

Tuesday, 14 December 2021, 2 pm

ZOOM Meeting

CoCoNUT TALK

Andrew Brock

Edinburgh Centre for Robotics, Scotland

Understanding when Deep Nets are trainable: Busting Batchnorm, Clipping Gradients, and Plotting Everything.

Under what conditions will a neural network train, and under what conditions will it train well? Years of experimentation have led the community to develop a fairly robust recipe book for training deep nets on common tasks, and a series of slightly delayed efforts have built up a reasonably deep understanding of the mechanisms underlying the success of these techniques. In this talk, I'll discuss our recent work on understanding and improving signal propagation in deep neural networks, with a focus on the process by which one might discover and visualize quantities of interest, use that knowledge to ground the development of new techniques in empirical understanding, and maybe land an ImageNet SOTA or two in the process.