

Deep learning with RBM (\equiv Deep Belief Networks)

→ good to extract features
 - Learn multiple layers of representation to make sense of images, sound, text


Ex. Handwritten digit recognition

1 2 3 4 5 6 7 8 9 0

↳ straight lines

↳ curves & lines

Dataset: 3,000 images, each 50×50 pixels \Rightarrow 2,500 pixels per img.



input (pixels)

h_1

h_2


h_3 (or output)

Phase I
 (Model this \equiv find \vec{w}_1)

Phase II

Phase III: Simple FFNN with backprop. (supervised learning because you know the labels)

! All phases are done indep.

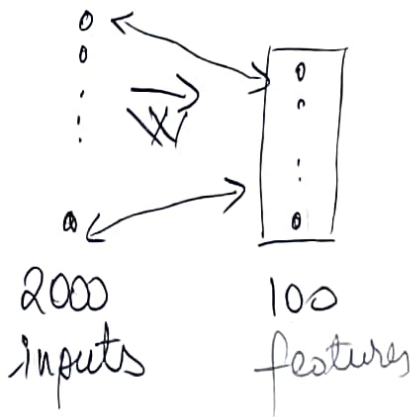
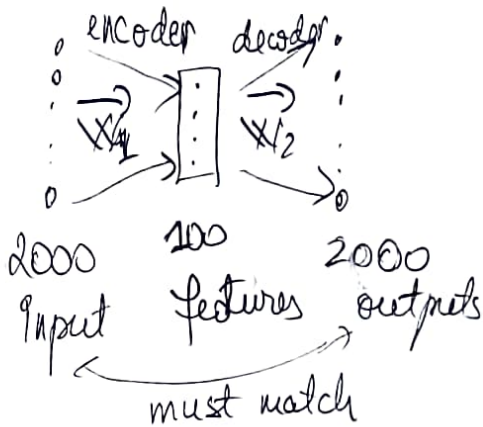


Feature extraction

Encoder-decoder vs. RBM

you learn two sets of weights

you learn only one set of weights



Ⓛ A NN with more than three layers is a deep net.

Ⓛ ≠ 12 layers, or even 1,000 layers

Ⓛ Deep learning - extracts features + finds patterns

