Multi-Objective Loss Balancing for Physics-Informed Deep Learning

Kraus Michael Anton and Rafael Bischof, ETH Zurich

Physics Informed Neural Networks (PINN) are algorithms from deep learning leveraging physical laws by including partial differential equations (PDE) together with a respective set of boundary and initial conditions (BC / IC) as penalty terms into their loss function. As the PDE, BC and IC loss function parts can significantly differ in magnitudes, due to their underlying physical units or stochasticity of initialisation, training of PINNs may suffer from severe convergence and efficiency problems, causing PINNs to stay beyond desirable approximation quality. In this work, we observe the significant role of correctly weighting the combination of multiple competitive loss functions for training PINNs effectively. To that end, we implement and evaluate different methods aiming at balancing the contributions of multiple terms of the PINNs loss function and their gradients. After review of three existing loss scaling approaches (Learning Rate Annealing, GradNorm as well as SoftAdapt), we propose a novel self-adaptive loss balancing of PINNs called ReLoBRaLo (Relative Loss Balancing with Random Lookback). Finally, the performance of ReLoBRaLo is compared and verified against these approaches by solving both forward as well as inverse problems on three benchmark PDEs for PINNs: Burgers' equation, Kirchhoff's plate bending equation and Helmholtz's equation. Our simulation studies show that ReLoBRaLo training is much faster and achieves higher accuracy than training PINNs with other balancing methods and hence is very effective and increases the sustainability of PINNs algorithms. The adaptability of ReLoBRaLo illustrates robustness across different PDE problem settings. The proposed method can also be employed to the wider class of penalized optimization problems, including PDE-constrained and Sobolev training apart from the studied PINNs examples.