

CRUNCH Seminars at Brown, Division of Applied Mathematics

Friday – March 20, 2020

DEEP VARIATIONAL INFORMATION BOTTLENECK

Liu Yang

We present a variational approximation to the information bottleneck of Tishby et al. (1999). This variational approach allows us to parameterize the information bottleneck model using a neural network and leverage the reparameterization trick for efficient training. We call this method “Deep Variational Information Bottleneck”, or Deep VIB. We show that models trained with the VIB objective outperform those that are trained with other forms of regularization, in terms of generalization performance and robustness to adversarial attack.