Explorations of Constrained Neural Networks

Neural networks are typically trained to minimize a loss function, but many applications require that the network's outputs satisfy additional constraints, such as bounds, smoothness, or physical laws. In this work, we explore how neural networks behave when we add soft constraints to the loss function. We begin with a simple regression task, training a neural network built on known mathematical functions (e.g., $y = x^2 - 3, y = cos(x)$) and apply constraints. We observe that constraint penalties can meaningfully improve the neural network's accuracy and that the choice we make on activation functions plays a critical role. We then turn to explore Physics-Informed Neural Networks (PINNs), which use PDE residuals as the constraints. We use a 1D heat equation to find out where and how much more the physics being enforced, using the collocation points, matter for accuracy, than the amount of data used for training. These observations serve as early steps toward understanding how to design constraints that make neural networks more reliable.

Yahriel Guel and Martine Ceberio
University of Texas at El Paso, 500 W. University
El Paso, TX 79968, USA,
e-mail: yiguel@miners.utep.edu