

Gradients in Deep Neural Networks at Initialization

Boris Hanin (Texas A&M University)

Time **Thursday, April 19, 2018 at 14:00**

Place **Campus Belval, room MNO 1.050**

Neural networks are finite (but typically high) dimensional families of non-linear functions whose practical uses and theoretical analysis is a field now called deep learning. I will give a brief introduction to deep learning and will focus on some probabilistic questions about the gradients of the function computed by neural networks at initialization (i.e. when the weights and biases are random).