Machine Learning

Yuh-Jye Lee

Data Science and Machine Intelligence Lab
National Chiao Tung University

February 21, 2017

Learning a Class from Examples

- Suppose we want to learn a class (concept) C.
- Example: "sport cars".
- Given a collection of cars, have people label them as sport cars (positive examples) or non-sport cars (negative examples).
- Our learning task: find a *description* that is shared by all of the positive examples and none of the negative examples
- Once we have this *description* for the concept C, we can
  - predict given a new unseen car, predict whether or not it is a sports car
  - describe/compress understand what people expect in a car.

Choosing an Input Representation

- Suppose that of all the features describing cars, we choose just two features: price and engine power.
- Let
  - x_1 represent the price (in USD)
  - x_2 represent the engine volume (in cm^3)
- Then each car is represented

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$$

- and its label y denotes its type

$$y = \begin{cases} 1 \text{ ; if } \mathbf{x} \text{ is a positive example} \\ \text{-1 ; if } \mathbf{x} \text{ is a negative example} \end{cases}$$

- Each example is represented by the pair $(\mathbf{x}, y)$
- A training set containing ℓ examples is represented by $\mathcal{X} = \{\mathbf{x}^t, y_t\}_{t=1}^{\ell}$

Plotting the Training Data

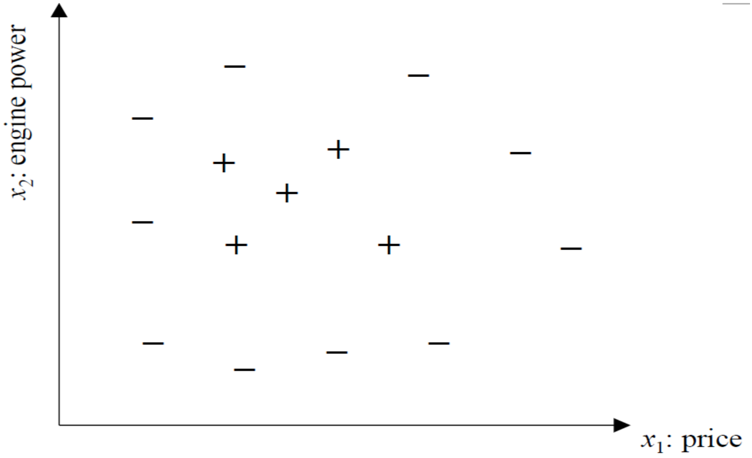

Figure: Training Data

Hypothesis Class

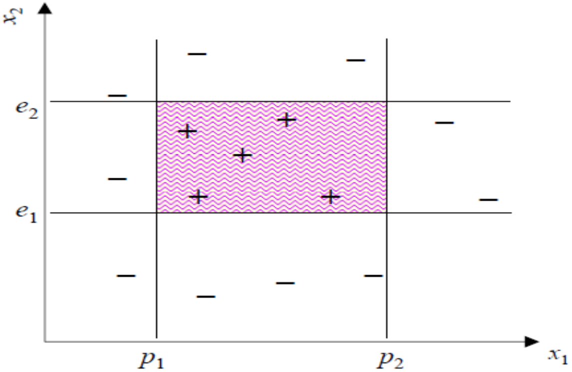

Figure: Suppose that we think that for a car to be a sports car, its price and its engine power should be in a certain range: $p_1 < price < p_2$ and $e_1 < engine < e_2$.

Concept Class

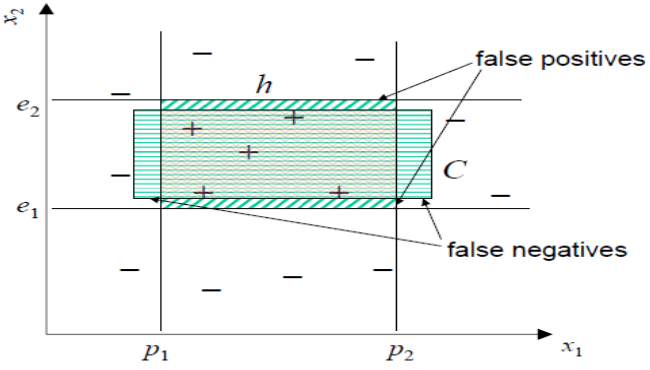

Figure: Suppose that the actual class is C task: find $h \in \mathcal{H}$ that is consistent (no training errors) with $\mathcal{X}$.

Choosing Hypothesis

- Empirical Error: proportion of training instances where predictions of h do not match the training set.

$$E(h \mid X) = \frac{1}{\ell} \sum_{t=1}^{\ell} \mathbf{1}\left(h(\mathbf{x^t}) \neq y_t\right)$$

- Each (p_1, p_2, e_1, e_2) defines a *hypothesis* $h \in \mathcal{H}$.
- We need to find the best one ...

Hypothesis Choice

Most specific?
Most general?

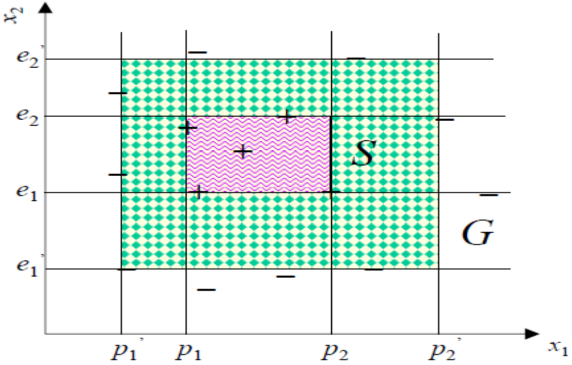

Figure: Most specific hypothesis S .Most general hypothesis G .

Consistent Hypothesis

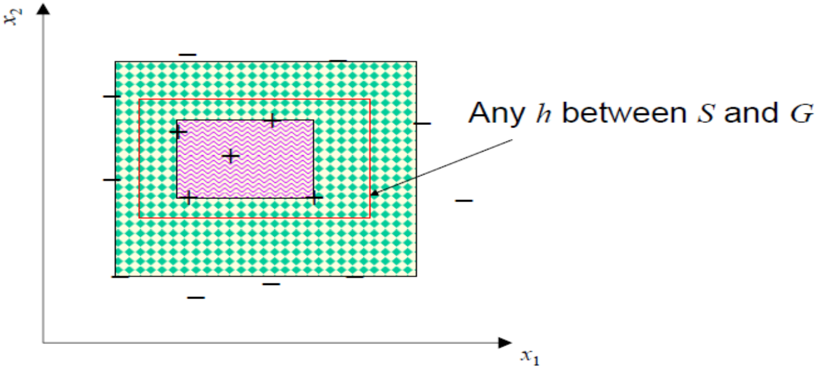

Figure: G and S define the boundaries of the Version Space.
The set of hypotheses more general than S and more specific than G forms the **Version Space** , the set of consistent hypotheses.

Now What?

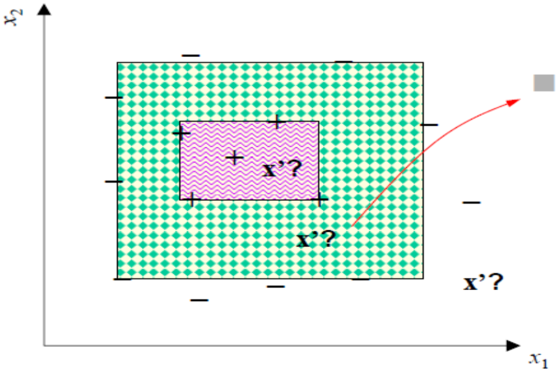

- *Using the average of S and G or just rejecting it to experts?*

Figure: How do we make prediction for a new $\mathbf{x}'$?

What is Machine Learning?

Representation + Optimization + Evaluation

- The most important reading assignment in my Machine Learning and Data Science and Machine Intelligence Lab at NCTU.
- Pedro Domingos *A few useful things to know about machine learning Communications of the ACM, Vol. 55 Issue 10, 78-87, October 2012 .*
- http://dl.acm.org/citation.cfm?id=2347755

The Master Algorithm

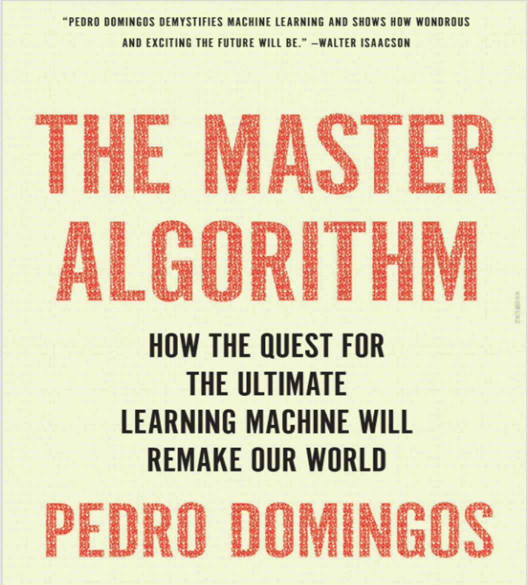

The Master Algorithm

The Basic Learning Concept

- Assumption: training instances are drawn from an unknown but fixed probability distribution $P(\mathbf{x}, y)$ independently.
- Our learning task:
  - Given a training set $S = \{(\mathbf{x}^1, y_1), (\mathbf{x}^2, y_2), \dots, (\mathbf{x}^\ell, y_\ell)\}$
  - We would like to construct a *rule*, $f(\mathbf{x})$ that can correctly predict the label y given unseen $\mathbf{x}$
  - If $f(\mathbf{x}) \neq y$ then we get some loss or penalty
  - For example: $\ell(f(\mathbf{x}), y) = \frac{1}{2}|f(\mathbf{x}) - y|$
- In the sport cars vs. non-sport cars example
  - **Represent** a car by two attributes: price and engine volumn
  - **Represent** the *hypothesis* (decision rule) by a *rectangle*
  - **Optimization**: Choose the hypothesis with the *smallest* error under certain *regularization*
  - **Evaluation**: Will our model work well in predicting the future examples? How good it will be? Is there any different setting to make it perform better?