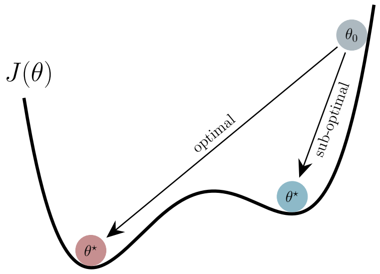

Project Thesis

Visual analysis of machine learning convergence

Supervised machine learning (ML) models typically have millions of parameter θ_i which are adjusted during training to minimize the so-called loss function $J(\boldsymbol{\theta})$[1]. The loss function quantifies the error between prediction $\hat{\mathbf{y}}(\boldsymbol{\theta})$ and target $\mathbf{y}$, effectively expressing high-dimensional information by a scalar value. Visually speaking, the training of ML models can be interpreted as the search for a minimum in the high-dimensional surface $J(\boldsymbol{\theta})$, often referred to as the *loss landscape*[2]. The shape of $J(\boldsymbol{\theta})$ determines the achieveable ML performance: convex curvature towards a minimum favors the training process while many non-optimal local minima may trap the model in non-optimal configuration.

This thesis aims at visualizing the loss surface for a given ML task under consideration of different loss functions, thus gaining deep insights into the dynamics of ML model training, ultimately answering the question, why some model configurations perform better than others. Visualizations in 3D can be achieved by dimensionality reduction methods[3] and provide qualitative understanding of the convergence behaviour. A collection of interactive visualizations using this approach can be found at https://losslandscape.com/.

The scope of this work covers the following tasks:

- Literature review of state of the art,
- Implementation of dimensionality reduction method presented in Li et al. (2018),
- Comparison of different model set-ups by means of visualized loss surface

Prerequisites:

- Demonstrated programming experience in Python and TensorFlow or PyTorch,
- Curiosity, excellent skills in independent work and communication.

Contact: M.Sc. Mathies Wedler (mathies.wedler@tuhh.de)

[1] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *nature*, *521*(7553), 436.

[2] Li, H., Xu, Z., Taylor, G., Studer, C., & Goldstein, T. (2018). Visualizing the loss landscape of neural nets, In *Proceedings of the 32nd international conference on neural information processing systems*, Montréal, Canada, Curran Associates Inc.

[3] Li et al., 2018.

Dynamics Group (M-14)
Hamburg University of Technology
www.tuhh.de/dyn