

Neural Networks part deux

Paul Gribble


Nov 15, 2012

backpropagation embellishments

- learning rate parameter
- momentum parameter
- local minima
- cross-validation
- online vs batch vs mini-batch learning
- weight decay


updating weights

steepest descent



the delta rule

conjugate gradient descent



fancy math

interpreting network weights

- input (many) -> hidden(few) -> output (fewer)
- hidden units represent input in a lower dimensional space
- examine hidden unit activity for presentation of specific (e.g. pared down) inputs
- visualize “receptive fields” of hidden units as weighting of their connections to inputs (what inputs do each hidden unit “care about”?)

Coursera (free) online courses

Neural Networks for Machine Learning

Geoff Hinton (Univ Toronto)

<https://www.coursera.org/course/neuralnets>

Machine Learning

Andrew Ng (Stanford Univ)

<https://www.coursera.org/course/ml>

sample code

XOR

- xor_aima.py
- xor.py
- xor_cg.py

try: learn the AND mapping

MNIST

- mnist_pybrain.py
- mnist_cg.py

try: change # of hidden units and explore how this affects network performance and generalization

assignment

- assignment6.py
- maybe too challenging?
- we will look at more code next week