M1 – Apprentissage

Michèle Sebag – Benoit Barbot
LRI – LSV

Sept. 2013

Where we are

Ast. series 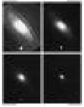

Pierre de Rosette 

World

Natural phenomenons

Human–related phenomenons

Data / Principles

Maths. Modelling

Common Sense

You are here

Where we are

Sc. data 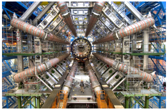

World

Natural phenomenons

Human–related phenomenons

Data / Principles

Maths. Modelling

Common Sense

You are here

Harnessing Big Data

Watson (IBM) defeats human champions at the quiz game Jeopardy (Feb. 11)

i	1	2	3	4	5	6	7	8	
1000^i	kilo	mega	giga	tera	peta	exa	zetta	yotta	bytes

- Google: 24 petabytes/day
- Facebook: 10 terabytes/day; Twitter: 7 terabytes/day
- Large Hadron Collider: 40 terabytes/seconds

Types of Machine Learning problems

WORLD – DATA – USER

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

Supervised Machine Learning

World $\rightarrow$ instance $\mathbf{x}_i$ $\rightarrow$ Oracle $\downarrow$ y_i

MNIST

Yann Le Cun, since end 80s

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

The 2005-2012 Visual Object Challenges

A. Zisserman, C. Williams, M. Everingham, L. v.d. Gool

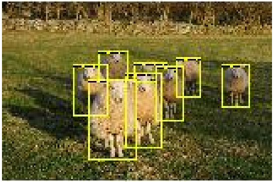

Supervised learning, notations

Input: set of $(\mathbf{x}, y)$

- An instance $\mathbf{x}$ e.g. set of pixels, $\mathbf{x} \in \mathbb{R}^D$
- A label y in $\{1, -1\}$ or $\{1, \ldots, K\}$ or $\mathbb{R}$

Supervised learning, notations

Input: set of $(\mathbf{x}, y)$

- An instance $\mathbf{x}$ e.g. set of pixels, $\mathbf{x} \in \mathbb{R}^D$
- A label y in $\{1, -1\}$ or $\{1, \ldots, K\}$ or $\mathbb{R}$

Pattern recognition

- Classification **Does the image contain the target concept ?**

$$h : \{ \text{ Images}\} \mapsto \{1, -1\}$$

- Detection **Does the pixel belong to the img of target concept?**

$$h : \{ \text{ Pixels in an image}\} \mapsto \{1, -1\}$$

- Segmentation

  Find contours of all instances of target concept in image

Unsupervised learning

Clustering

http://www.ofai.at/ elias.pampalk/music/

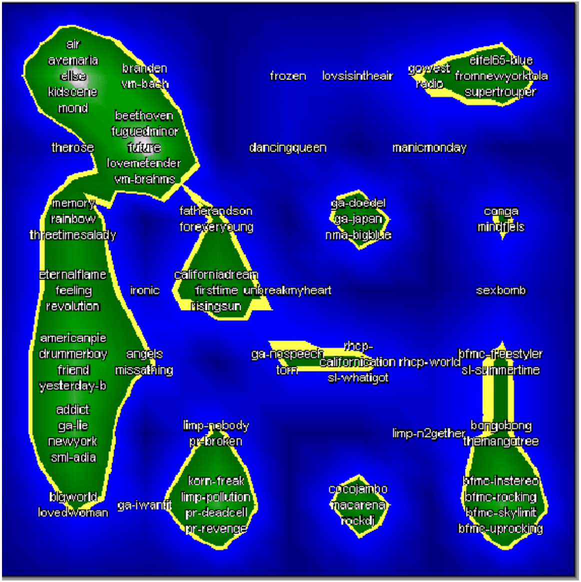

Unsupervised learning, issues

Hard or soft ?

- **Hard**: find a partition of the data
- **Soft**: estimate the distribution of the data as a mixture of components.

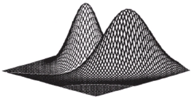

Parametric vs non Parametric ?

- **Parametric**: number K of clusters is known
- **Non-Parametric**: find K

  (wrapping a parametric clustering algorithm)

Unsupervised learning, 2

Collaborative Filtering

Netflix Challenge 2007-2008

Collaborative filtering, notations

Input

- A set of users n_u, ca 500,000
- A set of movies n_m, ca 18,000
- A $n_m \times n_u$ matrix: person, movie, rating
  Very sparse matrix: less than 1% filled...

Output

- Filling the matrix !

Collaborative filtering, notations

Input

- A set of users $\hfill n_u$, ca 500,000
- A set of movies $\hfill n_m$, ca 18,000
- A $n_m \times n_u$ matrix: person, movie, rating
  Very sparse matrix: less than 1% filled...

Output

- Filling the matrix !

Criterion

- (relative) mean square error
- ranking error

Reinforcement learning

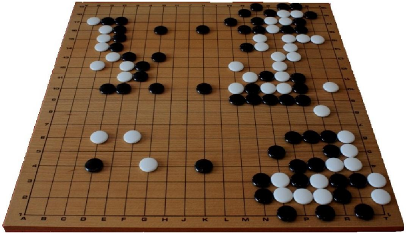

Reinforcement learning, notations

Notations

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

Goal

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

$$\text{Maximize } E[\pi] = \text{ Expected cumulative reward}$$

(detail later)

Some pointers

- My slides: http://tao.lri.fr/tiki-index.php?page=Courses
- Andrew Ng courses: http://ai.stanford.edu/~ang/courses.html
- PASCAL videos http://videolectures.net/pascal/
- Tutorials NIPS Neuro Information Processing Systems http://nips.cc/Conferences/2006/Media/
- About ML/DM http://hunch.net/

This course

WHO

- Michèle Sebag, machine learning — LRI
- Benoit Barbot, — LSV

WHAT

1. Introduction
2. Supervised Machine Learning
3. Unsupervised Machine Learning
4. Reinforcement Learning

WHERE: http://tao.lri.fr/tiki-index.php?page=Courses

Exam

Final:

- Questions
- Problems

Volunteers

- Some pointers are in the slides
  More ?

  here a paper or url
- Volunteers: read material, write one page, send it (sebag@lri.fr), oral presentation 5mn.

Overview

The roots of ML : AI

AI as search

AI and games

Promises?

What's new

Roots of AI

Bletchley

- Enigma cypher 1918-1945
- Some flaws/regularities
- Alan Turing (1912-1954) and Gordon Welchman: the Bombe
- Colossus

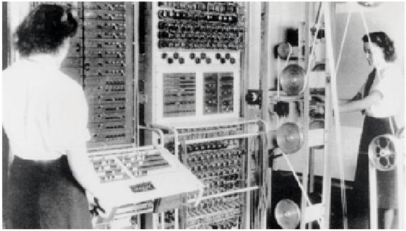

Dartmouth: when AI was coined

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.

An attempt will be made to find how to make machines use language, form abstraction and concepts ... and improve themselves.

Dartmouth: when AI was coined

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.

An attempt will be made to find how to make machines use language, form abstraction and concepts ... and improve themselves. John McCarthy, 1956

Before AI, the vision was there:

Alan Turing

Machine Learning, 1950

by (...) mimicking education, we should hope to modify the machine until it could be relied on to produce definite reactions to certain commands.

Before AI, the vision was there:

Alan Turing

Machine Learning, 1950

by (...) mimicking education, we should hope to modify the machine until it could be relied on to produce definite reactions to certain commands.

How ?

One could carry through the organization of an intelligent machine with only two interfering inputs, one for pleasure or reward, and the other for pain or punishment.

More ?

http://www.csee.umbc.edu/courses/471/papers/turing.pdf

The imitation game

The criterion:
Whether the machine could answer **questions in such a way that it will be extremely difficult to guess whether the answers are given by a man, or by the machine**

Critical issue
The extent we regard something as behaving in an intelligent manner is determined **as much by our own state of mind and training, as by the properties of the object under consideration**.

The imitation game, 2

A regret-like criterion

- Comparison to reference performance (oracle)
- More difficult task $\not\Rightarrow$ higher regret

Oracle = human being

- Social intelligence matters
- Weaknesses are OK.

Great expectations ! Promises

1955 : Logic Theorist Newell, Simon, Shaw, 1955

- Reading Principia Mathematica

  Whitehead and Russell, 1910-1913 ... an attempt to derive all mathematical truths from a well-defined set of axioms and inference rules in symbolic logic
- General Problem Solver Newell, Shaw, Simon, 1960

Within 10 years, a computer will

- be the world's chess champion
- prove an important theorem in maths
- compose good music
- set up the language for theoretical psychology

Overview

The roots of ML : AI

AI as search

AI and games

Promises?

What's new

Position of the problem

Symbols | **Operators**

numbers | $2 + 2 = 4$

concepts | $A,\ A \rightarrow B \models B$

Symbol manipulation

- Numbers and arithmetic operators $(+, \times, ...)$ — interpretation
  Arithmetics, Constraint Satisfaction
- Concepts, logical operators
  - Propositional — Inference, Constraint Satisfaction
  - Relational ($man(X), mortal(X), man(Socrates)$) — + unification
    Logic programming

Unification + Interpretation = Constraint Programming

Symbolic calculus: ingredients

Reasoning; navigate in a search tree

- States; tree nodes
- Navigation: select operators (= edges)

How

- Select promising operators
- Evaluate a state node
- Prune the search tree

Languages

IPL, Lisp, Prolog

- Lists
- Actions

Artificial Intelligence as Search

	Search space	Navigation	Criteria
Logic		+	
Expert Systems	+	+	
Games		+	+

Inference

Deduction

- **Modus ponens**

$$A, A \rightarrow B$$
$$\models B$$

- **Modus tollens**

$$\neg B, A \rightarrow B$$
$$\models \neg A$$

Comment

- Truth preserving $\models$
- Which déduction

More ?

http://homepages.math.uic.edu/ kauffman/Robbins.htm

Inference, 2

Induction

$\neg A, B$

(inference) $A \to B$

Inference, 2

Induction

$\neg A, B$

$$\text{(inference) } A \rightarrow B$$

Correlation & causality

- Many tuberculous people die in mountain regions
- Therefore ?

Inference, 3

Abduction

$B, \ A \rightarrow B$

(inférence) A

Inference, 3

Abduction

$B, \ A \rightarrow B$

(inférence) A

Multiple causes

- Drunk $\rightarrow$ Staggering. And you're staggering.
- Therefore ?

Halt

1972 : Winter AI

Dreyfus report

- Driving application: automatic translation
- Conjecture: cannot be syntactical (one must understand)
  - Paul is on the bus way; he is a friend; I push him (away)
  - Paul is on the bus way; he is an ennemy; I push him (under)

Discussion

- Everyone is able of deduction; the expert is able of reasoning
- The chinese room — Searle, “strong AI”
- Turing test

Artificial Intelligence as Search

	Search space	Navigation	Criteria
Logic		+	
Expert Systems	+	+	
Games		+	+

Expert systems

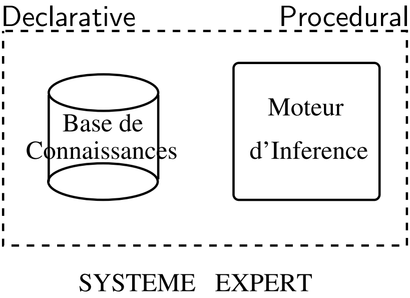

At the core of ES

- Knowledge base
  $A \rightarrow B$
- Inference engine
  $\models$, inference

Expert systems

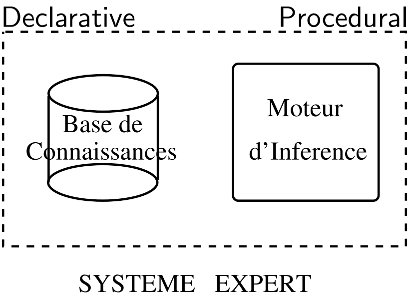

Forward chaining

Input Facts A

Output Diagnosis B

Backward chaining

Input B?

Output A? questions on facts leading to infer B

Lessons learned

Declarative style

- From orders/instructions to statements

Success

- Dendral: organic chemistry — Feigenbaum et al. 60s
- Mycin: medicine — Shortliffe, 76
- Molgen: molecular biology — Stefik, 81
- R1: computer assembly — McDermott, 82

Limitations

- Decreasing returns
- Human factors
- Statements... — but control needed

Bottleneck

Where do the knowledge base come from ?

Overview

The roots of ML : AI

AI as search

AI and games

Promises?

What's new

Artificial Intelligence as Search

	Search space	Navigation	Criteria
Logic		+	
Expert Systems	+	+	
Games		+	+

Why games ?

- Micro-worlds — finite number of states, actions
- Simple rules — known transitions (no simulator needed)
- Profound complexity — proof of principle of AI

MiniMax algorithm: brute force

backward induction; Nash equilibrium

The algorithm

1. Deploy the full game tree
2. Apply utility function to terminal states
3. Backward induction
   - On Max ply, assign max. payoff move
   - On Min ply, assign min. payoff move
4. At root, Max selects the move with max payoff.

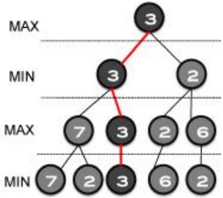

MiniMax algorithm: brute force

backward induction; Nash equilibrium

The algorithm

1. Deploy the full game tree
2. Apply utility function to terminal states
3. Backward induction
   - On Max ply, assign max. payoff move
   - On Min ply, assign min. payoff move
4. At root, Max selects the move with max payoff.

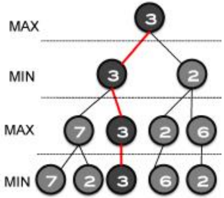

Comments

- Perfect play for deterministic, perfect information games
- Assumes perfect opponent
- Impractical: time and space complexity in $\mathcal{O}(b^d)$

Alpha-Beta: MiniMax with pruning

alphabeta(node, depth, α, β, Player)

- if depth = 0, return H(node)
- if Player = Max
  For each child node,

  $\alpha := max(\alpha, \text{alphabeta(child, depth-1,}\alpha, \beta, \text{not(Player)))}$

  if $\beta \leq \alpha$, cut *beta cut-off*
  return α

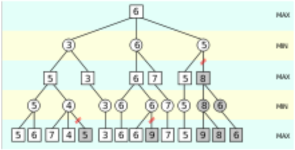

- if Player = Min
  For each child node,

  $\beta := min(\beta, \text{alphabeta(child, depth-1,}\alpha, \beta, \text{not(Player)))}$

  if $\beta \leq \alpha$, cut *alpha cut-off*
  return β

Alpha-Beta: MiniMax with pruning

Comments

- Pruning does not affect final result
- Good move ordering $\rightarrow$ complexity $\mathcal{O}(b^{\frac{d}{2}})$
- Same as $\sqrt{\text{branching factor}}$ for chess: $35 \rightarrow 6$.

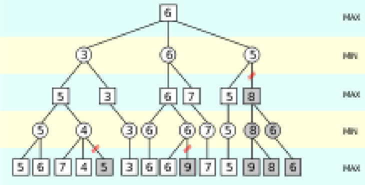

Chess: Deep Blue vs Kasparov

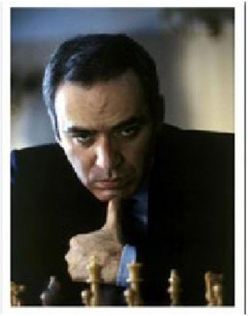

Ingredients

- Brute force; 200 million positions per second
- Look-ahead 12 plies
- Alpha-beta
- Tuning the heuristic function on a game archive
- Branching factor $b \sim 35$ good move ordering $b \sim 6$

Controversy

http://www.slideshare.net/toxygen/kasparov-vs-deep-blue

Dynamic programming

Principle

- Recursively decompose the problem in subproblems
- Solve and propagate

An example

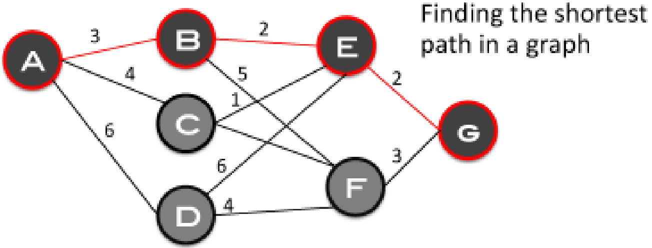

Dynamic programming & Learning

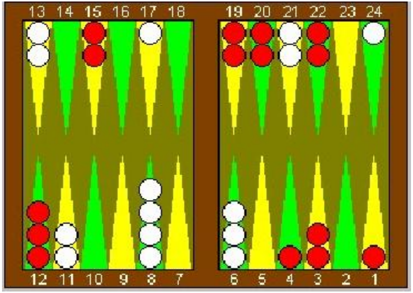

Backgammon Gerald Tesauro, 89-95

- State: raw description of a game (number of White or Black checkers at each location) $\mathbb{R}^D$
- Data: set of games
- A game: sequence of states $x_1, \ldots x_T$; value on last y_T: wins or loses

Dynamic programming & Learning

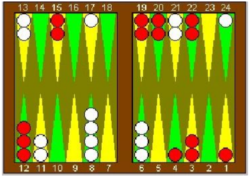

Learning

- Learned: $F : \mathbb{R}^D \mapsto [0,1]$ s.t.

$$\text{Minimize } |F(x_T) - y_T|; \qquad |F(x_\ell) - F(x_{\ell+1})|$$

- Search space: F is a neural net $\equiv w$ $\qquad \mathbb{R}^d$
- Learning rule $\qquad$ 200,000 games

$$\Delta w = \alpha(F(x_{\ell+1}) - F(x_\ell)) \sum_{k=1}^{\ell} \lambda^{\ell-k} \nabla w F(x_k)$$

Overview

La promesse (1960)

Within 10 years, a computer will

- be the world's chess champion
- prove an important theorem in maths
- compose good music
- set up the language for theoretical psychology

L'IA a beaucoup promis

The world's chess champion ?

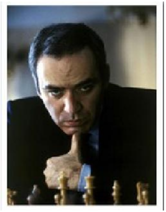

Discussion

Entre intelligence et force brute.

L'IA a beaucoup promis, 2

Prouver un théorème ?

The robot scientist

- Faits $\rightarrow$ Hypothèses
- Hypothèses $\rightarrow$ Expériences
- Expériences $\rightarrow$ Faits

- King R. D., Whelan, K. E., Jones, F. M., Reiser, P. G. K., Bryant, C. H., Muggleton, S., Kell, D. B. and Oliver, S. G. (2004) Functional genomic hypothesis generation and experimentation by a robot scientist. Nature 427 (6971) p247-252
- King R.D., Rowland J., Oliver S.G, Young M., Aubrey W., Byrne E., Liakata M., Markham M., Pir P., Soldatova L., Sparkes A., Whelan K.E., Clare A. (2009). The Automation of Science. Science 324 (5923): 85-89, 3rd April 2009

L'IA a beaucoup promis, 3

Composer de la bonne musique ?

Musac

L'IA a beaucoup promis, 4

Set up the language for theoretical psychology ?

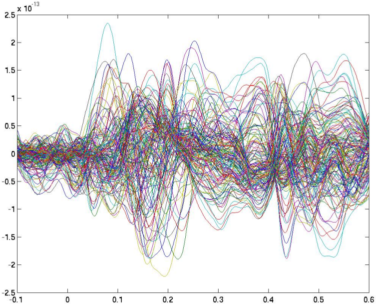

Neuro-imagerie – Interfaces Cerveau-Machine

L'IA a beaucoup promis, 4

Set up the language for theoretical psychology ?

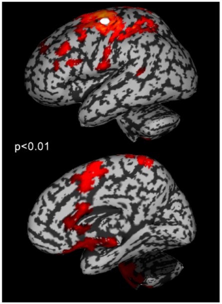

Test d'hypothèses multiples

http://videolectures.net/msht07_baillet_mht/

Overview

AI: The map and the territory

The 2005 DARPA Challenge

AI Agenda: What remains to be done — Thrun 2005

- Reasoning — 10%
- Dialog — 60%
- Perception — 90%

AI: Complete agent principles

Rolf Pfeiffer, Josh Bongard, Max Lungarella,
Jurgen Schmidhuber, Luc Steels, Pierre-Yves Oudeyer...

Situated cognition

Intelligence: not a goal, a means

brains are first and foremost control systems for embodied agents, and their most important job is to help such agents flourish.

Agent's goals: Intelligence is a means of

- Surviving
- Setting and completing self-driven tasks
- Completing prescribed tasks

What are the designer's goals ?

Research modes

Historical AI

- Identify sub-tasks
- Solve them

Bounded rationality

> 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.
>
> Herbert Simon, 1982

Lessons from 50 years

- We need descriptive knowledge: perceptual primitives, patterns, constraints, rules,
- We need control knowledge: policy, adaptation
- Knowledge can hardly be given: must be acquired
- We need interaction knowledge: retrieving new information, feedback

Meta-knowledge J. Pitrat, 2009

- Each goal, a new learning algorithm ?
- Problem reduction ? John Langford, http://hunch.net/

Artificial Intelligence

Search space — **ML**

- Representation — (Un) Supervised L.
- Patterns, Rules, Constraints (knowledge) — (Un) Supervised L., Data Mining
- Navigation policy — Reinforcement L.

Navigation

- Inference — Optimisation

Validation, control, feedback

- Criteria — Statistics

Questions

- Document: Perils and Promises of Big Data
  http://www.thinkbiganalytics.com/uploads/Aspen-Big_Data.pdf
- Quand les données disponibles augmentent
  qu'est-ce qui est différent ?
- Des limitations ?