

CSCI6506 – A little Machine Learning ‘101’

Instructor: Malcolm Heywood

September 7, 2011

1. What is Machine Learning (ML) and how does it relate to Artificial Intelligence (AI) as a whole?
2. How is a problem ‘posed’ to a ML algorithm to solve?
3. How do the basic ML genre differ?
 - *supervised* versus *unsupervised* versus *semi-supervised* versus *reinforcement* learning;
 - *filter* versus *wrapper* versus *embedded*;
4. What are the generic ML design questions and what implications do they have on the behavior of the resulting ML algorithm?
 - Goal versus stop (or termination) criteria – *cost function*;
 - Representation – describing *candidate solutions*;
 - Credit assignment – *exploration–exploitation tradeoff*;
5. What basic performance tradeoffs exist?
 - Representational biases;
 - Capacity for incorporating a priori knowledge?
 - Complexity–accuracy tradeoff;
 - the overlearning pathology;
 - use of validation data/ scenarios;
 - explicitly measuring complexity as a performance criterion.

- Scaling issues:
 - large *data sets*;
 - large *attribute spaces*;
 - large *state spaces*.
- What to learn against?
 - ***Layered Learning***: Structuring the task.
 - ***Competitive coevolution***: Not all exemplars/ opponents/ contexts are ‘created’ equal...
 - what has more impact on a candidate solution, the *learning algorithm* or the *environment* in which learning takes place?

6. Some ML holy grail,

- Problem decomposition;
- Credit assignment under delayed payoff;
- Establishing goals/ subtasks;
- Constructing useful training scenarios;
- Creativity: seeking out learning opportunities.