Average-case bounds for neural network robustness and generalization using path-based representations

Co-supervisors: Rémi Gribonval, Inria Lyon (remi.gribonval@inria.fr)
Laurent Jacques, Inria Lyon, and UCLouvain, Belgium (laurent.jacques@inria.fr)
Antoine Gonon, EPFL (antoine.gonon@epfl.ch)

Location: ENS de Lyon, in the Ockham team.

Profile: We are looking for a highly motivated student with a strong background in mathematics (machine learning, linear algebra, probability and statistics) and taste for computer science and Python programming. If the candidate is successful, this internship may be pursued by a PhD.

Keywords: Mathematics of deep learning, Lipschitz bounds, generalization bounds, measure concentration, directed acyclic graph, path-lifting

Context: Establishing as tight as possible bounds on neural network regularity properties (e.g., Lipschitz bounds) and generalization capacity is a key endeavour to improve their robustness, ensure their trustworthiness, harness their statistical significance, or reduce their carbon footprint.

Standard Lipschitz bounds based on products of layer-based norms for classical multilayer perceptrons (MLP) have recently been significantly sharpened thanks to the so-called *path-lifting* framework (Gonon et al., 2025, 2024), leading at the same time to their extension to DAG[1] ReLU[2] neural networks, a wide family covering most of the complex deep architectures used in today's practice. This is particularly interesting since, despite being of combinatorial dimension, the path-lifting is endowed with efficient computational tools allowing, e.g., to compute its ℓ^p-norms (which not only provide Lipschitz bounds but also reveal for example whether some paths in the network's computational graph are somehow negligible with respect to some others) with a single pass on the network.

Even the sharpest path-lifting based bounds, which are computable on an arbitrary trained DAG ReLU network, are worst-case in nature and remain *overly pessimistic* by several orders of magnitude.

While exactly computing the Lipschitz constant of multilayer perceptrons (MLP) is known to be NP-hard (Virmaux and Scaman, 2018), recent work has shown the potential of *average-case analyses* and measure concentration tools to obtain nearly-tight estimates for networks with i.i.d. random weights, corresponding exactly to the networks used at the initialization of the training procedure. This is, however, currently limited to random *linear* networks associated to matrix products (Mourrat, 2025) or to random *multilayer perceptrons* (Dirksen et al., 2025), which are far from current practice involving much more complex DAG ReLU architectures.

Objective: The goal of the internship is to design computable average case bounds for random DAG ReLU neural networks. Using the simplest shallow network architectures as a first guiding example, the intern will progressively explore the theoretical potential of the path-lifting formalism, with the primary aim of replacing worst-case bounds (e.g., based on Hölder-type inequalities) with average case ones exploiting randomness with concentration arguments. From a computational perspective, it will be interesting to explore how to more explicitly reveal interesting properties of an arbitrary (given) network via the ℓ^p-norms (or other, to be designed, efficiently computable functions) of its path-lifted representation. From a theoretical perspective, a first focus will be on Lipschitz bounds for i.i.d. networks, but extensions to other quantities, possibly involving the training dynamics or the statistical properties of the network, may be explored as a follow-up.

[1] directed acyclic graph

[2] rectified linear unit – one of the most standard nonlinearities used in deep learning.

References

Antoine Gonon, Nicolas Brisebarre, Elisa Riccietti, and Rémi Gribonval. A Rescaling-Invariant Lipschitz Bound Based on Path-Metrics for Modern ReLU Network Parameterizations. In *Proceedings of the 42nd International Conference on Machine Learning (ICML 2025)*, Vancouver (BC), Canada, July 2025. URL https://hal.science/hal-04584311.

Antoine Gonon, Nicolas Brisebarre, Elisa Riccietti, and Rémi Gribonval. A path-norm toolkit for modern networks: consequences, promises and challenges. In B. Kim, Y. Yue, S. Chaudhuri, K. Fragkiadaki, M. Khan, and Y. Sun, editors, *International Conference on Representation Learning*, volume 2024, pages 41108–41152, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/file/b3732a13897c4cea145c3bdece80de64-Paper-Conference.pdf.

Aladin Virmaux and Kevin Scaman. Lipschitz regularity of deep neural networks: analysis and efficient estimation. *Advances in Neural Information Processing Systems*, 31, 2018.

Jean-Christophe Mourrat. Operator $\ell^\infty \to \ell^\infty$ norm of products of random matrices, February 2025. URL http://arxiv.org/abs/2502.06711. arXiv:2502.06711 [math].

Sjoerd Dirksen, Patrick Finke, Paul Geuchen, Dominik Stöger, and Felix Voigtlaender. Near-optimal estimates for the ℓ^p-Lipschitz constants of deep random ReLU neural networks, June 2025. URL http://arxiv.org/abs/2506.19695. arXiv:2506.19695 [stat].