Phase Collapse in Neural Networks

Florentin Guth [1] John Zarka [1] Stéphane Mallat [2,3]

[1]DI, ENS, CNRS, PSL University, Paris, France [2]Collège de France, Paris, France [3]Flatiron Institute, New York, USA

Introduction

CNNs progressively increase linear separability of image classes while collapsing spatial dimensions:

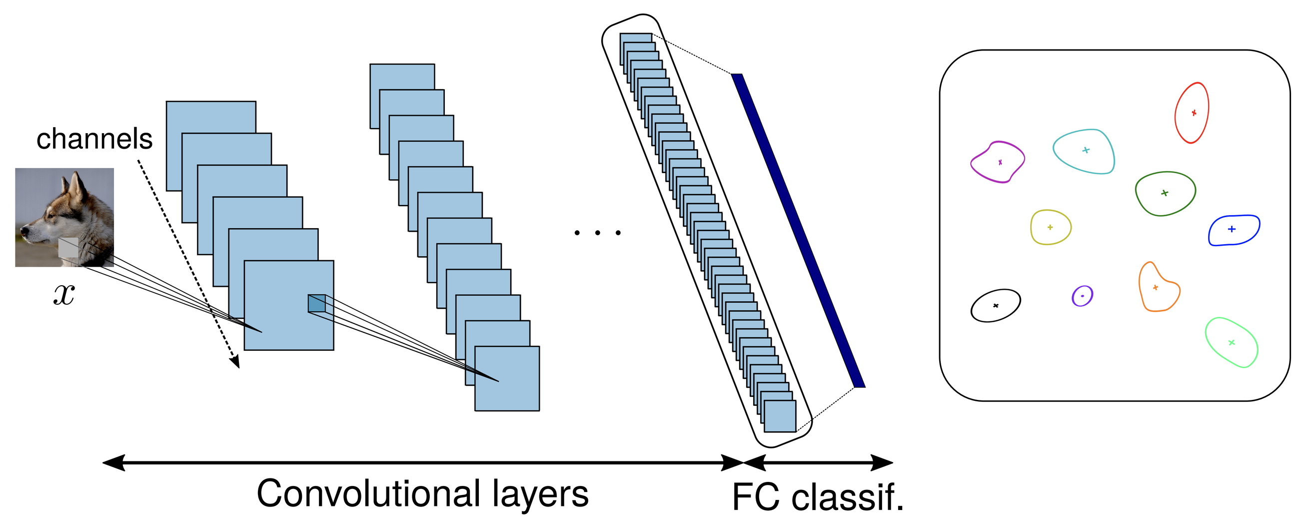

What is the role of the non-linearity and the network weights in this phenomenon?

Contributions:

- We show that a phase collapse of complex wavelet coefficients is both *sufficient* and *necessary* to reach ResNet-18 accuracy on ImageNet.
- Contrarily, iterated amplitude reductions, which threshold coefficients but preserve the phase, reach significantly lower accuracies.

⇒ **Provides an interpretation of the underlying functional role of ReLUs, otherwise hidden because of their flexibility.**

Small Translations as Phase Shifts

Complex wavelets diagonalize translations: a small translation τ of an image x is represented by a phase shift after a convolution with a complex wavelet ψ:

$$(\tau \cdot x) * \psi \approx e^{-i\xi \cdot \tau}(x * \psi)$$

with a relative error bounded by $\sigma|\tau|$.

ψ is a complex wavelet with center frequency ξ and spectral bandwidth σ:

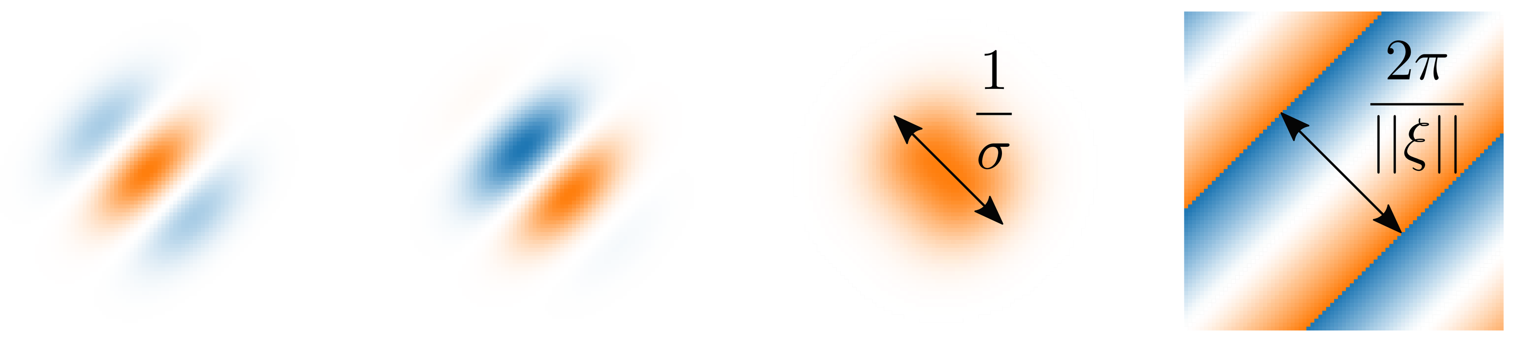

Real part Imaginary part Modulus Phase

For a class x invariant to translations, $x * \psi$ has a circularly-symmetric distribution because of phase shifts: $\forall \theta, \quad e^{i\theta} x \overset{d}{=} x$.

Class Mean Separation with Phase Collapses

Wavelet coefficients $x * \psi$ all have zero means for all translation-invariant classes. Linear classification cannot then use this information to discriminate such classes.

Phase collapse is defined as a complex modulus applied to wavelet coefficients. It can also be computed with real filters and ReLUs.

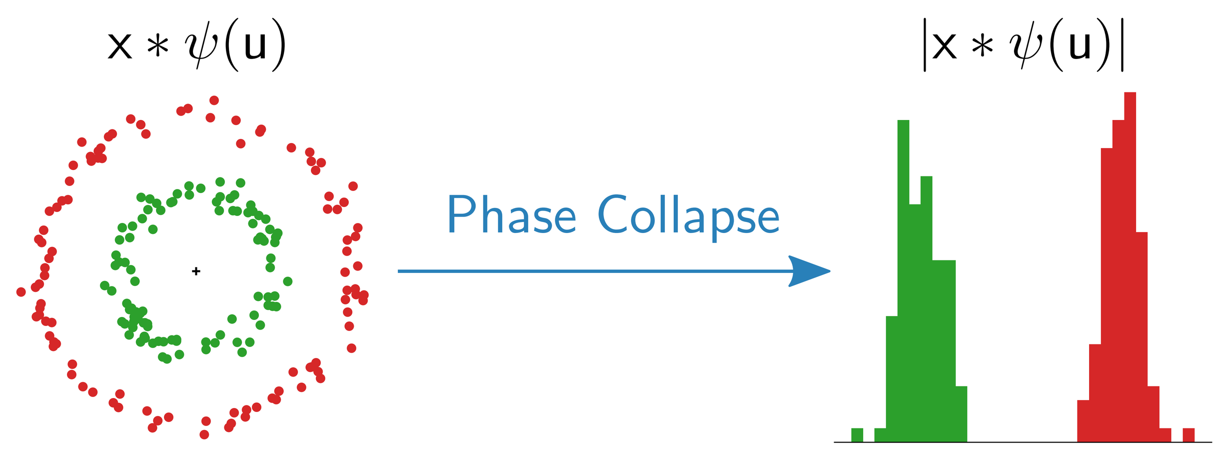

Phase collapse concentrates translation variability, which can separate class means.

Learned Scattering Network with Phase Collapses

We introduce an architecture which iterates **fixed** phase collapses $|W|$ and **learned** 1×1 convolutions $(P_j)_j$. There are **no biases**.

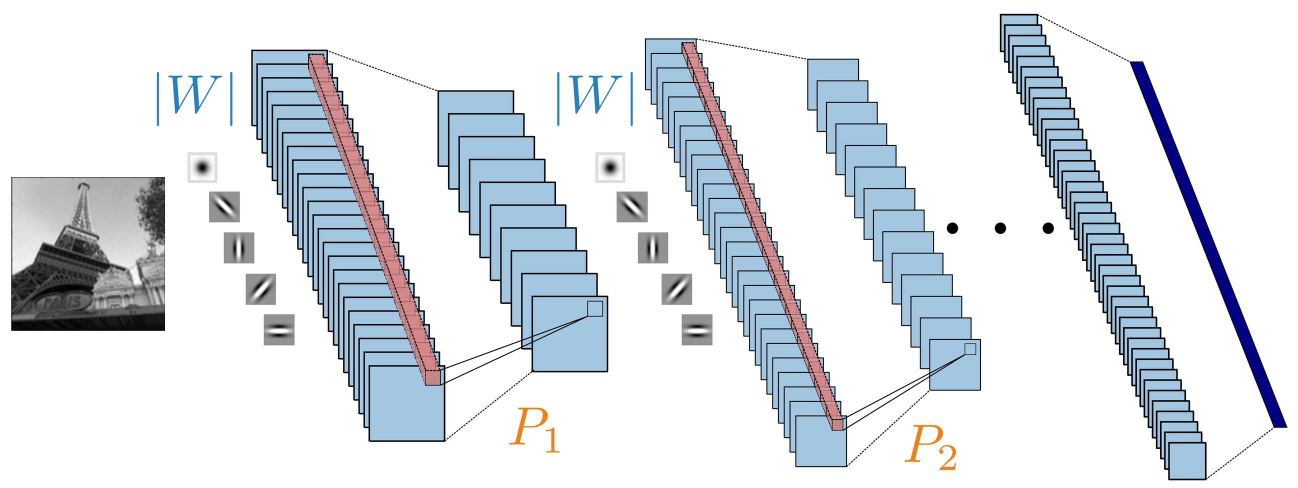

Skip-connections allow keeping some spatial information, which can be useful as input images are not fully stationary.

Reaches **ResNet-18 accuracy on ImageNet with only 12 layers**, largely outperforming the scattering baseline for which $(P_j)_j$ are predefined.

Error (%)		Scat	LScat	LScat + skip	ResNet
CIFAR-10		27.7	11.7	**7.7**	8.8
ImageNet	Top-5	54.1	15.2	11.0	**10.9**
	Top-1	73.0	35.9	**30.1**	30.2

ReLU and Phase Collapse

The ReLU can be decomposed into its even part, an absolute value with a dead-zone, and its odd part, a soft-thresholding:

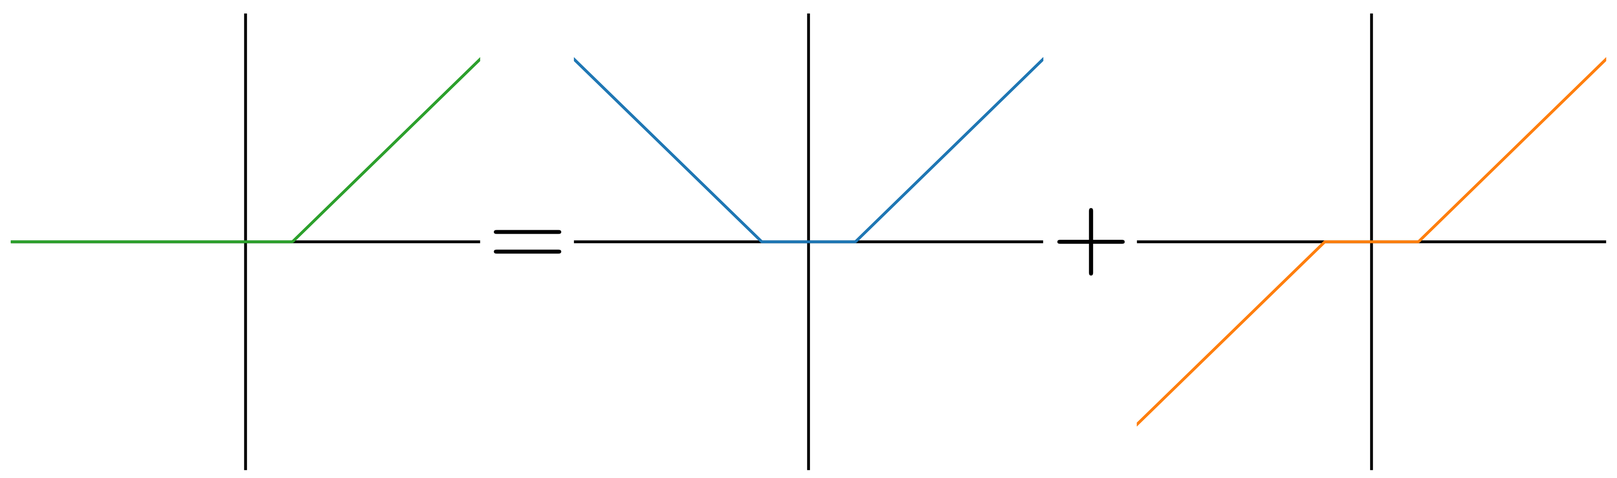

The absolute value collapses the sign (or phase over $\mathbb{C}$), which is in contrast preserved by a soft-thresholding. **Phase-preserving non-linearities cannot separate circularly-symmetric distributions, which keep zero means.**

Indeed, replacing the modulus with a complex soft-tresholding leads to a **three-fold increase in classification error**.

Error (%)	Scat	LScat + skip (cst)	LScat + skip (mod)	ResNet
CIFAR-10	27.7	22.5	**7.7**	8.8

Those results hold true when more general phase preserving are used. Similarly, replacing ReLUs with soft-thresholding within a standard ResNet architecture considerably degrades performances, while absolute values preserve high accuracies.

It is thus necessary for the non-linearity to act on the phase.

Conclusion

- ReLUs with biases can affect both the phase (absolute values) and the amplitude (soft-thresholdings) of network coefficients.
- Linear separation of classes results from acting on the phase rather than the amplitude.
- This can be constrained to collapsing the phase of wavelet coefficients.

⇒ **Phase collapses are both necessary and sufficient to linearly separate classes.**

Ongoing work: analyze the learned channel operators $(P_j)_j$ to design more efficient architectures with fewer parameters.

florentin.guth@ens.fr

john.zarka@ens.fr