


# Neural networks

Restricted Boltzmann machine - example

# RESTRICTED BOLTZMANN MACHINE

**Topics:** RBM, visible layer, hidden layer, energy function



Energy function: 
$$E(\mathbf{x}, \mathbf{h}) = -\mathbf{h}^\top \mathbf{W} \mathbf{x} - \mathbf{c}^\top \mathbf{x} - \mathbf{b}^\top \mathbf{h}$$

$$= -\sum_j \sum_k W_{j,k} h_j x_k - \sum_k c_k x_k - \sum_j b_j h_j$$

Distribution: 
$$p(\mathbf{x}, \mathbf{h}) = \exp(-E(\mathbf{x}, \mathbf{h})) / Z$$


← partition function  
(intractable)

# EXAMPLE OF DATA SET: MNIST

3	8	6	9	6	4	5	3	8	4	5	2	3	8	4	8
1	5	0	5	9	7	4	1	0	3	0	6	2	9	9	4
1	3	6	8	0	7	7	6	8	9	0	3	8	3	7	7
8	4	4	1	2	9	8	1	1	0	6	6	5	0	1	1
7	2	7	3	1	4	0	5	0	6	8	7	6	8	9	9
4	0	6	1	9	2	2	3	9	4	4	5	6	6	1	7
2	8	6	9	7	0	9	1	6	2	8	3	6	4	9	5
8	6	8	7	8	8	6	9	1	7	6	0	9	6	7	0

# FILTERS

(LAROCHELLE ET AL., JMLR2009)



# DEBUGGING

**Topics:** stochastic reconstruction, filters

- Unfortunately, we can't debug with a comparison with finite difference
- We instead rely on approximate “tricks”
  - ▶ we plot the average stochastic reconstruction  $\|\mathbf{x}^{(t)} - \tilde{\mathbf{x}}\|^2$  and see if it tends to decrease:
  - ▶ for inputs that correspond to image, we visualize the connection coming into each hidden unit as if it was an image
    - gives an idea of the type of visual feature each hidden unit detects
  - ▶ we can also try to approximate the partition function  $Z$  and see whether the (approximated) NLL decreases
    - On the Quantitative Analysis of Deep Belief Networks.  
Ruslan Salakhutdinov and Iain Murray, 2008