[x]x + b[x]$, it is clear that we can obtain the adapted template for class c as $a[x]_{c, \cdot} = \frac{df_{\Theta}(x)_c}{dx}$. We present below computed templates for one model, the LargeCNNmean. We remind that specific model descriptions are provided in Appendix B. All the other templates related to other topologies are provided in Appendix. We present in Fig. 8 different templates induced for one class with the input belonging to this same class, hence this template matching matches the class of interest. In the title of each subplot is provided the template matching output $\langle A[X_n]_{Y_n}, X_n \rangle$ as well as the bias $b[X_n]_{Y_n}$. One can see the adaptivity of the templates but also for the CIFAR10 case (on the right) the "denoising" of the background. Only the object of interest remains. We also present in Fig. 9 more specific templates for 4 different inputs. In particular we provide the templates of all the classes for a given input as well as the gram-matrix of those templates, representing their correlation. It is interesting to denote that the wrong class templates are not 0 neither white noise like but on the contrary inversely correlated w.r.t. the input and the correct template class. This phenomenon will be explained in details in the later section discussing the optimal template and their convergence. Finally, in Fig. 10 we provide the same type of analysis but for CIFAR10. Clearly the shape and class are less distinguishable in the templates, yet they are fitted to the input so that the template matching of the right class is indeed the maximum of all the classes.

Concerning the encoding of a given input x , a DNN find very close interpretation with standard signal processing tools. In fact, it encodes x via localization information and amplitudes as would be the case of Fourier transform with phase and amplitude. In fact, we denote by amplitude the result of $A[x]x$ and by phase the collection of regions at each layer in which the input belongs to. If we denote by $\Omega[x]^{(\ell)}$ the region

LargeCNNmean.pkl Inferred Template Examples, k=0 and values $((A[x_k]_{k \dots x}), b[x_k]_k)$



LargeCNNmean.pkl Inferred Template Examples, k=0 and values $((A[x_k]_{k \dots x}), b[x_k]_k)$




Figure 8: Left: Examples of Templates for class 0 given inputs X_n belonging to class 0 ($Y_n = 0$). As one can see, the class 0 templates fully adapt to their input for LargeCNNmean. Right: Templates on CIFAR10 associated to class plane. Interestingly the background is not captured and only the class shape is present.

at level ℓ in which $z^{(\ell)}$ belongs to, then a DNN representation of an input x is defined as follows.

Proposition 1. *A DNN encoding of an input x is define as the complementary couple of amplitude and phase defined as $(A[x]x, (\Omega[x]^{(\ell)})_{\ell=1}^L)$. This information is indeed enough to fully reconstruct an input as we will demonstrate later with deep neural network inversion and semi supervised applications.*

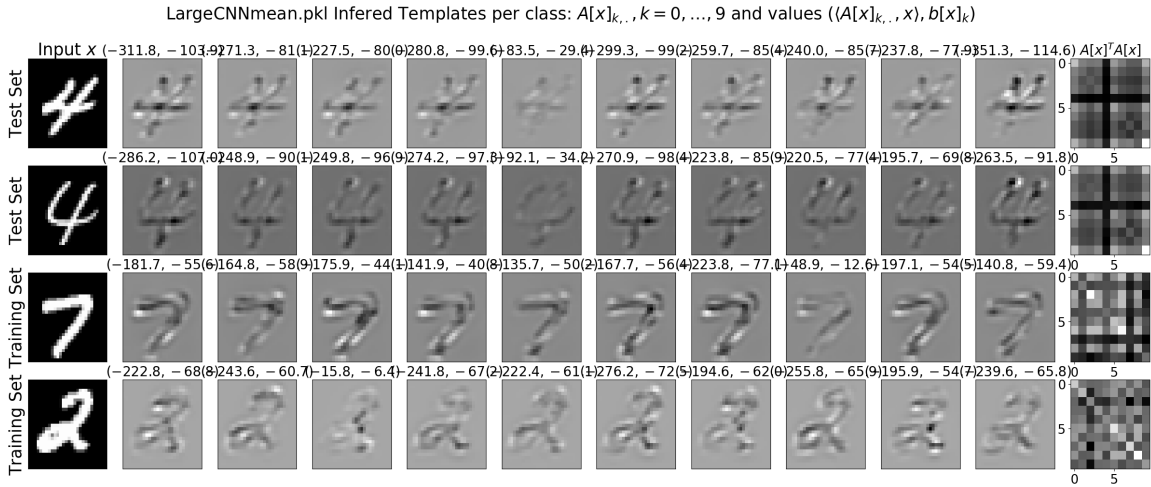


Figure 9: Depiction on MNIST for 4 inputs belonging to class 4, 4, 7, 2 of all the classes templates with classes from 0 (left) to 9 (right). Right column is the Gram matrix of the templates representing the correlation between templates of different classes for each input. In the subplot titles are the values of the template matching and the bias.

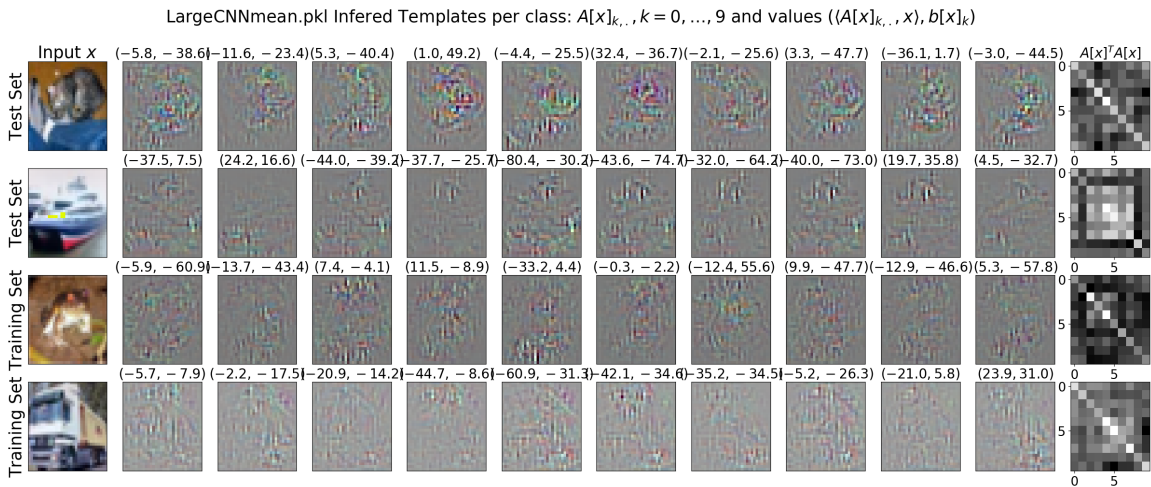


Figure 10: Depiction on CIFAR10 for 4 inputs belonging to class cat, boat, frog, truck, of all the classes templates with classes from 0 (left) to 9 (right). Right column is the Gram matrix of the templates representing the correlation between templates of different classes for each input. In the subplot titles are the values of the template matching and the bias.

4 Theoretical Results: Generalization, Optimal Learning, Memorization

In this section, we will leverage the derived spline framework to provide analytical understanding of current DNNs. To do so, we first demonstrate the impacts of regularization into the quest of generalization performances for DNNs. Through regularization we can obtain analytical optimal templates and thus provide a clear methodology to guide DNNs towards this optimum. Through this analysis, results on adversarial examples, DNN inversion and optimization schemes will be studied. Let first review the generalization problem for DNNs and the mathematical tools that can be leveraged.

4.1 What is Generalization for DNN and why Regularization is Key

As detailed in the introduction Section 1, generalization is the ability for the approximant f_Θ to reproduce the behavior of the true unknown functional f on new points not present in the training set \mathcal{D} . This being very general we now have to distinguish two important cases. First, when f represents the "law of nature" in the sense that it follows some fundamental unbreakable laws. Hence, given an input x , only one possible outcome $f(x)$ exists, fully determined by fundamental laws such as thermodynamics. A typical example would be to predict the internal energy of a system. The second case concerns machine learning. It is the one where f is "human", corresponding to human perceptions, a qualitative, normed interpretation of an input x . Typical example would be for x to be a pixel representation of a scene and the associated $f(x)$ the human associated label of the *main object of interest*. However, from individuals to others, the core definition of f might change, raising questions about the term generalization in computer vision. Due to this unclear definition, the only quantitative measure one has of generalization is the use of a test set on which f_Θ has not been trained. This test set is used to compare to predictions based on f_Θ with the known correct outputs of f . Since we focus on accuracy performance, we denote this loss by \mathcal{L}_{AC} standing for accuracy loss. As a result, one uses \mathcal{L}_{CE} the cross-entropy loss and \mathcal{D} to update the parameters Θ and \mathcal{L}_{AC} with \mathcal{D}_{test} for a quantitative measure of generalization of the trained approximant. Due to the discussed context, we have to consider generalization differently than simply maximizing the test set accuracy. In fact, we propose here to define generalization as a measure of performance consistency from the training set to the test set. Hence, even a less accurate model is favored if its ability do not vary from samples used for its training to new samples.

Definition 6. We define the generalization measure G of a trained network f_Θ by the average empirical difference of performance between training set and test set as

$$\mathcal{L}_G(\mathcal{D}, \mathcal{D}_{test}, f_\Theta) = d(\mathcal{L}_{AC}(\mathcal{D}, f_\Theta), \mathcal{L}_{AC}(\mathcal{D}_{test}, f_\Theta)), \quad (56)$$

with d a distance metric, \mathcal{L}_{AC} a generic performance metric linked with quality of the prediction, accuracy loss in our case

$$\begin{aligned} \mathcal{L}_{AC}(\mathcal{D}, f_\Theta) &= \frac{1}{\text{Card}(\mathcal{D})} \sum_{(X_n, Y_n) \in \mathcal{D}} \mathcal{L}_{AC}(Y_n, \hat{y}(X_n)) \\ \mathcal{L}_{AC}(\mathcal{D}_{test}, f_\Theta) &= \frac{1}{\text{Card}(\mathcal{D}_{test})} \sum_{(X_n, Y_n) \in \mathcal{D}_{test}} \mathcal{L}_{AC}(Y_n, \hat{y}(X_n)). \end{aligned}$$

As a result, a network performing similarly on train and test set is considered as optimal in term of

generalization of its underlying learned representation. This results in a systematic way to measure and learn topologies as we now define the overall objective.

Definition 7. *The optimal network given a finite training set \mathcal{D} and test set \mathcal{D}_{test} is defined as*

$$f_{\Theta^*}^* = \arg \min_{f_{\Theta}} \underbrace{\mathcal{L}_{AC}(\mathcal{D}, f_{\Theta})}_{\text{Empirical Risk Minimization: } \Theta^*} + \underbrace{\mathcal{L}_G(\mathcal{D}, \mathcal{D}_{test}, f_{\Theta})}_{\text{Structural Risk Minimization: } f^*} \quad (57)$$

This search of the optimal approximant can thus be done in a two-step process by first fixing a topology and then minimizing \mathcal{L}_{CE} synonym of minimization of \mathcal{L}_{AC} on the training set. Doing this over multiple topologies and then selecting the optimal network by search of the one with minimum generalization loss \mathcal{L}_G . Since this results in learning of tremendous possible models, one usually tries to find a way to translate \mathcal{L}_G into a differentiable loss that can be used on the training set. This usually takes the form of standard regularization such as Tikhonov, dropout and so on. Yet, those approaches can only impact the final parameters Θ , and thus have only a limited impact on the true generalization loss as opposed to topological changes in f . Nevertheless, we now develop in the following section precise analysis and results to link regularization with generalization, overfitting. This will also to understand the dataset memorization problem and what generalization actually means for DNNs. The next section will however build on those results and attempt to tackle this problem from a broader point of view via a systematic way to estimate \mathcal{L}_G prior learning hence allowing easier design search.

Given an approximant f_{Θ} , the search for best generalization performances is commonly interpreted as finding parameters Θ in a flat-minima region. A flat-minima region is a part of the parameter space associated with great generalization performance of the approximant. The term flatness is easily interpreted as follows. One seeks $\Theta + \epsilon$ to also belong to this region. Hence moving around Θ still produce great generalization leading to a flat generalization performance as $G(f_{\Theta}) \approx G(f_{\Theta+\epsilon})$. This is opposed to sharp-minima where $\|G(f_{\Theta}) - G(f_{\Theta+\epsilon})\| \gg \|\epsilon\|$. This analysis started long before current DNN outbreaks. Generalization, in addition of being associated to flat minima [Wolpert, 1994] is also mapped to complexity of networks [Hochreiter and Schmidhuber, 1995] which is linked with Kolmogorov complexity and Minimum Description Length. This comes from the fact that flat minima are associated to simpler networks which then leads to high generalization [Schmidhuber, 1994, Hochreiter and Schmidhuber, 1995]. However, standard analysis is hardly applied as the measure and definition of a DNN complexity is still not clear. Practical approaches aiming at guiding Θ towards flat-minima then took different forms. From one side, reduction of the number of degrees of freedom via weight sharing led to promising results [Nowlan and Hinton, 1992, Rumelhart and McClelland, 1986, Lang et al., 1990, Yann, 1987, LeCun et al., 1989] while being very general and model agnostic. Another approach uses early-stopping as in [Morgan and Bourlard, 1990, Weigend et al., 1990, Vapnik, 1992, Moody and Utans, 1994, Guyon et al., 1992] motivated by the famous point of inflexion of the testing error, first reducing till a breaking point where it increases. This point is the optimal to stop training as generalization error is minimized. Both methods require inside information and expertise. Thus the search for a more principled method possibly adaptive to any case led to regularization studies. To do so, penalization of complex networks was applied and led to great advances in the flat-minima search. This complexity based approach takes the form of Occam’s razor principle [Blumer et al., 1987] and was in practice applied via weight penalization [MacKay, 1996, Hinton, 1986, Hinton, 1987, Plaut et al., 1986, Williams, 1995]. For example, with norm based regularization. In fact, in the case of Tikhonov regularization, a very intuitive interpretation of the weights appear: **the weight amplitudes is equal or proportional to their er-**

ror derivative, a.k.a their importance. By making weights amplitude correlated to their role in the loss minimization, only the necessary one will remain nonzero hence simplifying the network through sparsity of the model/connections. Finally, input and/or weight noise applied during training has also found equivalences with regularization. It consists of perturbing the input or the current set of parameters with additive or multiplicative noise throughout the learning phase. This has recently taken the form of dropout, a multiplicative noise with Bernoulli variables [Srivastava et al., 2014, Gal and Ghahramani, 2016, Srivastava et al., 2014] randomly turning neurons or connections to 0 in DNNs. This concept goes back to synaptic noise [Murray and Edwards, 1993], and [Matsuoka, 1992] where generalization performances of neural networks trained with backpropagation is studied via introduction of noise to the input. It was then shown in 1995 [Bishop, 2008] that introduction of noise during learning is equivalent to a generalized Tikhonov regularization technique. More precisely, it has been shown that while additive noise provides an induced penalty term on the norm of the weights a la Tikhonov, multiplicative noise provides a weighting of this regularization based on the Fisher information of the weights [Li et al., 2016]. A probabilistic interpretation of dropout indeed demonstrates the push of the weights towards sparse solutions [Nalisnick et al., 2015]. Going further, an explicit regularization term is found from dropout and extended in [Wager et al., 2013]. Based on those approaches, we can now have the following intuitive explanation of why current topologies work so well:

- **Multi-Objective Regularization:** Introduction of noise during learning coupled with explicit norm based penalties on the weights
- **Weight-Sharing:** Convolutional topologies allow extremely efficient and smart weight sharing for perceptual tasks reducing the number of degrees of freedom while providing very high-dimensional mappings
- **Cross-Validation and Early-Stopping:** huge resources now allow fine search of topologies and hyper-parameters

We now provide in the next section theoretical results in the context of regularization. As we will see norm constraints on the free parameters, hence the templates for DNNs allow to obtain closed form optimal theoretical templates. From this, different results will be derived from adversarial example existence to dataset memorization and network inversion.

4.2 Learning Optimal Templates

In this section we study the learning of the DNN templates for two cases. First in the case of a loss function without regularization terms. Secondly when sparsity constraints is imposed. For both cases we study what are the optimal templates, their convergence and demonstrate the need for regularization. Regularization will be shown to make the problem of optimal template learning well-defined as well as being robust to poor weight initialization. Based on this, we then provide a methodology to quantify the quality of a given DNN topology and weight initialization schemes simply based on the induced templates and their potentials in the next Section ??.

4.2.1 Unregularized Learning Solution: Unstable Training

We present here the general analysis for any DNN using the cross-entropy loss function coupled with softmax activations. We denote by A_c the c^{th} template. There is no input conditioning ($A[x]$) as we aim at finding the explicit optimal form this template should have given the input x . Hence for now A_c is an generic template. The global loss function to be minimized is the negative cross-entropy between the true label Y_n and the estimation $\hat{y}(X_n)$. We remind that it is defined as

$$\mathcal{L}_{CE}(Y_n, \hat{y}(X_n)) = -(\langle A_{Y_n}, x \rangle + b_{Y_n}) + \log \left(\sum_{c=1}^C e^{\langle A_c, x \rangle + b_c} \right). \quad (58)$$

From this we apply standard iterative gradient based minimization procedures to seek the optimal templates A_c for a given input X_n and for each classes. We have

$$\begin{aligned} \frac{\partial \mathcal{L}_{CE}(Y_n, \hat{y}(X_n))}{\partial A_{Y_n}} &= -X_n + X_n \frac{e^{\langle A_{Y_n}, X_n \rangle + b_{Y_n}}}{\sum_{c=1}^C e^{\langle A_c, X_n \rangle + b_c}} \\ &= X_n (\hat{y}(X_n)_{Y_n} - 1), \end{aligned} \quad (59)$$

$$\begin{aligned} \frac{\partial \mathcal{L}_{CE}(Y_n, \hat{y}(X_n))}{\partial A_c} &= X_n \frac{e^{\langle A_c, X_n \rangle + b_c}}{\sum_{k=1}^C e^{\langle A_k, X_n \rangle + b_k}} \\ &= X_n \hat{y}(X_n)_c, \forall c \neq Y_n, \end{aligned} \quad (60)$$

leading to the following gradient update rule with learning rate λ

$$\begin{aligned} A_{Y_n}^{(t+1)} &= A_{Y_n}^{(t)} - \lambda X_n (\hat{y}_{Y_n}(X_n) - 1) \\ &= A_{Y_n}^{(t)} + \lambda X_n (1 - \hat{y}_{Y_n}(X_n)) \end{aligned} \quad (61)$$

$$A_c^{(t+1)} = A_c^{(t)} - \lambda X_n \hat{y}(X_n)_c, \forall c \neq Y_n. \quad (62)$$

We do not analyze the behavior for the biases b since they do not interfere with the optimal templates. It is clear that $\hat{y}(X_n)_c > 0$ as well as $(1 - \hat{y}_{Y_n}(X_n)) > 0$. This implies that the update rule adds the re-scaled input X_n for the correct template A_{Y_n} whereas for the other classes, X_n is subtracted. This way, by adding or subtracting the input to the templates, the template matching mapping can either increase or decrease. It becomes clear that the sensitivity to initialization is extreme as given a starting random template, it can only moves along the input X_n direction. Hence there are infinitely many optimal templates, one for each starting point. Moreover, the similarity between the templates and the input will also depend on this initialization. We depict this phenomenon in Fig. 11 on the left subplot. For the case of a structured templates as in practice it is defined as a composition of affine mappings with different internal parameters, it is clear that the update will push the template as close as possible to this optimal based on the ability of the mappings composition to produce it. With increasing number of free parameters and network complexity it is fair to assume that the induced update will be close to this optimum. We now dive into the regularized case.

4.2.2 Regularized Learning: Global Optimum, Robust, implies Dataset Memorization

By adding a regularization penalty to the loss function such as sparsity constraint with norm based loss, we can obtain analytical optimal templates as the optimization problem becomes well-defined. As we will see,

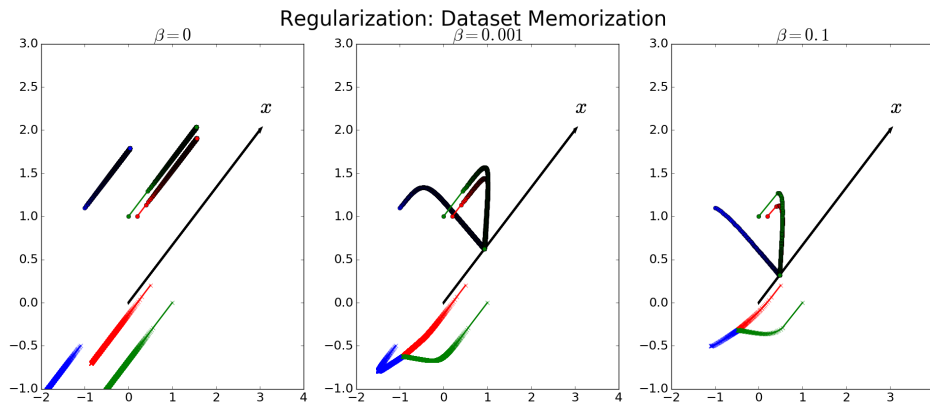


Figure 11: Tikhonov based regularization implies a push of the templates towards dataset memorization.

dataset memorization is the global optimum.

While memorization is often associated to overfitting and bad performances, we will see that this general statement is more ambiguous. By memorization, we denote the ability for DNNs template to become collinear to their input. In fact, the term memorization itself should be seen as a good ability of a DNN if it holds for arbitrary inputs from the manifold of interest. In this case, the DNN would be able to span all inputs of interests, effectively making it a basis of the training set, testing set and so on. We now study the impact of norm constraints on the optimal templates of DNNs with only assumption that all inputs X_n have same energy, in particular $\|X_n\|^2 = 1$.

Theorem 5. *In the case where all inputs have identity norm $\|X_n\| = 1, \forall n$ and assuming all templates denoted by $A[X_n]_c, c = 1, \dots, C$ have a norm constraints as $\sum_{c=1}^C \|A[X_n]_c\|^2 \leq K, \forall X_n$ then the unique globally optimal templates are*

$$A^*[X_n]_c = \begin{cases} \sqrt{\frac{C-1}{C}} K X_n, & \iff c = Y_n \\ -\sqrt{\frac{K}{C(C-1)}} X_n, & \text{else} \end{cases} \quad (63)$$

In order to highlight and provide intuition on the two derived results, we provide in Fig. 11 a simple example in 2 dimensional case. We initialize the templates randomly and set the templates of opposite classes to be equal a initialization to better see the update differences based on a same starting point. On the left part no Tikhonov regularization has been applied whereas the middle and right plots contains standard and strong regularization. We can see that with regularization, templates converge toward a rescaled X_n , with scale depending on the regularization parameter. This convergence speed is also dependent on this parameter. On the other hand for the no penalty case, a simple move of the templates can be seen as ill-posed and leads to very pool templates. Hence, dataset memorization can not be said to always be synonym to overfitting and thus poor performance as they might correspond to optimal templates. From the derived results, we now propose to analyze adversarial examples though the lenses of templates and optimal templates. As we will see, adversarial examples are natural and should not be fought, high sensitivity however should be controlled.

4.3 Adversarial Examples Are Natural, being Fooled is Not

In this section we propose a two-fold analysis of adversarial examples. Note that we focus here on model based adversarial examples which are optimized based on a trained DNN. Firstly we demonstrate that existence of adversarial noise is due to optimal templates. In particular we demonstrate that training with current loss function and condition can only lead to adversarial example existence. Secondly we study the sensitivity of DNN mappings in general including adversarial noise sensitivity. To do so we propose a simple methodology to compute the Lipschitz constant of any DNN layer via the corresponding LSO. By allowing explicit computation of the overall mapping sensitivity we further justify the need for sparsity constraints on the parameters. We first review briefly what are adversarial examples.

We remind that we denote a differentiable mapping $f_{\Theta} : \mathbb{R}^D \rightarrow \mathbb{R}^C$ that we call a pre-processing system producing a feature map e.g. SIFT, wavelets, DCN. In order to produce a class distribution \hat{y} a softmax nonlinearity is applied as presented in Section 1. We further denote as $X_n^{Y_n}$ an input image X_n of class Y_n . Adding to it a small optimized noise in order to fool the system to predict a wrong class $k, k \neq Y_n$ lead to the new input denoted $X_n^{Y_n \rightarrow k}$. In order to generate an adversarial example we use

$$X_n^{Y_n \rightarrow k} = X_n + \alpha \frac{d\hat{y}(X_n^{Y_n})_k}{dX_n^{Y_n}}. \quad (64)$$

It thus intuitively corresponds to pushing by an amplitude of $\alpha \in \mathbb{R}^+$ the natural input $X_n^{Y_n}$ with the wrong class template $A[X_n^{Y_n}]_k$. Clearly, the demonstration that a very small α leads to a complete change in the prediction $\hat{y}(X_n)$ implies that the joint system is not Lipschitz contractive since

$$\|\hat{y}(X_n^{Y_n}) - \hat{y}(X_n^{Y_n \rightarrow k})\| \gg \|X_n - X_n^{Y_n \rightarrow k}\| \quad (65)$$

$$\gg \alpha \left\| \frac{d\hat{y}(X_n)_k}{dX_n} \right\|. \quad (66)$$

We first review briefly previous work on adversarial examples. In order to become more robust to adversarial examples [Gu and Rigazio, 2014, Lyu et al., 2015] developed gradient regularization techniques by imposing a sparsity penalty term on the partial derivatives of the output the the deep nets layers w.r.t. their inputs. This is as we showed above on way to reduce the amplitude of the adversarial examples by imposing that $\left\| \frac{df_{\Theta}(X_n)}{dX_n} \right\| < \epsilon$. On the other hand, using adversarial examples during training has also been done showing better performances on the test set as in [Shaham et al., 2015]. However none of these methods provide a guarantee on the generalization of the technique for new unseen examples and adversarial examples. On the opposite in [Fawzi et al., 2015], has shown that for linear classifier, invariance can not be achieved in general. Similarly, in [Gu and Rigazio, 2014] the idea of being invariant is rejected stating that one can always engineer adversarial noise, however this happens to not be the case following the sufficient condition we propose. Another approach is proposed by defensive distillation [Papernot et al., 2016] which consists of artificially increasing the output of $f_{\Theta}(X_n)$ or the input to the softmax as stated which thus forces the prediction $\hat{y}(X_n)$ to become ϵ -close to a one hot representation and thus making the norm of the derivative w.r.t. input close to 0-norm, this is interesting has it exploits the vanishing gradient problem which has the same time pushed to 0 the margin to robustify the network against adversarial example and has been questioned in [Carlini and Wagner, 2016]. While competitive robustness has been achieved in [Gu and Rigazio, 2014] via the jacobian penalty term it was already hint that a deeper problem resides in the training procedures and the loss function as it is confirmed. In [Szegedy et al., 2013] additional weight decay on

the parameters is applied to the loss function thus reducing $\| \frac{df_{\Theta}(X_n)}{dX_n} \|$ leading to smaller adversarial energy without yet removing its presence.

4.3.1 Unregularized Optimal Templates Imply Adv. Noise

We study the impact of the templates concerning the presence of adversarial examples. When regularization is used during training, the optimal templates for any given input $X_n^{Y_n}$ are of the form $A[X_n^{Y_n}]_{Y_n} \propto X_n^{Y_n}$ and $A[X_n^{Y_n}]_c \propto -X_n^{Y_n}, c \neq Y_n$. Hence it is clear that adversarial examples based on model optimization can only scale down the input X_n hence $X_n^{Y_n \rightarrow c} = \kappa X_n^{Y_n}, \kappa < 1$. For the case of unregularized templates however, the input additive transformation will introduce the initial template. As we saw, the final template after learning results from the initialized templates plus a succession of updates pushing it by a rescaled version of X_n . Hence we denote $A[X_n^{Y_n}]_c = A[X_n^{Y_n}]_c^0 + \beta X_n^{Y_n}$. The noisy input thus corresponds to $X_n^{Y_n \rightarrow c} = X_n^{Y_n} + \kappa(A[X_n^{Y_n}]_c^0 - \beta X_n^{Y_n})$. As a result, it does not just scale down the input but add a random noise to the input. This random noise corresponding to part of the wrong class template, implies much higher unstabilities and can look like standard noise to us, as it is in practice.

We now present two properties on the gradient based update in the case of no regularization that will be used to demonstrate the rise of adversarial examples as part of the weight optimum convergence. This first one states that as the number of classes increase, the wrong class templates will be less and less updated in the opposite direction of the input making the initialization noise more present in the adversarial noise possible leading in the extreme case $X_n^{Y_n \rightarrow c} = X_n^{Y_n} + \kappa(A[X_n^{Y_n}]_c^0)$

Proposition 2. *The sum of the updates for the wrong classes add up the the opposite of the update of the correct class denoted as*

$$\begin{aligned} \sum_{c \neq Y_n} \frac{d\hat{y}(X_n^{Y_n})_c}{dX_n^{Y_n}} &= \sum_{c \neq Y_n} -X_n \hat{y}(X_n)_c \\ &= -X_n (1 - \hat{y}(X_n)_{Y_n}) \\ &= -\frac{d\hat{y}(X_n^{Y_n})_{Y_n}}{dX_n^{Y_n}}. \end{aligned}$$

We can thus see that as the number of classes grow the less the wrong class templates will move away from their initial point. We now present adversarial example specific analysis to quantify their impact on a given network using the tools and remarks developed in the previous section. Using the chain rule and the definition of adversarial examples defined in Eq. 64 we can measure the sensitivity of a network via analysis of the norm of $\frac{d\hat{y}(X_n)_c}{dX_n}$. We thus proceed to derive Lipschitz constant of DNNs in the next section.

4.3.2 Ensuring Adversarial Noise Robustness via Lipschitz Constant minimization: Contractive DNNs

In this section we describe the space contraction properties of DNNs and composition of LSOs in general. As we will see, deriving the exact formula for any given deep neural network is straightforward and will allow us to better understand what causes "chaotic" behaviors as seen with adversarial examples. Let first

remind that for differential mappings $f_{\Theta}^{(\ell)} : \mathbb{R}^{D^{(\ell-1)}} \rightarrow \mathbb{R}^{D^{(\ell)}}$ the Lipschitz constant K is equal to the infinite norm of the total derivative. Hence for LSOs $S[\mathbf{A}, \mathbf{b}, \Omega]$ differentiable almost everywhere we have

$$K \leq \max_{r=1, \dots, R^S} \|A_r\|^2, \quad (67)$$

We now briefly present the Lipschitz constant of the most used layers, namely the ReLU layer, the pooling layer and the affine transforms. Finally, we will conclude with the softmax nonlinearity, which is present in any classification framework.

Firstly we study the general case of the affine transforms.

Theorem 6. *For the affine mappings and nonlinearity layers we have*

$$\begin{aligned} K_W &= \|Wx_1 + b - (Wx_2 + b)\|^2 = \|Wx_1 - Wx_2\|^2 \\ &\leq \|W\|^2 \|x_1 - x_2\|^2 \end{aligned} \quad (68)$$

$$\begin{aligned} K_C &= \|Cx_1 + b - (Cx_2 + b)\|^2 = \|Cx_1 - Cx_2\|^2 \\ &\leq \|C\|^2 \|x_1 - x_2\|^2 \end{aligned} \quad (69)$$

$$\begin{aligned} K_{\sigma} &= \|A_{\sigma}[x_1]x_1 - A_{\sigma}[x_2]x_2\|^2 \leq \max_{r=1, \dots, R^S} \|A_{\sigma,r}\|^2 \|x_1 - x_2\|^2 \\ &\leq D^2 \|x_1 - x_2\|^2 \end{aligned} \quad (70)$$

$$\begin{aligned} K_{\rho} &= \|A_{\rho}[x_1]x_1 - A_{\rho}[x_2]x_2\|^2 \leq \max_{r=1, \dots, R^S} \|A_{\rho,r}\|^2 \|x_1 - x_2\|^2 \\ &\leq D^2 \|x_1 - x_2\|^2 \end{aligned} \quad (71)$$

with D representing the output dimension. Those translate into the norm of the weight for the FC and convolutional layers and upper bounded by the output dimension for ReLU, LReLU and max-pooling.

All the demonstration of the results are provided in the Appendix. We also present the softmax nonlinearity which is a strictly contractive operator. In fact, we have the following result. The softmax layer is strictly contractive with $K \leq \frac{C-1}{C^2}$. In general given a composition of affine spline operators with parameters $\mathbf{A}^{(\ell)}, \mathbf{b}^{(\ell)}$ for the ℓ^{th} operator, we have the Lipschitz constant of their composition defined as

$$\|f_{\Theta}(x) - f_{\Theta}(y)\|^2 \leq \left(\prod_{\ell=1}^L \max_{r=1, \dots, R} \|A_r^{(\ell)}\|^2 \right) \|x - y\|^2, \quad (72)$$

The composition of LSOs thus inherits this property regarding its Lipschitz constant. Based on the previously derived upper bounds we can thus analyze in general the regularity property of DNNs depending on their layer composition and special topologies or weight constraints. In particular, we propose to study the case of adversarial examples, a typical application of perturbation leading to unstable outputs. In fact, as presented in earlier section, adversarial examples represents an optimized perturbation introduced into an input before being fed into the DNN mapping. The large SNR implies that if the DNN output changes drastically, there is a clear regularity drawback for the mapping. As in practice this perturbation is able to

completely fool the network making it predict with very high accuracy the incorrect class. Hence, practical evidence demonstrate the lack of contractivity. Hence, based on the previous analysis, we see that two reasons exist.

Corollary 1. *Adversarial examples are caused by an "explosion" of the weight norms coupled with very high-dimensional mappings.*

In order to solve or lessen this effect two solutions appear. Firstly, regularization applied on the weights can reduce the norms hence the irregularity of the mapping. As seen before, this penalty term is also crucial for optimal template learning. However there is also a second way to prevent unstable outputs and it is via sparsity of the activation. This denotes the number of neurons firing after a ReLU or LReLU nonlinearity and can be easily measured given an input and a given layer as $\|A_{\sigma}^{(\ell)}[z^{(\ell-1)}]z^{(\ell-1)}\|_0$. This activation sparsity in fact can be upper bounded by a quantity smaller than the unconstrained $D^{(\ell)}$. For example, replacing the standard nonlinearity with one letting go through only the κ order statistics brings down the nonlinearity Lipschitz upper bound to κ . Another solution can be to impose an extra layer before the nonlinearity with aim to structure the input such that it can not be all positive, the worst case for the ReLU. This solution is discussed in details in Sec. ??.

5 Extension

In this section we propose to leverage the developed tools to propose some solution to current DNN drawbacks. This will consist of proposing a systematic way to ensure regularization and generalization measures via DNN inversion and input reconstruction. This will also allow the development of a generic semi-supervised and unsupervised strategy for DNNs. Secondly, we will study the impact of inhibitor connections to provide network stability, bias removal. One of the key concept of this part will consist of studying DNNs from a dual point of view : forward (template inference) pass and backward (reconstruction, learning) path. As we will see, adding the right connections can increase the forward sparsity whereas densifying the backward pass. Finally, we will develop a simple methodology to measure the quality and potential of untrained DNN topologies and weight initialization. This will find great application in topology search and automated DNN design as there is no longer need to train the network to obtain a qualitative measure. In all cases, we also provide experiments on MNIST, CIFAR10.

5.1 DNN Inversion: Input Reconstruction is Necessary for Generalization

Deep learning systems have made great strides recently in a wide range of difficult machine perception tasks. However, most systems are still trained in a fully supervised fashion requiring a large set of labeled data, which can be extremely tedious and costly to acquire. Hence, there is a great need to study the inversion problem of DNN such that *semi-supervised* algorithms can be used, leveraging both labeled and unlabeled data for learning and inference. Limited progress has been made on semi-supervised learning algorithms for deep neural networks [Rasmus et al., 2015, Salimans et al., 2016, Patel et al., 2015, Nguyen et al., 2016] and today’s methods suffer from a range of drawbacks, including training instability, lack of topology generalization, and computational complexity. Most importantly, there exists no universal methodology to equip any given deep net with an inversion scheme. In this section, we develop a universal methodology to invert a network allowing input reconstruction. This will allow for semi-supervised learning which can also be extended to unsupervised tasks with arbitrary DNN mappings. Our approach simply relies on the derived inverse mapping strategy of a deep network allowing to add an additional term to the loss function. This extra term will guide the weight updates such that information contained in unlabeled data are incorporated to the network. Our key insight is that the defined and general inverse function can be easily derived and computed; thus for unlabeled data points we can both compute and minimize the error between the input signal and the estimate provided by applying the inverse function to the network output without extra cost or change in the used model. The simplicity of this approach, coupled with its universal applicability promise to significantly advance the purview of semi-supervised and unsupervised learning. A series of experiments demonstrate that these modified networks have attain state-of-the-art performance in a range of semi-supervised learning tasks.

A major drawback to supervised learning is the need for a massive set of fully labeled training data. *Semi-supervised learning* relaxes this requirement by leaning Θ based on two datasets: a fully labeled set \mathcal{D} of N training data pairs and a "complementary" unlabeled set $\mathcal{D}_u := \{X_n, n = 1, \dots, N_u\}$ of N_u training inputs. Unlabelled training data is useful for learning, because the unlabelled inputs provide information on the statistical distribution of the data and will help to guide the learning of Θ to classify the supervised dataset as well as characterize the unlabeled samples present in \mathcal{D}_u . The lionshare of deep learning research has focused on supervised learning, because it has not been clear how to best incorporate unlabeled data points into the loss function to incorporate those unlabeled examples information in f_{Θ} . However, there has

been limited progress in a few directions which we now review.

When considering network inversion, standard approaches [?] The only DNN models will reconstruction ability are based on autoencoders as [?]. While more complex topologies have been used such as stacked convolutional autoencoder [?] there exists two main drawbacks. Firstly, the difficulty to train complex models in a stable manner even when using per layer optimization. Secondly, the difficulty to leverage supervised information into the learning of the representation. If one considers the problem of semi-supervised learning with deep neural networks, different methods have been developed. The improved *generative adversarial network* (GAN) technique [Salimans et al., 2016] couples two deep networks; a generative model that can create new signal samples on the fly and a discriminative network predicting the class of the labeled examples as well as the generated versus original nature of the input. Both neural nets are trained jointly as in typical GAN frameworks but the fact that the discriminator has to perform both tasks force it to incorporate the unlabeled signal information to distinguish them from the fake one generated by the generator. This enables the generator and the classifier to better model the data distribution by using the labeled examples to learn the class boundaries and the unlabeled examples to learn the distinction between “true” and “fake” signals. The main drawbacks to this approach include training instability of GANS [Arjovsky et al., 2017], its lack of portability to time series, and high-resolution images (e.g., Imagenet [Deng et al., 2009]) which so far are out-of-reach of GANs [Salimans et al., 2016] as well as the extra time and memory cost of training two deep networks. The *semi-supervised with ladder network* approach [Rasmus et al., 2015] employs a per-layer denoising reconstruction loss, which enables the system to be viewed as a stacked denoising autoencoder which is a standard and until now only way to tackle unsupervised tasks. By forcing the last denoising autoencoder to output an encoding describing the class distribution of the input, this deep unsupervised model is turned into a semi-supervised model. The main drawback from this methods relies in the lack of a clear path to generalize it to other network topologies, such as recurrent network or residual networks. Also, the per layer “greedy” reconstruction loss might be too restrictive unless correctly weighted pushing the need for a precise and large cross-validation of hyper-parameters. The probabilistic formulation of deep convolutional nets presented in [Patel et al., 2015, Nguyen et al., 2016] natively supports semi-supervised learning. The main drawback of this approach lies in its probabilistic nature requiring activation function to be ReLU and the DNN topology to be a DCN as well as inheriting standard difficulties of probabilistic graphical models in the context of large scale high dimensional dataset.

We propose a simple way to invert any piecewise differentiable mapping including DNNs. We provide the inverse mapping which requires no change in the current models and is computationally optimal as input reconstruction results in a backward pass as is used today for weight updates via backpropagation. This efficiency coupled with grounded mathematical motivation highlighting the need to reconstruct makes the approach a necessary step to improve DNNs performances. While tremendous applications would leverage this generalized inversion scheme, we will focus here one demonstrating the benefits of input reconstruction as a network regularized and provide semi-supervised experiments where we are able to reach state-of-the-art results with different neural network topologies. Highlighting that any supervised model can be kept as they are while simply changing the loss function should bring the proposed method to any known or to be known model in deep learning.

As demonstrated in previous sections, DNNs can be considered as composition of linear splines or be closely approximated as such. As a result, DNNs can be rewritten as a linear spline of the form

$$f_{\Theta}(x) = A[x]X_n + b[x], \tag{73}$$

where we denote by f_{Θ} a general DNN, x a generic input, $A[x], b[x]$ the spline parameters conditioned

on the input region. Based on this interpretation, DNN can be considered as template matching machines where $A[x]$ plays the role of an input adaptive template. To illustrate this point we provide for two common topologies the exact input-output mappings. For a standard deep convolutional neural network (DCN) with succession of convolutions, nonlinearities and pooling, one had

$$z_{CNN}^{(L)}(x) = W^{(L)} \underbrace{\prod_{\ell=L-1}^1 A_{\rho}^{(\ell)} A_{\sigma}^{(\ell)} C^{(\ell)}}_{\text{Template Matching}} x + W^{(L)} \underbrace{\sum_{\ell=1}^{L-1} \left(\prod_{j=L-1}^{\ell+1} A_{\rho}^{(j)} A_{\sigma}^{(j)} C^{(j)} \right) \left(A_{\rho}^{(\ell)} A_{\sigma}^{(\ell)} b^{(\ell)} \right)}_{\text{Bias}} + b^{(L)}.$$

(74)

where $z^{(\ell)}$ represents the latent representation at layer ℓ . Similarly for a deep residual network one has

$$z_{RES}^{(L)}(x) = W^{(L)} \underbrace{\left[\prod_{\ell=L-1}^1 \left(A_{\sigma,in}^{(\ell)} C_{in}^{(\ell)} + C_{out}^{(\ell)} \right) \right]}_{\text{Template Matching}} x + \underbrace{\sum_{\ell=L-1}^1 \left(\prod_{i=L-1}^{\ell+1} \left(A_{\sigma,in}^{(i)} C_{in}^{(i)} + C_{out}^{(i)} \right) \right) \left(A_{\sigma,in}^{(\ell)} b_{in}^{(\ell)} + b_{out}^{(\ell)} \right)}_{\text{Bias}} + b^{(L)}.$$

(75)

Based on those findings, and by imposing a simple L2 norm upper bound on the templates, it has been shown that optimal DNNs should be able to generate templates proportional to their input, positively for the belonging class and negatively for the others.

Theorem 7. *In the case where all inputs have identity norm $\|X_n\| = 1, \forall n$ and assuming all templates denoted by $A[X_n]_c, c = 1, \dots, C$ have a norm constraints as $\sum_{c=1}^C \|A[X_n]_c\|^2 \leq K, \forall X_n$ then the unique globally optimal templates are*

$$A^*[X_n]_c = \begin{cases} \sqrt{\frac{C-1}{C}} K X_n, & \iff c = Y_n \\ -\sqrt{\frac{K}{C(C-1)}} X_n, & \text{else} \end{cases} \quad (76)$$

As a result, We now leverage the analytical optimal DNN solution to demonstrate that reconstruction is indeed implied by such an optimum.

5.1.1 Optimal DNN Leads to Input Reconstruction

The study of reconstruction in the context of DNN differs from standard signal processing. In fact, for any given input, the number of templates is defined as the number of classes C which is in general much smaller than the number of input dimensions D . Hence, approaches based on the study of orthogonal or over-complete basis can not be applied. Also, composition of mappings is hardly studied in this context unless dealing with deep NMF (CITE) types of approaches. On the other hand, the templates or atoms are adapted to the input as opposed to NMF where they are optimized explicitly. Firstly, we define the inverse

of a DNN as

$$\begin{aligned}
f_{\Theta}^{-1}(X_n) &= A[X_n]^T (A[X_n]X_n + b[X_n]) \\
&= \sum_{c=1}^C (\langle A[X_n]_c, X_n \rangle + b[X_n]_c) A[X_n]_c.
\end{aligned} \tag{77}$$

Theorem 8. *In the case of no or negligible $A[X_n]^T b[X_n]$, optimal templates defined in the previous Theorem lead to exact reconstruction as we have by definition*

$$\sum_{c=1}^C \langle A[X_n]_c, x \rangle A[X_n]_c = (C-1)K \frac{1}{C(C-1)} \|X_n\|^2 X_n + K \frac{C-1}{C} \|X_n\|^2 X_n = X_n$$

Hence optimal templates imply perfect reconstruction.

However, exact reconstruction is in practice not optimal as from all the content present in X_n only a subpart is sufficient for the task at hand, there also is presence of noise and so on. In fact, the quantity $A[X_n]^T b[X_n]$ is negligible but not null representing this fact. Hence if we now consider the problem of reconstructing a noisy input, we can bridge this scheme to standard thresholding and denoising scheme such as wavelet thresholding. In particular we study now the case of Relu or LReLU activations and mean or max pooling. The ReLU with bias is equivalent to asymmetric soft thresholding. Hence we can see the inversion of DNNs are equivalent to a composition of soft-thresholding based denoising operators. With minimization of the reconstruction error, one then has an adaptive filter-bank able to span the dataset. While not being sufficient for generalization it is necessary as absence of templates adapted to an input is synonym of false induced representation. We now present a particular application of the derived reconstruction loss: semi-supervised.

5.1.2 Boundary Inversion Method, Implementation and Loss Function

We briefly describe how to apply the proposed strategy to a given task with arbitrary DNN. As exposed earlier all the needed changes happen in the loss functions by adding extra terms. As a result if automatic differentiation is used as in Theano [Bergstra et al., 2010], TensorFlow [Abadi et al., 2016] for example then it is sufficient to change the loss function and all the updates will be adapted via the change in the gradients for each of the parameters. The great efficiency of this inversion schemes is due to the following. As we have seen in the previous section, any deep network can be rewritten as a linear mapping. This leads to a simple derivation of a network inverse defined as f^{-1} that will be used to derive our unsupervised and semi-supervised loss function via

$$\begin{aligned}
f^{-1}(x, A[x], b[x]) &= A[x]^T (A[x]x + b[x]) \\
&= A[x]^T f(x; \Theta) \\
&= \frac{df(x; \Theta)}{dx}^T f(x; \Theta).
\end{aligned} \tag{78}$$

The main efficiency argument thus comes from the fact that

$$A[x] = \frac{df(x; \Theta)}{dx}, \quad (79)$$

allowing to very efficiently compute this matrix on any deep networks via differentiation as it would be done to back-propagate a gradient for example.

Interestingly for neural networks and many common frameworks s.a. wavelet thresholding, PCA, . . . , \mathcal{E} is considered as the reconstruction error as $(\frac{df(x)}{dx})^T f(x)$ is the definition of the inverse transform. In particular and for illustration purposes, we present in the table below some common frameworks for which \mathcal{E} represents exactly the reconstruction loss and thus Eq.?? is considered as the inverse transform. This

Table 3: Examples of frameworks with similar inverse transform definition.

	α_i	$f(x)_i$	loss
Sparse Coding	Learned	$\frac{\langle x, W_i \rangle}{\ W_i\ ^2}$	$\ x - \sum_i \alpha_i \frac{df(x)_i}{dx}\ ^2 + \lambda \ \alpha\ _1$
NMF	Learned	$\langle x, W_i \rangle$	$\ x - \sum_i \alpha_i \frac{df(x)_i}{dx}\ ^2$ s.t. $J_f(x) \geq 0$
PCA	$f(x)_i$	$\langle x, W_i \rangle$	$\ x - \sum_i \alpha_i \frac{df(x)_i}{dx}\ ^2$ s.t. $J_f(x)$ orthonormal
Soft Wavelet Thresh.	$f(x)_i$	$\max(\langle x, W_i \rangle - b_i, 0) \text{sign}(\langle x, W_i \rangle)$	$\ x - \sum_i \alpha_i \frac{df(x)_i}{dx}\ ^2$
Hard Wavelet Thresh.	$f(x)_i$	$1_{ \langle x, W_i \rangle - b_i > 0} \langle x, W_i \rangle$	$\ x - \sum_i \alpha_i \frac{df(x)_i}{dx}\ ^2$
Best Basis (WTA)	$f(x)_i$	$1_{i=\text{argmax}_i \frac{\langle x, W_i \rangle}{\ W_i\ ^2}} \langle x, W_i \rangle$	$\ x - \sum_i \alpha_i \frac{df(x)_i}{dx}\ ^2$
k-NN	1	$1_{i=\text{argmax}_{\langle x, W_i \rangle > -\ W_i\ ^2/2}} \langle x, W_i \rangle$	$\ x - \sum_i \alpha_i \frac{df(x)_i}{dx}\ ^2$

inversion scheme is often seen as an ill-posed problem. In fact, for the ReLU case for example, given an output activation, the negative values that have been filtered can not be reconstructed. However with the proposed method, the implied reconstruction is 0. This corresponds to reconstruction the input on the boundary of the region defined by the current ReLU activation. As one reconstruct with values further away from 0 into the negative side, as the reconstruction goes away from the region. Hence our propose scheme can be seen as an optimistic case where the input was assumed to lie on the boundaries of the region.

We now describe how to incorporate this loss for semi-supervised and unsupervised learning. We first define the R reconstruction loss as

$$R(X_n) = \|(\frac{df_{\Theta}(X_n)}{dX_n})^T f_{\Theta}(X_n) - X_n\|^2. \quad (80)$$

While we use the mean squared error, any other reconstruction loss which is differentiable can be used s.a. cosine similarity. We also introduce an additional "specialization" loss defined as the Shannon entropy of the prediction:

$$E(\hat{y}(X_n)) = - \sum_{c=1}^C \hat{y}(X_n)_c \log(\hat{y}(X_n)_c), \quad (81)$$

pushing the output distribution to have low entropy when minimized. The need for this loss is to make the unlabeled prediction with low-entropy a.k.a predicting a one-hot representation. As a result, we define our complete loss function as the combination of the standard cross entropy loss for labeled data denoted by $L_{CE}(Y_n, \hat{y}(X_n))$, the reconstruction loss and entropy loss as

$$\mathcal{L}(X_n, Y_n) = \alpha L_{CE}(Y_n, \hat{y}(X_n)) 1_{\{Y_n \neq \emptyset\}} + (1 - \alpha)[\beta R(X_n) + (1 - \beta)E(X_n)], \alpha, \beta \in [0, 1]^2, \quad (82)$$

The parameters α, β are introduced to form a convex combination of the losses, 2 of them being unsupervised, with α controlling the ratio between supervised and unsupervised loss and β the ration between the two unsupervised losses.

One can also use the presented loss to perform unsupervised tasks and clustering. In fact, we can see that by setting $\alpha = 0$ we are in a fully unsupervised framework, and, depending on the value of β , pushing the mapping f_{Θ} to produce a low-entropy, clustered, representation or rather being unconstrained and simply producing optimal reconstruction. Even in a fully unsupervised and reconstruction case ($\alpha = 0, \beta = 1$) the proposed framework is not similar to a deep-autoencoder for two main reasons. The first one lies in the fact that there is no greedy (per layer) reconstruction loss, only the final output is considered in the reconstruction loss. Secondly, while in both case there is parameter sharing, in our case there is also "activation" sharing which corresponds to the states (spline) that were used in the forward pass that will also be used for the backward one. In a deep autoencoder, the backward activation states are induced by the backward projection and will most likely not be equal to the forward ones.

5.1.3 Semi-Sup Experiments for State-of-the-art Performances across Topologies

We now present results of the approach on a semi-supervised tasks on the MNIST dataset where we are able to obtain state-of-the-art performances with different topologies showing the ability of the method to generalize to any topology as well as being competitive. MNIST is made of 70000 grayscale images of shape 28×28 which is split into a training set of 60000 images and a test set of 10000 images. We present results for the case with $N_L = 50$ which represents the number of samples from the training set which are labeled. All the others are unlabeled. In addition, 10 different topologies are tested to show the portability of the approach. The DNNs architecture details as well as training procedures are detailed below. Furthermore, we tested $\alpha \in \{0.7, 0.5, 0.3\}$ and found that better results were obtained on average with $\alpha = 0.7$ and thus present below all results with $\alpha = 0.7, \beta = 0.5$. Running the proposed semi-supervised scheme on MNIST led to the results presented in the table below. We are able to obtain better results than all the standard benchmarks but one. In particular, with $N_L = 50$ we are also able to outperform most method which use 100 labels. We used Theano and Lasagne libraries. Details on the learning procedure and the used topologies are provided in the next section, the code for reproducible results is in the attached materials. The column Sup1000 for MNIST corresponds to the accuracy after training of DNN using only supervised loss on 1000 data, showing the great impact of the proposed solution.

In addition we present in Fig. 12 the reconstruction for the case $N_L = 50$ as well as the test set accuracy in Fig. 13.

5.1.4 Extensions

Among the possible extensions, one can develop the reconstruction loss into a per layer reconstruction loss. Doing so, there is the possibility to weight each layer penalty bringing flexibility as well as meaningful reconstruction. Let define the per layer loss as

$$\mathcal{L}(X_n, Y_n) = \alpha L_{CE}(Y_n, \hat{y}(X_n)) 1_{\{Y_n \neq \emptyset\}} + \beta E(X_n) + \sum_{\ell=0}^{L-1} \gamma^{(\ell)} R^{(\ell)}(X_n), \quad (83)$$

N_L	50	100	Sup1000
SmallCNNmean	99.07 ,(0.7, 0.2)		94.9
SmallCNNmax	98.63,(0.7, 0.2)		95.0
SmallUCNN	98.85,(0.5, 0.2)		96.0
LargeCNNmean	98.63,(0.6, 0.5)		94.7
LargeCNNmax	98.79,(0.7, 0.5)		94.8
LargeUCNN	98.23,(0.5, 0.5)		96.1
Resnet2-32mean	99.11 ,(0.7, 0.2)		95.5
Resnet2-32max	99.14 ,(0.7, 0.2)		94.9
UResnet2-32	98.84,(0.7, 0.2)		95.6
Resnet3-16mean	98.67,(0.7, 0.2)		95.4
Resnet3-16max	98.56,(0.5, 0.5)		94.8
UResnet3-16	98.7,(0.7, 0.2)		95.5
Improved GAN [Salimans et al., 2016]	97.79 \pm 1.36	99.07 \pm 0.065	
Auxiliary Deep Generative Model [Maaløe et al., 2016]	-	99.04	
LadderNetwork [Rasmus et al., 2015]	-	98.94 \pm 0.37	
Skip Deep Generative Model [Maaløe et al., 2016]	-	98.68	
Virtual Adversarial [Miyato et al., 2015]	-	97.88	
catGAN [Springenberg, 2015]	-	98.61 \pm 0.28	
DGN [Kingma et al., 2014]	-	96.67 \pm 0.14	
DRMM [Nguyen et al., 2016]	78.27	86.59	
DRMM +NN penalty	77.9	87.72	
DRMM+KL penalty	97.54	98.64	
DRMM +KL+NN penalties	99.09	99.43	

Table 4: Test Error on MNIST for 50 and 100 labeled examples for the used networks as well as comparison with other methods. The column Sup1000 demonstrates the raw performance of the same networks trained only with the supervised loss with 1000 labels.




Figure 12: Reconstruction of the studied DNN models for 4 test set samples. In each subplot the columns from left to right correspond to: the original image, mean-pooling reconstruction, max-pooling reconstruction, inhibitor connections. Each group of subplot represents a specific topology being in clockwise order: LargeCNN, SmallCNN, Resnet2-32 and Resnet3-16.

with

$$R^{(\ell)}(X_n) = \left\| \left(\frac{df_{\Theta}(X_n)}{dz^{(\ell)}(X_n)} \right)^T f_{\Theta}(X_n) - z^{(\ell)}(X_n) \right\|^2. \quad (84)$$

Doing so, one can adopt a strategy in favor of high reconstruction objective for inner layers, close to the final latent representation $z^{(L)}$ and lessen the reconstruction cost for layers closer to the input X_n . In fact, inputs of standard dataset are usually noisy, with background, and the object of interest only contains a small energy w.r.t. the total energy of X_n . Another extension would be to update the weighting while performing




Figure 13: Er present in in this figures the test set accuracies during learning for all the studied topologies. In black are the DINN, and blue mean and max pooling. From left to right, the topologies are: SmallCNN, LargeCNN, Resnet3-16, Resnet2-32.

learning. Hence, if we denote by t the position in time such as the current epoch or batch, we now have the previous loss becoming

$$\mathcal{L}(X_n, Y_n; \Theta) = \alpha(t)L_{CE}(Y_n, \hat{y}(X_n))1_{\{Y_n \neq \emptyset\}} + \beta(t)E(X_n) + \sum_{\ell=0}^{L-1} \gamma^{(\ell)}(t)R^{(\ell)}(X_n). \quad (85)$$

One approach would be to impose some deterministic policy based on heuristic such as favoring reconstruction at the beginning to then switch to classification and entropy minimization. Finer approaches could rely on an explicit optimization schemes for those coefficients. One way to perform this, would be to optimize the loss weighting coefficients $\alpha, \beta, \gamma^{(\ell)}$ after each batch or epoch by backpropagation on the updates weights. Let define

$$\Theta(t+1) = \Theta(t) - \lambda \frac{dL(X_n, Y_n)}{d\Theta}, \quad (86)$$

representing an generic iterative update based on a given policy such as gradient descent. One can thus adopt the following update strategy for the hyper-parameters as

$$\gamma^{(\ell)}(t+1) = \gamma^{(\ell)}(t) - \frac{dL(X_n, Y_n; \Theta(t+1))}{d\gamma^{(\ell)}(t)}, \quad (87)$$

and so for all hyper-parameters. Another approach would be to use adversarial training to update those hyper-parameters where both update cooperate trying to accelerate learning.

6 Conclusion

We presented a natural reformulation of deep neural network as composition of adaptive partitioning splines and in general linear spline operators. By doing so we have been able to explicitly determine the optimal network weights, their impact for adversarial example, generalization, memorization. From this we built an intuitive and generic method to invert arbitrary networks giving rise to semi-supervised and unsupervised application. We obtained state-of-the-art performances on MNIST with CNNs and Resnets and provided supplemental experiments highlighting the ability of the introduced reconstruction error to regularize and improve generalization. We also proposed a simple criterion to judge the quality of a network and its initialization based on template analysis allowing fast topology search. Finally, by bridging many fields such as template matching, adaptive partitioning splines, and deep neural networks, we hope to allow further and deeper analysis of all the presented results and insights.

References

- [Abadi et al., 2016] Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., et al. (2016). Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*.
- [Arjovsky et al., 2017] Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein gan. *arXiv preprint arXiv:1701.07875*.
- [Atteia and Benbourhim, 1989] Atteia, M. and Benbourhim, M. (1989). Spline elastic manifolds. *Mathematical methods in computer aided geometric design*, pages 45–50.
- [Bajcsy and Kovačič, 1989] Bajcsy, R. and Kovačič, S. (1989). Multiresolution elastic matching. *Computer vision, graphics, and image processing*, 46(1):1–21.
- [Bastien et al., 2012] Bastien, F., Lamblin, P., Pascanu, R., Bergstra, J., Goodfellow, I., Bergeron, A., Bouchard, N., Warde-Farley, D., and Bengio, Y. (2012). Theano: new features and speed improvements. *arXiv preprint arXiv:1211.5590*.
- [Bengio et al., 2013] Bengio, Y., Boulanger-Lewandowski, N., and Pascanu, R. (2013). Advances in optimizing recurrent networks. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 8624–8628. IEEE.
- [Berger et al., 1994] Berger, J., Coifman, R. R., and Goldberg, M. J. (1994). Removing noise from music using local trigonometric bases and wavelet packets. *Journal of the Audio Engineering Society*, 42(10):808–818.
- [Bergstra et al., 2010] Bergstra, J., Breuleux, O., Bastien, F., Lamblin, P., Pascanu, R., Desjardins, G., Turian, J., Warde-Farley, D., and Bengio, Y. (2010). Theano: A cpu and gpu math compiler in python. In *Proc. 9th Python in Science Conf*, pages 1–7.
- [Bezhaev, 1988] Bezhaev, A. Y. (1988). Splines on manifolds. *Russian Journal of Numerical Analysis and Mathematical Modelling*, 3(4):287–300.
- [Bishop, 2008] Bishop, C. M. (2008). Training with noise is equivalent to tikhonov regularization. *Training*, 7(1).

- [Bloor and Wilson, 1990] Bloor, M. I. and Wilson, M. J. (1990). Representing pde surfaces in terms of b-splines. *Computer-Aided Design*, 22(6):324–331.
- [Blumer et al., 1987] Blumer, A., Ehrenfeucht, A., Haussler, D., and Warmuth, M. K. (1987). Occam’s razor. *Information processing letters*, 24(6):377–380.
- [Boyd and Xu, 2009] Boyd, J. P. and Xu, F. (2009). Divergence (runge phenomenon) for least-squares polynomial approximation on an equispaced grid and mock–chebyshev subset interpolation. *Applied Mathematics and Computation*, 210(1):158–168.
- [Breiman, 1993] Breiman, L. (1993). Hinging hyperplanes for regression, classification, and function approximation. *IEEE Transactions on Information Theory*, 39(3):999–1013.
- [Burr, 1981] Burr, D. J. (1981). Elastic matching of line drawings. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 3(6):708.
- [Carlini and Wagner, 2016] Carlini, N. and Wagner, D. (2016). Defensive distillation is not robust to adversarial examples. *arXiv preprint*.
- [Cheney, 1980] Cheney, E. W. (1980). *Approximation theory III*, volume 12. Academic Press New York.
- [Choromanska et al., 2015] Choromanska, A., Henaff, M., Mathieu, M., Arous, G. B., and LeCun, Y. (2015). The loss surfaces of multilayer networks. In *AISTATS*.
- [Chui, 1988] Chui, C. K. (1988). *Multivariate splines*. SIAM.
- [Chung et al., 2014] Chung, J., Gulcehre, C., Cho, K., and Bengio, Y. (2014). Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- [Coifman and Wickerhauser, 1992] Coifman, R. R. and Wickerhauser, M. V. (1992). Entropy-based algorithms for best basis selection. *IEEE Transactions on information theory*, 38(2):713–718.
- [Cybenko, 1989] Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems (MCSS)*, 2(4):303–314.
- [de Boor and DeVore, 1983] de Boor, C. and DeVore, R. (1983). Approximation by smooth multivariate splines. *Transactions of the American Mathematical Society*, 276(2):775–788.
- [de Brébisson and Vincent, 2015] de Brébisson, A. and Vincent, P. (2015). An exploration of softmax alternatives belonging to the spherical loss family. *arXiv preprint arXiv:1511.05042*.
- [Deng et al., 2009] Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. IEEE.
- [Duchi et al., 2011] Duchi, J., Hazan, E., and Singer, Y. (2011). Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159.
- [Fawzi et al., 2015] Fawzi, A., Fawzi, O., and Frossard, P. (2015). Analysis of classifiers’ robustness to adversarial perturbations. *arXiv preprint arXiv:1502.02590*.
- [Gal and Ghahramani, 2016] Gal, Y. and Ghahramani, Z. (2016). Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059.

- [Glorot et al., 2011] Glorot, X., Bordes, A., and Bengio, Y. (2011). Deep sparse rectifier neural networks. In *Aistats*, volume 15, page 275.
- [Graves, 2013] Graves, A. (2013). Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*.
- [Graves and Schmidhuber, 2005] Graves, A. and Schmidhuber, J. (2005). Framewise phoneme classification with bidirectional lstm networks. In *Neural Networks, 2005. IJCNN'05. Proceedings. 2005 IEEE International Joint Conference on*, volume 4, pages 2047–2052. IEEE.
- [Gu and Rigazio, 2014] Gu, S. and Rigazio, L. (2014). Towards deep neural network architectures robust to adversarial examples. *arXiv preprint arXiv:1412.5068*.
- [Gu et al., 2006] Gu, X., He, Y., and Qin, H. (2006). Manifold splines. *Graphical Models*, 68(3):237–254.
- [Guyon et al., 1992] Guyon, I., Vapnik, V., Boser, B., Bottou, L., and Solla, S. A. (1992). Structural risk minimization for character recognition. In *Advances in neural information processing systems*, pages 471–479.
- [Hannah and Dunson, 2013] Hannah, L. A. and Dunson, D. B. (2013). Multivariate convex regression with adaptive partitioning. *The Journal of Machine Learning Research*, 14(1):3261–3294.
- [He et al., 2016] 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.
- [Hecht-Nielsen et al., 1988] Hecht-Nielsen, R. et al. (1988). Theory of the backpropagation neural network. *Neural Networks*, 1(Supplement-1):445–448.
- [Hinton, 1986] Hinton, G. E. (1986). Learning distributed representations of concepts. In *Proceedings of the eighth annual conference of the cognitive science society*, volume 1, page 12. Amherst, MA.
- [Hinton, 1987] Hinton, G. E. (1987). Learning translation invariant recognition in a massively parallel networks. In *International Conference on Parallel Architectures and Languages Europe*, pages 1–13. Springer.
- [Hochreiter and Schmidhuber, 1995] Hochreiter, S. and Schmidhuber, J. (1995). Simplifying neural nets by discovering flat minima. In *Advances in neural information processing systems*, pages 529–536.
- [Hofer and Pottmann, 2004] Hofer, M. and Pottmann, H. (2004). Energy-minimizing splines in manifolds. *ACM Transactions on Graphics (TOG)*, 23(3):284–293.
- [Hornik et al., 1989] Hornik, K., Stinchcombe, M., and White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366.
- [Jayaraman et al., 2009] Jayaraman, S., Esakkirajan, S., and Veerakumar, T. (2009). Digital image processing tmh publication. *Year of Publication*.
- [Kim and De Araújo, 2007] Kim, H. Y. and De Araújo, S. A. (2007). Grayscale template-matching invariant to rotation, scale, translation, brightness and contrast. In *Pacific-Rim Symposium on Image and Video Technology*, pages 100–113. Springer.
- [Kingma and Ba, 2014] Kingma, D. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

- [Kingma et al., 2014] Kingma, D. P., Mohamed, S., Rezende, D. J., and Welling, M. (2014). Semi-supervised learning with deep generative models. In *Advances in Neural Information Processing Systems*, pages 3581–3589.
- [Korman et al., 2013] Korman, S., Reichman, D., Tsur, G., and Avidan, S. (2013). Fast-match: Fast affine template matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2331–2338.
- [Lang et al., 1990] Lang, K. J., Waibel, A. H., and Hinton, G. E. (1990). A time-delay neural network architecture for isolated word recognition. *Neural networks*, 3(1):23–43.
- [LeCun et al., 1989] LeCun, Y. et al. (1989). Generalization and network design strategies. *Connectionism in perspective*, pages 143–155.
- [LeCun et al., 1995] LeCun, Y., Jackel, L., Bottou, L., Cortes, C., Denker, J. S., Drucker, H., Guyon, I., Muller, U., Sackinger, E., Simard, P., et al. (1995). Learning algorithms for classification: A comparison on handwritten digit recognition. *Neural networks: the statistical mechanics perspective*, 261:276.
- [Li et al., 2016] Li, Y., Xu, R., and Liu, F. (2016). Whiteout: Gaussian adaptive regularization noise in deep neural networks. *arXiv preprint arXiv:1612.01490*.
- [Lyu et al., 2015] Lyu, C., Huang, K., and Liang, H.-N. (2015). A unified gradient regularization family for adversarial examples. In *Data Mining (ICDM), 2015 IEEE International Conference on*, pages 301–309. IEEE.
- [Maaløe et al., 2016] Maaløe, L., Sønderby, C. K., Sønderby, S. K., and Winther, O. (2016). Auxiliary deep generative models. *arXiv preprint arXiv:1602.05473*.
- [MacKay, 1996] MacKay, D. J. (1996). Bayesian methods for backpropagation networks. In *Models of neural networks III*, pages 211–254. Springer.
- [Magnani and Boyd, 2009] Magnani, A. and Boyd, S. P. (2009). Convex piecewise-linear fitting. *Optimization and Engineering*, 10(1):1–17.
- [Mallat, 1999] Mallat, S. (1999). *A wavelet tour of signal processing*. Academic press.
- [Mallat, 2008] Mallat, S. (2008). *A wavelet tour of signal processing: the sparse way*. Academic press.
- [Mallat, 2016] Mallat, S. (2016). Understanding deep convolutional networks. *Phil. Trans. R. Soc. A*, 374(2065):20150203.
- [Matsuoka, 1992] Matsuoka, K. (1992). Noise injection into inputs in back-propagation learning. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(3):436–440.
- [Meyer, 1993] Meyer, Y. (1993). Algorithms and applications. *SIAM, philadelphia*.
- [Miyato et al., 2015] Miyato, T., Maeda, S.-i., Koyama, M., Nakae, K., and Ishii, S. (2015). Distributional smoothing by virtual adversarial examples. *stat*, 1050:2.
- [Moody and Utans, 1994] Moody, J. and Utans, J. (1994). Architecture selection strategies for neural networks: Application to corporate bond rating prediction. In *Neural networks in the capital markets*, pages 277–300. John Wiley & Sons.

- [Morgan and Bourlard, 1990] Morgan, N. and Bourlard, H. (1990). Generalization and parameter estimation in feedforward nets: Some experiments. In *Advances in neural information processing systems*, pages 630–637.
- [Murray and Edwards, 1993] Murray, A. F. and Edwards, P. J. (1993). Synaptic weight noise during mlp learning enhances fault-tolerance, generalization and learning trajectory. In *Advances in neural information processing systems*, pages 491–498.
- [Nalisnick et al., 2015] Nalisnick, E., Anandkumar, A., and Smyth, P. (2015). A scale mixture perspective of multiplicative noise in neural networks. *arXiv preprint arXiv:1506.03208*.
- [Nguyen et al., 2016] Nguyen, T., Liu, W., Perez, E., Baraniuk, R. G., and Patel, A. B. (2016). Semi-supervised learning with the deep rendering mixture model. *arXiv preprint arXiv:1612.01942*.
- [Nishikawa, 1998] Nishikawa, H. (1998). Accurate piecewise linear continuous approximations to one-dimensional curves: Error estimates and algorithms.
- [Nowlan and Hinton, 1992] Nowlan, S. J. and Hinton, G. E. (1992). Simplifying neural networks by soft weight-sharing. *Neural computation*, 4(4):473–493.
- [Olshausen et al., 1996] Olshausen, B. A. et al. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609.
- [Papernot et al., 2016] Papernot, N., McDaniel, P., Wu, X., Jha, S., and Swami, A. (2016). Distillation as a defense to adversarial perturbations against deep neural networks. In *Security and Privacy (SP), 2016 IEEE Symposium on*, pages 582–597. IEEE.
- [Patel et al., 2015] Patel, A. B., Nguyen, T., and Baraniuk, R. G. (2015). A probabilistic theory of deep learning. *arXiv preprint arXiv:1504.00641*.
- [Peña, 2000] Peña, J. M. (2000). On the multivariate horner scheme. *SIAM journal on numerical analysis*, 37(4):1186–1197.
- [Plaut et al., 1986] Plaut, D. C. et al. (1986). Experiments on learning by back propagation.
- [Rasmus et al., 2015] Rasmus, A., Berglund, M., Honkala, M., Valpola, H., and Raiko, T. (2015). Semi-supervised learning with ladder networks. In *Advances in Neural Information Processing Systems*, pages 3546–3554.
- [Reinsch, 1967] Reinsch, C. H. (1967). Smoothing by spline functions. *Numerische mathematik*, 10(3):177–183.
- [Rister and Rubin, 2017] Rister, B. and Rubin, D. L. (2017). Piecewise convexity of artificial neural networks. *Neural Networks*, 94:34–45.
- [Rumelhart et al., 1988] Rumelhart, D. E., Hinton, G. E., Williams, R. J., et al. (1988). Learning representations by back-propagating errors. *Cognitive modeling*, 5(3):1.
- [Rumelhart and McClelland, 1986] Rumelhart, D. E. and McClelland, J. L. (1986). Parallel distributed processing: Explorations in the microstructure of cognition: Foundations (parallel distributed processing).
- [Salimans et al., 2016] Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. (2016). Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, pages 2226–2234.

- [Savel'ev, 1995] Savel'ev, I. V. (1995). Splines and manifolds. *Russian Mathematical Surveys*, 50(6):1306–1307.
- [Schmidhuber, 1994] Schmidhuber, J. (1994). Discovering problem solutions with low kolmogorov complexity and high generalization capability. In *Machine Learning: Proceedings of the Twelfth International Conference*. Citeseer.
- [Schoenberg, 1964] Schoenberg, I. J. (1964). On interpolation by spline functions and its minimal properties. In *On Approximation Theory/Über Approximationstheorie*, pages 109–129. Springer.
- [Schumaker, 2007] Schumaker, L. (2007). *Spline functions: basic theory*. Cambridge University Press.
- [Shaham et al., 2015] Shaham, U., Yamada, Y., and Negahban, S. (2015). Understanding adversarial training: Increasing local stability of neural nets through robust optimization. *arXiv preprint arXiv:1511.05432*.
- [Smith, 1985] Smith, G. D. (1985). *Numerical solution of partial differential equations: finite difference methods*. Oxford university press.
- [Springenberg, 2015] Springenberg, J. T. (2015). Unsupervised and semi-supervised learning with categorical generative adversarial networks. *arXiv preprint arXiv:1511.06390*.
- [Srivastava et al., 2014] Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958.
- [Szegedy et al., 2013] 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*.
- [Tieleman and Hinton, 2012] Tieleman, T. and Hinton, G. (2012). Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2):26–31.
- [Tropp, 2004] Tropp, J. A. (2004). Greed is good: Algorithmic results for sparse approximation. *IEEE Transactions on Information theory*, 50(10):2231–2242.
- [Uchida and Sakoe, 2005] Uchida, S. and Sakoe, H. (2005). A survey of elastic matching techniques for handwritten character recognition. *IEICE transactions on information and systems*, 88(8):1781–1790.
- [Vapnik, 1992] Vapnik, V. (1992). Principles of risk minimization for learning theory. In *Advances in neural information processing systems*, pages 831–838.
- [Veit et al., 2016] Veit, A., Wilber, M. J., and Belongie, S. (2016). Residual networks behave like ensembles of relatively shallow networks. In *Advances in Neural Information Processing Systems*, pages 550–558.
- [Wager et al., 2013] Wager, S., Wang, S., and Liang, P. S. (2013). Dropout training as adaptive regularization. In *Advances in neural information processing systems*, pages 351–359.
- [Weigend et al., 1990] Weigend, A. S., Huberman, B. A., and Rumelhart, D. E. (1990). Predicting the future: A connectionist approach. *International journal of neural systems*, 1(03):193–209.
- [Williams, 1995] Williams, P. M. (1995). Bayesian regularization and pruning using a laplace prior. *Neural computation*, 7(1):117–143.

- [Wolpert, 1994] Wolpert, D. H. (1994). Bayesian backpropagation over io functions rather than weights. In *Advances in neural information processing systems*, pages 200–207.
- [Xu et al., 2015] Xu, B., Wang, N., Chen, T., and Li, M. (2015). Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*.
- [Yann, 1987] Yann, L. (1987). *Modèles connexionnistes de l'apprentissage*. PhD thesis, These de Doctorat, Universite Paris 6.
- [Zeiler, 2012] Zeiler, M. D. (2012). Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.
- [Zhang et al., 2016] Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2016). Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*.
- [Zhang et al., 1997] Zhang, J., Yan, Y., and Lades, M. (1997). Face recognition: eigenface, elastic matching, and neural nets. *Proceedings of the IEEE*, 85(9):1423–1435.

A Extra Material

A.1 Spline Operator

In this section we first review the literature on splines in order to ease the introduction of multivariate spline functions which will be the building block of the spline operators. We then discuss properties of the defined mathematical objects s.t. the next section on the rewriting of deep neural networks via spline operators become intuitive.

A.1.1 Spline Functions[FINI]

Throughout this section, the main reference comes from the formidable monograph [Schoenberg, 1964] revisited ad extended in [Schumaker, 2007]. We first present standard univariate splines constructed via piecewise polynomial functions. Let's now present the beautiful development of spline functions.

We first define the space of *univariate polynomials of order m* as

$$\mathcal{P}_m = \{p : p(x) = \sum_{i=1}^m c_i x^{i-1}, c_1, \dots, c_m, x \in \mathbb{R}\}, \quad (88)$$

$$= \text{span}\{1, x, \dots, x^{m-1}\}, \quad (89)$$

note that the order m is equal to the number of degrees of freedom, and is the degree of the polynomial minus one. A specific element of this set denoted by p_m is thus fully determined by its specific coefficients $\mathbf{c} = (c_1, \dots, c_m) \in \mathbb{R}^m$.

Definition 8. we denote the polynomial of order m with parameters $\mathbf{c} \in \mathbb{R}^m$ by

$$\begin{aligned} p_m[\mathbf{c}] : \mathbb{R} &\rightarrow \mathbb{R} \\ x &\rightarrow \sum_{i=1}^m c_i x^{i-1}. \end{aligned} \quad (90)$$

While a polynomial $p_m[\mathbf{c}]$ acts with the same set of parameters across all its input space, it is possible to define a partition Ω of the input space on which "local" polynomials can act, hence depending on the region of an input, not necessarily the same polynomial with parameters will be used for the mapping, this defines piecewise polynomials. The set Ω being a collection of R regions ω_r of the input space forms a partition of $[a, b] \subset \mathbb{R}$ s.t.

$$\begin{aligned} \Omega &= \{\omega_r, i = 1, \dots, R\}, \\ \omega_i \cap \omega_j &= \emptyset, \forall i \neq j, \\ \cup_{r=1}^R \omega_r &= [a, b], \end{aligned} \quad (91)$$

where the partition can be extended to the real line by adding the elements $]-\infty, a[$ and $]b, \infty[$. From this,

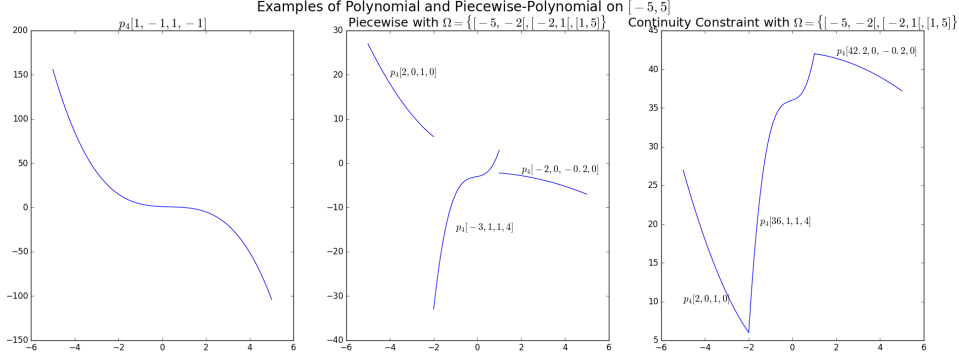


Figure 14: Illustrative examples of polynomial/piecewise-polynomial/continuous piecewise-polynomial.

we denote the space of piecewise polynomials of order $\mathbf{m} = (m_1, \dots, m_R) \in \mathbb{Z}^R$ as

$$\mathcal{PP}_{\mathbf{m}}(\Omega) = \{f : \forall x \in \omega_r, \exists p_r \in \mathcal{P}_{m_r} | f(x) = p_r(x), r = 1, \dots, R\}. \quad (92)$$

It is clear that the number of degrees of freedom is thus m_r for region r and thus for the piecewise polynomial it is equal to $\sum_{r=1}^R m_r$. We denote this collection of coefficients as \mathbf{c} and the coefficients specific to region r by $\mathbf{c}_r \in \mathbb{R}^{m_r}$.

Definition 9. We denote the piecewise polynomial of order \mathbf{m} with parameters $\mathbf{c} \in \mathbb{R}^{\sum_{r=1}^R m_r}$, a partition of $O \subset \mathbb{R}$ defined by Ω with $\text{Card}(\Omega) = R$ by

$$\begin{aligned} pp_{\mathbf{m}}[\mathbf{c}, \Omega] : \mathbb{R} &\rightarrow \mathbb{R} \\ x &\rightarrow \sum_{r=1}^R p_{m_r}[\mathbf{c}_r](x) \mathbf{1}_{\{x \in \omega_r\}} = \sum_{r=1}^R \left(\sum_{i=1}^{m_r} \mathbf{c}_{r,i} x^{i-1} \right) \mathbf{1}_{\{x \in \omega_r\}}. \end{aligned} \quad (93)$$

This defines a mapping with which different polynomials can be "activated" depending on the input location and the partition, hence the piecewise property.

We can now define the space of splines as a subset of this piecewise polynomial space by adding a "regularity" constraint between polynomials of neighbouring/adjacent regions. If we denote by $x_r, r = 0, \dots, R$ the "knots" of the partitions with $x_0 = a$ and $x_R = b$ then the space of polynomial splines is defined as

$$\begin{aligned} S(\mathcal{P}_m, \mathcal{M}, \Omega) = \{f : \forall x \in \omega_r, \exists p_r \in \mathcal{P}_m, f(x) = p_r(x), D^j p_r(x_r) = D^j p_{r+1}(x_r), \\ j = 0, \dots, m-1-M_i, r = 1, \dots, R\} \end{aligned} \quad (94)$$

with $\mathcal{M} = (M_1, \dots, M_K)$ the multiplicity vector which is the "smoothness" conditioning of the spline and D^j the differential operator of order j . Two simple examples are presented below :

$$\begin{aligned} S(\mathcal{P}_m, (m, \dots, m), \Omega) &= \mathcal{PP}_{(m, \dots, m)}(\Omega), \\ S(\mathcal{P}_m, (1, \dots, 1), \Omega) &= \mathcal{PP}_{(m, \dots, m)}(\Omega) \cap C^{m-2}([a, b]), \end{aligned}$$

the first is the least constraint spline while the latter is the most constraint one in term of boundary conditions.

A.1.2 Multivariate spline functions[FINI]

We present the multivariate polynomials and splines allowing to process multivariate inputs of dimension $K \geq 1$, with a slight difference from the literature [Schumaker, 2007, Schoenberg, 1964, Chui, 1988, de Boor and DeVore, 1983] where we allow non-rectangular regions for general spaces of dimension d , as most of the development of irregular grids focus on 2/3-dimensional spaces for PDE specific applications. We thus omit here the introduction to the known tensor multivariate splines as they are constructed with rectangular regions and thus will not be used in the later sections. From this, all the tools will be made clear for us to present the next section which consists of "adapting" the spline and piecewise multivariate polynomial terminology and notations for use in deep learning via the development of the spline operator, for mappings going to $\mathbb{R}^K, K > 1$.

Definition 10. *Let first define the multivariate integer set as*

$$\mathbb{N}^d = \{\alpha : \alpha = (\alpha_1, \dots, \alpha_d), \alpha_i \in \mathbb{N}, i = 1, \dots, d\}. \quad (95)$$

Using this notation, we denote the space of multivariate polynomials, given $\Lambda \subset \mathbb{N}^d$ as

$$\mathcal{P}_\Lambda^d = \text{span}\{x^\lambda : \lambda \in \Lambda\}, \quad (96)$$

$$= \{p : p(x) = \sum_{\lambda \in \Lambda} c_\lambda x^\lambda, c_\lambda \in \mathbb{R}\}, \quad (97)$$

where we denoted $x^\lambda = \prod_{i=1}^d x_i^{\lambda_i}$ with x_i the i^{th} input dimension of x , we also denote by \mathbf{c} the ordered collection of the c_λ . The collection Λ thus holds all the possible configuration of power for each of the input dimension and their configuration, as each element $\lambda \in \Lambda$ defines uniquely a combination of some power of the input dimensions. For example in the 3-dimensional setting, $\lambda = (2, 1, 0)$ leads to $x^\lambda = x_1^2 x_2$.

Definition 11. *We rewrite the multivariate polynomial acting on \mathbb{R}^d with parameters $\mathbf{c} \in \mathbb{R}^{\text{Card}(\Lambda)}$ and order $\Lambda \subset \mathbb{Z}_+^d$ by*

$$\begin{aligned} p_\Lambda^d[\mathbf{c}] : \mathbb{R}^d &\rightarrow \mathbb{R} \\ x &\rightarrow \sum_{\lambda \in \Lambda} c_\lambda x^\lambda. \end{aligned} \quad (98)$$

Note that a particular case occurs given a tuple $\mathbf{m} = (m_1, \dots, m_d)$ with the property that $\Lambda_{\mathbf{m}} = \{\alpha : 0 \leq \alpha_i \leq m_i, i = 1, \dots, d\} = \otimes_{i=1}^d \{0, \dots, m_i\}$, we also denote by Λ_m the case where $\mathbf{m} = (m, \dots, m)$. For example with $\mathbf{m} = (2, 2)$ we have the basis functions $\mathcal{P}_{\Lambda_2}^d = \{1, x, y, xy\}$ which is a bilinear polynomial. We now develop the piecewise version of the multivariate polynomial space. Given an arbitrary partition Ω of $O \subset \mathbb{R}^d$ and corresponding $\Lambda_r \subset \mathbb{N}^d$ we define a piecewise multivariate polynomial

$$\mathcal{PP}_\Lambda^d = \{f : \forall \omega \in \Omega, \forall x \in \omega, \exists p \in \mathcal{P}_{\Lambda_r}^d : f(x) = p(x)\}, \quad (99)$$

Definition 12. *We rewrite the piecewise multivariate polynomial acting on $O \subset \mathbb{R}^d$ with Ω a partition of O ,*

corresponding $\Lambda_r \subset \mathbb{Z}_+^d$ and with parameters $\mathbf{c} \in \mathbb{R}^{\sum_{r=1}^R \text{Card}(\Lambda_r)}$ by

$$pp_{\Lambda}^d[\mathbf{c}, \Omega] : \mathbb{R}^d \rightarrow \mathbb{R}$$

$$x \rightarrow \sum_{r=1}^R \left(\sum_{\lambda \in \Lambda_r} c_{r,\lambda} x^\lambda \right) 1_{\{x \in \omega_r\}}. \quad (100)$$

From this set of piecewise-polynomial functional, we can now define the space of splines by adding a smoothness constraints between neighboring regions. As in the univariate case we introduce \mathcal{M} the tuple of regularization coefficients forcing for each neighboring regions to have piecewise polynomials with same first derivatives, up to the order specified by the \mathcal{M} entry corresponding to it. As our work focus on two simple cases of regularization, we present the most and least constraint multivariate splines, with respectively $\mathcal{M} = (m, \dots, m) := \mathcal{M}_m$ and $\mathcal{M} = (1, \dots, 1) := \mathcal{M}_1$ as

$$\mathcal{S}(\mathcal{P}_{\Lambda_m}^d; \Omega; \mathcal{M}_1) = \mathcal{PP}_{\Lambda_m}^d(\Omega) \cap \mathcal{C}^{m-2}(O), \quad (101)$$

$$\mathcal{S}(\mathcal{P}_{\Lambda_m}^d; \Omega; \mathcal{M}_m) = \mathcal{PP}_{\Lambda_m}^d(\Omega). \quad (102)$$

where we remind that $\Lambda_m = \otimes_{i=1}^d \{0, \dots, m\}$. Since we have $\mathcal{S}(\mathcal{P}_{\Lambda_m}^d; \Omega; \mathcal{M}_1) \subset \mathcal{S}(\mathcal{P}_{\Lambda_m}^d; \Omega; \mathcal{M}_m)$ we now present the general formulation and results for the latter as other cases are "restricted" cases than can be derived from it, we thus now denote a multivariate spline simply by $\mathcal{S}(\mathcal{P}_{\Lambda_m}^d; \Omega)$. Given a spline $s^d \in \mathcal{S}(\mathcal{P}_{\Lambda_m}^d; \Omega)$ we denote by \mathbf{c} the parameters of the splines, namely, the coefficients of the m -order polynomials for each region.

Definition 13. We denote the multivariate spline for the case $s \in \mathcal{S}(\mathcal{P}_{\Lambda_m}^d; \Omega)$ with parameters $\mathbf{c} \in \mathbb{R}^{m \times \text{Card}(\Omega)}$ given a partition Ω of $O \subset \mathbb{R}^d$ by

$$s[\mathbf{c}, \mathcal{P}_{\Lambda_m}^d, \Omega] : \mathbb{R}^d \rightarrow \mathbb{R} \quad (103)$$

$$x \rightarrow s[\mathbf{c}, \mathcal{P}_{\Lambda_m}^d, \Omega](x) = pp_{\Lambda_m}^d[\mathbf{c}, \Omega](x) \quad (104)$$

where this definition is based on the use of \mathcal{M}_m as defined above.

If we denote by $\mathcal{J}[\Omega_k][x]$ the index of the region in which x belongs according to the Ω_k partitioning as

$$\mathcal{J}[\Omega_k] : \mathbb{R}^d \rightarrow \{1, \dots, \text{Card}(\Omega_k)\} \quad (105)$$

$$x \rightarrow \sum_{i=1}^{\text{Card}(\Omega_k)} i * 1_{\{x \in \omega_{k,i}\}}, \quad (106)$$

where $\omega_{k,i}$ is the i^{th} region of Ω_k . It is clear that given an input x . We can thus simplify notations by introducing this region selection operator coupled with the multivariate spline now defined as

$$s[\mathbf{c}, \mathcal{P}_{\Lambda_m}^d, \Omega] : \mathbb{R}^d \rightarrow \mathbb{R} \quad (107)$$

$$x \rightarrow s[\mathbf{c}, \mathcal{P}_{\Lambda_m}^d]_{\mathcal{J}[\Omega]}(x). \quad (108)$$

Linear Multivariate Splines[FINI]

We describe briefly a special case in which all the local mappings are linear. It will become of importance in the next section for the introduction of affine spline operators and neural networks. We first consider a special multivariate polynomial defined with total order m and denoted by $\mathcal{P}_{|m|}^d$ where $|m|$ denotes the total order property as opposed to the standard polynomials of order m . It is explicitly defined as the space

Definition 14.

$$\mathcal{P}_{|m|}^d = \text{span}\left\{\prod_{i=1}^d x_i^{\alpha_i}, \alpha \in \mathbb{Z}_+^d, |\alpha| < |m|\right\}, \quad (109)$$

with $m \in \mathbb{N}$ and $|\alpha| = \sum_{i=1}^d \alpha_i$.

As an example, if we consider the 3D space with $x = (x_1, x_2, x_3)$ the space of linear polynomials as

$$\mathcal{P}_{|2|}^3 = \text{span}\{1, x_1, x_2, x_3\}, \text{ 3D polynomial of total order 2}$$

as opposed to the nonlinear but multi-linear case of

$$\mathcal{P}_2^3 = \text{span}\{1, x_1, x_2, x_3, x_1x_2, x_1x_3, x_2x_3\}, \text{ 3D polynomial of order 2.}$$

As we will focus on linear cases in the remaining of the study we now simplify notations for the latter.

Definition 15. A linear K dimensional polynomial $p \in \mathcal{P}_{|2|}^K$ with coefficients $\mathbf{a} \in \mathbb{R}^K$, $b \in \mathbb{R}$ is denoted as

$$p^K[\mathbf{a}, b](x) = \langle \mathbf{a}, x \rangle + b. \quad (110)$$

From this linear polynomial we define the linear piecewise polynomials according to the last section notations. Given a partition of $O \subset \mathbb{R}^d$ defined as $\Omega = \{\omega_i, i = 1, \dots, I\}$ s.t. $\cup_{i=1}^I \omega_i = O$ and $\omega_i \cap \omega_j = \emptyset, \forall i \neq j$, the space of linear piecewise polynomials is defined as

$$\mathcal{PP}_{|2|}^d(\Omega) = \{f : \forall x \in \omega_i, \exists p_i \in \mathcal{P}_{|2|}^d | f(x) = p_i(x), i = 1, \dots, K\}. \quad (111)$$

From this, splines are defined in a similar fashion as in the last section where we have

$$\mathcal{S}(\mathcal{P}_{|2|}^d; \Omega; \mathcal{M}_1) = \mathcal{PP}_{|2|}^d(\Omega) \cap \mathcal{C}^{m-2}(O), \quad (112)$$

$$\mathcal{S}(\mathcal{P}_{|2|}^d; \Omega; \mathcal{M}_2) = \mathcal{PP}_{|2|}^d(\Omega). \quad (113)$$

We now focus for the case of total order 2 and we omit the constraint \mathcal{M}_m as we present the general case, any other one is a restriction of the coefficients to fulfill the smoothness boundary condition.

Definition 16. We denote multivariate linear splines the case with total order $|2|$ denoted by $\mathcal{S}(\mathcal{P}_{|2|}^d; \Omega)$ given the partition Ω of $O \subset \mathbb{R}^d$ with $\text{Card}(\Omega) = R$ and with parameters $(\mathbf{a}_r, b_r) \in \mathbb{R}^d \times \mathbb{R}, r = 1, \dots, R$ by using the last proposition as

$$s[(\mathbf{a}_r, b_r)_{r=1}^R; \mathcal{P}_{|2|}^d, \Omega](x) = \sum_{r=1}^R (\langle \mathbf{a}_r, x \rangle + b_r) 1_{\{x \in \omega_r\}} \quad (114)$$

$$= \mathbf{a}[x]^T x + b[x], \quad (115)$$

where a_r represent the slope and b_r the intercept for each region. The input dependant selection is abbreviated via

$$\mathbf{a}[x] = \mathbf{a}_{\mathcal{J}[\Omega](x)} \text{ and } b[x] = b_{\mathcal{J}[\Omega](x)}. \quad (116)$$

From now on, the term multivariate is dropped as the mappings will be explicit. We define a **local linear spline** function as a special case where the support of \mathbf{a}_r is constrained, in the sense that some dimensions are constraint to be 0. This forces \mathbf{a}_r to only act on a sub-part of the input x , thus the local property where locality is again in the dimension domain as opposed to the input domain. This "0-constraint" is denoted by the collection of indices on which we enforce it om each of the \mathbf{a}_r :

$$\Gamma(\mathbf{a}_r) := \{i \in \{1, \dots, d\} : [\mathbf{a}_r]_i \equiv 0\}, \quad (117)$$

where $[\mathbf{a}_r]_i$ denotes the i^{th} dimension of \mathbf{a}_r and $[\mathbf{a}_r]_i \equiv 0$ represents the presence of the 0-constraint for the given dimension. We also denote by $\Gamma^C(\mathbf{a}_r)$ the unconstrained part where the C upperscript stands for contrapose. In fact, those two collections are complementary w.r.t to the list of indices and thus given one the other is uniquely defined as $\Gamma^C(\mathbf{a}_r) = \{1, \dots, D\} \setminus \Gamma(\mathbf{a}_r)$ and vice-versa. Finally, for further precision, we call a **uniform local linear spline** the case where $\Gamma(\mathbf{a}_r) = \Gamma(\mathbf{a}_p), \forall r \neq p$, thus, the dimensions of x on which the linear transform acts does not depend on the partition, they are "shared". This constraint allows one to control the way the mapping "sensitivity" to the input space dimensions, in fact by setting the constraints one can ensure that some changes in the input for those constraint dimensions will not affect the output. In fact, given a uniform local linear spline with constraint $\Gamma(\mathbf{a}_1)$ we have

$$\forall \epsilon \in \mathbb{R}^d, [\epsilon]_i = 0, i \in \Gamma(\mathbf{a}_1), s[(\mathbf{a}_r, b_r)_{r=1}^R; \Omega](x) = s[(\mathbf{a}_r, b_r)_{r=1}^R; \Omega](x + \epsilon) \quad (118)$$

$$\mathcal{S} \left[\left(s[\mathbf{c}_k, \mathcal{P}_{\Lambda_{m_k}}^d, \Omega_k] \right)_{k=1}^K \right] (x) = \sum_{\alpha \in \alpha} \begin{bmatrix} s[\mathbf{c}_1, \mathcal{P}_{\Lambda_{m_1}}^d, \Omega_1](x) \\ \vdots \\ s[\mathbf{c}_K, \mathcal{P}_{\Lambda_{m_K}}^d, \Omega_K](x) \end{bmatrix} \mathbf{1}_{\{x \in \omega_\alpha\}} \quad (119)$$

$$= \sum_{\alpha \in \alpha} \begin{bmatrix} \sum_{\lambda \in \Lambda_{m_1}} c_{\alpha_1, \lambda} x^\lambda \\ \vdots \\ \sum_{\lambda \in \Lambda_{m_K}} c_{\alpha_K, \lambda} x^\lambda \end{bmatrix} \mathbf{1}_{\{x \in \omega_\alpha\}} \quad (120)$$

A.2 Dataset Memorization Proof

Theorem 9. Assuming all templates denoted by $A[X_n]_c, c = 1, \dots, C$ have a norm constraints as $\sum_{c=1}^C \|A[X_n]_c\|^2 \leq K, \forall X_n$ and that all the inputs have identity norm $\|X_n\| = 1, \forall x$ then the unique optimal templates are

$$A^*[X_n]_c = \begin{cases} \sqrt{\frac{C-1}{C}} K X_n, & \iff c = Y_n \\ -\sqrt{\frac{K}{C(C-1)}} X_n, & \text{else} \end{cases} \quad (121)$$

Proof. We aim at minimizing the cross-entropy loss function for a given input X_n belonging to class Y_n , we

also have the constraint $\sum_{c=1}^C \|A[X_n]_c\|^2 \leq K$. The loss function is thus convex on a convex set, it is thus sufficient to find a extremum point. We denote the augmented loss function with the Lagrange multiplier as

$$l(A[X_n]_1, \dots, A[X_n]_C, \lambda) = -\langle A[X_n]_{Y_n}, X_n \rangle + \log \left(\sum_{c=1}^C e^{\langle A[X_n]_c, X_n \rangle} \right) - \lambda \left(\sum_{c=1}^C \|A[X_n]_c\|^2 - K \right).$$

The sufficient KKT conditions are thus

$$\begin{aligned} \frac{dl}{dA[X_n]_1} &= -1_{\{Y_n=1\}}x + \frac{e^{\langle A[X_n]_1, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}}x - 2\lambda A[X_n]_1 = \underline{0} \\ &\vdots \\ \frac{dl}{dA[X_n]_C} &= -1_{\{Y_n=1\}}x + \frac{e^{\langle A[X_n]_C, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}}x - 2\lambda A[X_n]_C = \underline{0} \\ \frac{\partial l}{\partial \lambda} &= K - \sum_{c=1}^C \|A[X_n]_c\|^2 = 0 \end{aligned}$$

We first proceed by identifying λ as follows

$$\begin{aligned} \left. \begin{array}{l} \frac{dl}{dA[X_n]_1} = \underline{0} \\ \vdots \\ \frac{dl}{dA[X_n]_C} = \underline{0} \end{array} \right\} &\implies \sum_{c=1}^C A[X_n]_c^T \frac{dl}{dA[X_n]_c} = 0 \\ &\implies -\langle A[X_n]_{Y_n}, x \rangle + \sum_{c=1}^C \frac{e^{\langle A[X_n]_c, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} \langle A[X_n]_c, x \rangle - 2\lambda \sum_{c=1}^C \|A[X_n]_c\|^2 = 0 \\ &\implies \lambda = \frac{1}{2K} \left(\sum_{c=1}^C \frac{e^{\langle A[X_n]_c, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} \langle A[X_n]_c, x \rangle - \langle A[X_n]_{Y_n}, x \rangle \right) \end{aligned}$$

Now we plug λ in $\frac{dl}{dA[X_n]_k}, \forall k = 1, \dots, C$

$$\begin{aligned} \frac{dl}{dA[X_n]_k} &= -1_{\{Y_n=k\}}x + \frac{e^{\langle A[X_n]_k, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}}x - 2\lambda A[X_n]_k \\ &= -1_{\{Y_n=k\}}x + \frac{e^{\langle A[X_n]_k, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}}x - \frac{1}{K} \sum_{c=1}^C \frac{e^{\langle A[X_n]_c, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} \langle A[X_n]_c, x \rangle A[X_n]_k \\ &\quad + \frac{1}{K} \langle A[X_n]_{Y_n}, x \rangle A[X_n]_k \end{aligned}$$

we now leverage the fact that $A[X_n]_i = A[X_n]_j, \forall i, j \neq Y_n$ to simplify notations

$$\begin{aligned}
\frac{dl}{dA[X_n]_k} &= \left(\frac{e^{\langle A[X_n]_k, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} - 1_{\{k=Y_n\}} \right) x - \frac{C-1}{K} \frac{e^{\langle A[X_n]_i, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} \langle A[X_n]_i, x \rangle A[X_n]_k \\
&\quad + \frac{1}{K} \left(1 - \frac{e^{\langle A[X_n]_{Y_n}, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} \right) \langle A[X_n]_{Y_n}, x \rangle A[X_n]_k \\
&= \left(\frac{e^{\langle A[X_n]_k, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} - 1_{\{k=Y_n\}} \right) x - \frac{C-1}{K} \frac{e^{\langle A[X_n]_i, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} \langle A[X_n]_i, x \rangle A[X_n]_k \\
&\quad + \frac{C-1}{K} \frac{e^{\langle A[X_n]_i, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} \langle A[X_n]_{Y_n}, x \rangle A[X_n]_k \\
&= \left(\frac{e^{\langle A[X_n]_k, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} - 1_{\{k=Y_n\}} \right) x + \frac{C-1}{K} \frac{e^{\langle A[X_n]_i, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} \langle A[X_n]_{Y_n} - A[X_n]_i, x \rangle A[X_n]_k
\end{aligned}$$

We now proceed by using the proposed optimal solutions $A^*[X_n]_c, c = 1, \dots, C$ and demonstrate that it leads to an extremum point which by nature of the problem is the global optimum. We denote by i any index different from Y_n , first case $k = Y_n$:

$$\begin{aligned}
\frac{dl}{dA[X_n]_{Y_n}} &= - \left(1 - \frac{e^{\langle A[X_n]_{Y_n}, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} \right) x + \frac{C-1}{K} \frac{e^{\langle A[X_n]_i, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} \langle A[X_n]_{Y_n} - A[X_n]_i, x \rangle A[X_n]_{Y_n} \\
&= - \frac{C-1}{K} \frac{e^{\langle A[X_n]_i, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} x + \frac{C-1}{K} \frac{e^{\langle A[X_n]_i, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} \langle A[X_n]_{Y_n} - A[X_n]_i, x \rangle A[X_n]_{Y_n} \\
&= - \frac{C-1}{K} \frac{e^{\langle A[X_n]_i, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} (-x + \langle A[X_n]_{Y_n} - A[X_n]_i, x \rangle A[X_n]_{Y_n}) \\
&= - \frac{C-1}{K} \frac{e^{\langle A[X_n]_i, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} \left(-x + \langle \sqrt{\frac{C-1}{C}} K X_n + \sqrt{\frac{K}{C(C-1)}} X_n, x \rangle \sqrt{\frac{C-1}{C}} K X_n \right) \\
&= \frac{(C-1)}{K} \frac{e^{\langle A[X_n]_i, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} (-X_n + \|X_n\|^2 X_n) \\
&= 0
\end{aligned}$$

Other cases $k \neq Y_n$

$$\begin{aligned}
\frac{dl}{dA[X_n]_i} &= \frac{e^{\langle A[X_n]_i, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} x + \frac{C-1}{K} \frac{e^{\langle A[X_n]_i, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} \langle A[X_n]_{Y_n} - A[X_n]_i, x \rangle A[X_n]_i \\
&= \frac{e^{\langle A[X_n]_i, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} \left(x + \frac{C-1}{K} \langle A[X_n]_{Y_n} - A[X_n]_i, x \rangle A[X_n]_i \right) \\
&= \frac{e^{\langle A[X_n]_i, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} \left(x - \frac{C-1}{K} \left\langle \sqrt{\frac{C-1}{C}} K X_n + \sqrt{\frac{K}{C(C-1)}} X_n, x \right\rangle \sqrt{\frac{K}{C(C-1)}} X_n \right) \\
&= \frac{e^{\langle A[X_n]_i, x \rangle}}{\sum_{c=1}^C e^{\langle A[X_n]_c, x \rangle}} (X_n - \|X_n\|^2 X_n) \\
&= 0
\end{aligned}$$

□

A.3 Conditions for local to global inference

As shown in the classification tasks, having the global inference property is not always synonym of better accuracy. In fact, it is not because one network is able to produce optimal templates according to its constraint topology that the resulting mapping is "better" than a local inference done on an "unconstraint" mapping. We believe that better network conditioning would allow to have global inference and be "optimal" for classification tasks across topologies. We remind that the used constraints are only sufficient conditions and thus can be replaced with many others. Overall, we also believe that the induced convex property that is paired with the global template inference property could also be leveraged during optimization to obtained faster and smarter training as for example is the case with sparse coding.

As we saw, a spline operator made of convex independent multivariate splines can have a input region selection or inference easily done making it input adaptive agnostic of the final space partition. Since a deep neural network is a composition of spline operators it is interesting to study the conditions for this inference to see if its locally optimal inference can become a global optimal inference.

We now study the conditions in order for one to pull this maximization process outside of the inner layer which would transform this greedy per layer maximization a global maximization step as for example in the case of two layers

$$\operatorname{argmax}_{\Phi^{(1)}} \langle \mathcal{S}^{(2)}(\mathcal{S}^{(1)}(x)), 1 \rangle = \operatorname{argmax}_{\Phi^{(1)}} \langle \mathcal{S}^{(1)}(x), 1 \rangle.$$

In order to analyze this possibility we first remind that in the case of a composition of affine spline operators, we can always rewrite the inner layers mappings as an affine transform, thus we now present the following result.

Theorem 10. *Given a spline operator $\mathcal{S}[\Phi, \Omega]$ made of independent multivariate convex affine splines, we*

have

$$\begin{aligned} \operatorname{argmax}_{\phi \in \Phi} \langle BW\phi(y) + b, 1 \rangle &= \operatorname{argmax}_{\phi \in \Phi} \langle \phi(y), 1 \rangle \\ &= \Phi[y], \end{aligned} \quad (122)$$

if and only if $\sum_d (W_{d,k} + b_d) \kappa_d^B [W\phi(y) + b] > 0, \forall k$ and increases w.r.t. $\phi(y)$.

Proof. We seek to prove the equality of the argmax for every input of the layer y ,

$$\begin{aligned} \operatorname{argmax}_{\phi \in \Phi} \langle BW\phi(y) + b, 1 \rangle &= \operatorname{argmax}_{[\phi_1, \dots, \phi_K] \in \mathcal{C}[\phi_1[\cdot], \dots, \phi_K[\cdot]]} \langle BW\phi(y) + b, 1 \rangle \\ &= \operatorname{argmax}_{[\phi_1, \dots, \phi_K] \in \mathcal{C}[\phi_1[\cdot], \dots, \phi_K[\cdot]]} \left\langle \begin{bmatrix} \sum_d B[W\phi(y) + b]_{1,d} \sum_i W_{d,i} \phi_i(y) + \sum_d B[W\phi(y) + b]_{1,d} b_d \\ \vdots \\ \sum_d B[W\phi(y) + b]_{K,d} \sum_i W_{d,i} \phi_i(y) + \sum_d B[W\phi(y) + b]_{K,d} b_d \end{bmatrix}, 1 \right\rangle \\ &= \operatorname{argmax}_{[\phi_1, \dots, \phi_K] \in \mathcal{C}[\phi_1[\cdot], \dots, \phi_K[\cdot]]} \sum_k \sum_d B[W\phi(y) + b]_{k,d} \sum_i W_{d,i} \phi_i(y) + \sum_k \sum_d B[W\phi(y) + b]_{k,d} b_d \\ &= \operatorname{argmax}_{[\phi_1, \dots, \phi_K] \in \mathcal{C}[\phi_1[\cdot], \dots, \phi_K[\cdot]]} \sum_i \phi_i(y) \sum_d W_{d,i} \sum_k B[W\phi(y) + b]_{k,d} + \sum_d b_d \sum_k B[W\phi(y) + b]_{k,d} \\ &= \operatorname{argmax}_{[\phi_1, \dots, \phi_K] \in \mathcal{C}[\phi_1[\cdot], \dots, \phi_K[\cdot]]} \sum_i \phi_i(y) \alpha_i [W\phi(y) + b] + \beta [W\phi(y) + b] \\ &= \begin{bmatrix} \operatorname{argmax}_{\phi_1 \in \phi_1[\cdot]} \sum_i \phi_i(y) \alpha_i [W\phi(y) + b] + \beta [W\phi(y) + b] \\ \vdots \\ \operatorname{argmax}_{\phi_K \in \phi_K[\cdot]} \sum_i \phi_i(y) \alpha_i [W\phi(y) + b] + \beta [W\phi(y) + b] \end{bmatrix} \\ &= \begin{bmatrix} \operatorname{argmax}_{\phi_1 \in \phi_1[\cdot]} \phi_1(y) \\ \vdots \\ \operatorname{argmax}_{\phi_K \in \phi_K[\cdot]} \phi_K(y) \end{bmatrix} \\ &= \operatorname{argmax}_{\Phi \in \Phi[\cdot]} \langle \Phi(y), 1 \rangle \\ &= \Phi[y] \end{aligned}$$

□

Corollary 2. *In order to a deep neural network to have globally optimal inference, we have the following sufficient conditions*

- *Unconstrained first layer filters and bias*
- *Positive filters and nonnegative bias for inner-layers, strictly increasing nonlinearities, last layer should be a fc-layer.*

Proposition 3. *We this mentioned properties, one also has the following property*

$$\begin{aligned} \operatorname{argmax}_{\phi \in \Phi} (BW\phi(y) + b)_k &= \operatorname{argmax}_{\phi \in \Phi} \langle \phi(y), 1 \rangle \\ &= \Phi[y], \forall k \end{aligned} \quad (123)$$

Hence the local inference leads to the same spline as the one maximizing each output neuron of the network.

A.4 Space Contraction and Adversarial Examples

We present the softmax nonlinearity case which is as opposed to the intuition a strictly contractive operator. In fact, we have the following result.

Theorem 11. *The softmax layer is strictly contractive with $K = \frac{D-1}{D^2}$*

Proof. We now from that

$$\|f(x) - f(y)\|_2^2 \leq \max_x \|Df(x)\|_F^2 \|x - y\|_2^2, \quad (124)$$

thus we now analyze $\max_x \|Df(x)\|_F^2$.

$$\max_{p \in \Delta_D} \|Df(p)\|_F^2 = \max_{p \in \Delta_D} \sum_{i=1}^D \sum_{j=1, j \neq i}^D p_i^2 p_j^2 + \sum_{i=1}^D p_i^2 (1 - p_i)^2$$

where we used Δ_D the simplex of dimension D defined as

$$\Delta_D = \{x \in \mathbb{R}^{D+1} \mid \sum_i x_i = 1\}. \quad (125)$$

The Lagrangian is the augmented loss function with the constrain as

$$\mathcal{L}(p) = \sum_{i=1}^D \sum_{j=1, j \neq i}^D p_i^2 p_j^2 + \sum_{i=1}^D p_i^2 (1 - p_i)^2 + \lambda \left(\sum_{i=1}^D p_i - 1 \right),$$

we now seek the stationary points

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial p_k} &= 4p_k \sum_{j=1, j \neq k}^D p_j^2 + 2p_k(1-p_k)^2 - 2(1-p_k)p_k^2 + \lambda \\
&= 4p_k \sum_{j=1, j \neq k}^D p_j^2 + 2p_k - 4p_k^2 + 2p_k^3 - 2p_k^2 + 2p_k^3 + \lambda \\
&= 4p_k \sum_{j=1, j \neq k}^D p_j^2 + 2p_k - 6p_k^2 + 4p_k^3 + \lambda \\
&= 4p_k \sum_{j=1}^D p_j^2 + 2p_k - 6p_k^2 + \lambda, \quad \forall k \\
\frac{\partial \mathcal{L}}{\partial \lambda} &= \sum_{i=1}^D p_i - 1
\end{aligned}$$

we now seek to solve the system $\nabla \mathcal{L} = 0$. Since we have $\frac{\partial \mathcal{L}}{\partial p_k} = 0 \forall k$, it is clear that $\sum_{k=1}^D \frac{\partial \mathcal{L}}{\partial p_k} = 0$ leading to

$$\begin{aligned}
&\sum_{k=1}^D \frac{\partial \mathcal{L}}{\partial p_k} = 0 \\
\Rightarrow \sum_{k=1}^D \left(4p_k \sum_{j=1}^D p_j^2 + 2p_k - 6p_k^2 + \lambda \right) &= 0 \\
\Rightarrow 4 \sum_{j=1}^D p_j^2 + 2 - 6 \sum_{j=1}^D p_j^2 + D\lambda &= 0 \\
\Rightarrow D\lambda &= 2 \sum_{j=1}^D p_j^2 - 2 \\
\Rightarrow \lambda &= \frac{2}{D} \left(\sum_{j=1}^D p_j^2 - 1 \right).
\end{aligned}$$

Now plugin this into the gradient of \mathcal{L} , we have the updated system of equations

$$\begin{aligned}
4p_k \sum_{j=1}^D p_j^2 + 2p_k - 6p_k^2 + \\
\frac{2D}{D} \left(\sum_{j=1}^D p_j^2 - 1 \right) &= 0, \quad \forall k \\
\sum_{i=1}^D p_i &= 1
\end{aligned}$$

and thus in vector form leads to

$$\mathbf{p}(2\|\mathbf{p}\|_2^2 + 1) - 3\mathbf{p} \bullet \mathbf{p} = (1 - \|\mathbf{p}\|_2^2)\text{vec}(1/D),$$

where \bullet denotes the element-wise product. It is clear that \mathbf{p} must be constant across its dimensions, with the constraint this leads to $p = \text{vec}(1/D)$. It is indeed a maximum as on the boundary of the domain $f = 0$ and at the point we have

$$\begin{aligned} f(\text{vec}(1/D)) &= \sum_{i=1}^D \sum_{j=1, j \neq i}^D \frac{1}{D^4} + \sum_{i=1}^D \frac{1}{D^2} \left(1 - \frac{1}{D}\right)^2 \\ &= \frac{D-1}{D^3} + \frac{1}{D} \left(1 - \frac{1}{D}\right)^2 \\ &= \frac{D-1}{D^3} + \frac{1}{D} - \frac{2}{D^2} + \frac{1}{D^3} \\ &= \frac{D-1 + D^2 - 2D + 1}{D^3} \\ &= \frac{D(D-1)}{D^3} \\ &= \frac{D-1}{D^2} \end{aligned}$$

□

A.5 Function Approximation, Orbits, Class Separation, Generalization and Activation Graph

In this section, we develop some simple formulations of invariant learning and orbits in order to provide on way to define generalization for deep learning.

A.5.1 Function Approximation and Orbits for Invariant Learning

Firstly, as a deep network is a composition of affine spline operators, hence a spline operators itself, we analyze what is the function that is approximated. For the classification case, the objective is to learn the mapping that predicts a density distribution representation the probability of class belonging for the input. This means that if the target class is k we aim at learning

$$f(x) = \mathbf{e}_k, \forall x \in \mathcal{X}_k, \tag{126}$$

where \mathcal{X}_k represents the collection of data belonging to class k . It is clear from this formulation that f aims at learning orbits, namely the manifold of class k . Given that in general the input space is high dimensional, and the training set finite, the main challenge one has to face for generalization is to be able to perform efficient interpolation given new test points. Using splines for manifold learning and interpolation has been used for example in [Hofer and Pottmann, 2004, Savel'ev, 1995, Atteia and Benbourhim, 1989, Bezhaev, 1988, Gu et al., 2006] for low dimensional cases. In fact, the interpolation of splines is flexible enough to learn locally sharp functions via small partition of the input space, and yet allow for robust interpolation.

During the learning phase, by changing the weights regions of the input space are learned along with the per-region mappings. As a result during learning, back-propagation will guide the partition s.a. Eq. 126 is fulfilled. At test time, the regions are fixed and in order to generalize well, it is sufficient that the region in which the input belongs has the affine mapping corresponding to the right class.

We base the following analysis on [Mallat, 2016] in which orbits and invariance learning is brought in the context of deep learning, especially convolutional neural networks.

Firstly, for notation and simplification we call y_x the label associated to a given input x which belongs to some space $\mathcal{X} \subset \mathbb{R}^d$. First, all given inputs $x, x' \in \mathcal{X}^2$ are separated independently from the class belonging as

$$\|x - x'\| > 0, \forall y_x, y_{x'}, \quad (127)$$

but as this input is fed into a deep neural to be transformed into a succession of $z^{(1)}, z^{(2)}, \dots, z^{(L)}$, we aim at those representation to somehow regroup together when considering inputs from the same class, and conversely, for difference classes, these representation should not "colide". We represent this separation as

$$y_x \neq y_{x'} \implies \|z^{(\ell)}[x] - z^{(\ell)}[x']\| > 0, \forall \ell = 1, \dots, L \quad (128)$$

on the other hand, for same class inputs, we aim at having

$$y_x = y_{x'} \implies \|z^{(\ell)}[x] - z^{(\ell)}[x']\| > \|z^{(\ell+1)}[x] - z^{(\ell+1)}[x']\|, \forall \ell = 1, \dots, L - 1 \quad (129)$$

which is a soft condition requiring that in the limit of depth, same class input collide together. We denote by $D^{(\ell)}[x, x']$ the operator representing this separation as

$$D^{(\ell)}[x, x'] = \|z^{(\ell)}[x] - z^{(\ell)}[x']\|. \quad (130)$$

It is clear that the separation measure D is an indicator of the modeling power of deep neural networks. From this, the analysis proposed in [Mallat, 2016] relies on the aggregation or invariance to group actions while moving through the inner layers. Given a set \mathcal{X} , such as the space of images of the world, and a group G of operators acting on \mathcal{X} , an orbit of $x \in \mathcal{X}$ is defined as the set of all points

$$G.x = \{g.x : g \in G\}. \quad (131)$$

For example, G can be the group of rotation operators acting on images, and thus given an image, its orbit is the set of all the rotated version of this image. In general, this invariance learning is restricted to local symmetries. In fact, if we take the example of MNIST dataset, being invariant to small rotation is beneficial, yet, the "global" rotation group will bring a 1 to become a 7 or a 6 to become a 9. As a result, being invariant to the action group of rotations is globally detrimental but locally beneficial, as it is more generally for various transformations.

Via the presented work, we can now postulate on the way these orbits are approximated. In fact, we remind that a deep neural networks, given an input, produces an affine transformation to produce its output via the adapted template matching. Given one layer we simplify as an affine transformation followed by a nonlinearity:

$$z^{(\ell)} = \mathbf{A}^{(\ell)}(Wz^{(\ell-1)} + b^{(\ell)}), \quad (132)$$

we have that given two inputs $z_1^{(\ell-1)}$ and $z_2^{(\ell-1)}$, their respective output "energy" is for standard deep net-

works

$$\langle z_1^{(\ell)}, 1 \rangle = \max_A \langle A^{(\ell)}(W z_1^{(\ell-1)} + b^{(\ell)}), 1 \rangle, \quad (133)$$

$$\langle z_2^{(\ell)}, 1 \rangle = \max_A \langle A^{(\ell)}(W z_2^{(\ell-1)} + b^{(\ell)}), 1 \rangle. \quad (134)$$

Since the affine transforms are usually convolutions with filters of small sizes, it is likely that those two quantities will be close. However, analyzing their separateness leads to insightful results. We now consider the case of using the ReLU nonlinearity where we thus remind that the possible matrices A are diagonal with diagonal values belonging to $\{0, 1\}$.

$$\begin{aligned} D^{(\ell)}[x_1, x_2] &= \|A_1^{(\ell)}(W^{(\ell)} z_1^{(\ell-1)} + b) - A_2^{(\ell)}(W^{(\ell)} z_2^{(\ell-1)} + b)\|^2 \\ &= \|A_1^{(\ell)}(W^{(\ell)} z_1^{(\ell-1)} + b)\|^2 + \|A_2^{(\ell)}(W^{(\ell)} z_2^{(\ell-1)} + b)\|^2 \\ &\quad - 2\langle A_1^{(\ell)}(W^{(\ell)} z_1^{(\ell-1)} + b), A_2^{(\ell)}(W^{(\ell)} z_2^{(\ell-1)} + b) \rangle \\ &= \|A_1^{(\ell)}(W^{(\ell)} z_1^{(\ell-1)} + b)\|^2 + \|A_2^{(\ell)}(W^{(\ell)} z_2^{(\ell-1)} + b)\|^2 \\ &\quad - 2(W^{(\ell)} z_1^{(\ell-1)} + b)^T A_1^{(\ell)T} A_2^{(\ell)}(W^{(\ell)} z_2^{(\ell-1)} + b). \end{aligned}$$

While the first two terms simply represent the amount of energy going through the layer, the last term is the "pivot" as it is directly link to the "similarity" of the two inputs. In fact, for the ReLU and LReLU it is easy to see that $A_1^{(\ell)T} A_2^{(\ell)}$ is a diagonal matrix that represents the region associated with the intersection of the two input regions. In fact, to illustrate this, with the ReLU, the d^{th} element of the diagonal is 1 if the corresponding dimension of the affine transform is greater than 0. Now given two inputs, their affine transforms produce two features maps with associated ReLU as just described. An element-wise multiplication of those induced A matrices will "leave" a 1 in the d^{th} dimension if and only if both features maps have their d^{th} dimension greater than 0. As a result, it is clear that

$$- 2(W^{(\ell)} z_1^{(\ell-1)} + b)^T A_1^{(\ell)T} A_2^{(\ell)}(W^{(\ell)} z_2^{(\ell-1)} + b) \propto m(\omega_1 \cap \omega_2), \quad (135)$$

where ω_1 and ω_2 correspond to the regions in which $W^{(\ell)} z_1^{(\ell-1)} + b$ and respectively $W^{(\ell)} z_2^{(\ell-1)} + b$ belong to. As a result, we can see that as this measure of space intersection $m(\omega_1 \cap \omega_2)$ grows, as $D^{(\ell)}[x_1, x_2]$ diminishes. From this, it is interesting to note that the succession of $D^{(\ell)}[x_1, x_2], \ell = 1, \dots, L$ can be summarized by the succession of $m(\omega_1^{(\ell)} \cap \omega_2^{(\ell)})$, a.k.a the measure of the space intersection induced after all the affine transformations. If we now go back to the definition of orbit and invariance learning per layer, it is clear that it translates to learning a layer s.a. $m(\omega_1^{(\ell)} \cap \omega_2^{(\ell)})$ is proportional to the action group "locality" applied between x_1 and x_2 .

Also, on the first layers, the number of points from the training set belonging to each region might be very small as we remind that the number of possible region for a d -dimensional ReLU/LReLU is 2^d , and in practice $d \gg 784$ which is much obviously much smaller than any possible dataset. However, layers after layers, the dimension eventually goes down till ultimately being equal to K the number of classes to predict. Thus, it is likely that

$$m(\omega_1^{(1)}, \omega_2^{(1)}) \gg 0, \forall y_1, y_2,$$

thus all the points are "scattered" across the possible regions. From that, only remains the task to collide together points of the same class via this succession of spline operators.

For the network to generalize, it is now enough that for a new observation x we have $D^{(\ell)}[x, x_1] < D^{(\ell)}[x, x_2]$ for some ℓ and for all x_2 of the wrong classes. As this occur with greater ℓ as the amount of "generalization" required increases since it relates directly to the amount of composition needed to disentangle the hierarchy of group actions. If this occur at $\ell = 1$ for all possible points, this translates to have a task linearly separable in the first place. In practice we note that this measure $m(\omega_1^{(\ell)} \cap \omega_2^{(\ell)})$ can be easily defined via

$$m(\omega_1^{(\ell)} \cap \omega_2^{(\ell)}) = \|A_1^{(\ell)T} A_2^{(\ell)}\|. \quad (136)$$

A.5.2 Activation Graph, Paths

We now present some interesting results concerning the core structure of the activated regions of the affine spline operators that form deep neural networks. In particular we will see some properties when considering the following graph given an ordered collection of affine spline operators $\left(\mathcal{S}^{(\ell)}[\mathbf{A}^{(\ell)}, \mathbf{b}^{(\ell)}, \mathbf{\Omega}^{(\ell)}]\right)_{\ell=1}^L$ as

$$\mathcal{G}\left[\left(\mathcal{S}^{(\ell)}[\mathbf{A}^{(\ell)}, \mathbf{b}^{(\ell)}, \mathbf{\Omega}^{(\ell)}]\right)_{\ell=1}^L\right] = (V, E(x)) \quad (137)$$

with

$$V = \cup_{\ell=1}^L \mathbf{\Omega}^{(\ell)}, \quad (138)$$

$$E(x) = \{(\omega_{\alpha^*}^{(\ell)}, \omega_{\alpha^*}^{(\ell+1)})_{\ell=1}^{L-1}\}, \quad (139)$$

where it is clear that V the set of vertex corresponds to all the possible regions of all the possible spline operators while the set of edges which is input dependant corresponds to the active region linked from one spline operator to the next one. This input dependency is natural as we recall that for each spline function, depending on the region in which its input lie, an specific affine function is used to produce the output which in turn is fed into the next affine spline operator. This per spline region selection also characterize the input, and can be used to compare different input signals. Intuitively, if given two observations x_1 and x_2 the sequence of produced regions is the same, the two input x_1 and x_2 line in the same parametric line of $z^{(\ell)}$.

Proposition 4. *The possible regions that can take an input when considering each region of each spline as a node gives rise to a bipartite graph. Thus, $(V, E(x))$ is always a bipartite graph $\forall x$.*

It is clear that since the connectivity only go from one spline possible region to the other, we can create the bipartite graph by considering the two following sets of nodes

$$V_1 = \cup_{\ell=1, \ell \text{ odd}}^L \mathbf{\Omega}^{(\ell)},$$

$$V_2 = \cup_{\ell=1, \ell \text{ even}}^L \mathbf{\Omega}^{(\ell)},$$

and thus is clear that for any input x , we have the two partitions V_1 and V_2 as no edges can link two nodes of V_1 nor two nodes of V_2 .

Definition 17. *The activation graph (AG) corresponding to a signal x is the collection of active sub-regions among all the splines defining the deep network, and thus it is $E(x)$. Its root is an output and its leaves are at the input level.*




Figure 15: Examples of images tested for MNIST on 4 different classes Evolution of the sub-regions distances over the layers.

Note that the introduced bipartite graph $(V, E(x))$ via the selection regions of the affine spline operators is different from the derived neural network paths from [Choromanska et al., 2015, Nguyen et al., 2016]. In those works, the created graph has nodes all the possible neurons of any layers, and the edges are between neurons that fired, as those work only apply in the case of max-pooling and ReLU. As a result, the result AG has no interesting structure and is in general not usable for as the number of edges is extremely large. On the other hand, with the proposed region associated graph, we believe that much better results can be derived either in term of input characterization, outlier detection or generally structure behaving of the spline operators. By the natural sparsity we also reduce the computational overhead and are able to better visualize, interpret the AG.

A.5.3 Experiments

In order to highlight the aspect of "separation" between classes and invariance learning we provide some simple experiments on the MNIST dataset. We use three fully trained deep neural networks, the SmallCNN, LargeCNN and Resnet3 topologies. We then take a sample from the test set made of 25 examples of 4 classes. We present the distance matrix at each level of those networks before and after training. We used to compute the distance the definition introduced in 136. We present the used images in Fig. 15 and the corresponding distance matrices and results in Fig. 16,17,18 respectively for the SmallCNN, LargeCNN and resnet.

The class corresponding to the digit 1 is very interesting to analyze as all the given observations can be mapped to the same orbit just via a global rotation matrix as opposed to all the other classes provided here. As a result we can see the need for "depth" being null for this class, translated as a distance matrix almost optimal at the first layer output in Fig. 17.

SmallCNNmax Distance Through Layers Init/Trained

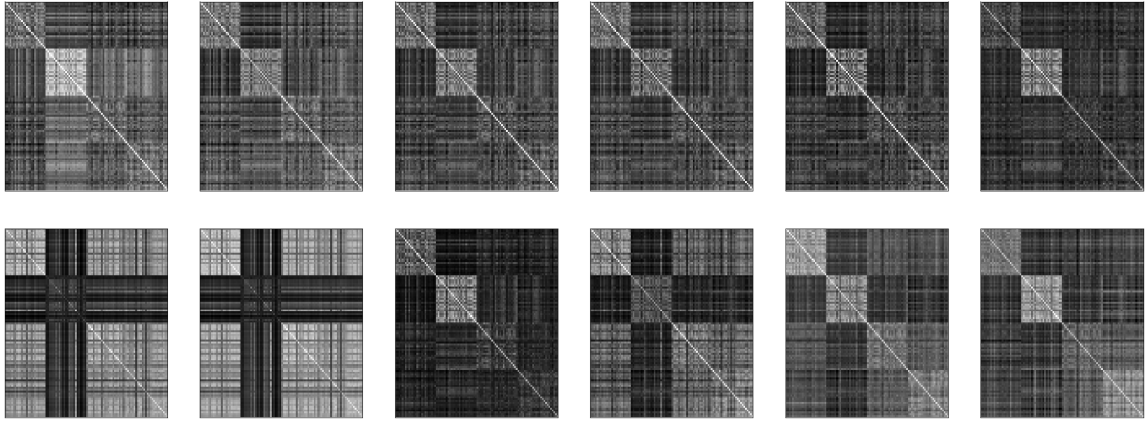


Figure 16: SmallCNNmean Evolution of the sub-regions distances over the layers.

LargeCNNmax Distance Through Layers Init/Trained




Figure 17: LargeCNNmax Evolution of the sub-regions distances over the layers. ditto figures 15 16.

Resnet3mean Distance Through Layers Init/Trained




Figure 18: Resnet3mean Evolution of the sub-regions distances over the layers.

B Dataset and Model Description

B.1 CIFAR10

Table 5: Mapping from integer to class label.

0	airplane
1	automobile
2	bird
3	cat
4	deer
5	dog
6	frog
7	horse
8	ship
9	truck

B.2 Networks Description and Training Details

We remind that for all the networks topologies, the mean/max describe the mode of the pooling layer whereas the presence of W transcribes to wavelets a.k.a the convex version of the network where the convolutional layers and nonlinearity layers are replaced with the one proposed in this work, also, F and NF represent the Fixed and NonFixed version of the newly introduced convolutional layer as described in its section. We present in the Table below the two topologies, Small and Large.

Table 6: Deep CNN Architectures

SmallCNN (131512)	LargeCNN ()
Input	Input
Conv (32,5,5) full	Conv (96,3,3) same
Pool (2,2)	Conv (96,3,3) full
Conv (64,3,3) valid	Conv (96,3,3) full
Conv (64,3,3) full	Pool (2,2)
Pool (2,2)	Conv (192,3,3) valid
Conv (128,3,3) valid	Conv (192,3,3) full
Conv (10,1,1)	Conv (192,3,3) valid
MeanPool (6,6)	Pool (2,2)
SoftMax	Conv (192,3,3) valid
	Conv (192,1,1) valid
	Conv (10,1,1) valid
	MeanPool(6,6)
	SoftMax

Note that for the standard cases (no W) the nonlinearities are always leaky rectifiers, for the convex cases (W) the first convolutional layer is left unconstrained and with a leaky rectifier in order to fulfill the template matching theorem.

Given this residual block, we define a simple Resnet (as opposed to wide Resnets) which topology depends on a parameter $n > 0$ defining the number of blocks per stages. The full Resnet has 3 stages of "widening" given by the following table.

For all topologies, Adam optimizer [Kingma and Ba, 2014] is used with a learning rate decay starting at 0.005 or 0.0005 for the Large topology. No regularization is applied nor batch normalization in order to better highlight the impact of the proposed changes. No shuffling is performed between epochs and the only normalization done is to have each observation with 0 mean and infinite norm equals to 1. For CIFAR10 this normalization is not done per channel (RGB) but across them. All the weights initialization as kept as default given the lasagne classes.

Table 7: Residual Architectures

Resnet3 ()	Resnet5 ()
Input	Input
Conv (16,3,3) valid	Conv (16,3,3) valid
Block (0)	Block (0)
Block (0)	Block (0)
Block (0)	Block (0)
Block (1)	Block (1)
Block (0)	Block (0)
Block (0)	Block (0)
Block (0)	Block (0)
Block (1)	Block (1)
Block (0)	Block (0)
Block (0)	Block (0)
GlobalMeanPool	Block (0)
SoftMax	Block (0)
	GlobalMeanPool
	SoftMax