A historical perspective on Machine Learning

(on the occasion of the 25th Benelearn)

Luc De Raedt

A historical perspective on Machine Learning

(on the occasion of the 25th Benelearn)

Luc De Raedt

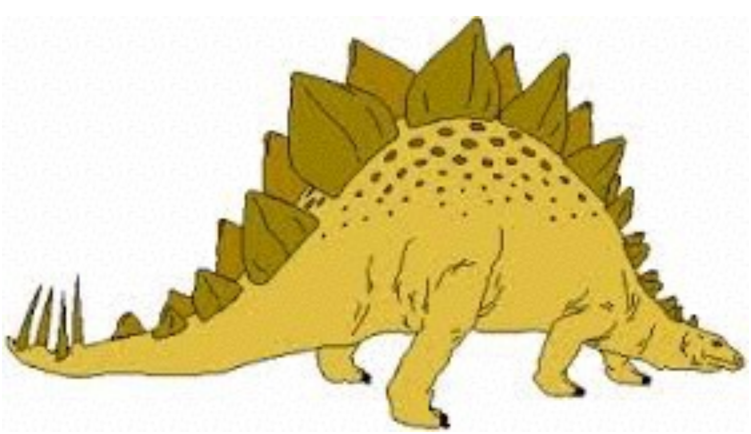

A historical perspective on Machine Learning

(on the occasion of the 25th Benelearn)

Luc De Raedt

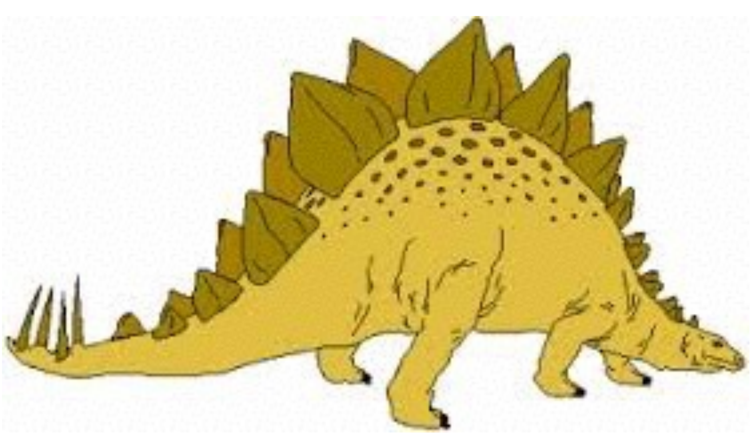

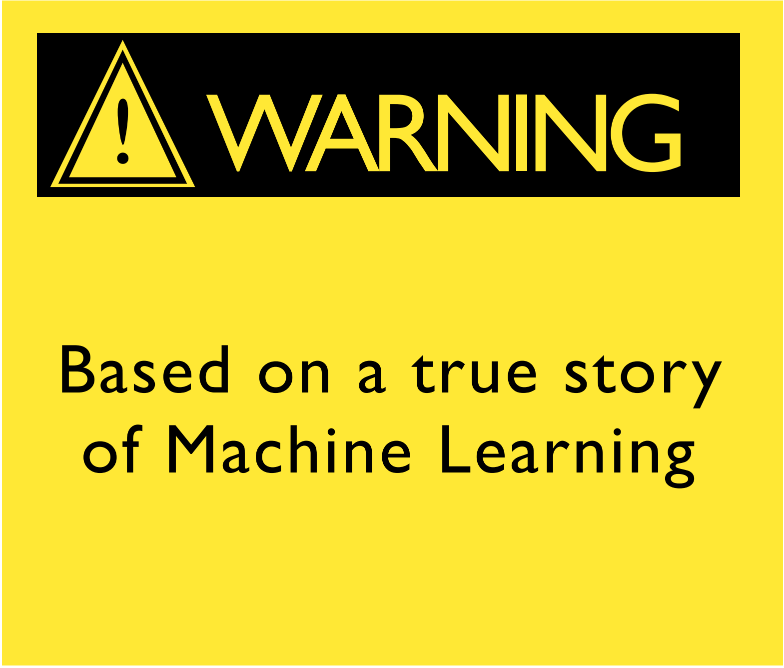

A historical perspective on Machine Learning

(on the occasion of the 25th Benelearn)

Luc De Raedt

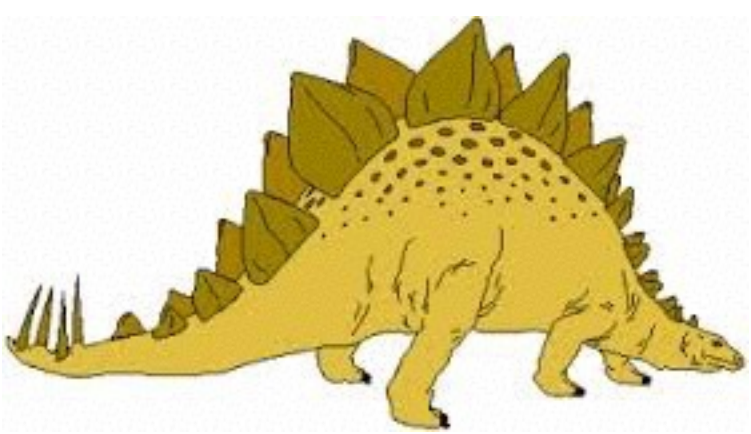

WARNING

Based on a true story of Machine Learning

SCIENTIFIC

ADVISORY

PERSONAL PERSPECTIVE

Machine Learning: an AI approach

Machine Learning: an AI approach

Ryszard Michalski …

Machine Learning: an AI approach

Ryszard Michalski, Tom Mitchell, Jaime Carbonell

Machine Learning: an AI approach

Ryszard Michalski, Tom Mitchell, Jaime Carbonell

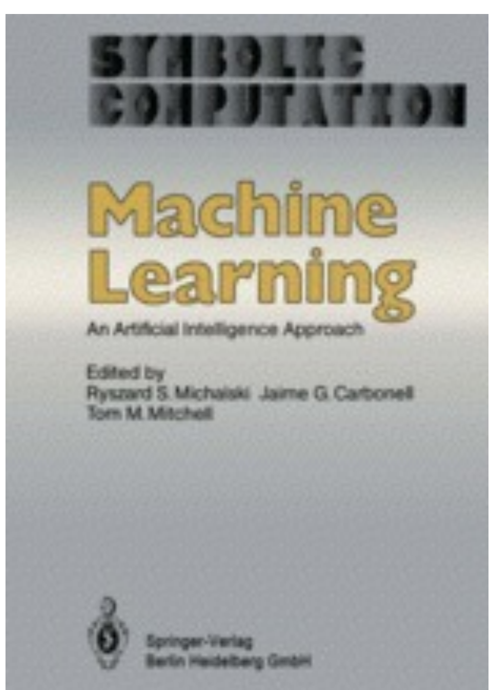

1983

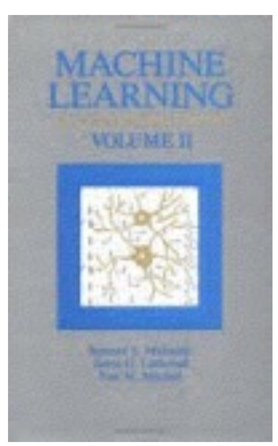

1986

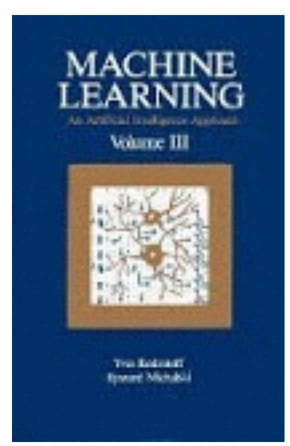

1990

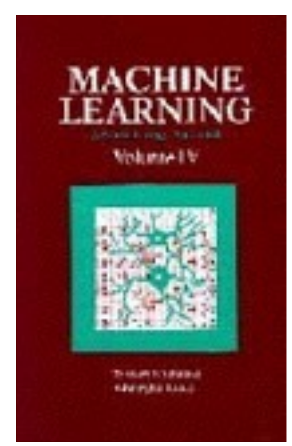

1994

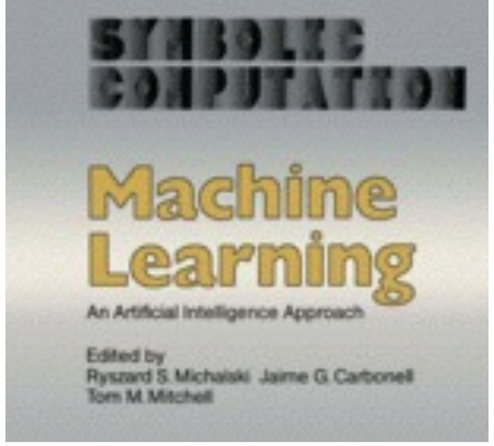

1983

Preface

The ability to learn is one of the most fundamental attributes of intelligent behavior. Consequently, progress in the theory and computer modeling of learning processes is of great significance to fields concerned with understanding intelligence. Such fields include cognitive science, artificial intelligence, information science, pattern recognition, psychology, education, epistemology, philosophy, and related disciplines.

The recent observance of the silver anniversary of artificial intelligence has been heralded by a surge of interest in machine learning—both in building models of human learning and in understanding how machines might be endowed with the ability to learn. This renewed interest has spawned many new research projects and resulted in an increase in related scientific activities. In the summer of 1980, the First Machine Learning Workshop was held at Carnegie-Mellon University in Pittsburgh. In the same year, three consecutive issues of the *International Journal of Policy Analysis and Information Systems* were specially devoted to machine learning (No. 2, 3 and 4, 1980). In the spring of 1981, a special issue of the *SIGART Newsletter* No. 76 reviewed current research projects in the field.

Before 1980 —Handbook of AI 1981 overview

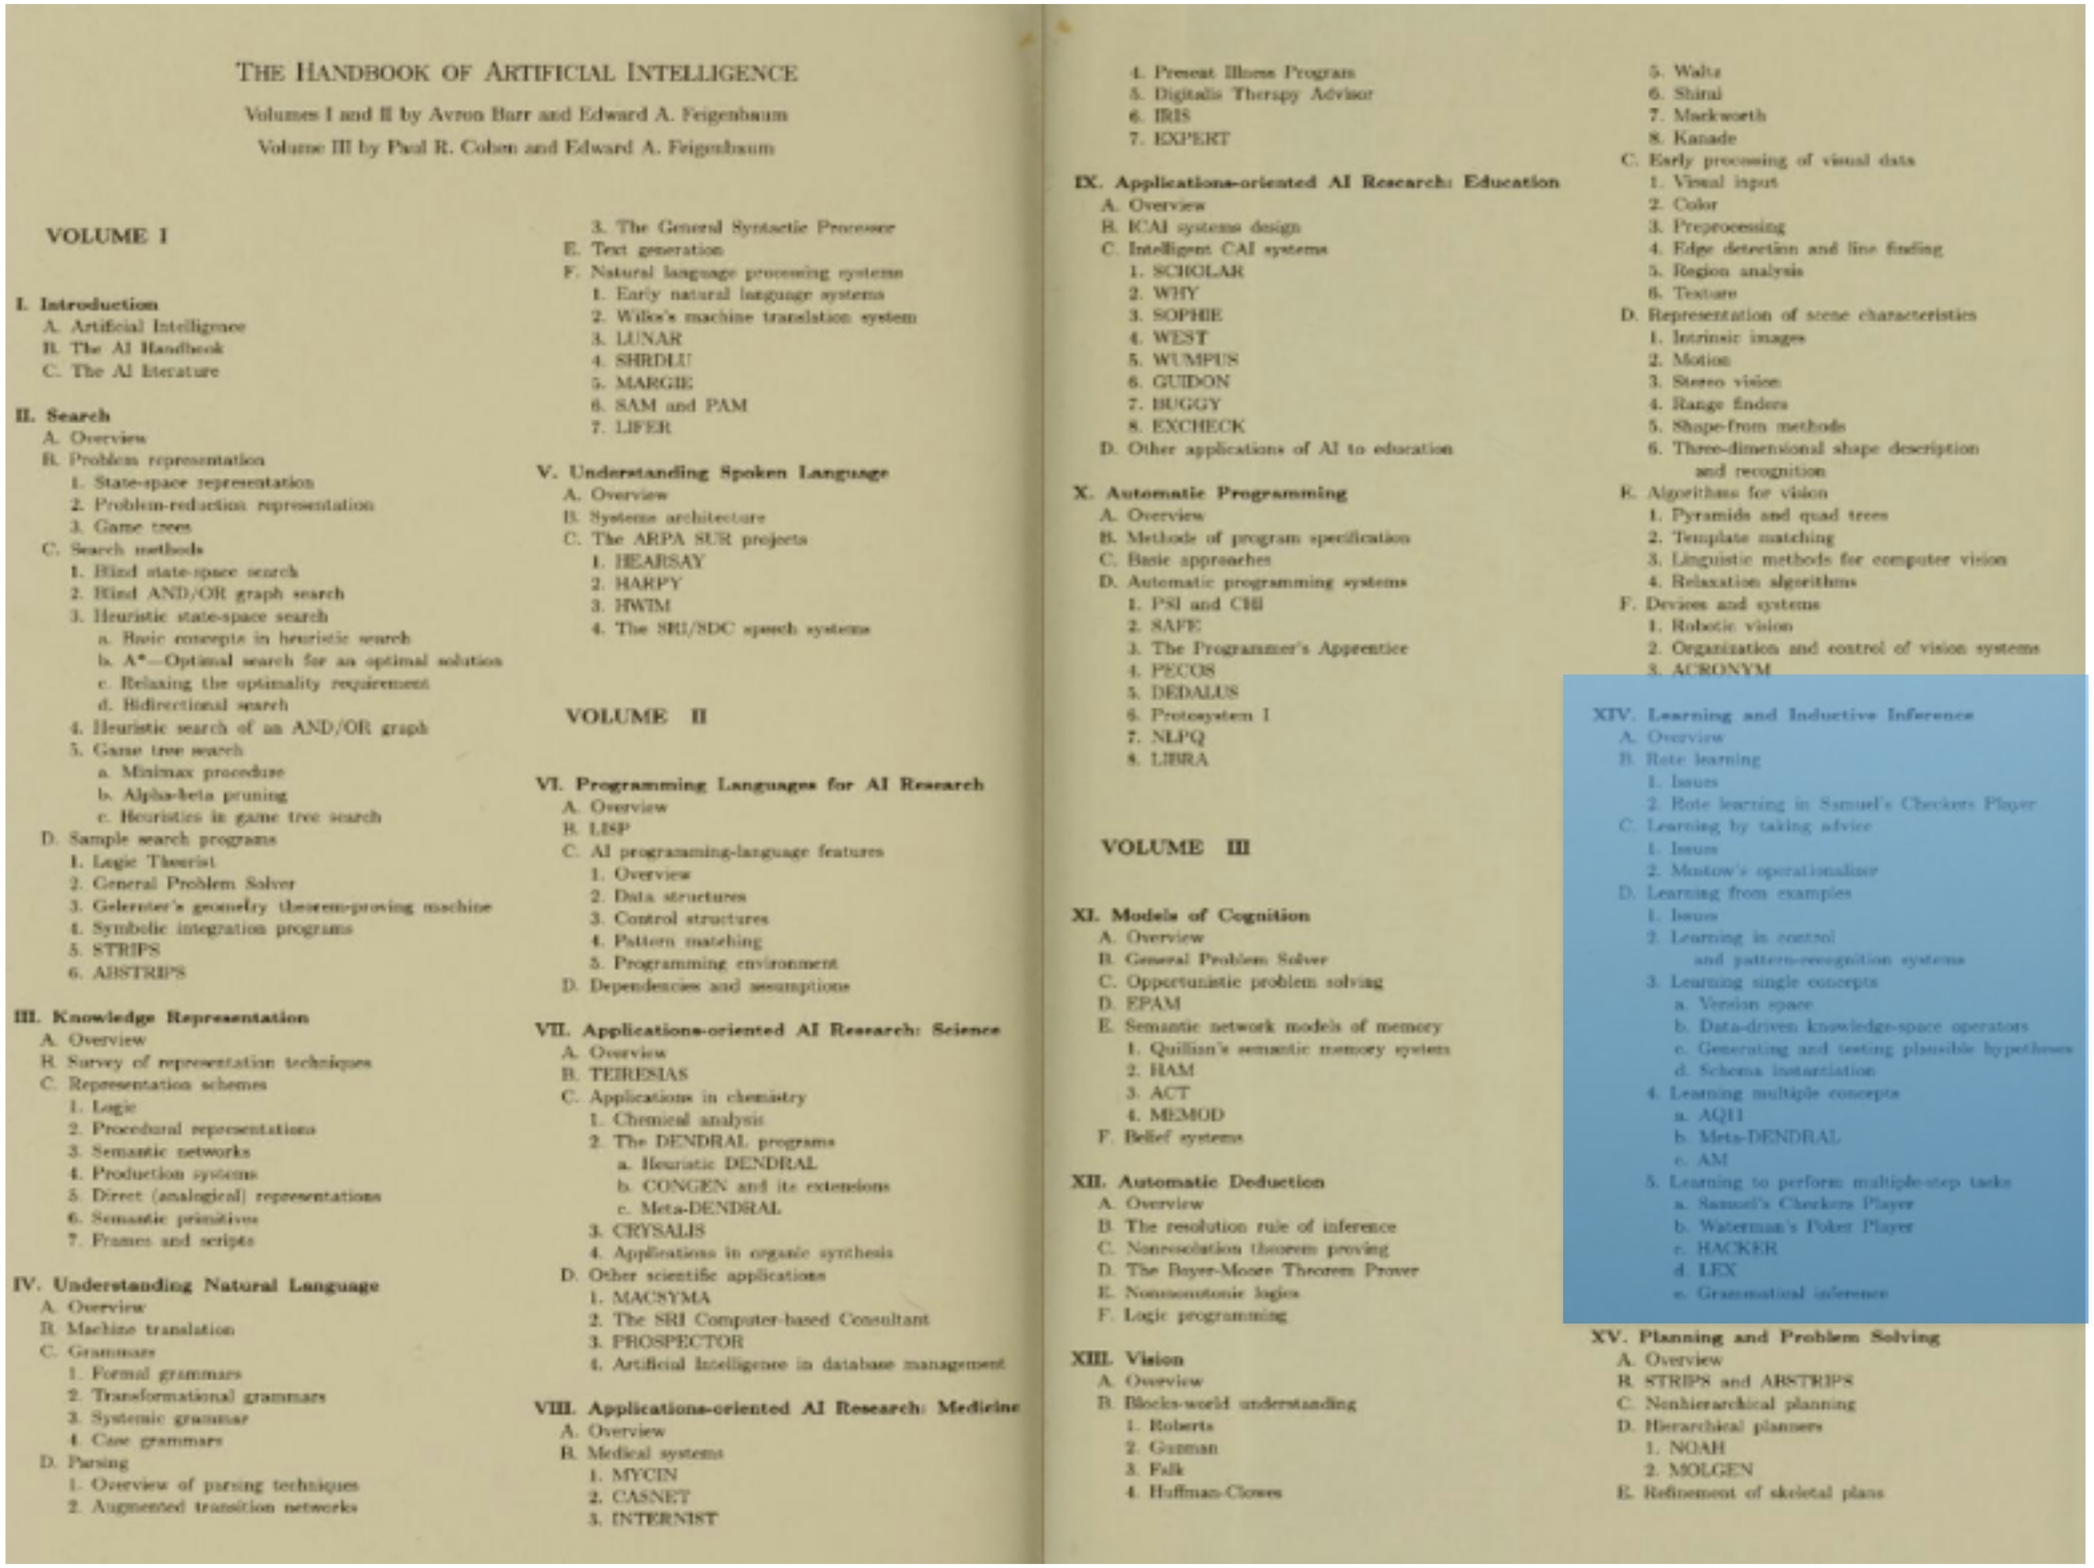

THE HANDBOOK OF ARTIFICIAL INTELLIGENCE

Volumes I and II by Avron Barr and Edward A. Feigenbaum

Volume III by Paul R. Cohen and Edward A. Feigenbaum

VOLUME I

I. Introduction
- A. Artificial Intelligence
- B. The AI Handbook
- C. The AI literature

II. Search
- A. Overview
- B. Problem representation
  - 1. State-space representation
  - 2. Problem-reduction representation
  - 3. Game trees
- C. Search methods
  - 1. Blind state-space search
  - 2. Blind AND/OR graph search
  - 3. Heuristic state-space search
    - a. Basic concepts in heuristic search
    - b. A*—Optimal search for an optimal solution
    - c. Relaxing the optimality requirement
    - d. Bidirectional search
  - 4. Heuristic search of an AND/OR graph
  - 5. Game tree search
    - a. Minimax procedure
    - b. Alpha-beta pruning
    - c. Heuristics in game tree search
- D. Sample search programs
  - 1. Logic Theorist
  - 2. General Problem Solver
  - 3. Gelernter's geometry theorem-proving machine
  - 4. Symbolic integration programs
  - 5. STRIPS
  - 6. ABSTRIPS

III. Knowledge Representation
- A. Overview
- B. Survey of representation techniques
- C. Representation schemes
  - 1. Logic
  - 2. Procedural representations
  - 3. Semantic networks
  - 4. Production systems
  - 5. Direct (analogical) representations
  - 6. Semantic primitives
  - 7. Frames and scripts

IV. Understanding Natural Language
- A. Overview
- B. Machine translation
- C. Grammars
  - 1. Formal grammars
  - 2. Transformational grammars
  - 3. Systemic grammar
  - 4. Case grammars
- D. Parsing
  - 1. Overview of parsing techniques
  - 2. Augmented transition networks
  - 3. The General Syntactic Processor
- E. Text generation
- F. Natural language processing systems
  - 1. Early natural language systems
  - 2. Wilks's machine translation system
  - 3. LUNAR
  - 4. SHRDLU
  - 5. MARGIE
  - 6. SAM and PAM
  - 7. LIFER

V. Understanding Spoken Language
- A. Overview
- B. Systems architecture
- C. The ARPA SUR projects
  - 1. HEARSAY
  - 2. HARPY
  - 3. HWIM
  - 4. The SRI/SDC speech systems

VOLUME II

VI. Programming Languages for AI Research
- A. Overview
- B. LISP
- C. AI programming-language features
  - 1. Overview
  - 2. Data structures
  - 3. Control structures
  - 4. Pattern matching
  - 5. Programming environment
- D. Dependencies and assumptions

VII. Applications-oriented AI Research: Science
- A. Overview
- B. TEIRESIAS
- C. Applications in chemistry
  - 1. Chemical analysis
  - 2. The DENDRAL programs
    - a. Heuristic DENDRAL
    - b. CONGEN and its extensions
    - c. Meta-DENDRAL
  - 3. CRYSALIS
  - 4. Applications in organic synthesis
- D. Other scientific applications
  - 1. MACSYMA
  - 2. The SRI Computer-based Consultant
  - 3. PROSPECTOR
  - 4. Artificial Intelligence in database management

VIII. Applications-oriented AI Research: Medicine
- A. Overview
- B. Medical systems
  - 1. MYCIN
  - 2. CASNET
  - 3. INTERNIST
  - 4. Present Illness Program
  - 5. Digitalis Therapy Advisor
  - 6. IRIS
  - 7. EXPERT

IX. Applications-oriented AI Research: Education
- A. Overview
- B. ICAI systems design
- C. Intelligent CAI systems
  - 1. SCHOLAR
  - 2. WHY
  - 3. SOPHIE
  - 4. WEST
  - 5. WUMPUS
  - 6. GUIDON
  - 7. BUGGY
  - 8. EXCHECK
- D. Other applications of AI to education

X. Automatic Programming
- A. Overview
- B. Methods of program specification
- C. Basic approaches
- D. Automatic programming systems
  - 1. PSI and CHI
  - 2. SAFE
  - 3. The Programmer's Apprentice
  - 4. PECOS
  - 5. DEDALUS
  - 6. Protosystem I
  - 7. NLPQ
  - 8. LIBRA

VOLUME III

XI. Models of Cognition
- A. Overview
- B. General Problem Solver
- C. Opportunistic problem solving
- D. EPAM
- E. Semantic network models of memory
  - 1. Quillian's semantic memory system
  - 2. HAM
  - 3. ACT
  - 4. MEMOD
- F. Belief systems

XII. Automatic Deduction
- A. Overview
- B. The resolution rule of inference
- C. Nonresolution theorem proving
- D. The Boyer-Moore Theorem Prover
- E. Nonmonotonic logics
- F. Logic programming

XIII. Vision
- A. Overview
- B. Blocks-world understanding
  - 1. Roberts
  - 2. Guzman
  - 3. Falk
  - 4. Huffman-Clowes
  - 5. Waltz
  - 6. Shirai
  - 7. Mackworth
  - 8. Kanade
- C. Early processing of visual data
  - 1. Visual input
  - 2. Color
  - 3. Preprocessing
  - 4. Edge detection and line finding
  - 5. Region analysis
  - 6. Texture
- D. Representation of scene characteristics
  - 1. Intrinsic images
  - 2. Motion
  - 3. Stereo vision
  - 4. Range finders
  - 5. Shape-from methods
  - 6. Three-dimensional shape description and recognition
- E. Algorithms for vision
  - 1. Pyramids and quad trees
  - 2. Template matching
  - 3. Linguistic methods for computer vision
  - 4. Relaxation algorithms
- F. Devices and systems
  - 1. Robotic vision
  - 2. Organization and control of vision systems
  - 3. ACRONYM

XIV. Learning and Inductive Inference
- A. Overview
- B. Rote learning
  - 1. Issues
  - 2. Rote learning in Samuel's Checkers Player
- C. Learning by taking advice
  - 1. Issues
  - 2. Mostow's operationalizer
- D. Learning from examples
  - 1. Issues
  - 2. Learning in control and pattern-recognition systems
  - 3. Learning single concepts
    - a. Version space
    - b. Data-driven knowledge-space operators
    - c. Generating and testing plausible hypotheses
    - d. Schema instantiation
  - 4. Learning multiple concepts
    - a. AQ11
    - b. Meta-DENDRAL
    - c. AM
  - 5. Learning to perform multiple-step tasks
    - a. Samuel's Checkers Player
    - b. Waterman's Poker Player
    - c. HACKER
    - d. LEX
    - e. Grammatical inference

XV. Planning and Problem Solving
- A. Overview
- B. STRIPS and ABSTRIPS
- C. Nonhierarchical planning
- D. Hierarchical planners
  - 1. NOAH
  - 2. MOLGEN
- E. Refinement of skeletal plans

https://archive.org/details/handbookofartific02barr

Before 1980 —Handbook of AI 1981 overview

7. BUGGY
8. EXCHECK

D. Other applications of AI to education

X. Automatic Programming

A. Overview
B. Methods of program specification
C. Basic approaches
D. Automatic programming systems
1. PSI and CHI
2. SAFE
3. The Programmer's Apprentice
4. PECOS
5. DEDALUS
6. Protosystem I
7. NLPQ
8. LIBRA

VOLUME III

XI. Models of Cognition

A. Overview
B. General Problem Solver
C. Opportunistic problem solving
D. EPAM
E. Semantic network models of memory
1. Quillian's semantic memory system
2. HAM
3. ACT
4. MEMOD
F. Belief systems

XII. Automatic Deduction

A. Overview
B. The resolution rule of inference
C. Nonresolution theorem proving
D. The Boyer-Moore Theorem Prover
E. Nonmonotonic logics
F. Logic programming

XIII. Vision

A. Overview
B. Blocks-world understanding
1. Roberts
2. Guzman
3. Falk
4. Huffman-Clowes

4. Range finders
5. Shape-from methods
6. Three-dimensional shape description

1. Pyramids and quad trees
2. Template matching
3. Linguistic methods for computer vision
4. Relaxation algorithms

F. Devices and systems
1. Robotic vision
2. Organization and control of vision systems
3. ACRONYM

XIV. Learning and Inductive Inference

A. Overview
B. Rote learning
1. Issues
2. Rote learning in Samuel's Checkers Player
C. Learning by taking advice
1. Issues
2. Mostow's operationalizer
D. Learning from examples
1. Issues
2. Learning in control and pattern-recognition systems
3. Learning single concepts
a. Version space
b. Data-driven knowledge-space operators
c. Generating and testing plausible hypotheses
d. Schema instantiation
4. Learning multiple concepts
a. AQ11
b. Meta-DENDRAL
c. AM
5. Learning to perform multiple-step tasks
a. Samuel's Checkers Player
b. Waterman's Poker Player
c. HACKER
d. LEX
e. Grammatical inference

XV. Planning and Problem Solving

A. Overview
B. STRIPS and ABSTRIPS
C. Nonhierarchical planning
D. Hierarchical planners
1. NOAH
2. MOLGEN
E. Refinement of skeletal plans

https://archive.org/details/handbookofartific02barr

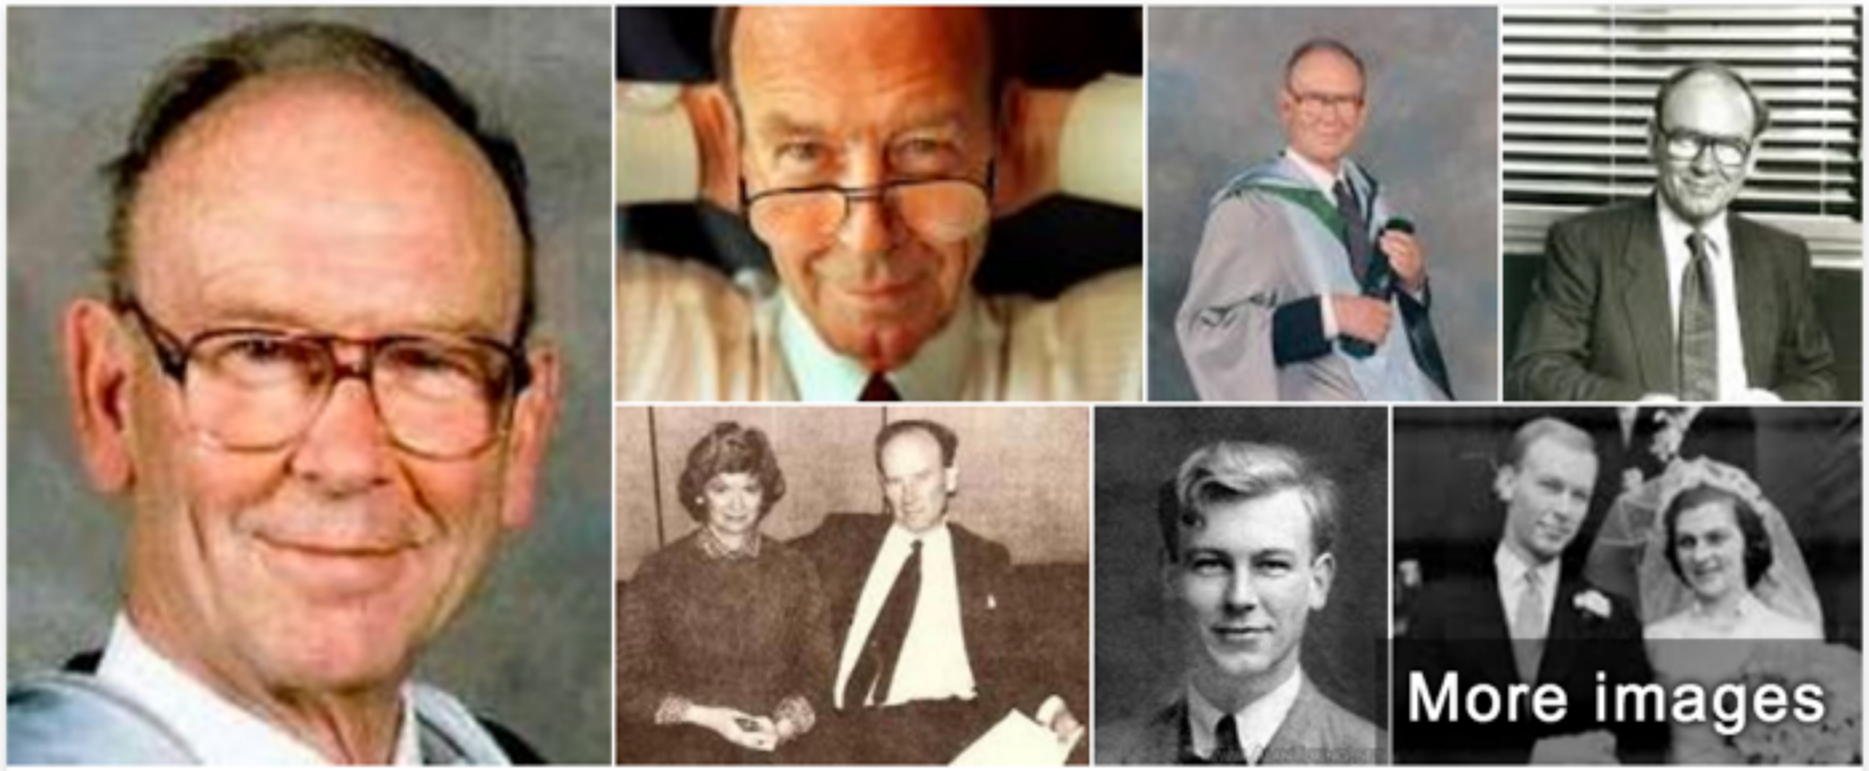

Donald Michie

Donald Michie was a British researcher in artificial intelligence. During World War II, Michie worked for the Government Code and Cypher School at Bletchley Park, contributing to the effort to solve "Tunny," a German teleprinter cipher. Wikipedia

Born: November 11, 1923, Yangon, Myanmar (Burma)

Died: July 7, 2007, England, United Kingdom

Children: Jonathan Michie, Susan Michie

Books: Donald Michie: Machine Intelligence, Biology and More, more

Education: Rugby School, Balliol College

Awards: IJCAI Award for Research Excellence

Menace (Michie 63)

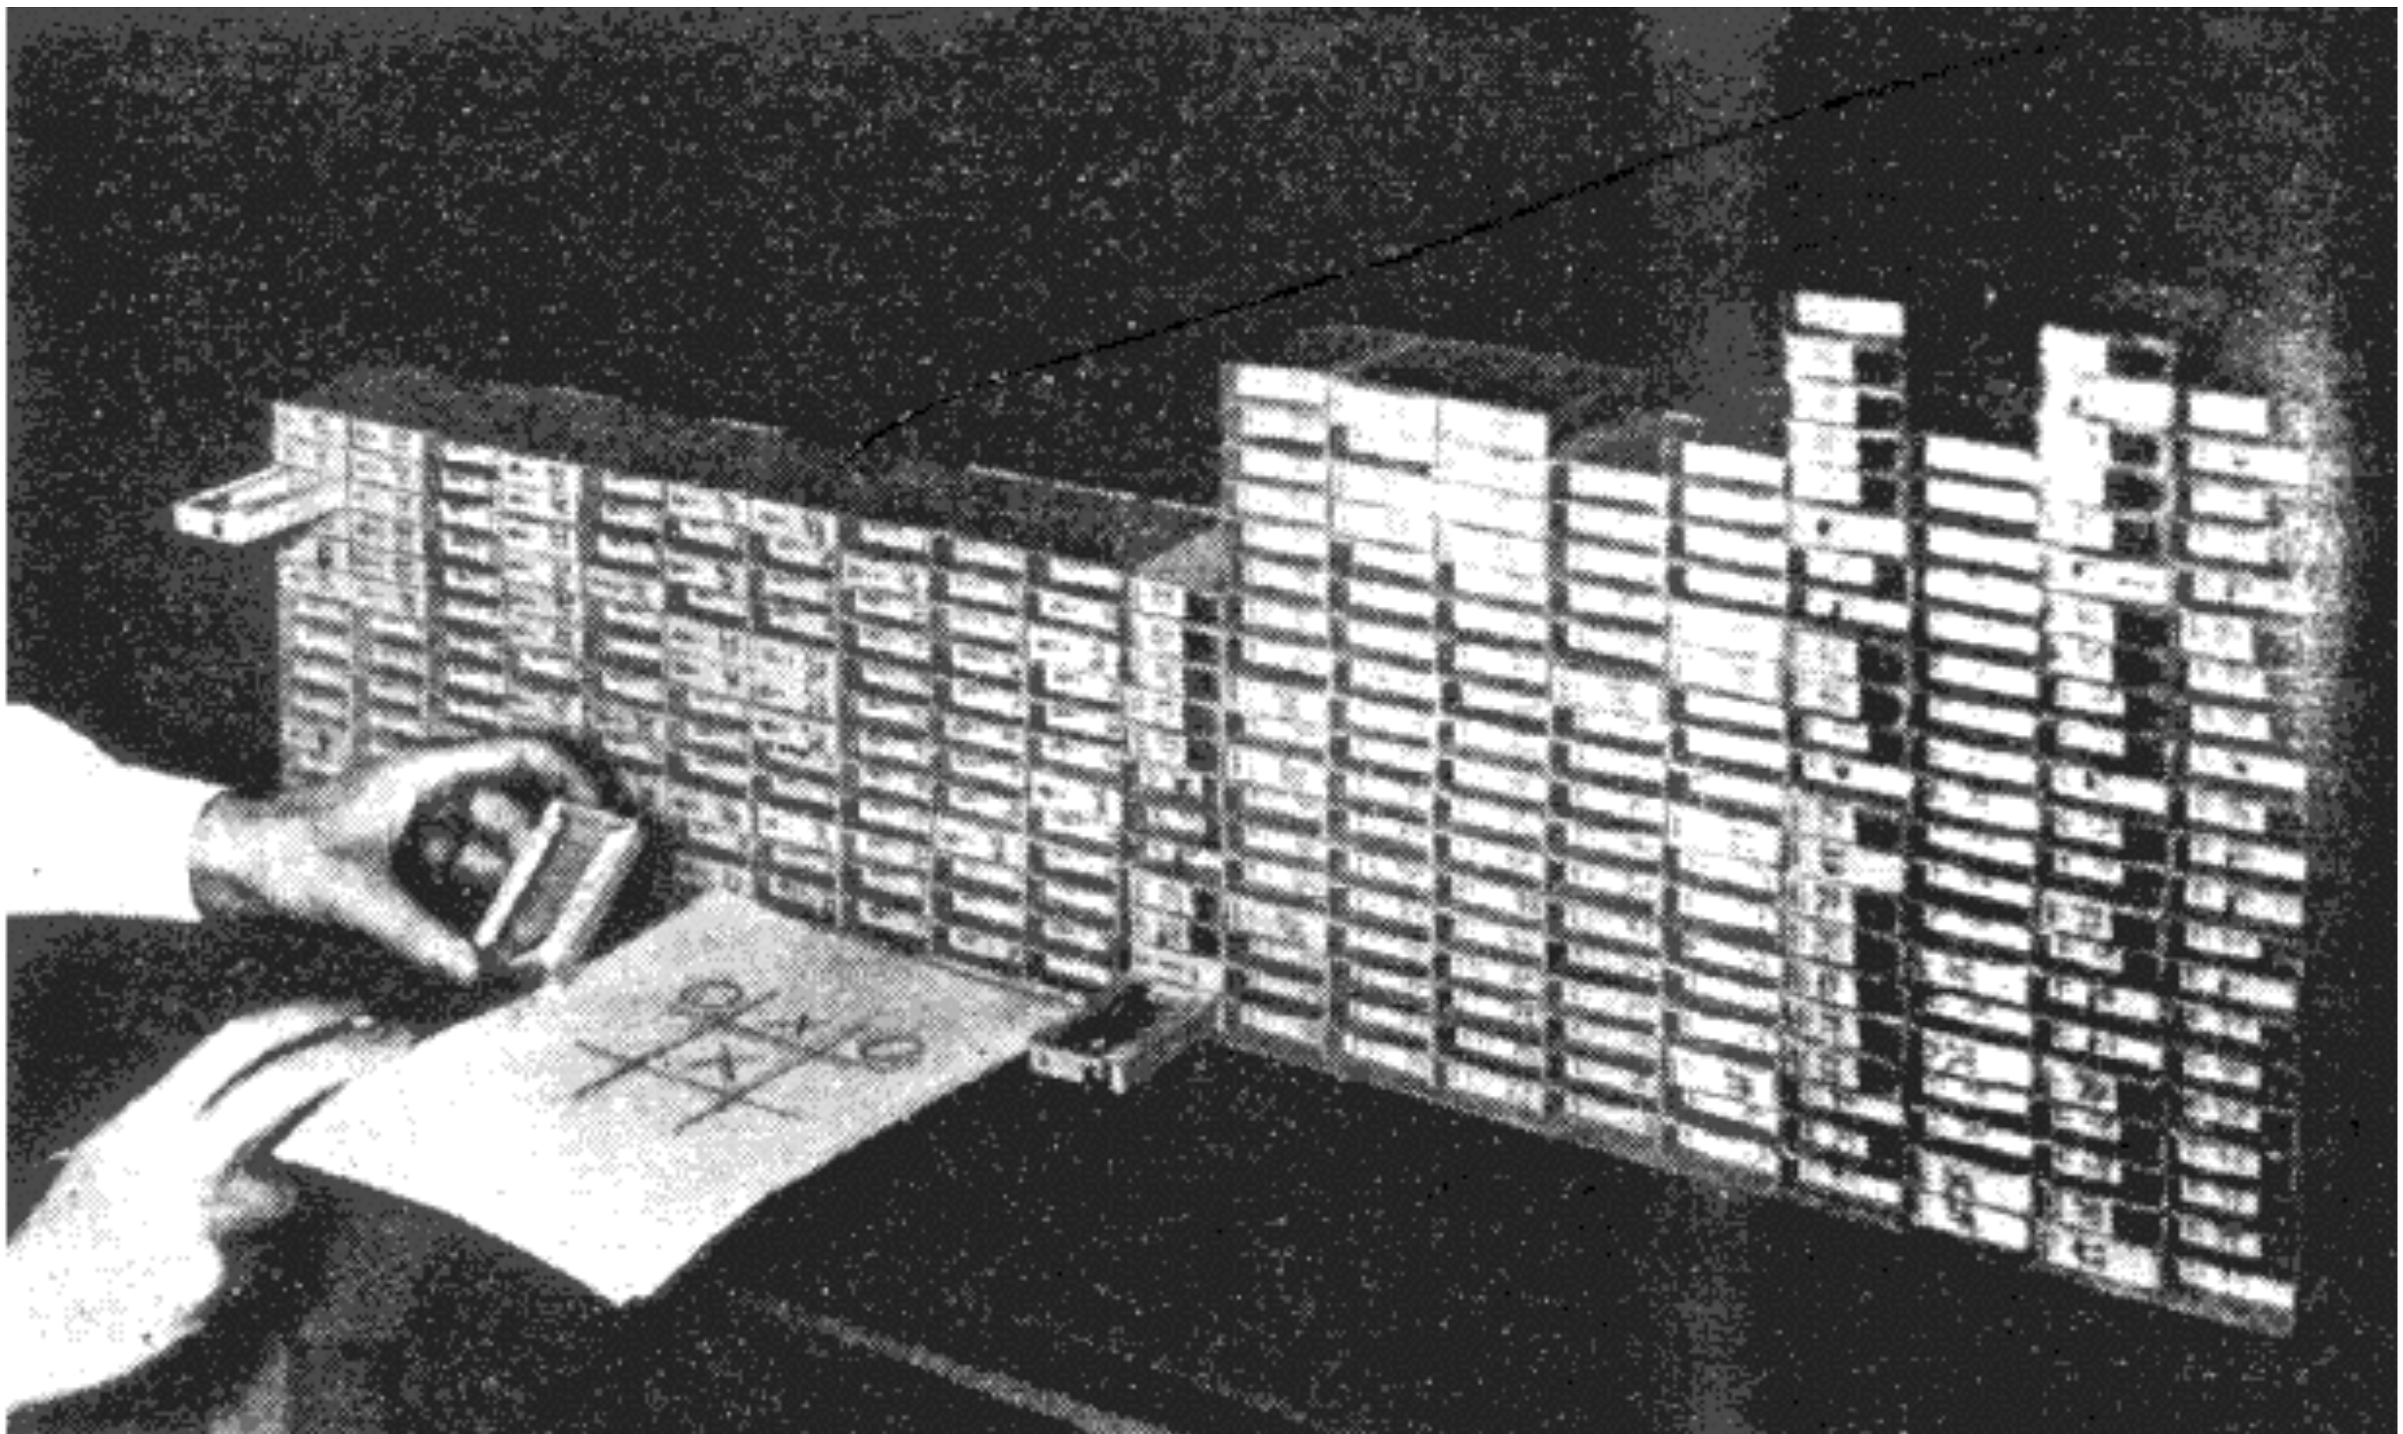

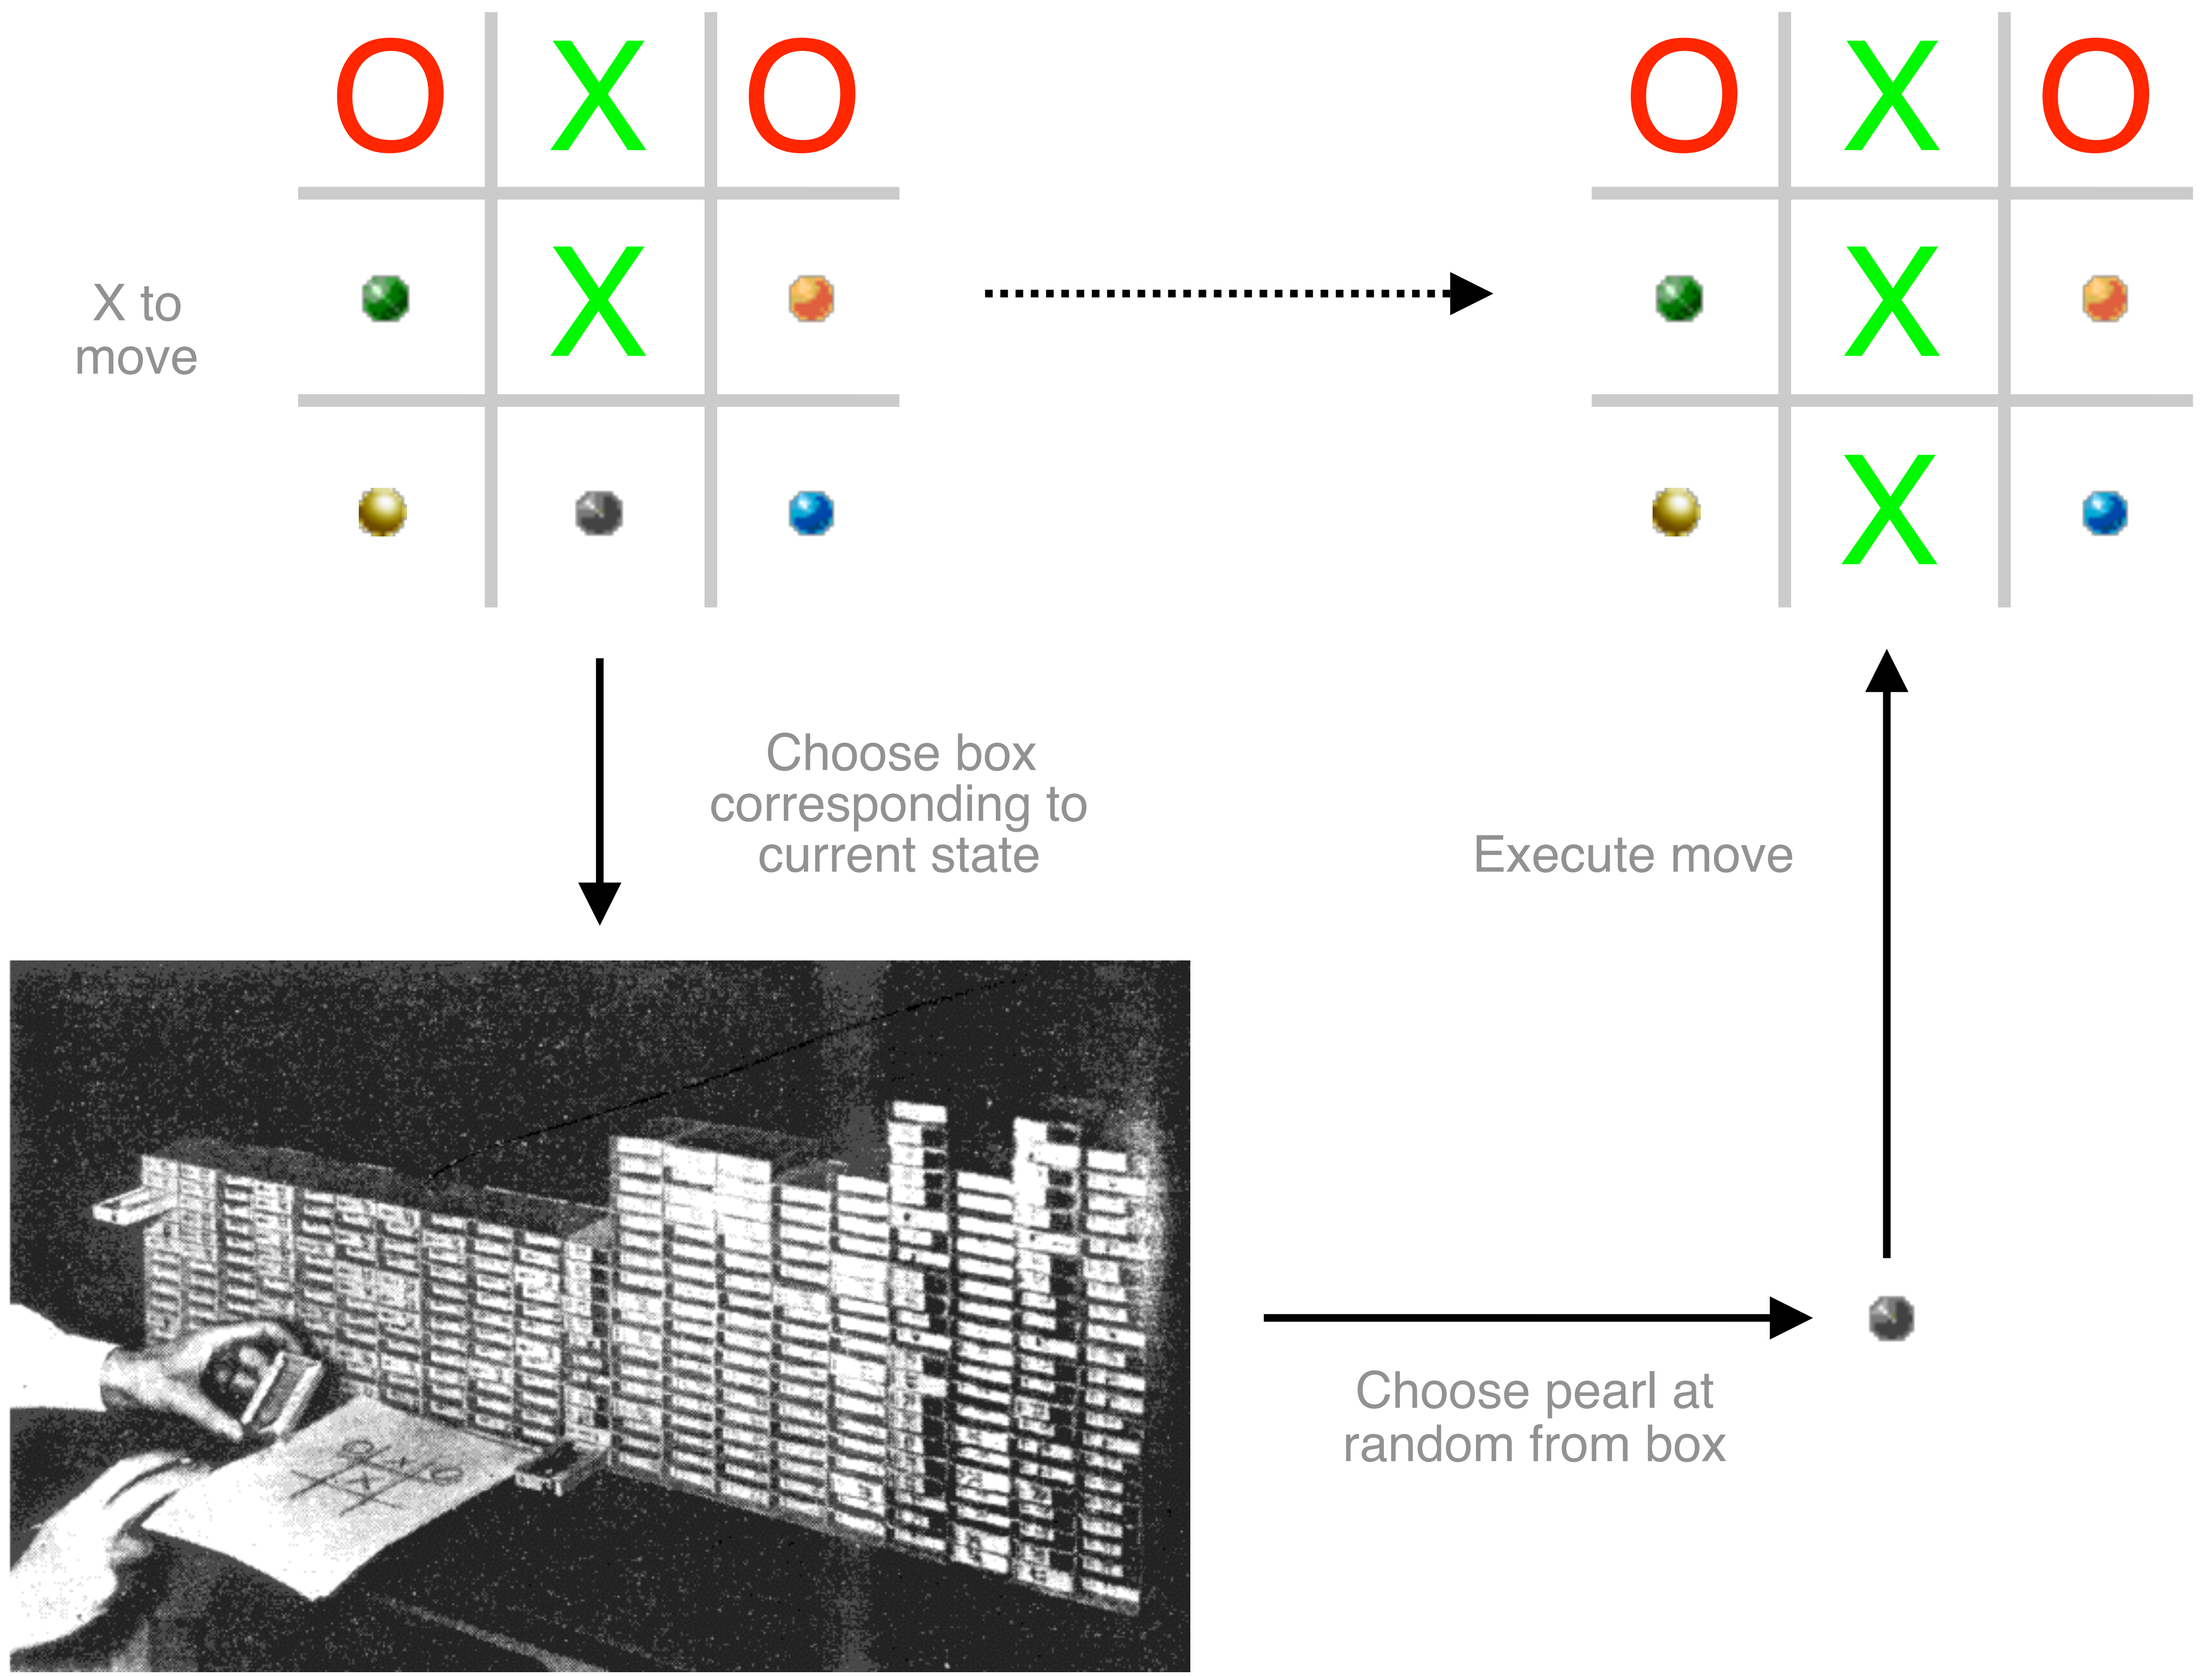
X to move
Choose box corresponding to current state
Execute move
Choose pearl at random from box

Menace (Michie 1963)

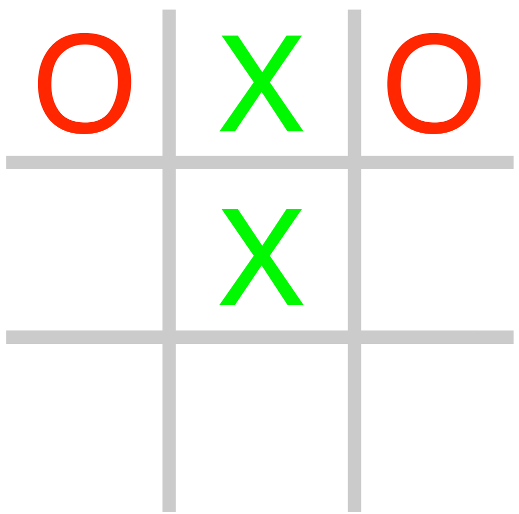

Learns Tic-Tac-Toe

Hardware:

287 Boxes
(1 for each state)

Pearls in 9 colors
(1 color per position)

Play principle:

Choose box corresponding to current state

Choose pearl at random from box

Play corresponding move

Learning algorithm:

Game lost -> retain all pearls used
(*negative reword - reinforcement)*

Game won -> for each select pearl, add a pearl of the same color to box
(*positive reward - reinforcement*)

BOXES (1968)

https://www.youtube.com/watch?v=qF2fFMrNUCQ

- basis of reinforcement learning

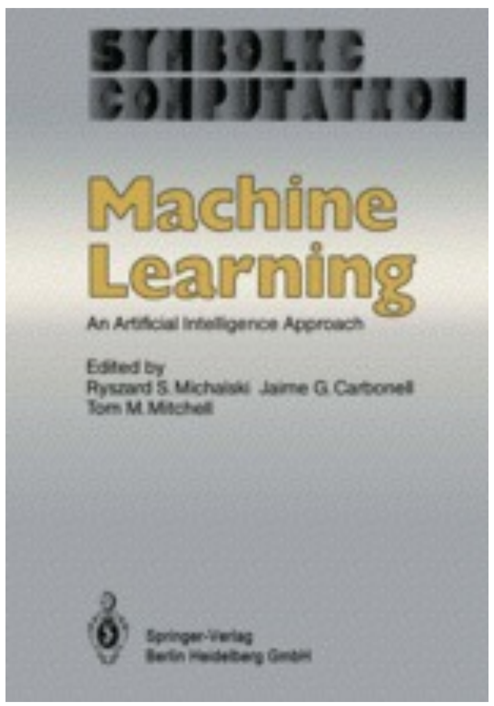
SYMBOLIC
COMPUTATION
Machine
Learning
An Artificial Intelligence Approach
Edited by
Ryszard S. Michalski Jaime G. Carbonell
Tom M. Mitchell
Springer-Verlag
Berlin Heidelberg GmbH

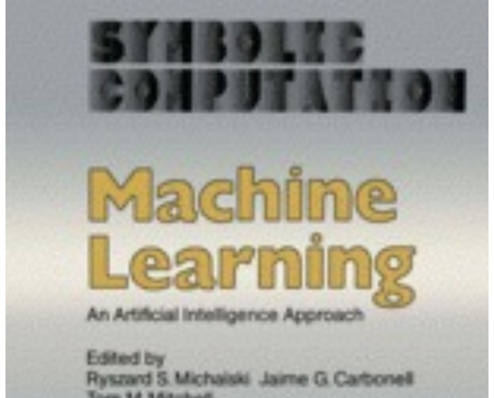

Why should machines learn ?

Herbert Simon (1916-2001)

Turing Award 1975, Nobel prize Economics 1978

The only partially satisfactory definition I've been able to find is that learning is any change in a system that allows it to perform better the second time on repetition of the same task or on another task drawn from the same population. The change should be more or less irreversible—not irreversible in the sense that you can't unlearn (although that sometimes is hard, especially unlearning bad habits) but irreversible in that the learning doesn't go away rapidly and autonomously. *Learning denotes changes in the system that are adaptive in the sense that they enable the system to do the same task or tasks drawn from the same population more efficiently and more effectively the next time.*

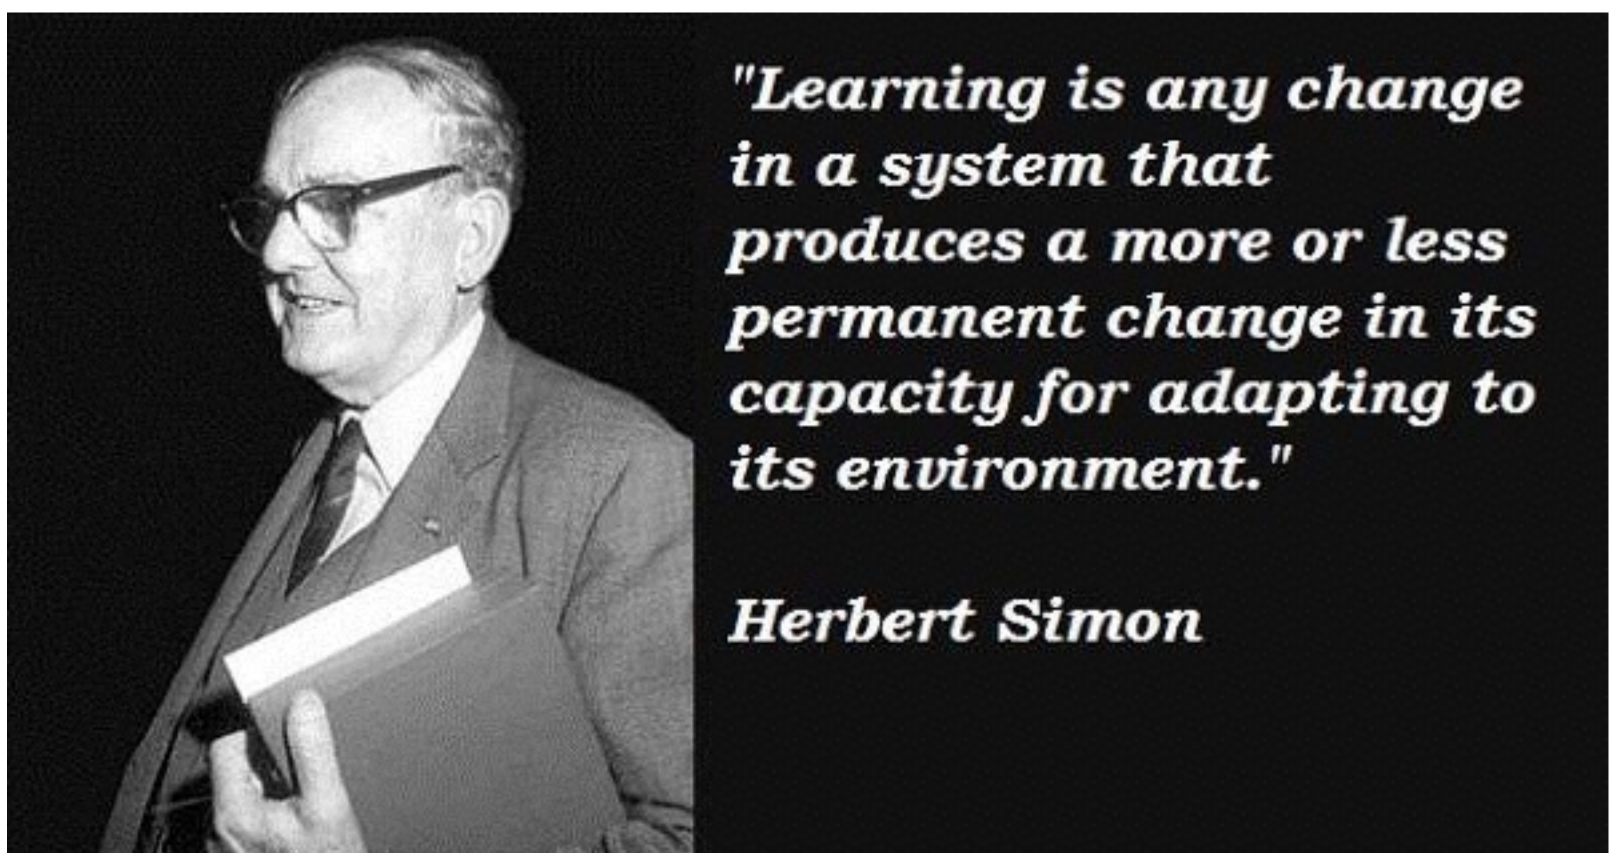

1. I would give a very high priority to research aimed at simulating, and thereby understanding, human learning. It may be objected that such research is not AI but cognitive psychology or cognitive science or something else. I don't really care what it is called; it is of the greatest importance that we deepen our understanding of human learning, and the AI community possesses a large share of the talent that can advance us toward this goal.
2. I would give a high priority, also, to basic research aimed at understanding why human learning is so slow and inefficient, and correspondingly, at examining the possibility that machine learning schemes can be devised that will avoid, for machines as well as people, some of the tediousness of learning.
3. I would give a high priority to research on the natural language interface between computer systems and human users. Again, it does not matter whether you call it research on learning or research on understanding. We do want systems, particularly in the knowledge engineering area, in which we don't have to know the internal language or representation in order to interact with them. This is especially true, as I have just argued, if the systems are to be cumulative over many years.
4. I think there is an important place for research on programming from incomplete instructions (automatic programming), which is not unrelated to the preceding item. Giving instructions to a skilled programmer is different from writing the program yourself—else why hire the programmer? It is a very important research question to ask whether we can get the computer to be the skilled programmer.
5. My final priority is research on discovery programs—programs that discover new things. We may regard discovery itself as a form of learning, but in addition we will want to give a discovery system learning capabilities because we will want it to preserve and to be able to use all the new things it finds.

1980 … 1986

- First workshops on Machine Learning (first conference in 1993)
- Focus on AI and Cognitive Science paradigm
- Focus on SYMBOLIC Methods, on HUMAN like learning, on AUTOMATED DISCOVERY
- IJCAI 85 in LA had 3000 academic participants (10 000 with industry included?) These were the days of expert systems
- No role for SUBSYMBOLIC methods / NEURAL NETS
- NIPS would start in 1986, with the revival of Neural Networks (Parallel Distributed Processing / Connectionism — Rumelhart and McClelland)
- https://www.youtube.com/watch?v=ilP4aPDTBPE (1989)

From the Dartmouth 1956 proposal

The following are some aspects of the artificial intelligence problem:

1. **Automatic Computers** : If a machine can do a job, then an automatic calculator can be programmed to simulate the machine. The speeds and memory capacities of present computers may be insufficient to simulate many of the higher functions of the human brain, but the major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.

2. **How Can a Computer be Programmed to Use a Language :** It may be speculated that a large part of human thought consists of manipulating words according to rules of reasoning and rules of conjecture. From this point of view, forming a generalization consists of admitting a new word and some rules whereby sentences containing it imply and are implied by others. This idea has never been very precisely formulated nor have examples been worked out.

3. **Neuron Nets :** How can a set of (hypothetical) neurons be arranged so as to form concepts. Considerable theoretical and experimental work has been done on this problem by Uttley, Rashevsky and his group, Farley and Clark, Pitts and McCulloch, Minsky, Rochester and Holland, and others. Partial results have been obtained but the problem needs more theoretical work.

4. **Theory of the Size of a Calculation :** If we are given a well-defined problem (one for which it is possible to test mechanically whether or not a proposed answer is a valid answer) one way of solving it is to try all possible answers in order. This method is inefficient, and to exclude it one must have some criterion for efficiency of calculation. Some consideration will show that to get a measure of the efficiency of a calculation it is necessary to have on hand a method of measuring the complexity of calculating devices which in turn can be done if one has a theory of the complexity of functions. Some partial results on this problem have been obtained by Shannon, and also by McCarthy.

5. **Self-Improvement** Probably a truly intelligent machine will carry out activities which may best be described as self-improvement. Some schemes for doing this have been proposed and are worth further study. It seems likely that this question can be studied abstractly as well.

6. **Abstractions** A number of types of ``abstraction'' can be distinctly defined and several others less distinctly. A direct attempt to classify these and to describe machine methods of forming abstractions from sensory and other data would seem worthwhile.

7. **Randomness and Creativity** A fairly attractive and yet clearly incomplete conjecture is that the difference between creative thinking and unimaginative competent thinking lies in the injection of a some randomness. The randomness must be guided by intuition to be efficient. In other words, the educated guess or the hunch include controlled randomness in otherwise orderly thinking.

1986

- The Machine Learning Journal was founded by **Pat Langley,** Ryszard Michalski, Jaime Carbonnell and Tom M. Mitchell
  - find an own venue for ML research …
  - same focus / bias initially (cf. Langley, MLJ 2011)
- The 1st European Working Session on ML was organised in Orsay by Yves Kodratoff

Video
0;58-3:08 +
22:25-25:43 ? +
44:05-45

Ryszard S. Michalski

Ryszard S. Michalski was a Polish-American computer scientist. Michalski was Professor at George Mason University and a pioneer in the field of machine learning.
Wikipedia

Born: May 7, 1937, Kalush, Ukraine

Died: September 20, 2007, Fairfax, Virginia, United States

Fields: Machine learning, Artificial intelligence

1989

- The 1st KDD workshop was organised at IJCAI 1989 in Detroit, attended by 67 participants (among which most of the key players in ML and KDD …)

- Panel with Ross Quinlan, Pat Langley, and Larry Kerschberg

- Donald Michie predicts that ``*The next area that is going to explode is the use of machine learning tools as a component of large scale data analysis''* (AI Week, March 15, 1990)

- First KDD conference 1995

Call for Participation:
IJCAI-89 Workshop on Knowledge Discovery in Databases
Sunday, August 20 (tentative), Detroit MI, USA

The growth in the amount of available databases far outstrips the growth of corresponding knowledge. This creates both a need and an opportunity for extracting knowledge from databases. Many recent results have been reported on extracting different kinds of knowledge from databases, including diagnostic rules, drug side effects, classes of stars, rules for expert systems, and rules for semantic query optimization.

Knowledge discovery in databases poses many interesting problems, especially when databases are large. Such databases are usually accompanied by substantial domain knowledge which can significantly facilitate discovery. Access to large databases is expensive - hence the need for sampling and other statistical methods. Finally, knowledge discovery in databases can benefit from many available tools and techniques from several different fields including expert systems, machine learning, intelligent databases, knowledge acquisition, and statistics.

Topics of interest include:

- Discovery and use of approximate rules
- Knowledge-based discovery methods
- Integration of knowledge-based and statistical methods
- Efficient heuristic algorithms for discovery
- Automatic knowledge acquisition
- Construction of expert systems from data
- Discovery in medical and scientific data
- Bias for human understandability of discovered knowledge
- Learning query optimization rules and integrity constraints
- Knowledge discovery as a threat to database security and privacy

“Donald Michie”

12:20 / 21:20 / 31:54

Benelearn 1991

- 1st Benelearn in Leuven, 70 participants, 10 talks, lunches paid by FWO, Invited talks by Yves Kodratoff and Katharina Morik
- 1992, Amsterdam (Van Someren)
- 1993, Brussels (Van de Velde)
- 1994, Rotterdam (Bioch) - 27 presentations !
- 1995, Brussels (ULB)
- 1996, Maastricht
- 1997, Tilburg
- 1998, Wageningen …

Learning and Knowledge

- Explicit goal to learn new “knowledge”, focus on results that are “understandable”
- Explicit goal to reason with that knowledge (eg. in problem solving)
- Explicit goal to learn rich representations, to learn for use in expert systems
- Cf. e.g. Dietterich, MLJ 86, Michalski’s trains etc.

Machine learning

- Was initially broadening its scope from purely symbolic and knowledge based to
  - probabilistic methods
  - reinforcement learning
  - case-based reasoning and instance based learning
  - problem solving …
  - was getting a diverse and open-minded field !

cf Langley, MLJ 2011

ML as an experimental science

- The title of an MLJ paper by Kibler and Langley in 1988; see also Langley MLJ 2011
- Presented also as keynote as EWSL 88
- Point of view that ML systems should not be just systems that do something, but should be evaluated according to scientific principles, through setting up experiments in a systematic way
- UCI Database
- Side-effect : focus on tasks that are easy to evaluate, on particular datasets, on classification and regression …
- Another side-effect: you have to beat the competition …

cf Langley, MLJ 2011

The Experimental Study of Machine Learning

Pat Langley (Langley@ptolemy.arc.nasa.gov)
AI Research Branch, Mail Stop 244–17, NASA Ames Research Center, Moffett Field, CA 94035 USA

Dennis Kibler (Kibler@ics.uci.edu)
Department of Information & Computer Science, University of California, Irvine, CA 92717 USA

learning behaviors.

As normally defined, an *experiment* involves systematically varying one or more *independent* variables and examining their effect on some *dependent* variables. Thus, a machine learning experiment requires more than a single observation of a system's behavior; it requires a number of observations made under different conditions. In each case, one must measure some aspect of the system's behavior for comparison across the different conditions.

6 phases :

1. Formulating Hypotheses
2. Design experiments and select Samples
3. Running experiments and compile results
4. Test hypotheses
5. Explain unexpected results
6. Report

Introduction of SVMs

- Around 1992-95 by Vapnik, Cortes et al.
- Enormous boost in performance
- Principled theory, interesting mathematics coming from a new community (physics, optimisation…)
- But also had a profound influence on the nature of machine learning
- Side-effect — shift of focus of ML, towards optimisation, math and Linear Algebra …
- Side-effect — change of the field …

ICML 2005

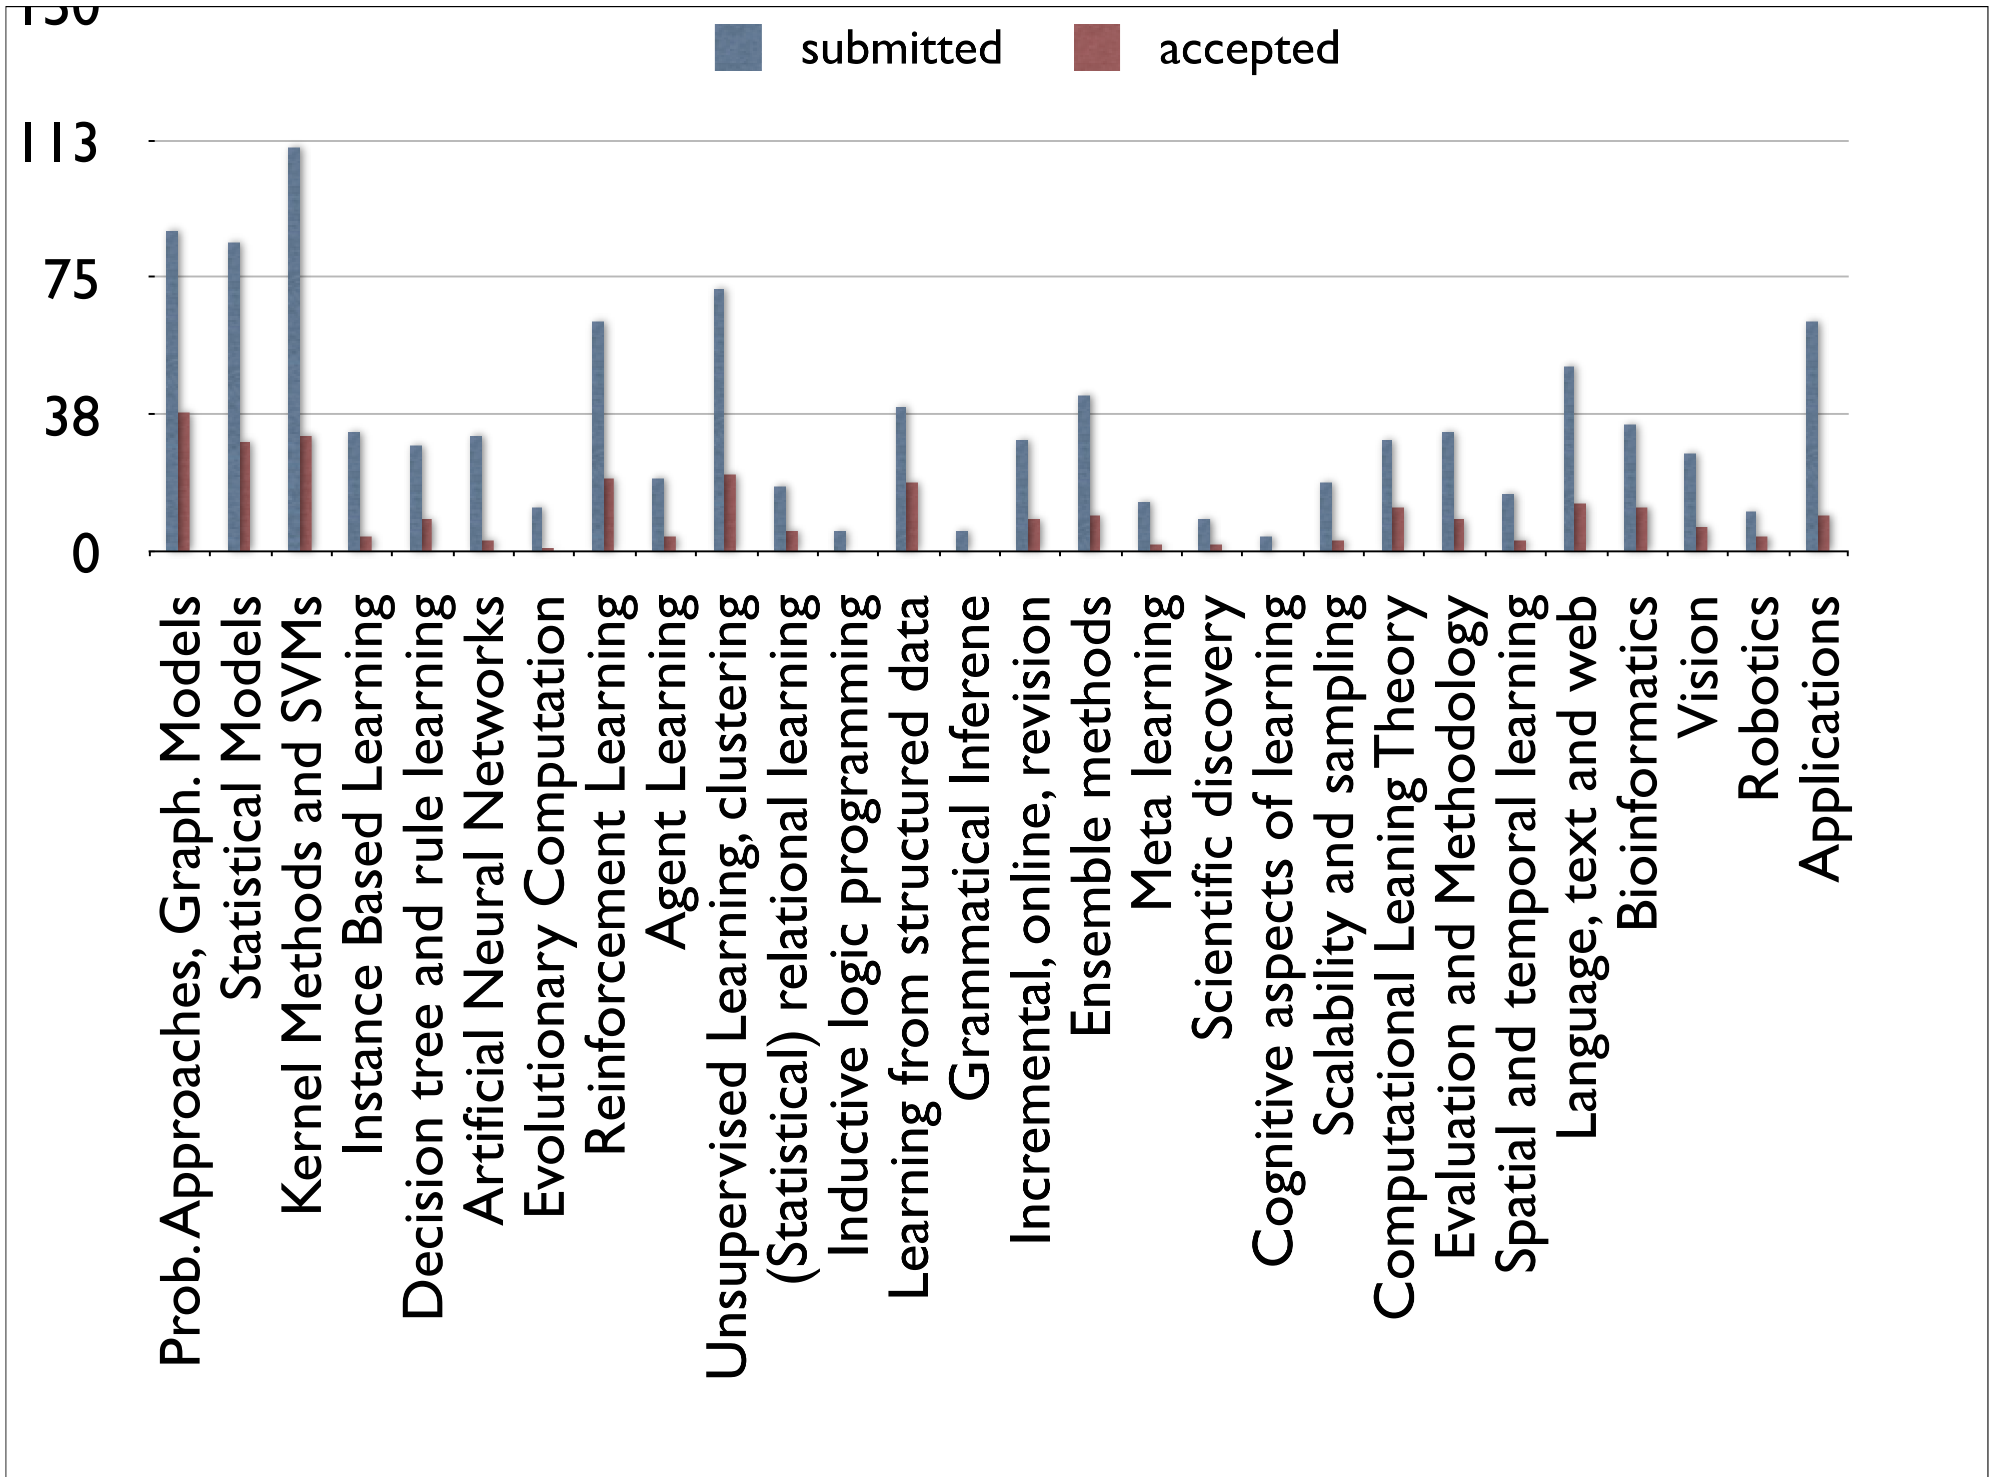

150
113
75
38
0
submitted
accepted
Prob. Approaches, Graph. Models
Statistical Models
Kernel Methods and SVMs
Instance Based Learning
Decision tree and rule learning
Artificial Neural Networks
Evolutionary Computation
Reinforcement Learning
Agent Learning
Unsupervised Learning, clustering
(Statistical) relational learning
Inductive logic programming
Learning from structured data
Grammatical Inferene
Incremental, online, revision
Ensemble methods
Meta learning
Scientific discovery
Cognitive aspects of learning
Scalability and sampling
Computational Leaning Theory
Evaluation and Methodology
Spatial and temporal learning
Language, text and web
Bioinformatics
Vision
Robotics
Applications

ICML 2005

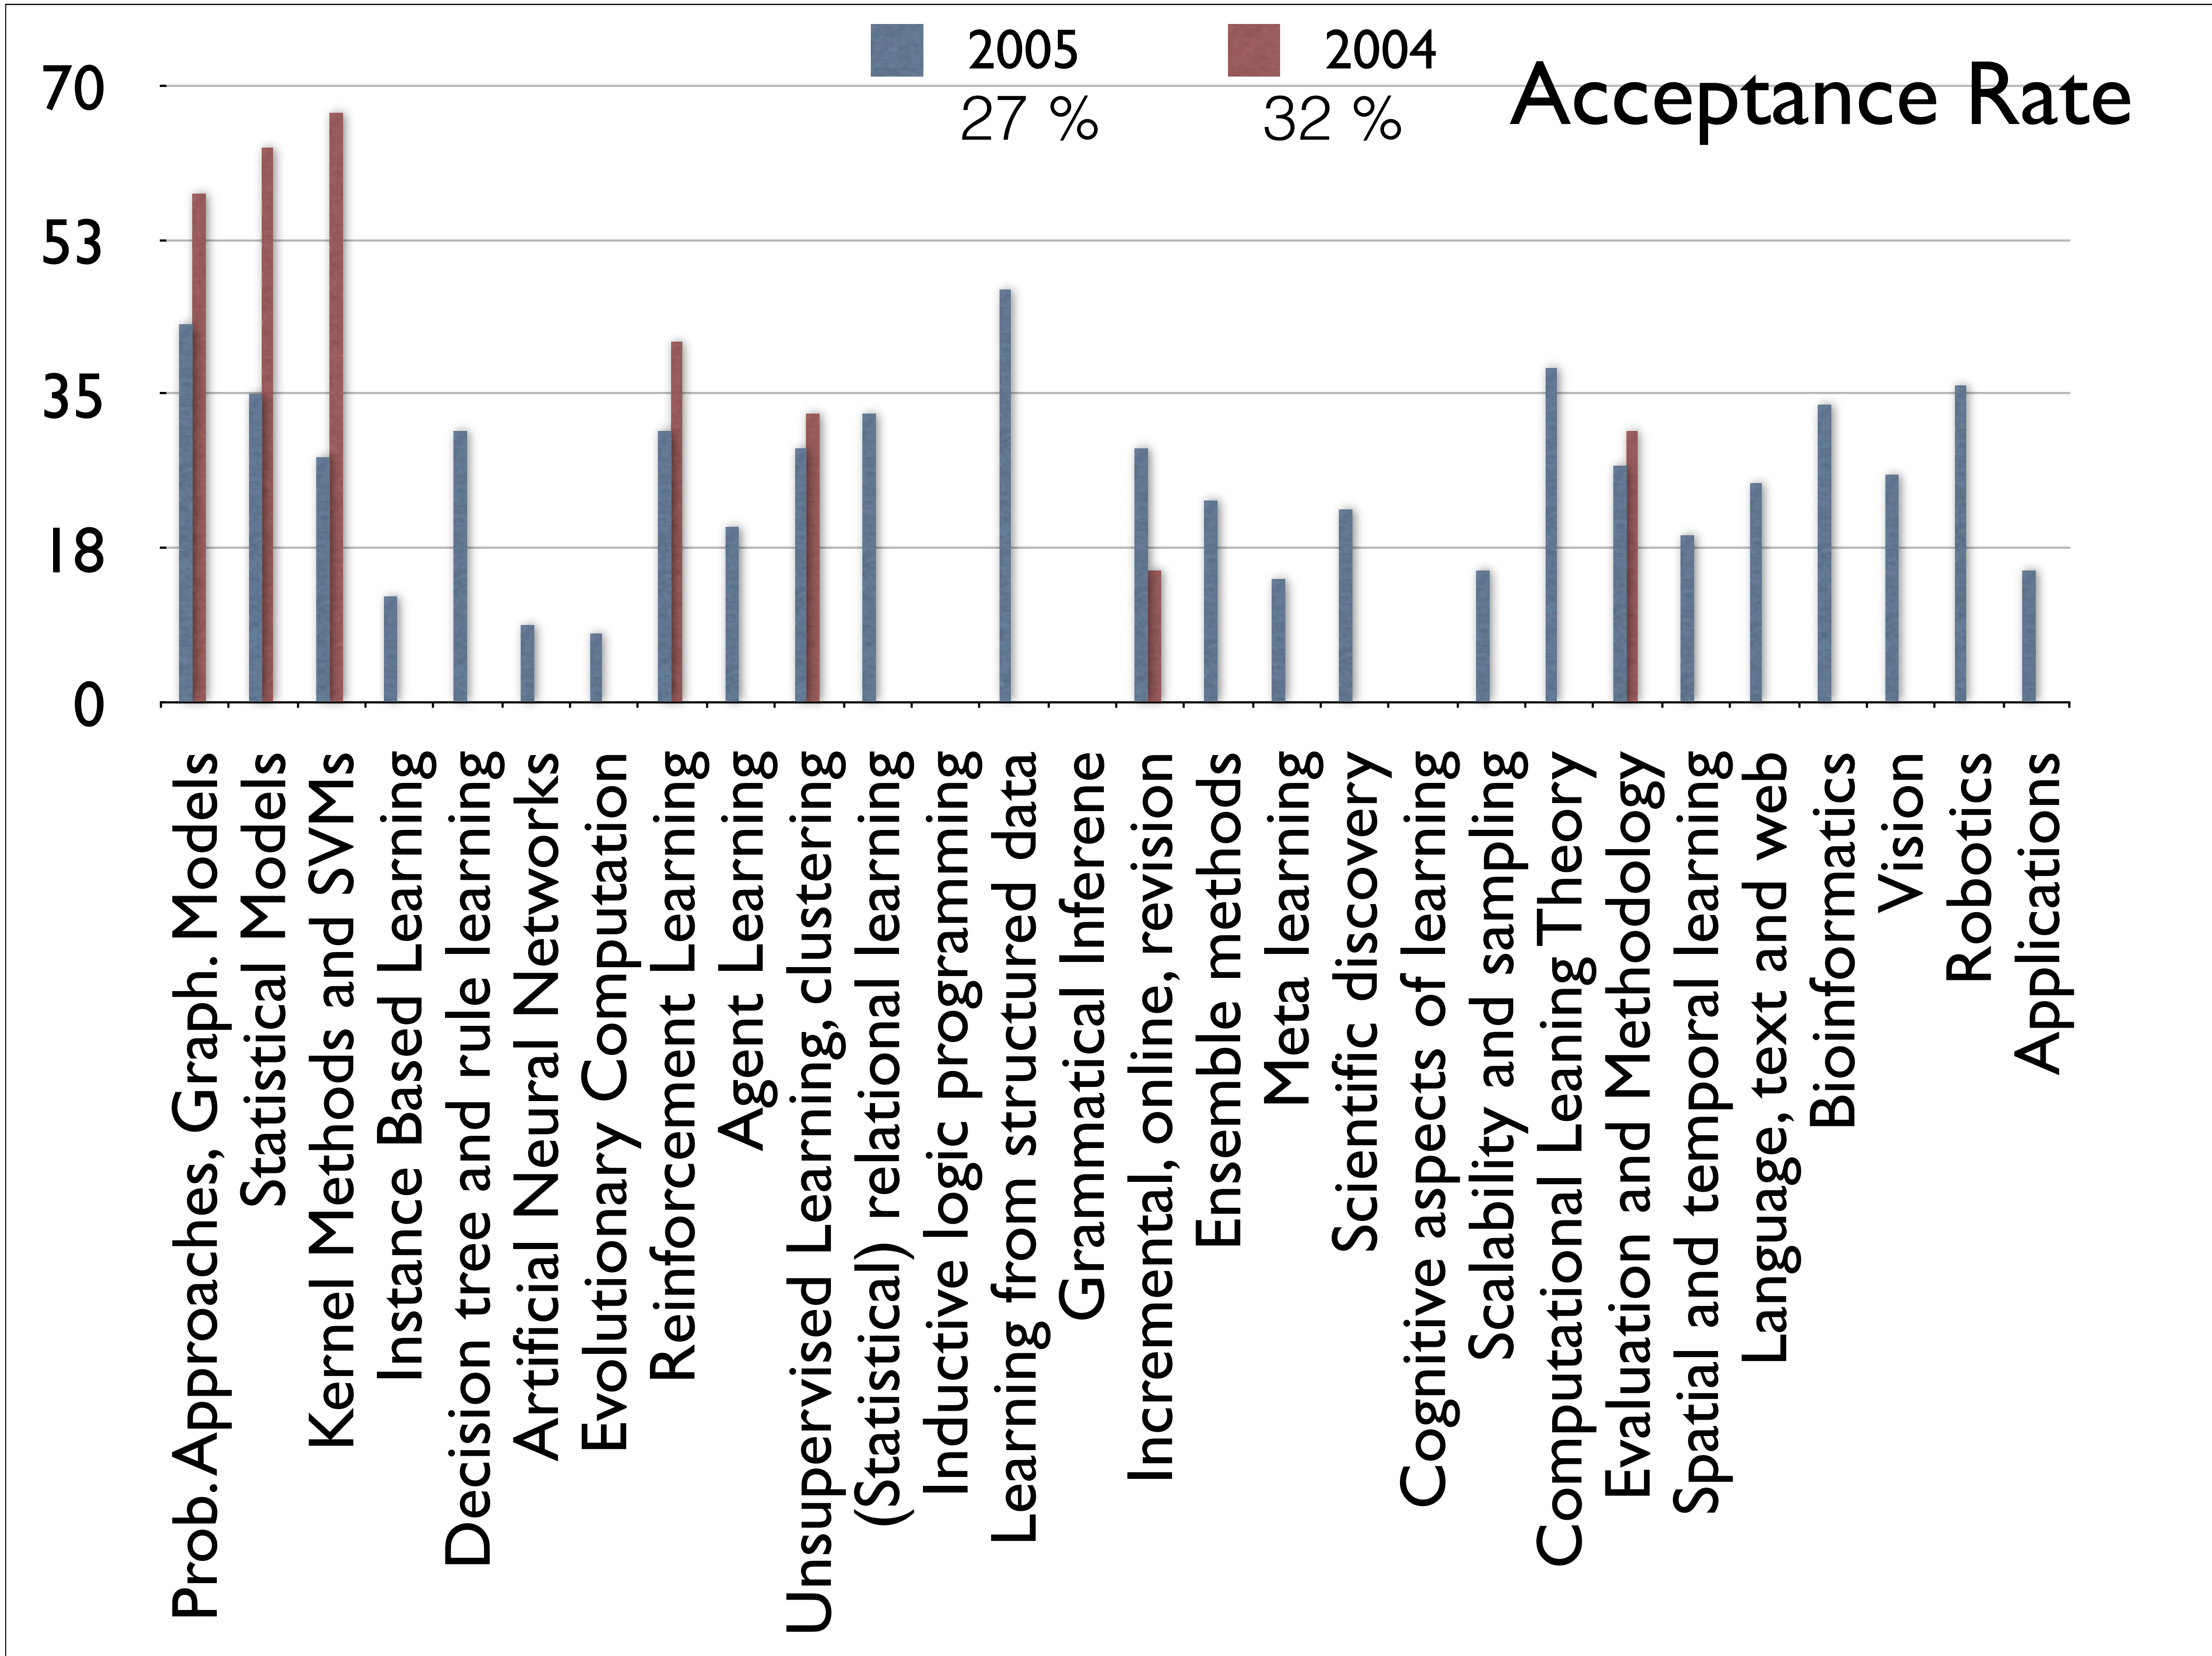
2005
27 %
2004
32 %
Acceptance Rate
70
53
35
18
0
Prob. Approaches, Graph. Models
Statistical Models
Kernel Methods and SVMs
Instance Based Learning
Decision tree and rule learning
Artificial Neural Networks
Evolutionary Computation
Reinforcement Learning
Agent Learning
Unsupervised Learning, clustering
(Statistical) relational learning
Inductive logic programming
Learning from structured data
Grammatical Inferene
Incremental, online, revision
Ensemble methods
Meta learning
Scientific discovery
Cognitive aspects of learning
Scalability and sampling
Computational Leaning Theory
Evaluation and Methodology
Spatial and temporal learning
Language, text and web
Bioinformatics
Vision
Robotics
Applications

Observations

- The social aspects of science (of ML?)
- Fields evolve, have biases, communities are dynamic, split, merge …
- Quite important to retain identity, to remain broad enough, yet coherent enough, …
- cf. The structure of scientific revolutions, Thomas Kuhn

Evolution …

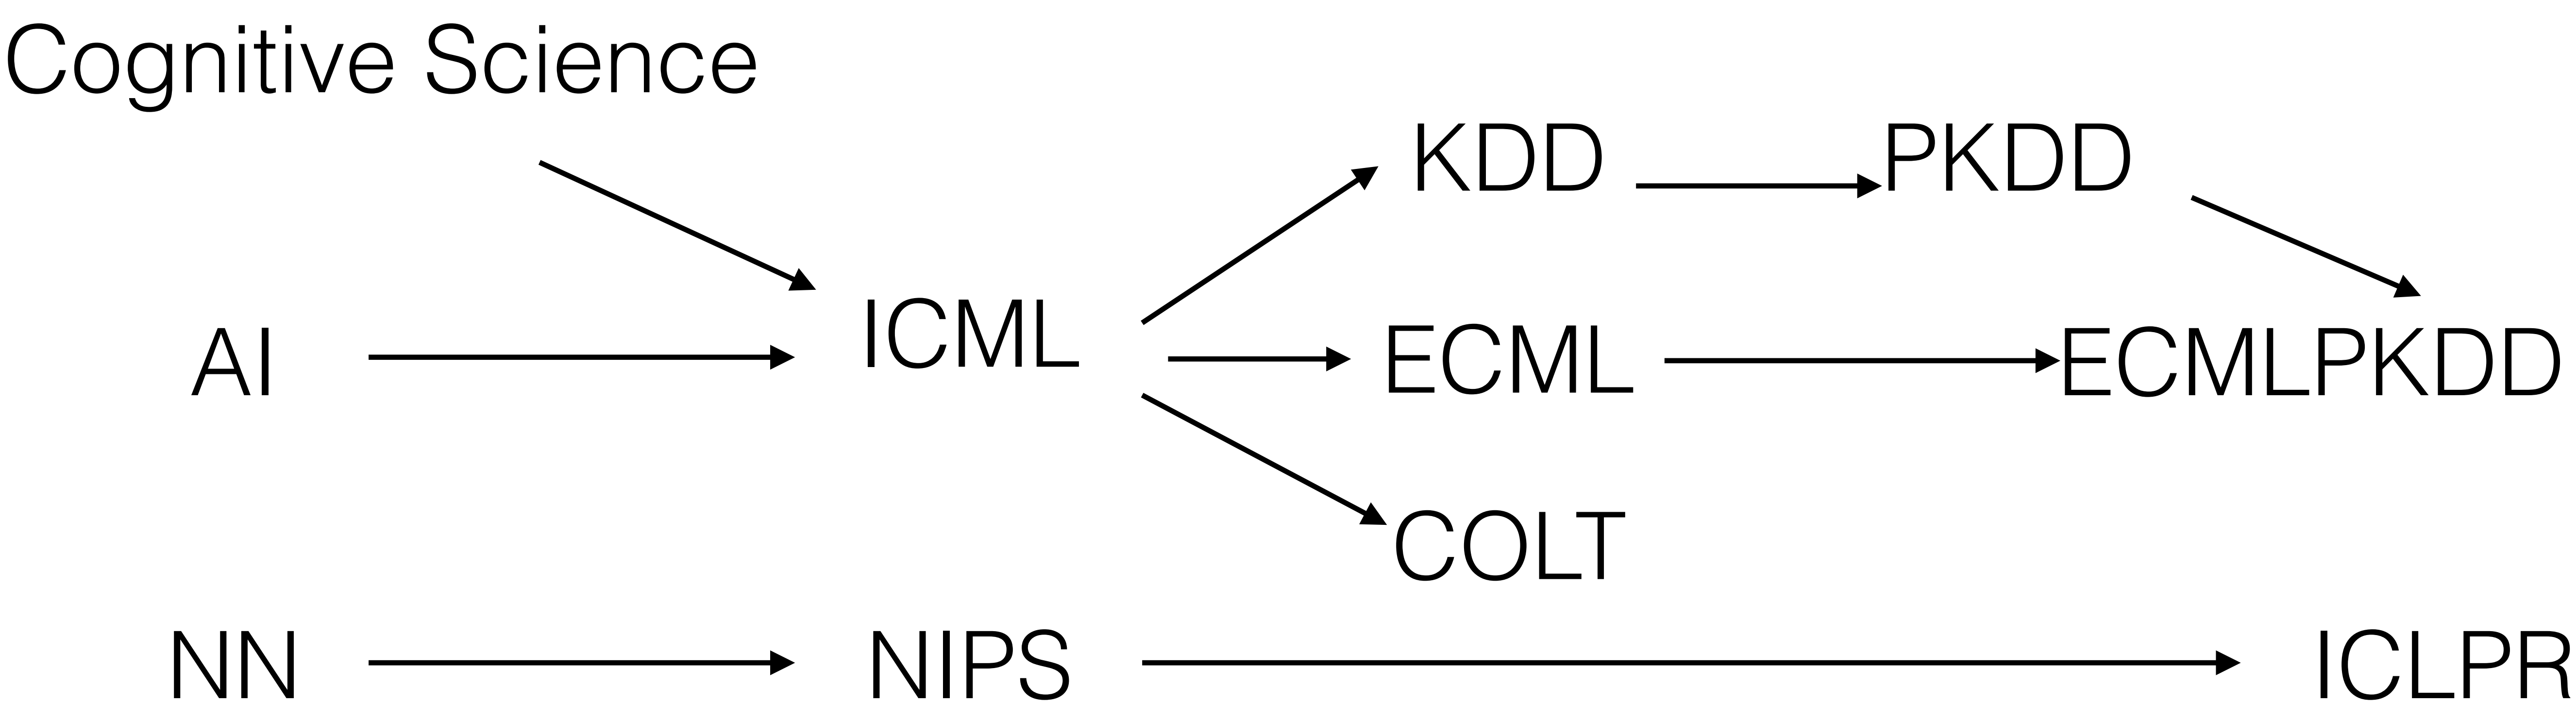
Cognitive Science
KDD
PKDD
AI
ICML
ECML
ECMLPKDD
COLT
NN
NIPS
ICLPR

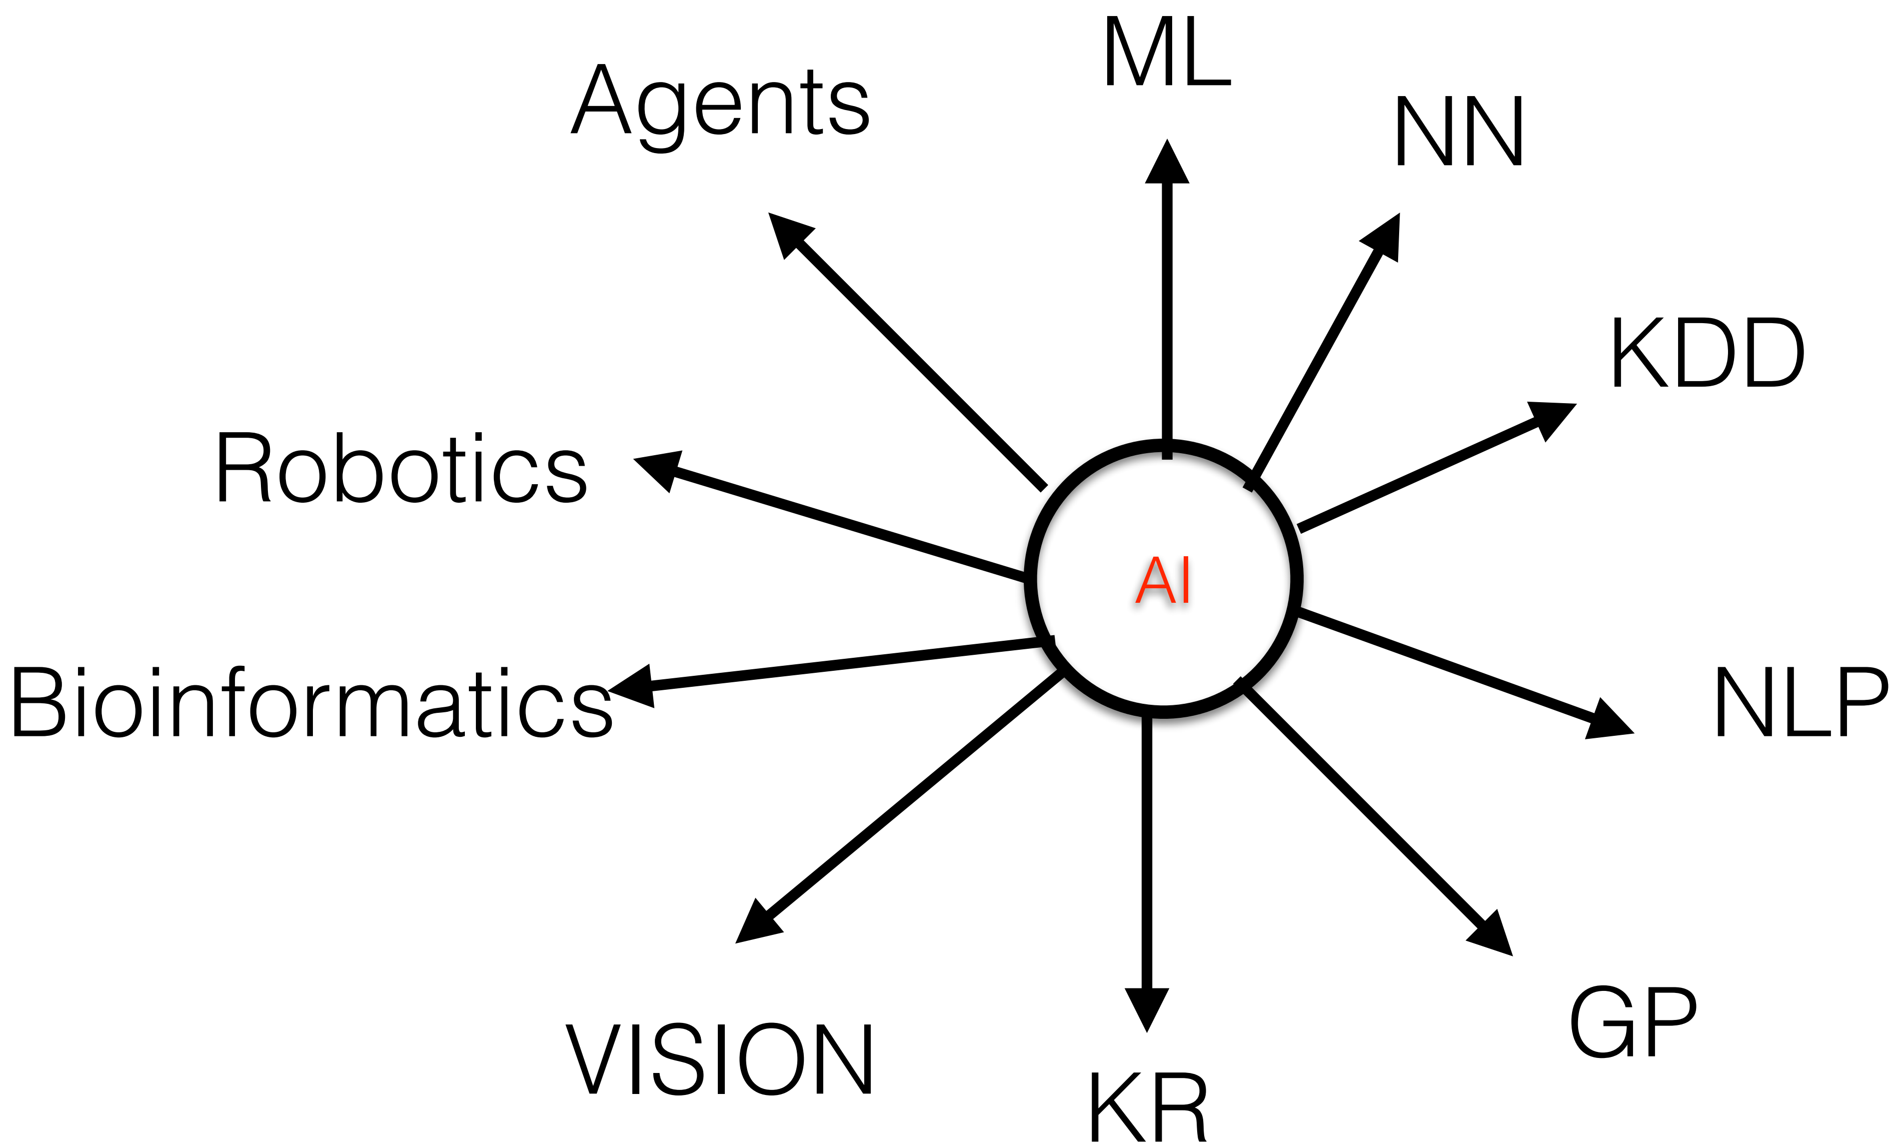

Agents
ML
NN
KDD
Robotics
AI
Bioinformatics
NLP
VISION
KR
GP

ML today

- Enormous progress made, impressive applications, used in many other fields as the enabling technology

- A healthy field

- Attracting loads of attention

- Unreasonable expectations ? a bit like AI ?

- Big data / data science … splitting off ?

ML today

- Diversity could be better ? are not we converging too fast ? More exploration would be useful ?
    - in terms of methodology
    - in terms of tasks
- Let's play more with the problem set up ?
- Links to AI, to human learning, to reasoning ?

What is next ?

- Get AI more into the picture …
- Machine learning : an AI approach ?

Can we automate Data Science / Machine Learning ?

- The robot scientist (Ross King et al. Nature 2004)

- Can we apply that idea to data science/machine learning itself ?
- One possible solution to the lack of data scientists today

The KDD process

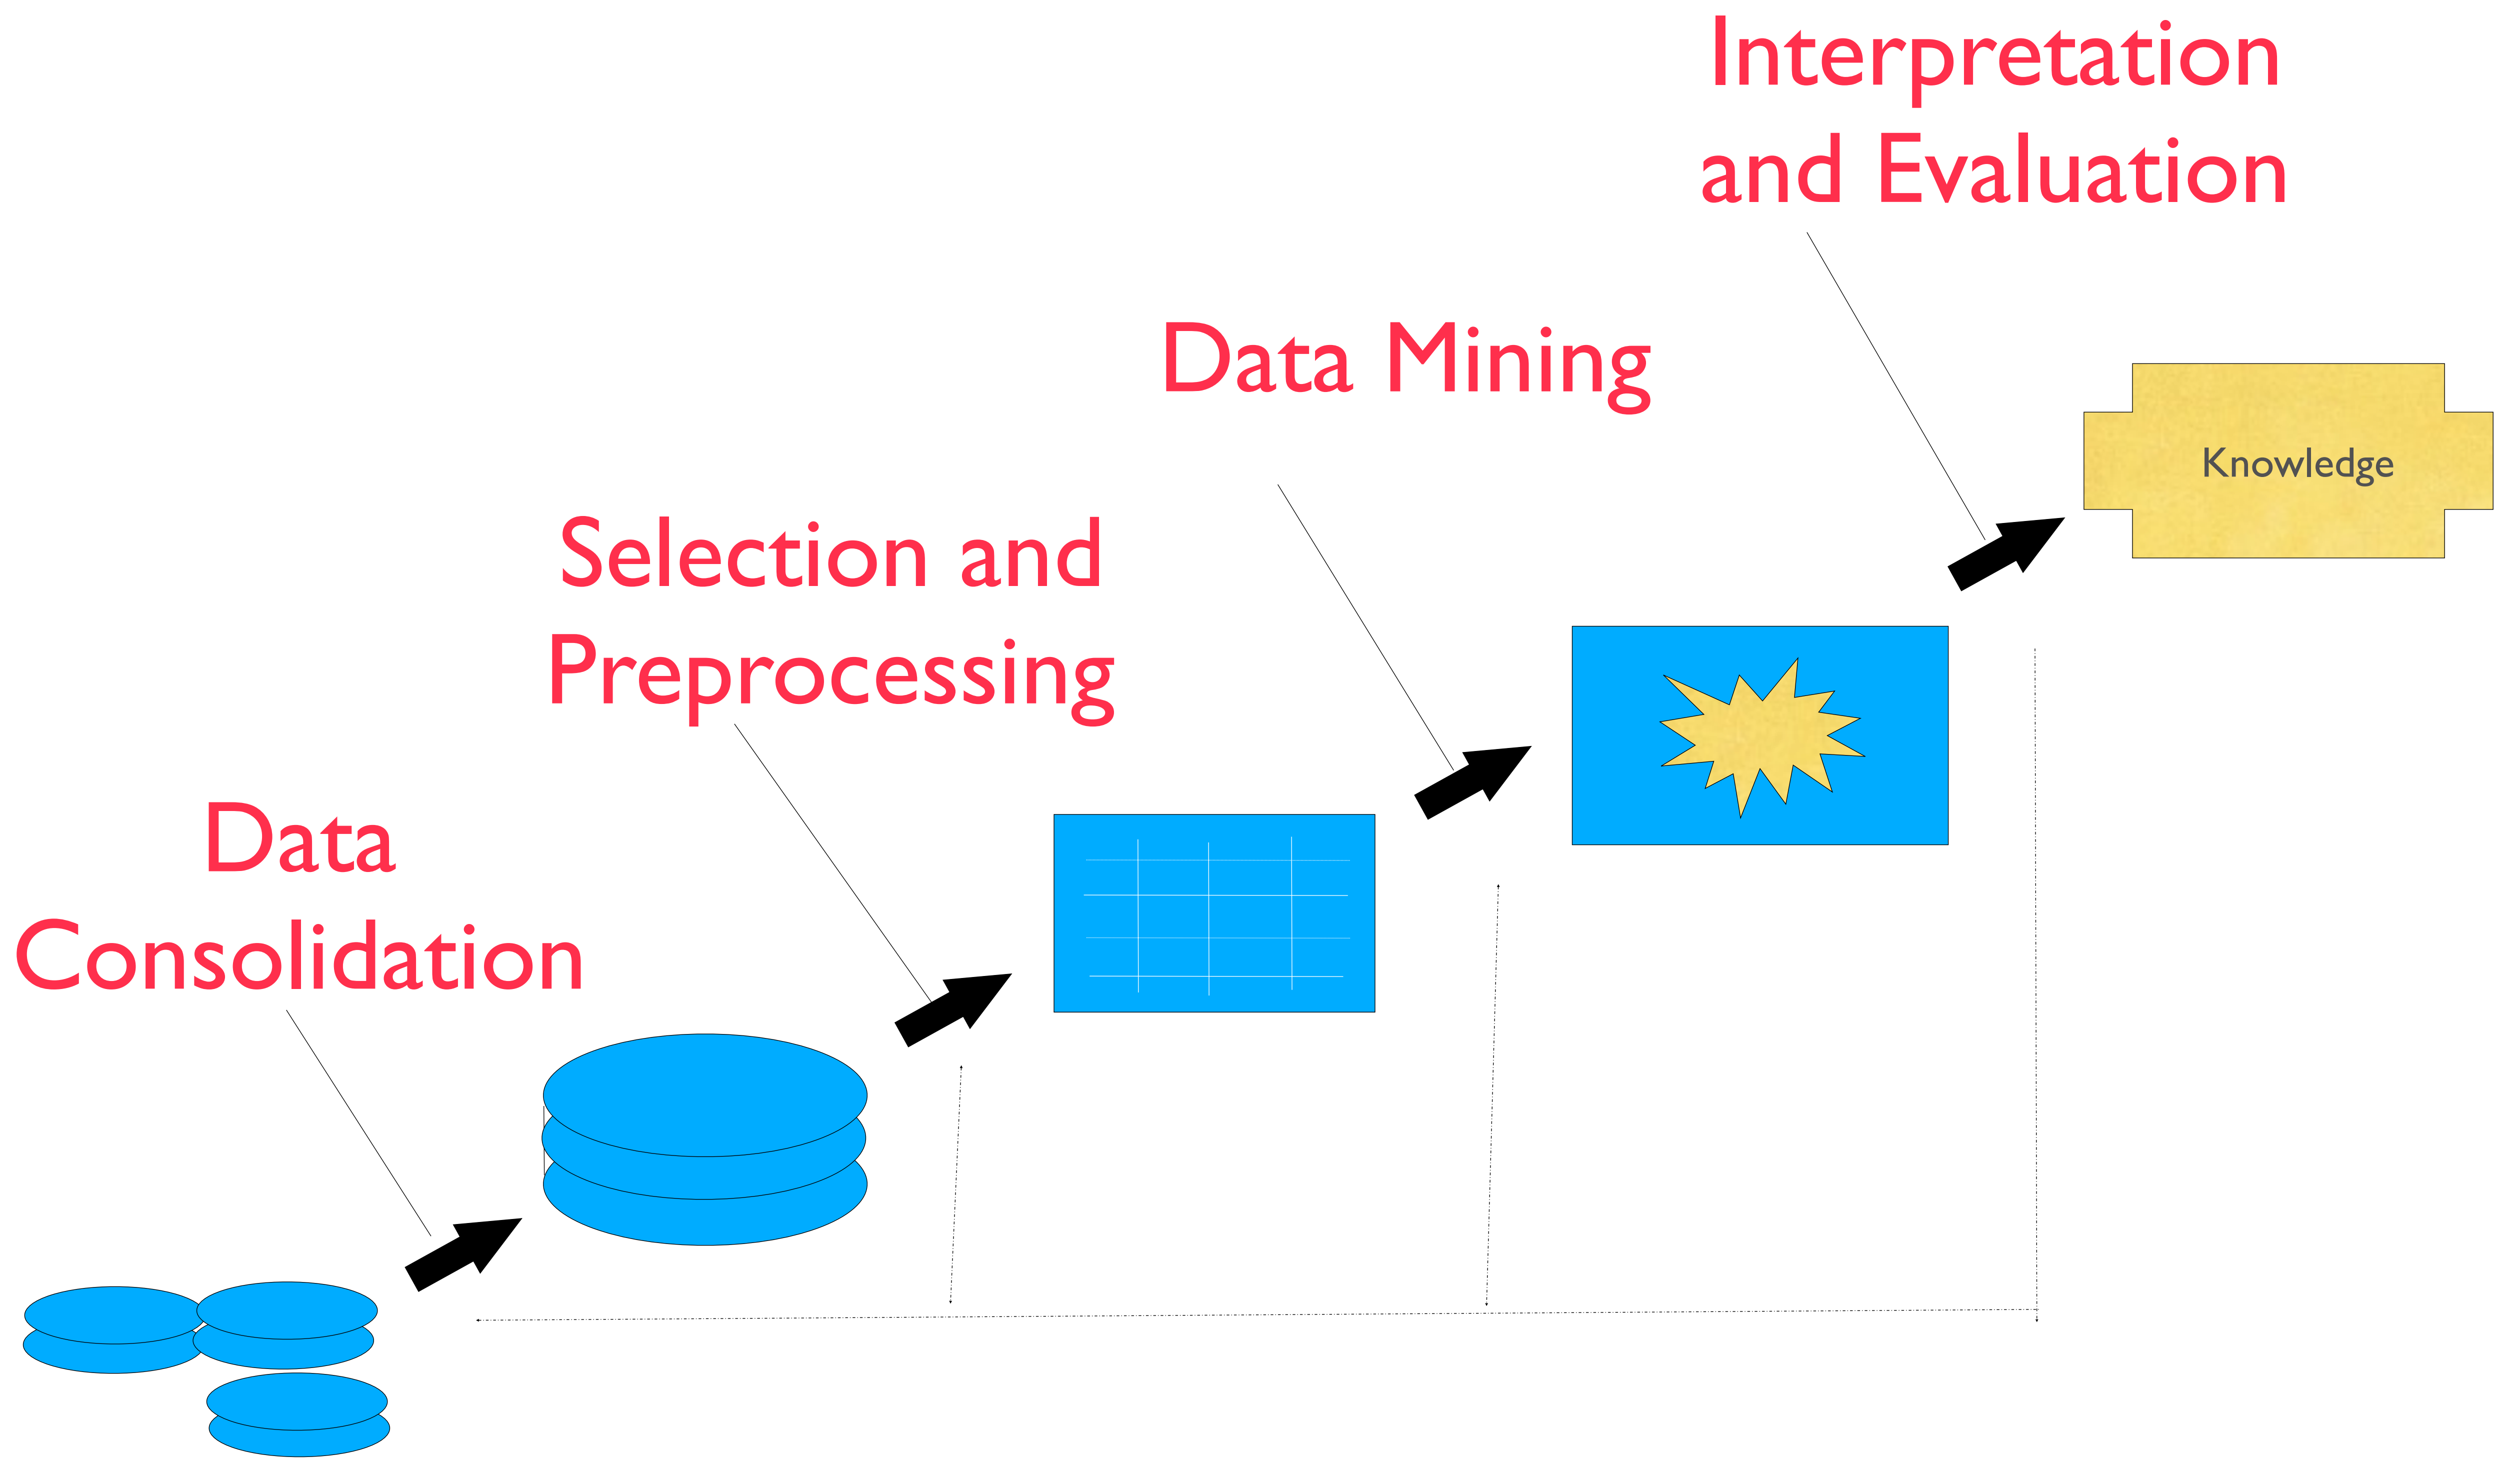

Advanced ERC Grant

Synthesising Inductive Data Models

Data Model

Inductive Model

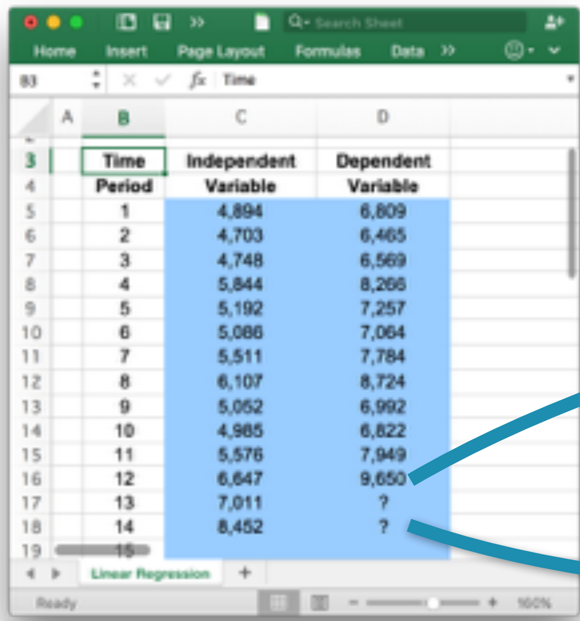

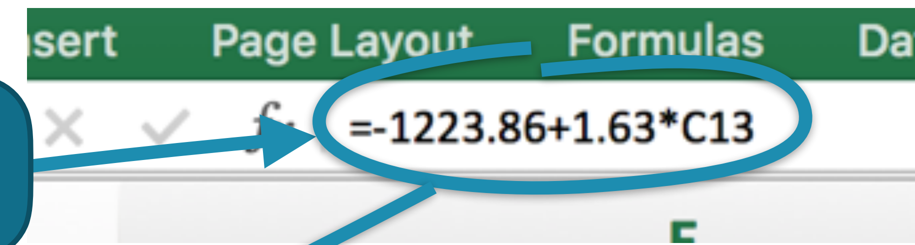

Discover patterns and rules present in a Data Model

+

Apply patterns to make predictions and support decisions

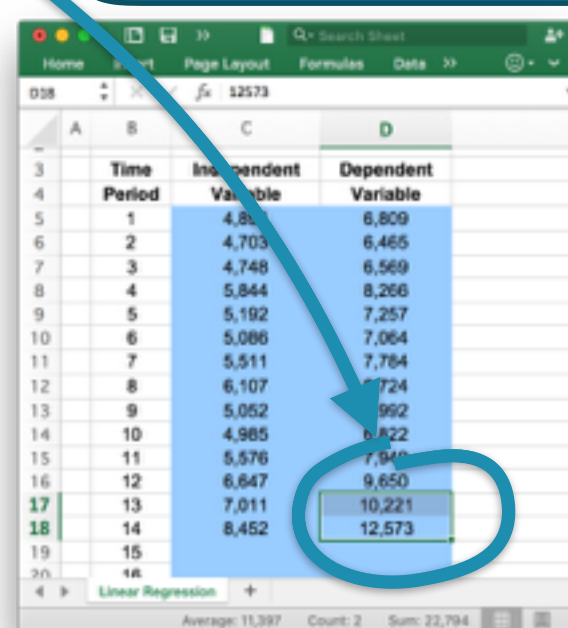

1. The synthesis system “learns the learning task”. It identifies the right learning tasks and learns appropriate IMs
2. The system may need to restructure the data set before IM synthesis can start
3. A unifying IDM language for a set of core patterns and models will be developed

DTAI Declaratieve Talen en Artificiële Intelligentie

erc

KU LEUVEN

Thanks

BTW: we are hiring PhD students and post-docs !