

A Renormalization Group Approach To Universality of Learning in Deep Neural Networks

G. Peraza Coppola^{(1), (2)}, M. Helias^{(1), (3)}, and Z. Ringel⁽⁴⁾

⁽¹⁾ Institute for Advanced Simulation (IAS-6), Computational and Systems Neuroscience, Jülich Research Centre, Jülich, Germany ⁽²⁾ RWTH Aachen University, Aachen, Germany ⁽³⁾ Department of Physics, RWTH Aachen University, Aachen, Germany ⁽⁴⁾ The Recah Institute of Physics, The Hebrew University of Jerusalem, Jerusalem, Israel

The search for universality—the idea that diverse systems can exhibit similar large-scale behavior independent of microscopic details—has long guided theoretical physics. Recently, this notion has begun to influence machine learning, where researchers seek unifying principles that describe how complex models behave across architectures, datasets, and training regimes. Just as physical systems near criticality that display universal scaling laws, learning systems may exhibit analogous patterns that transcend their implementation details. Such systems are often characterized by self-similarity, meaning that their structure or behavior looks similar across different scales or resolutions. They often display power-law correlations and universal behavior, making them amenable to analysis through the renormalization group framework.




Figure 1: Predicted neural scaling laws by Renormalization Group approach. Expected training loss per sample $\langle \mathcal{L} \rangle / P$ is plotted as a function of the training dataset size P . Dashed-black line represents the prediction for the Gaussian Process (no perturbation). Colored lines represent the prediction for systems with different perturbation strengths (feature learning regime).

Interestingly, power-law statistics and weaker forms of universality also appear in empirical datasets and modern deep learning models, suggesting that renormalization concepts may provide valuable insights into their structure and dynamics. We provide a theoretical framework based on the renormalization group to study self-similarity and its breakdown in the learning curves of a class of weakly non-linear (non-lazy) neural networks trained on data with power-law spectra. We demonstrate that features typically overlooked in conventional analyses—such as discrete spectral structure and the absence of translation invariance—lead to significant deviations from standard renormalization group behavior, both quantitatively and qualitatively. Notably, we show that the traditional notion of scaling dimensions is naturally replaced by that of scaling intervals. Despite these modifications, our approach preserves key aspects of the renormalization group formalism: it allows for the classification of perturbations into relevant and irrelevant directions and reveals a form of asymptotic universality at large data scales, governed by a Gaussian-Process-like ultraviolet fixed point.