



Neural Networks and Deep Learning: Deep Learning Theory

Nicolas Thome


Conservatoire National des Arts et Métiers (Cnam)
Département Informatique

Deep Learning Era




IMAGENET



- ▶ Deep Learning: huge impact in terms of experimental results



A group of people shopping at an outdoor market.
There are many vegetables at the fruit stand.



Understanding Deep Learning





IMAGENET



- ▶ **BUT: formal understanding still limited**
 - ▶ Optimization: non-convex problem
 - ▶ Model: ability to untangle manifold
 - ▶ Robustness to over-fitting & generalization
 - ▶ Stability, uncertainty estimate

Non-Convex Optimization

- ▶ One of the main historical shortcomings of deep neural networks
- ▶ In practice, not really an issue with modern neural networks, WHY?
- ▶ Some preliminary answer elements:
 - ▶ In high dimensional space random functions tend to have few local minima but many saddle points [Dauphin et al., 2014]
 - ▶ Empirically, gradient descent methods manage to escape [Goodfellow and Vinyals, 2015] saddle points




Result for shallow
linear autoencoder²

Non-Convex Optimization

- ▶ WHY non-convex optimization is not a major practical issue for deep learning?
- ▶ Some preliminary answer elements:
 - ▶ Most of local minima have about the same objective value [Haeffele and Vidal, 2015, Choromanska et al., 2014]

(Cartoon of
Dauphin et al 2014's
worldview)




Deep Learning and Generalization

- ▶ Rademacher complexity: capacity of a model to fit random label :

$$\mathcal{R}_n(\mathcal{H}) = E_{\sigma} \left[\sup_{h \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n \sigma_i h(x_i) \right]$$

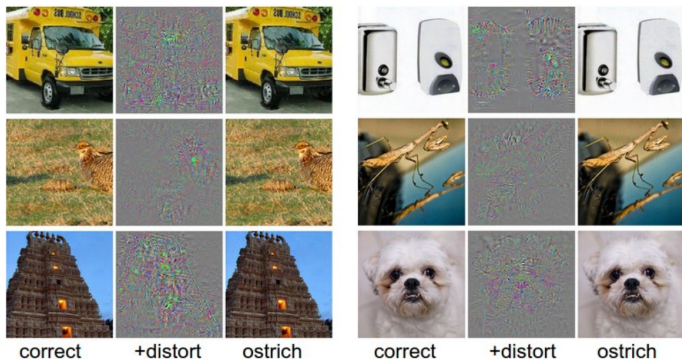
- ▶ Rethinking generalization [Zhang et al., 2017], ICLR



- ▶ Deep models easily fits random labels !!
- ▶ $\mathcal{R}_n(\mathcal{H}) \approx 1 \Rightarrow$ no theoretical guarantee on generalization performances
- ▶ Classical learning theory insufficient to explain the good generalization behavior of deep models


Deep Learning (DL) & Stability

- ▶ Deep Models not necessarily robust to input variations
- ▶ Deep Models do not naturally capture uncertainty
- ▶ Ex: Adversarial Examples



Deep Learning (DL) & Uncertainty: Problem

Softmax output in deep neural network \neq confidence (uncertainty) measure!



- ▶ Often wrong prediction \leftrightarrow unjustified high confidence
- ▶ Uncertainty however crucial in major applicative domains:
 - ▶ Healthcare
 - ▶ Autonomous driving
 - ▶ Nuclear

Formal theory explaining deep learning success: infancy

- ▶ **Optimization:** preliminary results for non-convex functions [Dauphin et al., 2014, Choromanska et al., 2014, Goodfellow and Vinyals, 2015, Haeffele and Vidal, 2015]
- ▶ **Regularization:** to be established
- ▶ **Stability:** studies under signal processing perspective [Bruna and Mallat, 2013], kernel methods [Bietti and Mairal, 2017]
- ▶ **Uncertainty:** preliminary connections between Bayesian models and dropout [Gal and Ghahramani, 2016]

TO BE CONTINUED ...

References I



Bietti, A. and Mairal, J. (2017).

Group invariance and stability to deformations of deep convolutional representations.
CoRR, [abs/1706.03078](#).



Bruna, J. and Mallat, S. (2013).

Invariant scattering convolution networks.
IEEE Trans. Pattern Anal. Mach. Intell., 35(8):1872–1886.



Choromanska, A., Henaff, M., Mathieu, M., Arous, G. B., and LeCun, Y. (2014).

The loss surface of multilayer networks.
CoRR, [abs/1412.0233](#).



Dauphin, Y., Pascanu, R., Gülçehre, Ç., Cho, K., Ganguli, S., and Bengio, Y. (2014).

Identifying and attacking the saddle point problem in high-dimensional non-convex optimization.
CoRR, [abs/1406.2572](#).



Gal, Y. and Ghahramani, Z. (2016).

Dropout as a bayesian approximation: Representing model uncertainty in deep learning.
In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML'16, pages 1050–1059. JMLR.org.



Goodfellow, I. J. and Vinyals, O. (2015).

Qualitatively characterizing neural network optimization problems.
In *ICLR*.



Haeffele, B. D. and Vidal, R. (2015).

Global optimality in tensor factorization, deep learning, and beyond.
CoRR, [abs/1506.07540](#).

References II



Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2017).
Understanding deep learning requires rethinking generalization.