An Introduction to Machine Learning

Fabio A. González Ph.D.

Depto. de Ing. de Sistemas e Industrial
Universidad Nacional de Colombia, Bogotá

February 5, 2008

Content

1. Patterns and Generalization
   - Generalizing from patterns
   - Overfitting/ Overlearning
2. Learning Problems
   - Supervised
   - Non-supervised
   - Active
   - On-line
3. Learning Techniques

What is a pattern?

- Data regularities
- Data relationships
- Redundancy
- Generative model

Learning a Boolean function

x_1	x_2	f_1	f_2	...	f_{16}
0	0	0	0	...	1
0	1	0	0	...	1
1	0	0	0	...	1
1	1	0	1	...	1

- How many Boolean functions of n variables are?
- How many candidate functions are removed by a sample?
- Is it possible to generalize?

Inductive bias

- In general, the learning problem is *ill-posed* (more than one possible solution for the same particular problem, solution sensitive to small changes on the problem)
- It is necessary to make additional assumptions about the kind of pattern that we want to learn
- **Hypothesis space**: set of valid patterns that can be learnt by the algorithm

What is a good pattern?

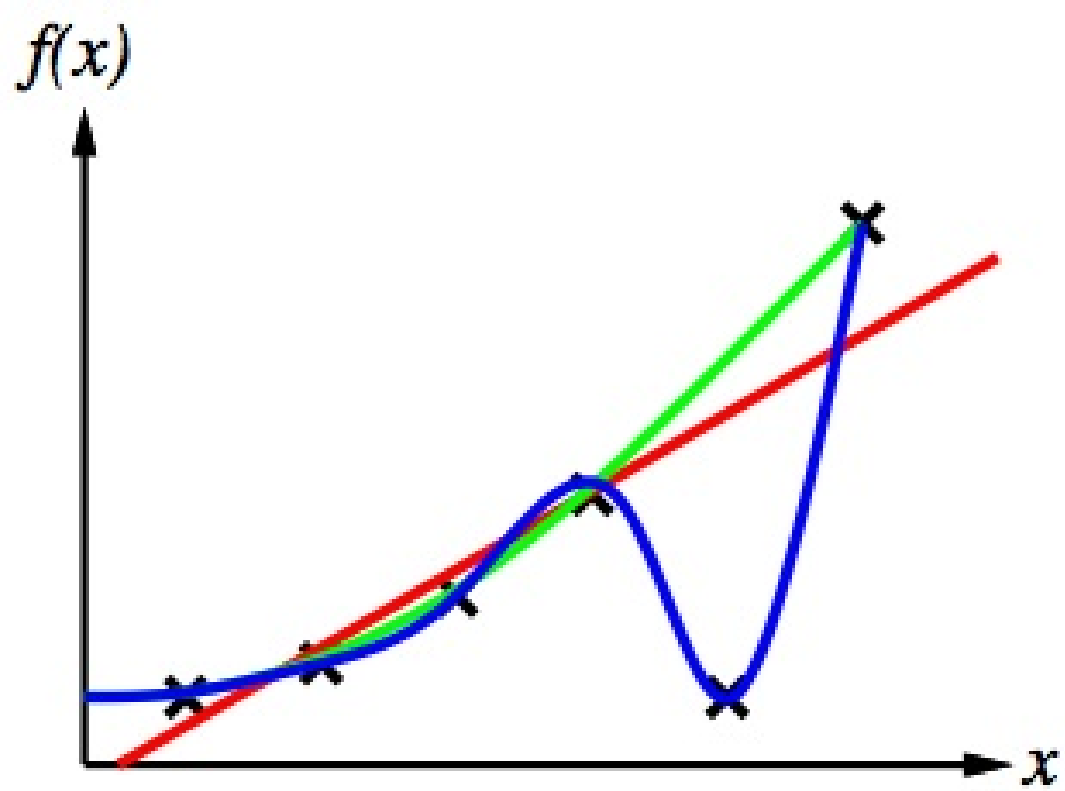

What is a good pattern?

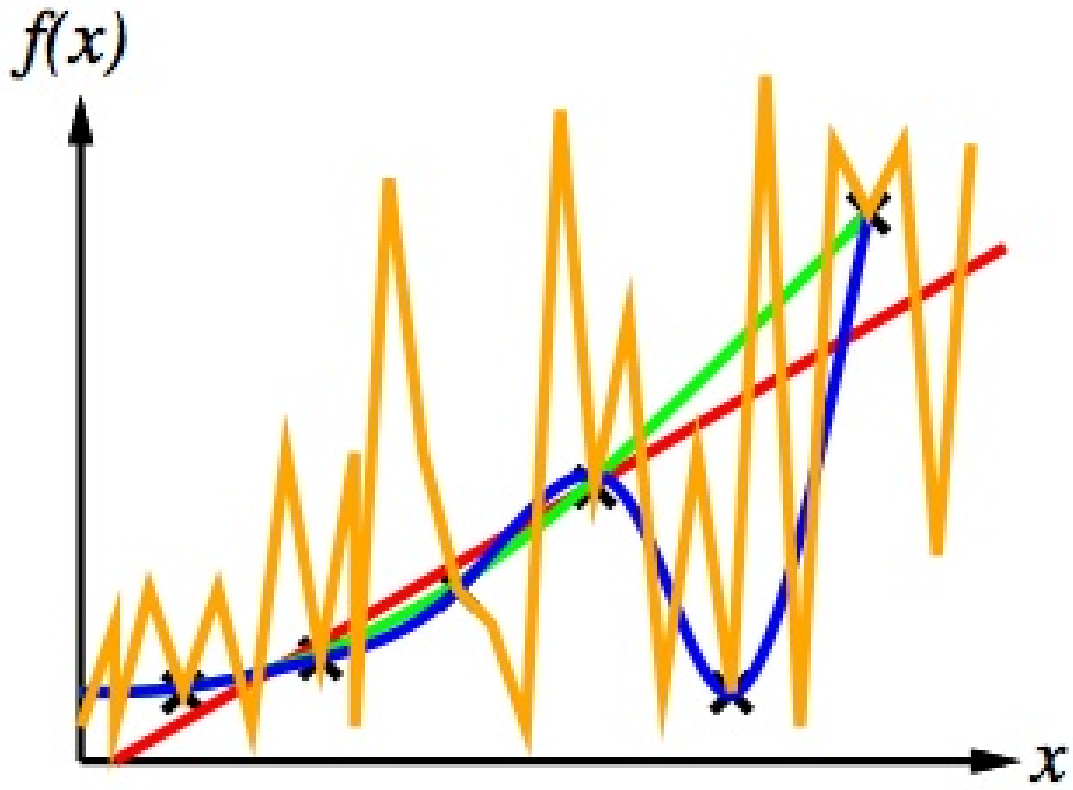

Occam's Razor

from Wikipedia:

Occam's razor (also spelled Ockham's razor) is a principle attributed to the 14th-century English logician and Franciscan friar William of Ockham. The principle states that the explanation of any phenomenon should make as few assumptions as possible, eliminating, or "shaving off", those that make no difference in the observable predictions of the explanatory hypothesis or theory. The principle is often expressed in Latin as the lex parsimoniae (law of succinctness or parsimony).

"All things being equal, the simplest solution tends to be the best one."

Types

- Supervised learning
- Non-supervised learning
- Semi-supervised learning
- Active learning
- On-line learning

Supervised learning

- **Fundamental problem**: to find a function that relates a set of inputs with a set of outputs
- Typical problems:
  - Classification
  - Regression

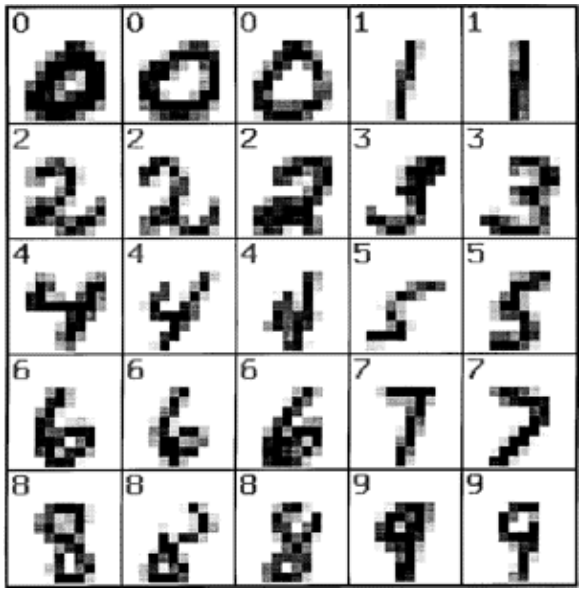

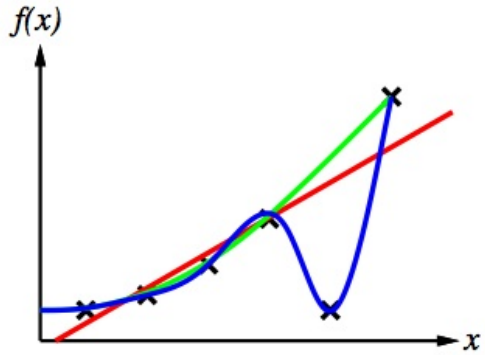

Non-supervised learning

- There are not labels for the training samples
- **Fundamental problem**: to find the subjacent structure of a training data set
- Typical problems: clustering, data compression
- Some samples may have labels, in that case it is called semi-supervised learning

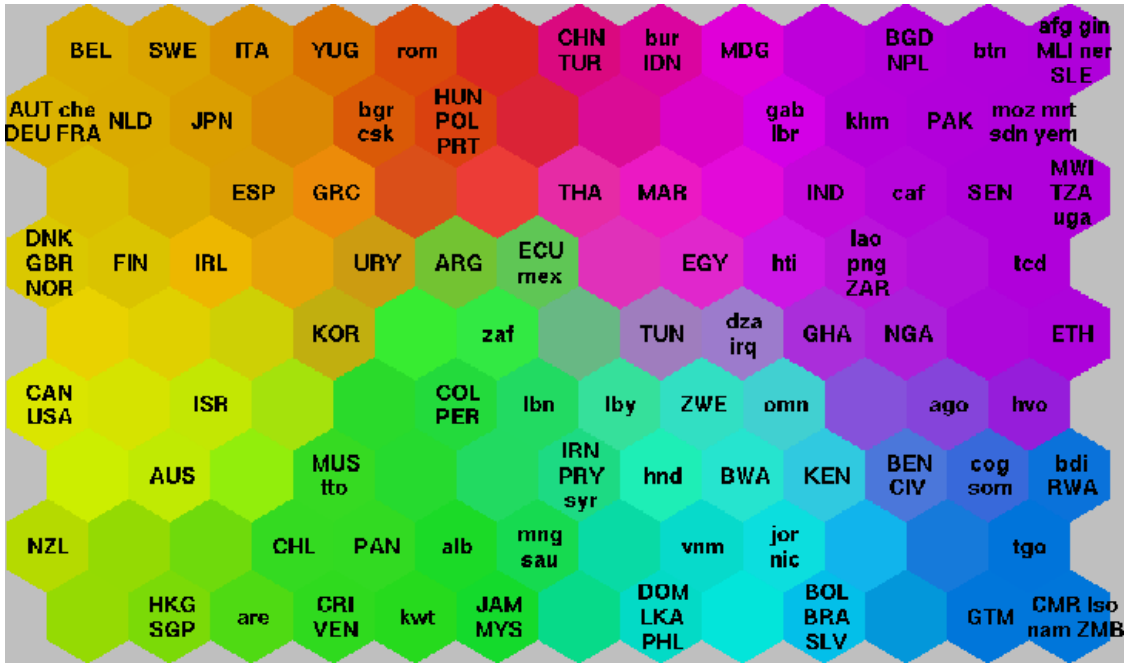

Active/reinforcing learning

- Generally, it happens in the context of an agent acting in an environment
- The agent is not told whether it has make the right decision or not
- The agent is punished or rewarded (not necessarily in an immediate way)
- **Fundamental problem**: to define a policy that allows to maximize the positive stimulus (reward)

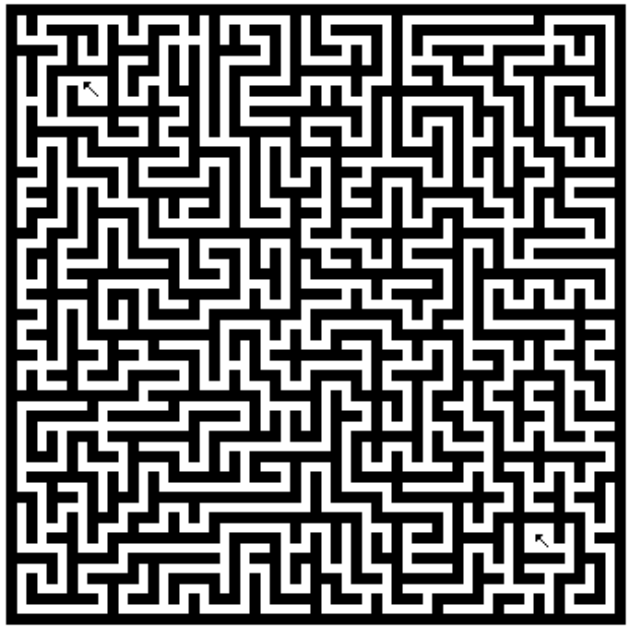

On-line learning

- Only one pass through the data
  - big data volume
  - real time
- It may be supervised or unsupervised
- **Fundamental problem**: to extract the maximum information from data with minimum number of passes

Representative techniques

- Computational
  - Decision trees
  - Nearest-neighbor classification
  - Graph-based clustering
  - Association rules
- Statistical
  - Multivariate regression
  - Linear discriminant analysis
  - Bayesian decision theory
  - Bayesian networks
  - K-means
- Computational-Statistical
  - SVM
  - AdaBoost
- Bio-inspired
  - Neural networks
  - Genetic algorithms
  - Artificial immune systems

Alpaydin, E. 2004 Introduction to Machine Learning (Adaptive Computation and Machine Learning). The MIT Press. (Cap 1,2)