Contribution ID: **46** Type: **not specified**

Neural Networks Loss Landscape Convergence in Hessian Low-Dimensional Space

Tuesday, 20 May 2025 13:52 (12 minutes)

Understanding how a neural network's loss landscape changes as we add more training data is important for efficient training. Although larger datasets reshape this high-dimensional surface, the point when extra data stop making a big difference is unclear. We show that near a local minimum the loss landscape stabilizes once the dataset exceeds a certain size. To study this, we project the full parameter space onto a lower-dimensional subspace formed by the Hessian's top eigenvectors, highlighting the most important curvature directions. Within this subspace, we use Monte Carlo sampling to estimate how the loss changes more precisely. Experiments on standard image classification tasks demonstrate that our low-dimensional analysis pinpoints when the landscape stops evolving, offering practical guidance for balancing training cost with performance improvements.

Primary author: НИКИТИН, Артем

Presenter: НИКИТИН, Артем

Session Classification: 20-Машинное обучение и нейросети

Track Classification: Машинное обучение и нейросети