Neural Networks Loss Landscape Convergence in Different Low-Dimensional Spaces

Tem Nikitin Nikita Kiselev

Moscow Institute of Physics and Technology

April 1, 2025

Problem Statement

- Neural networks achieve higher accuracy with larger datasets.
- However, computational resources limit training possibilities.
- Key question: **When does additional data stop reshaping the loss landscape?**

Proposed Method

- Analyze loss landscape via Hessian eigenvectors.
- Project high-dimensional parameter space onto lower-dimensional space.
- Utilize Monte-Carlo methods to approximate landscape changes.

Main Idea

- **Hessian-based projections** identify primary curvature directions.
- Calculate a Δ-function to quantify changes when increasing dataset size:

$$\Delta_k = \int (L_k(w) - L_{k-1}(w))^2 p(w) dw$$

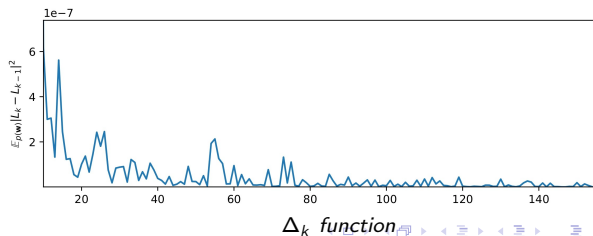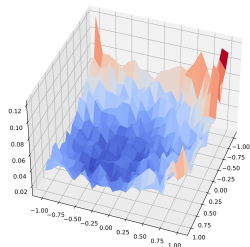- Identify the threshold at which further data does not significantly alter landscape.

Neural Networks Loss Landscape Convergence in Different Low-Dimensional Spaces

Goal: Measure how the loss function changes as the training set size grows:

$$\Delta_k = \mathbb{E}_{p(\mathbf{w})}\Big(\mathcal{L}_k(\mathbf{w}) - \mathcal{L}_{k-1}(\mathbf{w})\Big)^2.$$

Method:

- **Monte Carlo:** Generate points near the minimum according to $p(\mathbf{w})$ and average the differences.
- **Hessian Eigenvectors:** Use directions with the largest eigenvalues to focus on key curvature components.

Δ_k function

Experiment Results

- Experiments conducted on MNIST and Fashion-MNIST datasets.
- Low-dimensional visualization clearly illustrates stabilization.
- Verified theoretical bounds empirically.

Practical Implications

- Enables identification of minimal viable dataset size.
- Reduces computational cost significantly.
- Provides guidelines for efficient data collection.

References

Soydaner, D. *Optimization algorithms for deep learning*, 2020.

Wu et al. *Loss landscape and generalization*, 2017.

Petersen, Pedersen. *The matrix cookbook*, 2012.

Deng, L. *MNIST database*, 2012.

Xiao et al. *Fashion-MNIST*, 2017.

Kiselev, Grabovoy. *Unraveling the Hessian*, 2024.

Sagun et al. *Hessian analysis of networks*, 2018.