A Tour of Machine Learning

Michèle Sebag

TAO

Dec. 5th, 2011

Examples

- Cheques
- Spam
- Robot
- Helicopter
- Netflix
- Playing Go
- Google

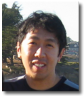

http://ai.stanford.edu/~ang/courses.html

Reading cheques

LeCun et al. 1990

MNIST: The drosophila of ML

Fig. 4. Size-normalized examples from the MNIST database.

Classification

Spam – Phishing – Scam

Classification, Outlier detection

The 2005 Darpa Challenge

Thrun, Burgard and Fox 2005

Autonomous vehicle Stanley – Terrains

Robots

Kolter, Abbeel, Ng 08; Saxena, Driemeyer, Ng 09

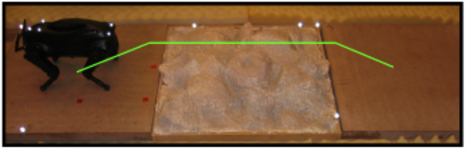

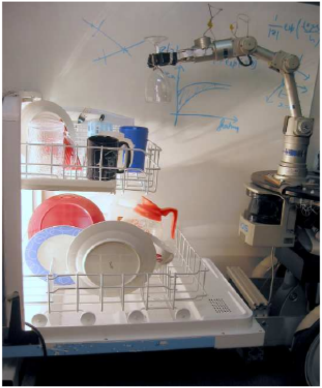

Reinforcement learning

Classification

Robots, 2

Toussaint et al. 2010

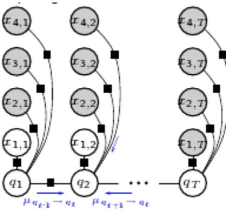

(a) Factor graph modelling the variable interactions

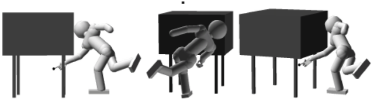

(b) Behaviour of the 39-DOF Humanoid:
Reaching goal under Balance and Collision constraints

Bayesian Inference for Motion Control and Planning

Go as AI Challenge

Gelly Wang 07; Teytaud et al. 2008-2011

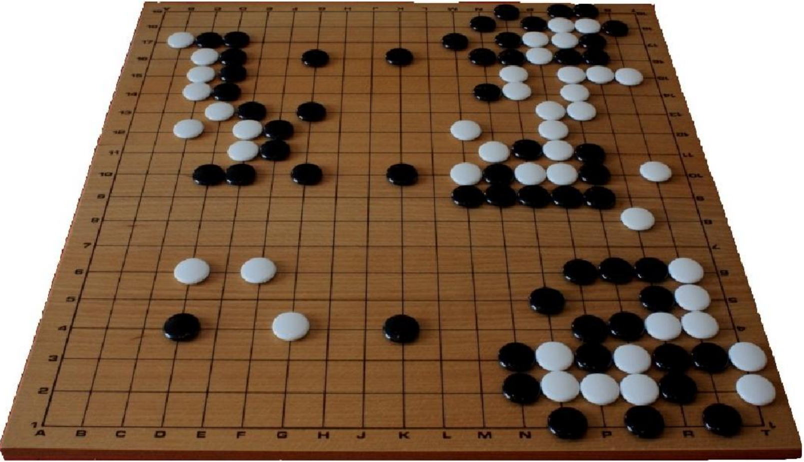

Reinforcement Learning, Monte-Carlo Tree Search

Netflix Challenge 2007-2008

Collaborative Filtering

The power of big data

- Now-casting — outbreak of flu
- Public relations $>>$ Advertizing

Sparrow, Science 11

In view of Dartmouth 1956 agenda

We propose a study of artificial intelligence [..]. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.

Where we are

Ast. series 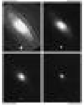

Pierre de Rosette 

World

Natural phenomenons

Human–related phenomenons

Data / Principles

Maths. Modelling

Common Sense

You are here

Data

Example

- row : example/ case
- column : feature/variables/attribute
- attribute : class/label

age	employme	education	edun	marital	...	job	relation	race	gender	hour	country	wealth
					...							
39	State_gov	Bachelors	13	Never_mar	...	Adm_cleric	Not_in_fam	White	Male	40	United_Sta	poor
51	Self_emp_	Bachelors	13	Married	...	Exec_man	Husband	White	Male	13	United_Sta	poor
39	Private	HS_grad	9	Divorced	...	Handlers_c	Not_in_fam	White	Male	40	United_Sta	poor
54	Private	11th	7	Married	...	Handlers_c	Husband	Black	Male	40	United_Sta	poor
28	Private	Bachelors	13	Married	...	Prof_speci	Wife	Black	Female	40	Cuba	poor
38	Private	Masters	14	Married	...	Exec_man	Wife	White	Female	40	United_Sta	poor
50	Private	9th	5	Married_sp	...	Other_serv	Not_in_fam	Black	Female	16	Jamaica	poor
52	Self_emp_	HS_grad	9	Married	...	Exec_man	Husband	White	Male	45	United_Sta	rich
31	Private	Masters	14	Never_mar	...	Prof_speci	Not_in_fam	White	Female	50	United_Sta	rich
42	Private	Bachelors	13	Married	...	Exec_man	Husband	White	Male	40	United_Sta	rich
37	Private	Some_coll	10	Married	...	Exec_man	Husband	Black	Male	80	United_Sta	rich
30	State_gov	Bachelors	13	Married	...	Prof_speci	Husband	Asian	Male	40	India	rich
24	Private	Bachelors	13	Never_mar	...	Adm_cleric	Own_child	White	Female	30	United_Sta	poor
33	Private	Assoc_acc	12	Never_mar	...	Sales	Not_in_fam	Black	Male	50	United_Sta	poor
41	Private	Assoc_voc	11	Married	...	Craft_repai	Husband	Asian	Male	40	*MissingVa	rich
34	Private	7th_8th	4	Married	...	Transport_	Husband	Amer_India	Male	45	Mexico	poor
26	Self_emp_	HS_grad	9	Never_mar	...	Farming_fi	Own_child	White	Male	35	United_Sta	poor
33	Private	HS_grad	9	Never_mar	...	Machine_c	Unmarried	White	Male	40	United_Sta	poor
38	Private	11th	7	Married	...	Sales	Husband	White	Male	50	United_Sta	poor
44	Self_emp_	Masters	14	Divorced	...	Exec_man	Unmarried	White	Female	45	United_Sta	rich
41	Private	Doctorate	16	Married	...	Prof_speci	Husband	White	Male	60	United_Sta	rich
:	:	:	:	:	:	:	:	:	:	:	:	:

Instance space $\mathcal{X}$

- Propositionnal : $\mathcal{X} \equiv \mathbb{R}^d$
- Structured : sequential, spatio-temporal, relational.

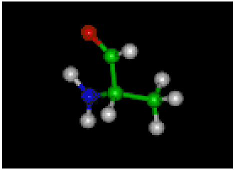

aminoacid

Types of Machine Learning problems

WORLD – DATA – USER

Observations	+ Target	+ Rewards
Understand Code	Predict Classification/Regression	Decide Policy
Unsupervised LEARNING	Supervised LEARNING	Reinforcement LEARNING

Unsupervised Learning

Example: a bag of songs

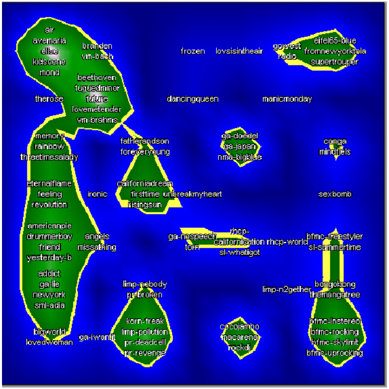

Find categories/characterization

Find names for sets of things

From observations to codes

What's known

- Indexing
- Compression

What's new

- Accessible to humans

Find codes with meanings

Position of the problem

Given Data, structure (distance, model space)

Find Code and its performance

Minimum Description Length

$$\text{Minimize } (\text{Adequacy}(\text{Data}, \text{Code}) + \text{Complexity}(\text{Code}))$$

What is difficult

- Impossibility thm
  scale-invariance, richness, consistency are incompatible
- Distances are elusive — curse of dimensionality

Unsupervised Learning

- Crunching data
- Finding correlations
- "Telling stories"
- Assessing causality

Causation and Prediction Challenge, Guyon et al. 10

Ultimately

- Make predictions good enough
- Build cases
- Take decisions

Visualization

Maps of cancer in Spain

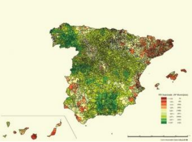

Breast

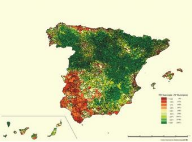

Lungs

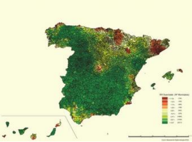

Stomach

http://www.elpais.com/articulo/sociedad/contaminacion/industrial/multiplica/tumores/Cataluna/Huelva/Asturias/elpepusoc/20070831elpepisoc_2/Tes

Types of Machine Learning problems

WORLD – DATA – USER

Observations	+ Target	+ Rewards
Understand Code	Predict Classification/Regression	Decide Policy
Unsupervised LEARNING	Supervised LEARNING	Reinforcement LEARNING

Supervised Learning

Context

$$\text{World} \rightarrow \text{instance } \mathbf{x}_i \rightarrow \begin{array}{c}\text{Oracle}\\ \downarrow \\ y_i\end{array}$$

Input:	Training set $\mathcal{E} = \{(\mathbf{x_i}, y_i), i = 1 \ldots n, x_i \in \mathcal{X}, y_i \in \mathcal{Y}\}$
Output:	Hypothesis $h : \mathcal{X} \mapsto \mathcal{Y}$
Criterion:	few mistakes (details later)

First task

- Propose a criterion $\mathcal{L}$
  - Consistency

When number n of examples goes to ∞
and the target concept h^* is in $\mathcal{H}$
Algorithm finds $\hat{h}_n$, with

$$lim_{n \to \infty} h_n = h^*$$

  - Convergence speed

$$||h^* - h_n|| = \mathcal{O}(1/\ln n), \mathcal{O}(1/\sqrt{n}), \mathcal{O}(1/n), ...\mathcal{O}(2^{-n})$$

Second task

- Optimize $\mathcal{L}$
  - $+$ Convex optimization: guarantees, reproducibility
  - (...) ML has suffered from an acute convexivitis epidemy

Le Cun et al. 07

H. Simon, 58:

In complex real-world situations, optimization becomes approximate optimization since the description of the real-world is radically simplified until reduced to a degree of complication that the decision maker can handle.
Satisficing seeks simplification in a somewhat different direction, retaining more of the detail of the real-world situation, but settling for a satisfactory, rather than approximate-best, decision.

What is the point ?

Underfitting Overfitting

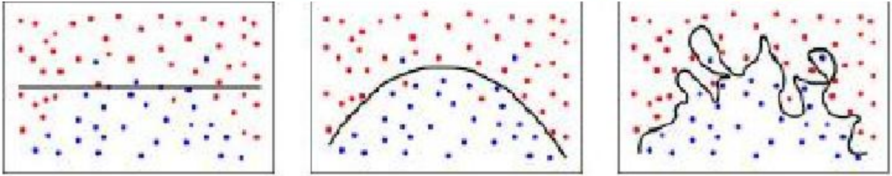

The point is not to be perfect on the training set

What is the point ?

Underfitting Overfitting

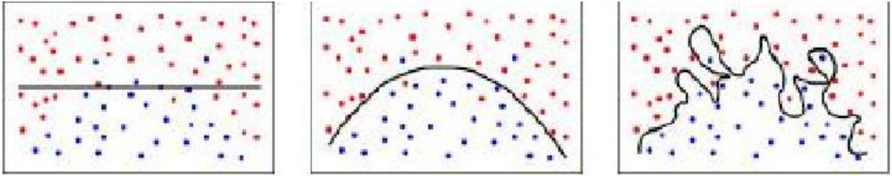

The point is not to be perfect on the training set

The villain: overfitting

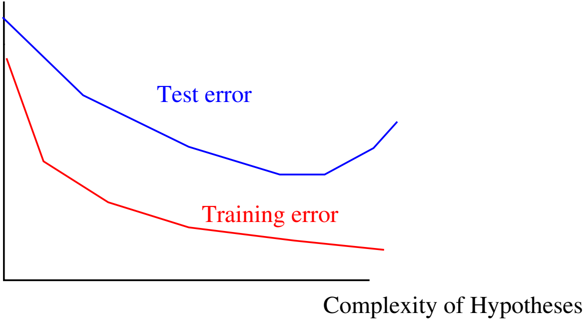

What is the point ?

Prediction good on future instances

Necessary condition:

Future instances must be similar to training instances

"identically distributed"

Minimize (cost of) errors $\ell(y, h(x)) \geq 0$

not all mistakes are equal.

Error: Find upper bounds

Vapnik 92, 95

Minimize expectation of error cost

$$\text{Minimize } E[\ell(y, h(x))] = \int_{X \times Y} \ell(y, h(x)) p(x, y) dx\, dy$$

Vapnik 92, 95

Minimize expectation of error cost

$$\text{Minimize } E[\ell(y, h(x))] = \int_{X \times Y} \ell(y, h(x)) p(x, y) dx\, dy$$

Principle

Si h "is well-behaved" on $\mathcal{E}$, and h is "sufficiently regular" h will be well-behaved in expectation.

$$E[F] \leq \frac{\sum_{i=1}^{n} F(x_i)}{n} + c(F, n)$$

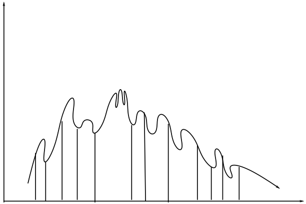

Minimize upper bounds

If x_i iid

Then

$$\text{Generalization error} < \text{Empirical error} + \text{Penalty term}$$

Find

$$h^* = \text{argmin }_h Fit(h, Data) + Penalty(h)$$

Minimize upper bounds

If x_i iid

Then

$$\text{Generalization error} < \text{Empirical error} + \text{Penalty term}$$

Find

$$h^* = \text{argmin }_h Fit(h, Data) + Penalty(h)$$

Designing penalty/regularization term

- Some guarantees
- Incorporate priors
- A tractable optimization problem

Supervised ML as Methodology

Phases

1. Collect data — expert, DB
2. Clean data — stat, expert
3. Select data — stat, expert
4. Data Mining / Machine Learning
   - Description — *what is in data ?*
   - Prediction — *Decide for one example*
   - Agregate — *Take a global decision*
5. Visualisation — chm
6. Evaluation — stat, chm
7. Collect new data — expert, stat

Trends

Extend scopes

- Active Learning: collect useful data
- Transfert/Multi-task learning: relax iid assumption

Prior knowledge structured spaces

- In the feature space Kernels
- In the regularization term

Big data

- Who does control the data ?
- When does brute force win ?

Types of Machine Learning problems

WORLD – DATA – USER

Observations	+ Target	+ Rewards
Understand Code	Predict Classification/Regression	Decide Policy
Unsupervised LEARNING	Supervised LEARNING	Reinforcement LEARNING

Reinforcement Learning

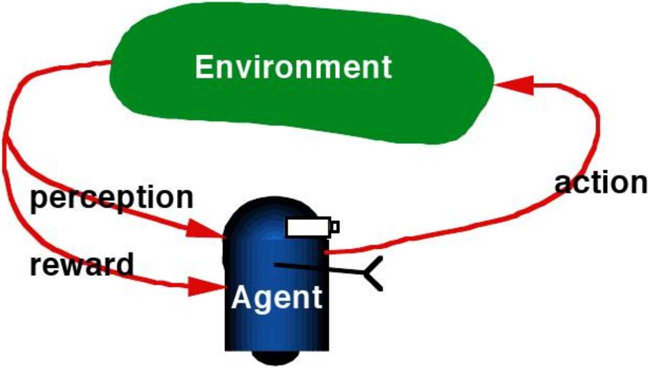

Context

- Agent temporally (and spatially) situated
- Learns and plans
- To act on the (stochastic, uncertain) environment
- To maximize cumulative reward

Sutton Barto 98; Singh 05

Init

World is unknown *Model of the world*

Some actions, in some states, yield rewards, possibly delayed, with some probability.

Output

$$\text{Policy} = \text{strategy} = (\text{State} \rightarrow \text{Action})$$

Goal: Find policy π^* maximizing in expectation

Sum of (discounted) rewards collected using π starting in s_0

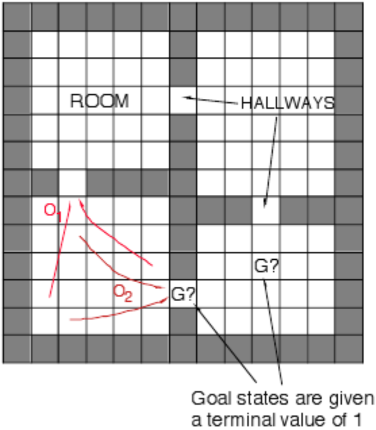

4 rooms

4 hallways

4 unreliable primitive actions

Fail 33% of the time

8 multi-step options

(to each room's 2 hallways)

Given goal location, quickly plan shortest route

All rewards zero

$\gamma = .9$

Reinforcement learning

Of several responses made to the same situation, those which are accompanied or closely followed by satisfaction to the animal will – others things being equal – be more firmly connected with the situation, so that when it recurs, they will more likely to recur;

those which are accompanied or closely followed by discomfort to the animal will – others things being equal – have their connection with the situation weakened, so that when it recurs, they will less likely to recur;

the greater the satisfaction or discomfort, the greater the strengthening or weakening of the link.
Thorndike, 1911.

Formalization

Given

- State space $\mathcal{S}$
- Action space $\mathcal{A}$
- Transition function $p(s, a, s') \mapsto [0, 1]$
- Reward $r(s)$

Find $\pi : \mathcal{S} \mapsto \mathcal{A}$

$$\text{Maximize } E[\pi] = \sum_{s_{t+1} \sim p(s_t, \pi(s_t))} \gamma^{t+1} r(s_{t+1})$$

Tasks

Three interdependent goals

- Learn a world model (p, r)
- Through experimenting
- Exploration vs exploitation dilemma

Issues

- Sparing trials; Inverse Optimal Control
- Sparing observations: Learning descriptions
- Load balancing

Applications

Classical applications

1. Games
2. Control, Robotics
3. Planning, scheduling — OR

New applications

- Whenever several interdependent classifications are needed
- Lifelong learning: self-$*$ systems — Autonomic Computing

Challenges

ML: A new programming language

- Design programs with learning primitives
- Reduction of ML problems — Langford et al. 08
- Verification ?

ML: between data acquisition and HPC

- giga, tera, peta, exa, yottabites
- GPU — Schmidhuber et al. 10