# Algebro-geometric study of loss landscapes

Kathlén Kohn, Sandra Di Rocco

A fundamental goal in the theory of deep learning is to explain why the optimization of the loss function of a neural network does not seem to be affected by the presence of non-global local minima. Many papers have addressed this issue by studying the *static landscape* of the loss function [1, 4, 7, 11]. More recently, researchers have begun to search for explanations based on the *dynamics* of optimization [2, 9]. We believe however that the study of the *static* properties of the loss function, i.e. the structure of its critical locus, is not settled. Even in the case of linear networks, the existing literature paints a purely analytical picture of the loss, and provides no explanation as to *why* such architectures exhibit no bad local minima. A complete understanding of the critical locus should be a prerequisite for investigating the dynamics of the optimization.

The goal of this project is to use methods from algebraic geometry to investigate the static properties of the loss function of neural networks with different activation functions, such as linear, polynomial or ReLU activations. In these cases, the *neuromanifold* (the space of functions parameterized by a network with a fixed architecture) is a (semi-)algebraic set. For instance, in the case of linear activation functions, the neuromanifold is a determinantal variety, a classical object of study in algebraic geometry. Hence, in the above cases, the optimization of the loss function is an optimization over a (semi-)algebraic set. The number and nature of the critical points of such an optimization is governed by several concepts recently studied in algebraic geometry: *ED discriminants* [6], *bottlenecks* [5], *Voronoi varieties* [3]. Still, a full understanding of the critical locus of the loss function of neural networks using these algebraic notions is widely open. Partial results have been obtained for linear [10] and polynomial [8] activations. A possible unified framework to enhance our current understanding could be provided by introducing and analyzing *signature discriminants* as the algebraic locus where the signature of the critical points changes.

A candidate should have the necessary background in Algebraic Geometry, introductory ML theory and demonstrate computational skills (e.g., familiarity with Python or Macaulay2).

## References


[1] P. Baldi, K. Hornik: *Neural networks and principal component analysis: Learning from examples without local minima*, Neural Networks **2** (1989) 53–58

[2] L. Chizat, F. Bach: *On the Global Convergence of Gradient Descent for Over-parameterized Models using Optimal Transport*, arXiv: 1805.09545

[3] D. Cifuentes, K. Ranestad, B. Sturmfels, M. Weinstein: *Voronoi Cells of Varieties*, arXiv: 1811.08395

[4] A. Choromanska, M. Henaff, M. Mathieu, G. B. Arous, Y. LeCun: *The Loss Surfaces of Multilayer Networks*, Proceedings of the 18th International Conference on Artificial Intelligence and Statistics, AISTATS 2015

[5] S. Di Rocco, D. Eklund, M. Weinstein: *The bottleneck degree of algebraic varieties*, arXiv: 1904.04502

[6] J. Draisma, E. Horobet, G. Ottaviani, B. Sturmfels, R. Thomas: *The Euclidean distance degree of an algebraic variety*, Foundations of Computational Mathematics **16** (2016) 99–149

[7] K. Kawaguchi: *Deep Learning without Poor Local Minima*, CoRR, arXiv: 1605.07110

[8] J. Kileel, M. Trager, J. Bruna: *On the Expressive Power of Deep Polynomial Neural Networks*, arXiv: 1905.12207

[9] S. Mei, A. Montanari, P.-M. Nguyen: *A Mean Field View of the Landscape of Two-Layers Neural Networks*, arXiv: 1804.06561

[10] M. Trager, K. Kohn, J. Bruna: *Pure and Spurious Critical Points: a Geometric Study of Linear Networks*, arXiv: 1910.01671

[11] L. Venturi, A. S. Bandeira, J. Bruna: *Spurious valleys in two-layer neural network optimization landscapes*, arXiv: 1802.06384