Learning how the mind learns

- Sanjit Singh Batra

- Diagnosis of diseases like malaria, osteoporosis etc. with common physical and physiological symptoms is well known.
- Consider ADHD or Alzheimer's disease. There are no reliable methods to diagnose these diseases till date.
- Why not? Because through blood tests we understand what happens to the liver in the case of typhoid. But here, here we don't have enough understanding of the anatomy of the brain to figure out what “changes” in the case of Alzheimer's.

Can we do something better?

- Yes!

The Idea!

- Given the data of so many diagnosed patients, and the increasing feasibility of methods that help us “peek” into the brain, like fMRI(functional Magnetic Resonance Imaging) and such, we can build classifiers not only to diagnose such diseases, but also to predict their onset.

The Experiment

- Data from patients diagnosed with Alzheimer's and ADHD were considered in two separate cases.
- A Deep Neural Network were trained with the extent to which they had the disease as:-

  0(No disease), 0.5(suspected) or 1(confirmed).

- The network had 30 neurons in 2 hidden layers, 50 neurons in the top-level layer, and 3 labels. Input vectors were 60-D after PCA. The 'structure' was :

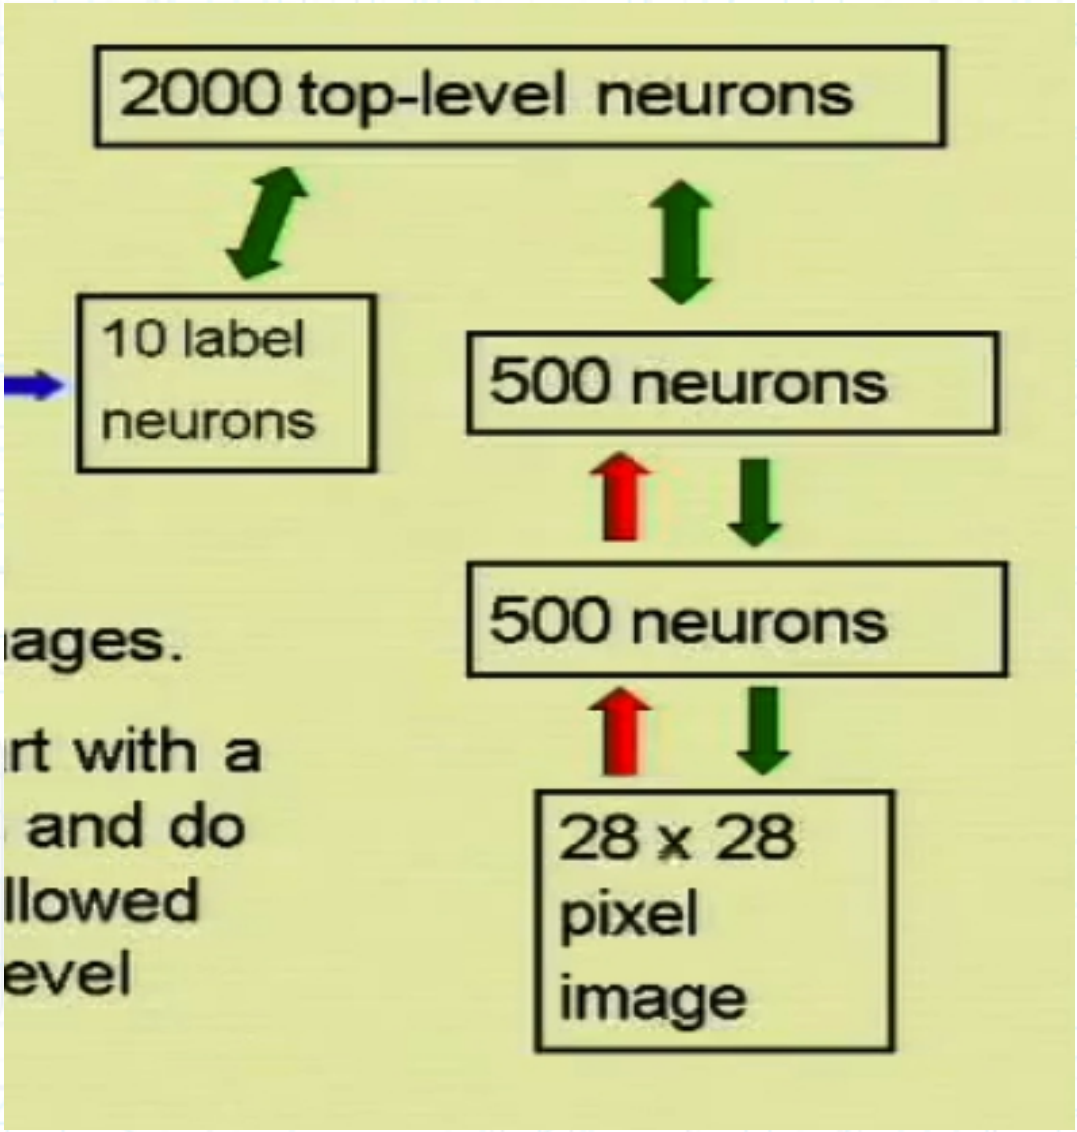

Results

Dataset:	SVM with PCA	SVM with PCA using Kernel Optimization	Deep Neural Networks
Alzheimer's	64.1%	67.9%	78.71%
ADHD	68.2%	69.1%	79.39%

Discussion

- SVMs maximize minimum error. What Deep Nets do is quite different. As the network learns its weights, there is a notion of energy of the system which is being minized. Ravines are carved in this landscape, one corresponding to each class. So, despite large variability in fMRI scans even for the same patient,

  deeper ravines means better classifiers.

A little adventure into 'Mind Reading'

- A simple idea, not so simple in the recent past!
- Apply unsupervised learning methods to fMRI scans of individuals as they observe objects in a highly “controlled” environment.
- Now, when they think, their thoughts can be visually reconstructed! Seems like magic...!

- Sample Experiment: Deep Neural Networks have a special ability of being able to “hallucinate”.

  To do this we essentially perform an upward pass and then a downward pass in the network.

  We train on a set of images of **cars**(same orientation) and try to see that once trained, what is the reconstruction of a different image of the object. Here are some very interesting pictures:

Results

- What was being reconstructed:
- 


- Reconstruction or “hallucination”
- 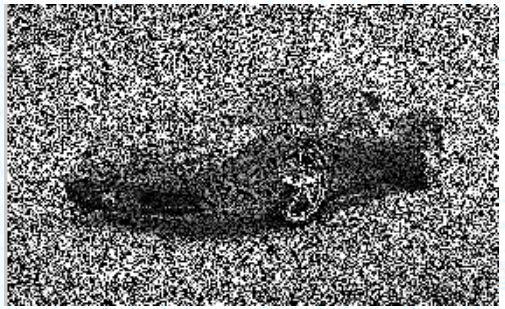
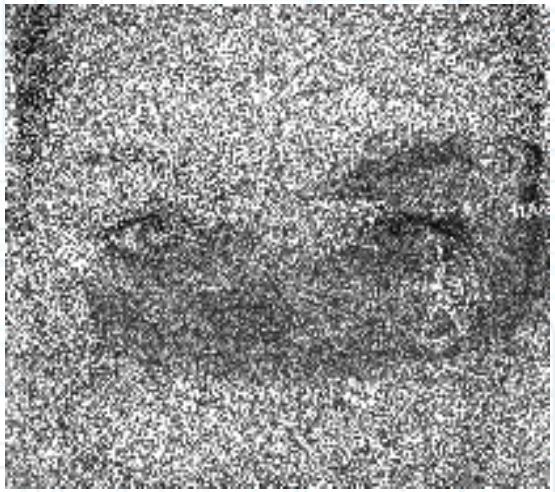

*Not seen while training.

Le Fin.