

Neural networks

Deep learning - deep autoencoder

DEEP AUTOENCODER

Topics: deep autoencoder

- Pre-training can be used to initialize a deep autoencoder

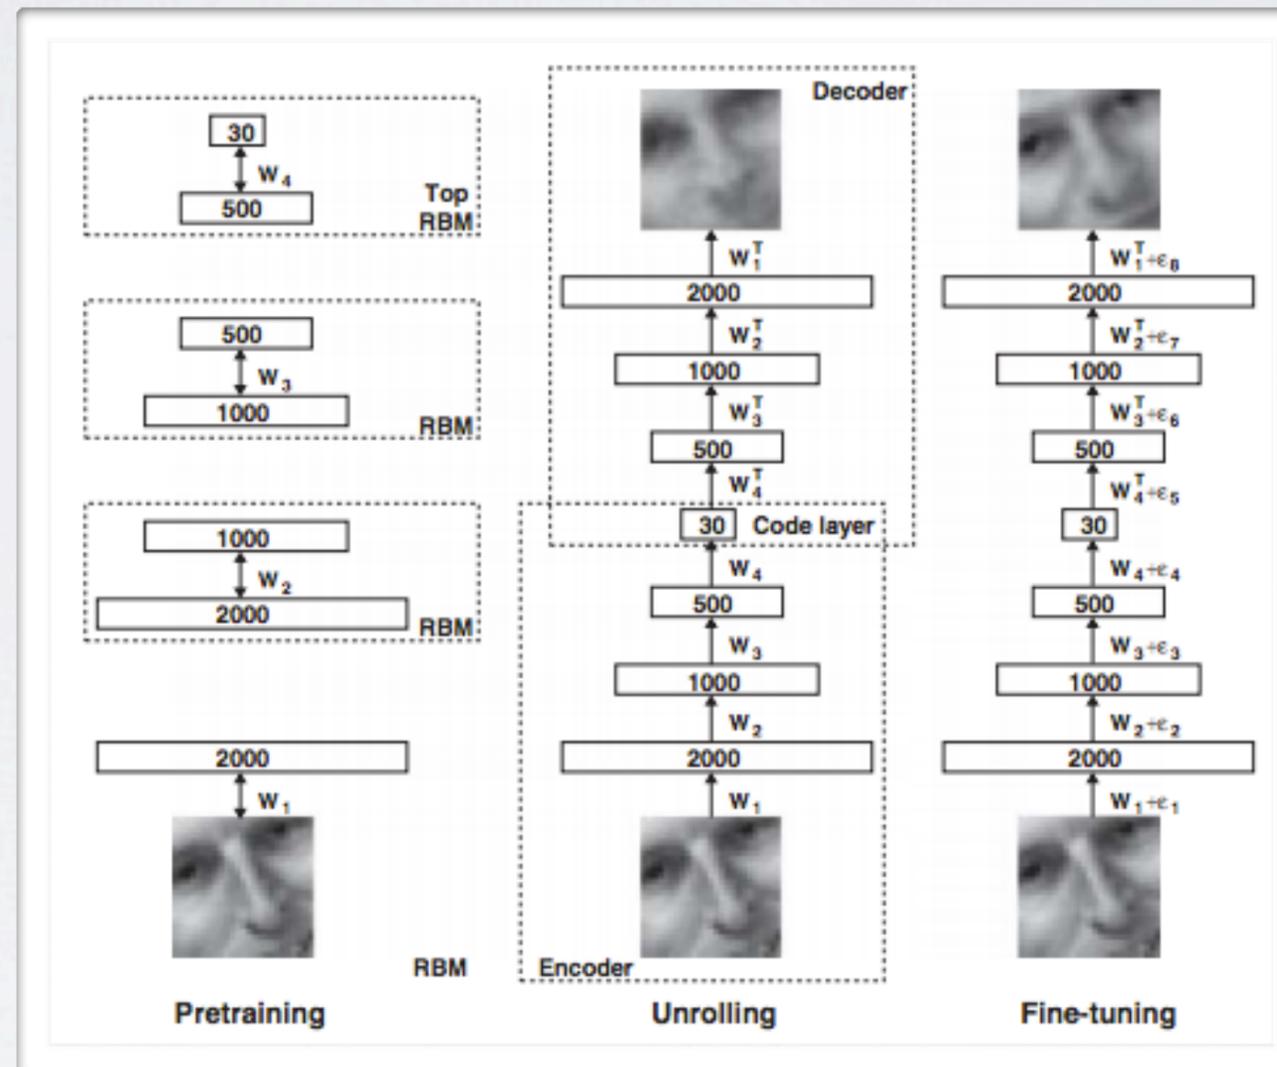
- ▶ This is an example of a situation where underfitting is an issue

- perhaps surprisingly, pre-training initializes the optimization problem in a region with better local optima of training objective

- ▶ Each RBM used to initialize parameters both in encoder and decoder (“unrolling”)

- ▶ Better optimization algorithms can also help

- Deep learning via Hessian-free optimization. James Martens, 2010



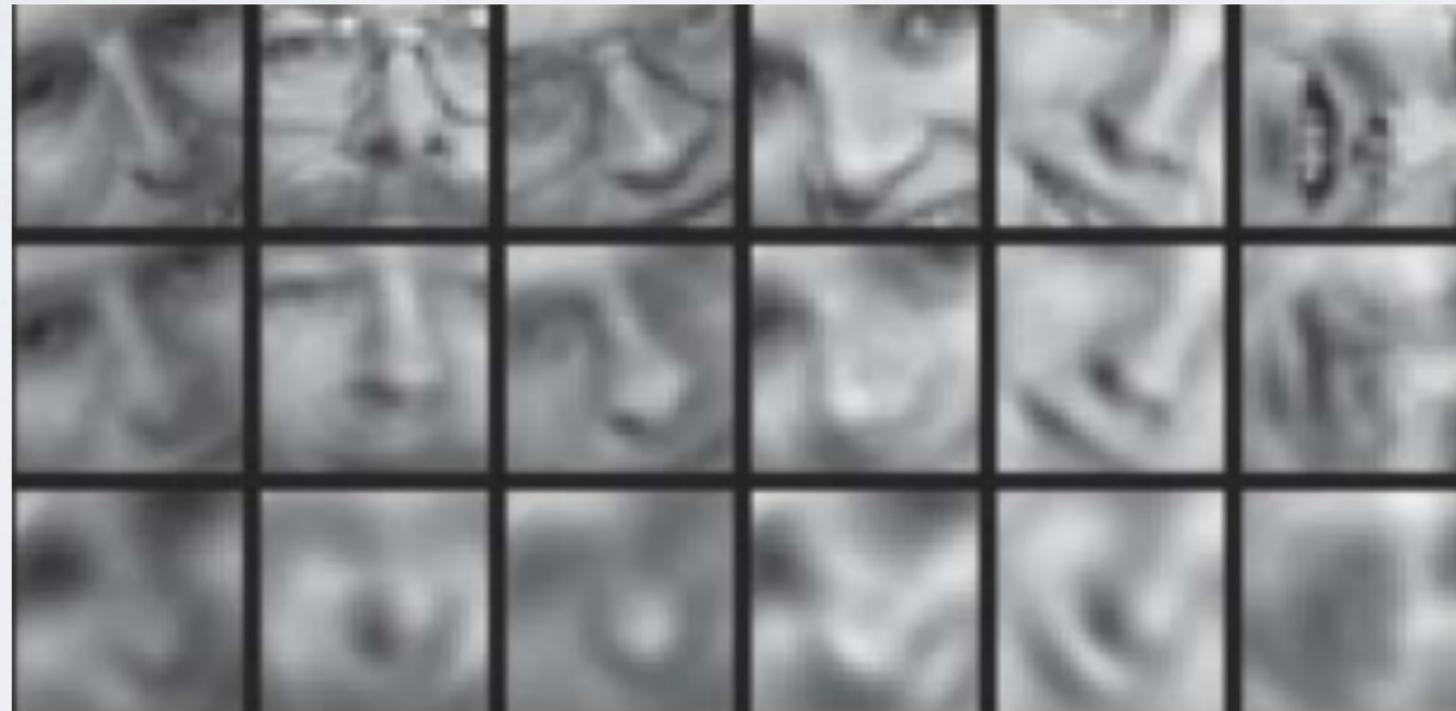
From Hinton and Salakhutdinov, Science, 2006

DEEP AUTOENCODER

Topics: deep autoencoder

- Can be used to reduce the dimensionality of the data
 - will have better reconstruction than a single layer network (i.e. PCA)

Original data



Deep autoencoder
reconstruction

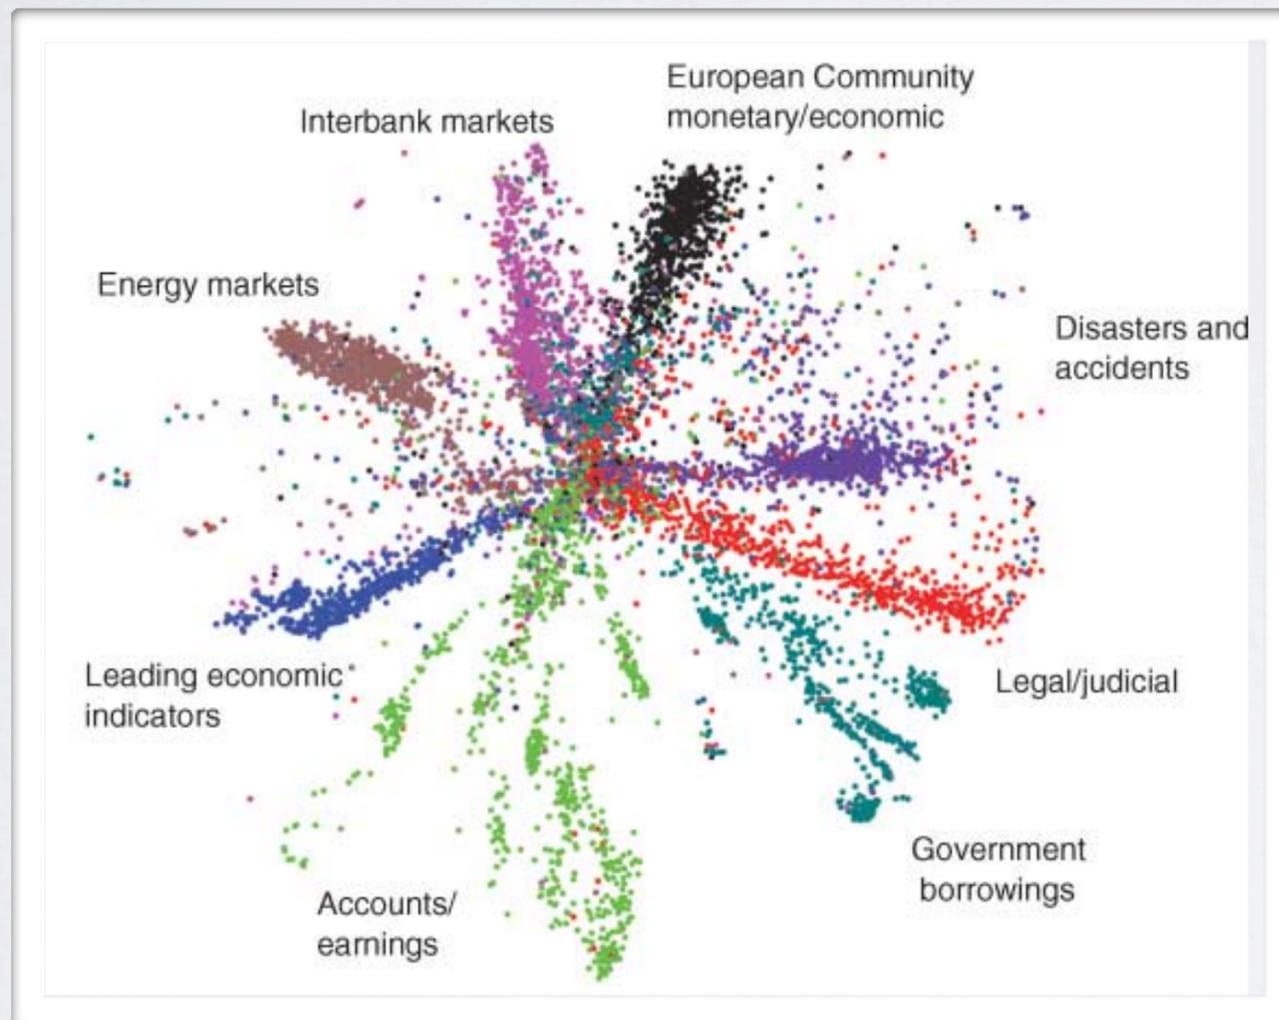
PCA reconstruction

From Hinton and Salakhutdinov, Science, 2006

DEEP AUTOENCODER

Topics: deep autoencoder

- If we reduce to 2D, we can visualize the data (e.g. a collection of document)



From Hinton and Salakhutdinov, Science, 2006