INTRODUCTION TO

PAC LEARNING

by: *alpaydin@boun.edu.tr*
http://www.cmpe.boun.edu.tr/~ethem/i2ml

Postedited by*: R. Basili*

Learning a Class from Examples

- Class C of a "family car"
  - Prediction: Is car x a "family car"?
  - Knowledge extraction: What do people expect from a family car?
- Output:

  Positive (+) and negative (–) examples

- Input representation:

 x_1: price, x_2 : engine power

Training set $\mathcal{X}$

$$\mathcal{X} = \{\boldsymbol{x}^t, r^t\}_{t=1}^N$$

label

$$\boldsymbol{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

$$r = \begin{cases} 1 \text{ if } \boldsymbol{x} \text{ is positive} \\ 0 \text{ if } \boldsymbol{x} \text{ is negative} \end{cases}$$

x_2 : Engine power

x_2^t

x_1^t

x_1: Price

Class *C*

$$\left(p_1 \leq \text{price} \leq p_2\right) \text{AND} \left(e_1 \leq \text{engine power} \leq e_2\right)$$

In general we do not know
C(x).

Hypothesis class $\mathcal{H}$

$$h(\boldsymbol{x}) = \begin{cases} 1 \text{ if } h \text{ classifies } \boldsymbol{x} \text{ as positive} \\ 0 \text{ if } h \text{ classifies } \boldsymbol{x} \text{ as negative} \end{cases}$$

Empirical error:

$$E(h \mid \mathcal{X}) = \sum_{t=1}^{N} 1\left(h\left(\boldsymbol{x}^t\right) \neq r^t\right)$$

S, G, and the Version Space

most specific hypothesis, S

most general hypothesis, G

$h \in \mathcal{H}$, between S and G is consistent and make up the version space

(Mitchell, 1997)

Probably Approximately Correct (PAC) Learning

- How many training examples are needed so that the tightest rectangle S which will constitute our hypothesis, will probably be approximately correct?

  - We want to be confident (above a level) that
  - The error probability is bounded by some value

Probably Approximately Correct (PAC) Learning

- How many training examples are needed so that the tightest rectangle S which will constitute our hypothesis, will be probably approximately correct?

  - We want to be confident (above a level) that
  - The error probability is bounded by some value

Probably Approximately Correct (PAC) Learning

- In PAC learning, given a class *C* and examples drawn from some unknown but fixed distribution *p(x)*, we want to find the number of examples *N*, such that with probability at least *1 - δ*, *h* has error at most *ε* ?

  (Blumer et al., 1989)

$$P(\ C\Delta h \leq \varepsilon\ ) \geq \delta$$

  where *CΔh* is (area of the) "the region of difference between *C* and *h*", and *δ>0, ε>0.*

PAC Learning

- How many training examples *m* should we have, such that with probability at least *1 - δ*, h has error at most *ε* ?

(Blumer et al., 1989)

- Let prob. of a + ex. in each strip be at most *ε/4*
- Pr that random ex. misses a strip: *1- ε/4*
- Pr that *m* random instances miss a strip:

$$(1 - ε/4)^m$$

- Pr that *m* random instances instances miss 4 strips: $4(1 - ε/4)^m$
- We want $1-4(1 - ε/4)^m ≥ 1- δ$ or $4(1 - ε/4)^m ≤ δ$
- Using $1-x ≤ e^{-x}$ it follows that

$$4e^{-εm/4} ≤ δ \quad OR$$

- Divide by *4,* take *ln...* and show that
$$m ≥ (4/ε)ln(4/δ)$$

Probably Approximately Correct (PAC) Learning

How many training examples m should we have, such that with probability at least 1 - δ, h has error at most ε ?

(Blumer et al., 1989)

$$m \geq (4/\varepsilon)ln(4/\delta)$$

- *m* increases slowly with $1/\varepsilon$ and $1/\delta$

- Say ε=1% with confidence 95%, pick $m \geq 1752$
- Say ε=10% with confidence 95%, pick $m \geq 175$

VC (Vapnik-Chervonenkis) Dimension

- *N* points can be labeled in 2^N ways as +/–

- $\mathcal{H}$ shatters *N* if there exists a set of *N* points such that $h \in \mathcal{H}$ is consistent with all of these possible labels:
  - Denoted as: $VC(\mathcal{H}) = N$
  - Measures the capacity of *H*

- Any learning problem definable by *N* examples can be learned with no error by a hypothesis drawn from *H*

What is the VC dimension of axis-aligned rectangles?

Formal Definition

The VC Dimension

Definition: the VC dimension of a set of functions $H = \{h(\mathbf{x}, \alpha)\}$ is d if and only if there exists a set of points $\{x^i\}_{i=1}^d$ such that these points can be labeled in all 2^d possible configurations, and for each labeling, a member of set H can be found which correctly assigns those labels, but that no set $\{x^i\}_{i=1}^q$ exists where $q > d$ satisfying this property.

VC (Vapnik-Chervonenkis) Dimension

- $\mathcal{H}$ shatters N if there
  exists N points and $h \in \mathcal{H}$ such that
  h is consistent for any labelings
  of those N points.

- VC($axis\ aligned\ rectangles$) = 4

VC (Vapnik-Chervonenkis) Dimension

- *What does this say about using rectangles as our hypothesis class?*

- VC dimension is pessimistic: in general we do not need to worry about all possible labelings

- It is important to remember that one can choose the arrangement of points in the space, but then the hypothesis must be consistent with all possible labelings of those fixed points.

Examples

α

$x \longrightarrow$ f $\longrightarrow y^{est}$

$f(x, w) = sign(x.w)$

- denotes +1

∘ denotes -1

Examples

α

$x \longrightarrow$ f $\longrightarrow y^{est}$

$f(x, w, b) = sign(x.w + b)$

• denotes +1

○ denotes -1

Shattering

- Question: Can the following *f* shatter the following points?

$$f(x, w) = sign(x.w)$$

■ Answer: Yes. There are four possible training set types to consider:

w=(0,1)   w=(-2,3)   w=(2,-3)   w=(0,-1)

VC dim of linear classifiers in m-dimensions

If input space is *m-dimensional* and if **f** is *sign(**w**.**x**-b)*, what is the VC-dimension?

$$h=m+1$$

- Lines in 2D can shatter 3 points
- Planes in 3D space can shatter 4 points
- ...

Examples

$$\alpha$$

$$x \longrightarrow \boxed{f} \longrightarrow y^{est}$$

$$f(x,b) = sign(x.x - b)$$

- denotes +1

- denotes -1

Shattering

- Question: Can the following *f* shatter the following points?

$$f(x,b) = sign(x.x-b)$$

■Answer: Yes. Hence, the VC dimension of circles on the origin is at least 2.

- Note that if we pick two points at the same distance to the origin, they cannot be shattered. But we are interested if all possible labellings of some n-points can be shattered.

- How about 3 for circles on the origin (Can you find 3 points such that all possible labellings can be shattered?)?

Model Selection & Generalization

- Learning is an ill-posed problem; data is not sufficient to find a unique solution

- The need for inductive bias, assumptions about $\mathcal{H}$

- **Generalization:** How well a model performs on new data

- Different machines have different amounts of "power".
  Tradeoff between:
  - More power: Can model more complex classifiers but might overfit.
  - Less power: Not going to overfit, but restricted in what it can model.
  - **Overfitting**: $\mathcal{H}$ more complex than C or f
  - **Underfitting**: $\mathcal{H}$ less complex than C or f

Triple Trade-Off

- There is a trade-off between three factors (Dietterich, 2003):
  1. Complexity of $\mathcal{H}$, $c\,(\mathcal{H})$,
  2. Training set size, N,
  3. Generalization error, E, on new data

- □ As $N\uparrow$, $E\downarrow$
- □ As $c\,(\mathcal{H})\uparrow$, first $E\downarrow$ and then $E\uparrow$

Why care about "complexity"?

- We need a quantitative measure of complexity in order to be able to relate the training error (which we can observe) and the test error (that we'd like to optimize)

Complexity

- "Complexity" is a measure of a set of classifiers, not any specific (fixed) classifier

- • Many possible measures
  - degrees of freedom
  - description length
  - Vapnik-Chervonenkis (VC) dimension
  - etc.

Expected and Empirical error

$$
\hat{\mathcal{E}}_n(i) = \frac{1}{n} \sum_{t=1}^{n} \overbrace{\mathsf{Loss}(y_t, h_i(\mathbf{x}_t))}^{=\, 0,\, 1} = \text{ empirical error of } h_i(\mathbf{x})
$$

$$
\mathcal{E}(i) = E_{(\mathbf{x},y) \sim P} \big\{ \mathsf{Loss}(y, h_i(\mathbf{x})) \big\} = \text{ expected error of } h_i(\mathbf{x})
$$

Learning and VC-dimension

- Let d_{VC} be the VC-dimension of our set of classifiers F.

  Theorem: With probability at least $1 - \delta$ over the choice of the training set, for all $h \in F$

  $$\mathcal{E}(h) \leq \hat{\mathcal{E}}_n(h) + \epsilon(n, d_{VC}, \delta)$$

  where

  $$\epsilon(n, d_{VC}, \delta) = \sqrt{\frac{d_{VC}(\log(2n/d_{VC}) + 1) + \log(1/(4\delta))}{n}}$$

Model selection

- We try to find the model with the best balance of complexity and the fit to the training data
- Ideally, we would select a model from a nested sequence of models of increasing complexity (VC-dimension)

  Model 1 d_1
  Model 2 d_2
  Model 3 d_3

  where $d_1 \leq d_2 \leq d_3 \leq \ldots$

- The model selection criterion is: find the model class that achieves the lowest upper *bound* on the expected loss

  Expected error $\leq$ Training error $+$ Complexity penalty

Structural risk minimization cont'd

- We choose the model class F_i that minimizes the upper bound on the expected error:

$$\mathcal{E}(\hat{h}_i) \leq \hat{\mathcal{E}}_n(\hat{h}_i) + \sqrt{\frac{d_i(\log(2n/d_i) + 1) + \log(1/(4\delta))}{n}}$$

where $\hat{h}_i$ is the best classifier from F_i selected on the basis of the training set.
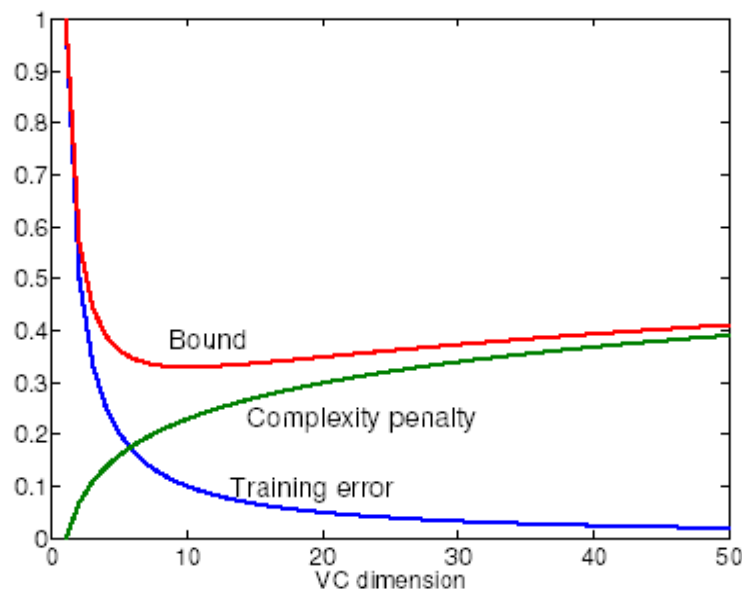
Structural risk minimization: example

linear

2^{nd} order polynomial

4^{th} order polynomial

8^{th} order polynomial

Structural risk minimization: example cont'd

- Number of training examples $n = 50$, confidence parameter $\delta = 0.05$.

Model	d_{VC}	Empirical fit	$\epsilon(n, d_{VC}, \delta)$
1^{st} order	3	0.06	0.5501
2^{nd} order	6	0.06	0.6999
4^{th} order	15	0.04	0.9494
8^{th} order	45	0.02	1.2849

- Structural risk minimization would select the simplest (linear) model in this case.

Summary: a learning machine

- A learning machine *f* takes an input *x* and transforms it, somehow using weights α, into a predicted output y^{est} = +/- 1

α is some vector of adjustable parameters

α

x ⟶ f ⟶ y^{est}
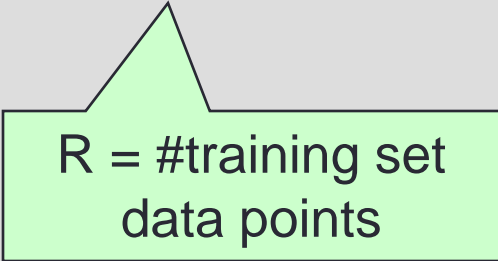
Back to test and empirical (training) error

- Given some machine **f**
- Define:

$$R(\alpha) = \text{TESTERR}\ (\alpha) = E\left[\frac{1}{2}\left|y - f(x,\alpha)\right|\right] = \begin{array}{l}\text{Probability of}\\ \text{Misclassification}\end{array}$$

$$R^{emp}(\alpha) = \text{TRAINERR}\ (\alpha) = \frac{1}{R}\sum_{k=1}^{R}\frac{1}{2}\left|y_k - f(x_k,\alpha)\right| = \begin{array}{l}\text{Fraction Training}\\ \text{Set misclassified}\end{array}$$

R = #training set data points

Vapnik-Chervonenkis dimension

$$\text{TESTERR}(\alpha) = E\left[\frac{1}{2}\left|y - f(x, \alpha)\right|\right] \qquad \text{TRAINERR}(\alpha) = \frac{1}{R}\sum_{k=1}^{R}\frac{1}{2}\left|y_k - f(x_k, \alpha)\right|$$
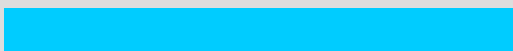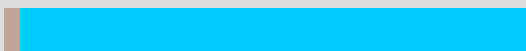
- Given some machine **f**, let *h* be its VC dimension (*h* does not depend on the choice of training set)
- **Let R be the number of training examples**
- Vapnik showed that with probability 1-η

$$\text{TESTERR}(\alpha) \leq \text{TRAINERR}(\alpha) + \sqrt{\frac{h(\log(2R/h) + 1) - \log(\eta/4)}{R}}$$

This gives us a way to estimate the error on future data based only on the training error and the VC-dimension of **f**

VC-dimension as measure of complexity

$$\text{TESTERR}\ (\alpha) \leq \text{TRAINERR}\ (\alpha) + \sqrt{\frac{h(\log(2R/h)+1)-\log(\eta/4)}{R}}$$

i	f_i	TRAINERR	VC-Conf	Probable upper bound on TESTERR	Choice
1	f_1				
2	f_2				
3	f_3				⊠
4	f_4				
5	f_5				
6	f_6				

Using VC-dimensionality

People have worked hard to find VC-dimension for..

- Decision Trees
- Perceptrons
- Neural Nets
- Decision Lists
- Support Vector Machines
- And many many more

All with the goals of

1. Understanding which learning machines are more or less powerful under which circumstances
2. Using Structural Risk Minimization for to choose the best learning machine

Alternatives to VC-dim-based model selection

Cross Validation

- To estimate generalization error, we need data unseen during training. We split the data as
  - Training set (50%)
  - Validation set (25%)
  - Test (publication) set (25%)
- Resampling when there is few data

Alternatives to VC-dim-based model selection

- What could we do instead of the scheme below?
    1. Cross-validation

i	f_i	TRAINERR	10-FOLD-CV-ERR	Choice
1	f_1			
2	f_2			
3	f_3			⊠
4	f_4			
5	f_5			
6	f_6			

Extra Comments

- An excellent tutorial on VC-dimension and Support Vector Machines

    C.J.C. Burges. A tutorial on support vector machines for pattern recognition. Data Mining and Knowledge Discovery, 2(2):955-974, 1998. http://citeseer.nj.nec.com/burges98tutorial.html

What you should know

- Definition of PAC learning
- The definition of a learning machine: $f(x, \alpha)$
- The definition of Shattering
- Be able to work through simple examples of shattering
- The definition of VC-dimension
- Be able to work through simple examples of VC-dimension
- Structural Risk Minimization for model selection
- Awareness of other model selection methods