NICOLAS LE ROUX, Université de Montréal, Québec, Canada
Representational Power of Restricted Boltzmann Machines and Deep Belief Networks

Deep Belief Networks (DBN) are generative neural network models with many layers of hidden causal variables, recently introduced by Hinton et al., along with a greedy layer-wise unsupervised learning algorithm. The building block of a DBN is a probabilistic model called a Restricted Boltzmann Machine (RBM), used to represent one layer of the model. Restricted Boltzmann Machines are interesting because inference is easy in them, and because they have been successfully used as building blocks for training deeper models.

We show that RBMs are universal approximators of discrete distributions. A first theorem shows that adding hidden units yields improved modeling power, while a second theorem shows that an RBM can model any discrete distribution.

We then study the question of whether DBNs with more layers are strictly more powerful in terms of representational power. This suggests another criterion for DBNs, obtained by considering that the top layer can perfectly fit its input.