Multifractal Structure of Neural Network Loss Landscapes

Gergő Török, Ferenc A. Bartha
University of Szeged, Szeged, Hungary

Understanding the geometry of neural-network loss landscapes remains a central problem in the mathematical study of deep learning. Standard diagnostics such as gradients, Hessians, and sharpness describe local behavior, but practical training is inherently finite-scale: stochastic optimization moves through parameter space using non-negligible updates, and the loss may respond differently across directions and perturbation scales. This work develops a finite-scale multifractal framework for studying this geometry along neural-network training trajectories.

The main contribution is a directional probing method that converts second-order finite-scale loss responses into scale-dependent empirical measures over parameter-space directions. At each checkpoint, the loss is evaluated under symmetric perturbations in randomly sampled, tensor-normalized directions. After normalization across directions, the resulting response distribution is analyzed using moment-based multifractal quantities, including empirical partition functions, mass-exponent profiles $\widehat{\tau}(q)$, Chhabra–Jensen-type descriptors, per-direction scaling exponents, and the scalar moment-scaling variability measure M_τ. These quantities provide finite-scale diagnostics of anisotropy and heterogeneous scaling, rather than assuming an exact asymptotic fractal structure.

The framework is applied to a ResNet-18 trained on CIFAR-10. Across training checkpoints, the measured multiscale geometry changes systematically: moment-scaling variability decreases, directional scaling exponents become more concentrated, and the estimated multifractal profiles become less dispersed. These trends occur together with improvements in training performance and are compared with curvature-based Edge-of-Stability diagnostics. The results suggest that successful training is accompanied not only by descent in the empirical loss, but also by a measurable regularization of finite-scale directional geometry.

Overall, this work provides a quantitative way to study loss landscapes beyond purely infinitesimal curvature. By combining finite-scale directional perturbations with multifractal moment analysis, it offers a compact description of how anisotropy, scale dependence, and stability-related curvature evolve during deep-network training. The proposed diagnostics may be useful for comparing optimization schedules, identifying geometric changes along training trajectories, and developing more refined mathematical descriptions of deep-learning dynamics.