Contribution ID: **1**   Type: **not specified**

Dropout: A Simple Way to Prevent Neural Networks from

Deep neural nets with a large number of parameters are very powerful machine learning systems. However, overtting is a serious problem in such networks. Large networks are also slow to use, making it dicult to deal with overtting by combining the predictions of many dierent large neural nets at test time. Dropout is a technique for addressing this problem. The key idea is to randomly drop units (along with their connections) from the neural network during training. This prevents units from co-adapting too much. During training, dropout samples from an exponential number of dierent \thinned" networks. At test time, it is easy to approximate the eect of averaging the predictions of all these thinned networks by simply using a single unthinned network that has smaller weights. This signicantly reduces overtting and gives major improvements over other regularization methods. We show that dropout improves the performance of neural networks on supervised learning tasks in vision, speech recognition, document classication and computational biology, obtaining state-of-the-art results on many benchmark data sets.

Primary authors: Mr HINTON, Georey (University of Toronto); Mr SUTSKEVER, Iyer

Presenter: Mr HINTON, Georey (University of Toronto)

Track Classification: Machine Learning