Case-Based Reasoning and the Upswing of AI

Agnar Aamodt[1] and Enric Plaza[2]

[1] Department of Computer Science, NTNU, Trondheim, Norway
agnar@ntnu.no
[2] IIIA – CSIC, Campus U.A.B., Bellaterra, Catalonia (Spain)
enric@iiia.csic.es

The history and evolution of AI has shaped Case-Based Reasoning (CBR) research and applications. We are currently living in an upswing of AI. To what degree does that mean an upswing of CBR as well? And what buttons should we push in order to increase the influence of CBR within the current AI summer, and beyond?

Artificial Intelligence as a scientific field was established at a Dartmouth College seminar in 1956, but ever since ancient times the idea of thinking as a formal and mechanistic process has occupied people's minds. After the firing of the starting shot in 56, the field has experienced both summers and winters, including two serious AI winters up to now. The causes behind these shifts in seasons have been subject to substantial discussion, and a compelling question of course is what to learn from this.

Case-Based Reasoning has had its own development history within the broader AI field. The grouping of AI methods into data-driven AI and knowledge-based AI [1] is also a familiar distinction in CBR. Recent trends in AI has clearly favoured the data-driven methods, and the well-known successes of Deep Neural Networks is a justification for that. But in order to widen the scope of AI methods, and be able to address a wider range of problems and applications, there are reasons to believe that a stronger knowledge-based influence will be needed in the years to come. Several authors have claimed we should look beyond the current upswing of AI, some have argued for methods inspired by human cognition, and others for a need to revitalize symbol-processing based on explicit knowledge representation. Pat Langley started the Cognitive Systems Movement [5], aimed at getting AI back to its roots of studying artefacts that explore the full range of human intelligence. The Artificial General Intelligence initiative addresses thinking machines with full human capabilities and beyond [3]. A focus on symbolic AI and knowledge representation issues has been strongly advocated by Hector Levesque [6], who also warns us to not to be blinded by short-term successes of particular methods.

Initiatives such as these are important to be aware of when we discuss future paths for CBR, and AI more generally. Moreover, the upswing of AI has created in the public, media, and decision makers a great confusion as to what AI is, where numerous concepts — AI, robots, ML, deep learning, and big data, together with the "smart" adjective before almost anything — are conflated and used interchangeably.

So, where is CBR in the overall AI landscape today? Does it live its own life alongside other main subareas, or are there sufficient similarities at the foundational level to group CBR with other methods? With a focus on machine learning, a division of the ML field into five "tribes" has been suggested [2], within which one such tribe is the "Analogizers", united by their reliance on similarity assessment as the basis for learning. It is a diverse tribe covering analogical reasoning, instance-based methods, and support vector machines. For

Invited talk, extended abstract

each of the five tribes a unifying ‘master algorithm’ is proposed, and some people may fall off their chairs when kernel machines is assigned as the unifying method for this tribe. Anyway, views like this may trigger discussions that will lead to a better understanding of CBR in relation to other AI methods.

Given the growing interest in cognitive foundations of AI, we recall the notions of System 1 and System 2 in human cognition presented by Kahneman [4]. System 1 is a model of human memory capable of ‘fast thinking’, basically performing recognition of new inputs and responding intuitively, while System 2 models the deliberate, explicit reasoning performed by humans. An important issue to discuss is how they could be related to an integrated view of CBR encompassing both data-driven and knowledge-intensive processes.

All these considerations open up some important future challenges and opportunities for CBR, including: How to interpret the revitalized cognitive turn in the paradigm of CBR? How can data-driven CBR be competitive with current ML developments? Can CBR offer a new kind of synergy of data-driven and knowledge-intensive approaches for AI?

References

1. 'Barcelona Declaration for the Proper Development and Use of Artificial Intelligence in Europe', http://www.bdebate.org/sites/default/files/barcelona-declaration_v7-1-eng.pdf, last accessed 2017/5/3.
2. Domingos, Pedro: The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World. Basic Books, New York (2015).
3. Goertzel, Ben: Artificial General Intelligence: Concept, State of the Art, and Future Prospects. *Journal of Artificial General Intelligence*, 5(1) (Dec 2014)
4. Kahneman, Daniel: Thinking, Fast and Slow. Penguin Books (2011).
5. Langley, Pat: The cognitive systems paradigm. *Advances in Cognitive Systems Journal* 1, 3-13 (2012).
6. Levesque, Hector: On our best behavior. *Artificial Intelligence*, 212(1), 27-35 (2014).

CBR and the Upswing of AI

Enric Plaza
Institut d'Investigació en Intel·ligència Artificial
CSIC

Agnar Aamodt
Department of Computer Science
NTNU

ICCBR 2017 Invited talk

Our slot

- Why did **we** become CBR-ers?
- A bit of AI history – through changing seasons
- CBR history snippets
- The current success of AI – and prospects ahead
  - Data-driven vs. Knowledge-based AI
- How will/may CBR contribute to future AI?
- Discussion

- So: Why, oh why, did ***we*** choose CBR?

BARCELONA DECLARATION FOR THE PROPER DEVELOPMENT AND USAGE OF ARTIFICIAL INTELLIGENCE IN EUROPE

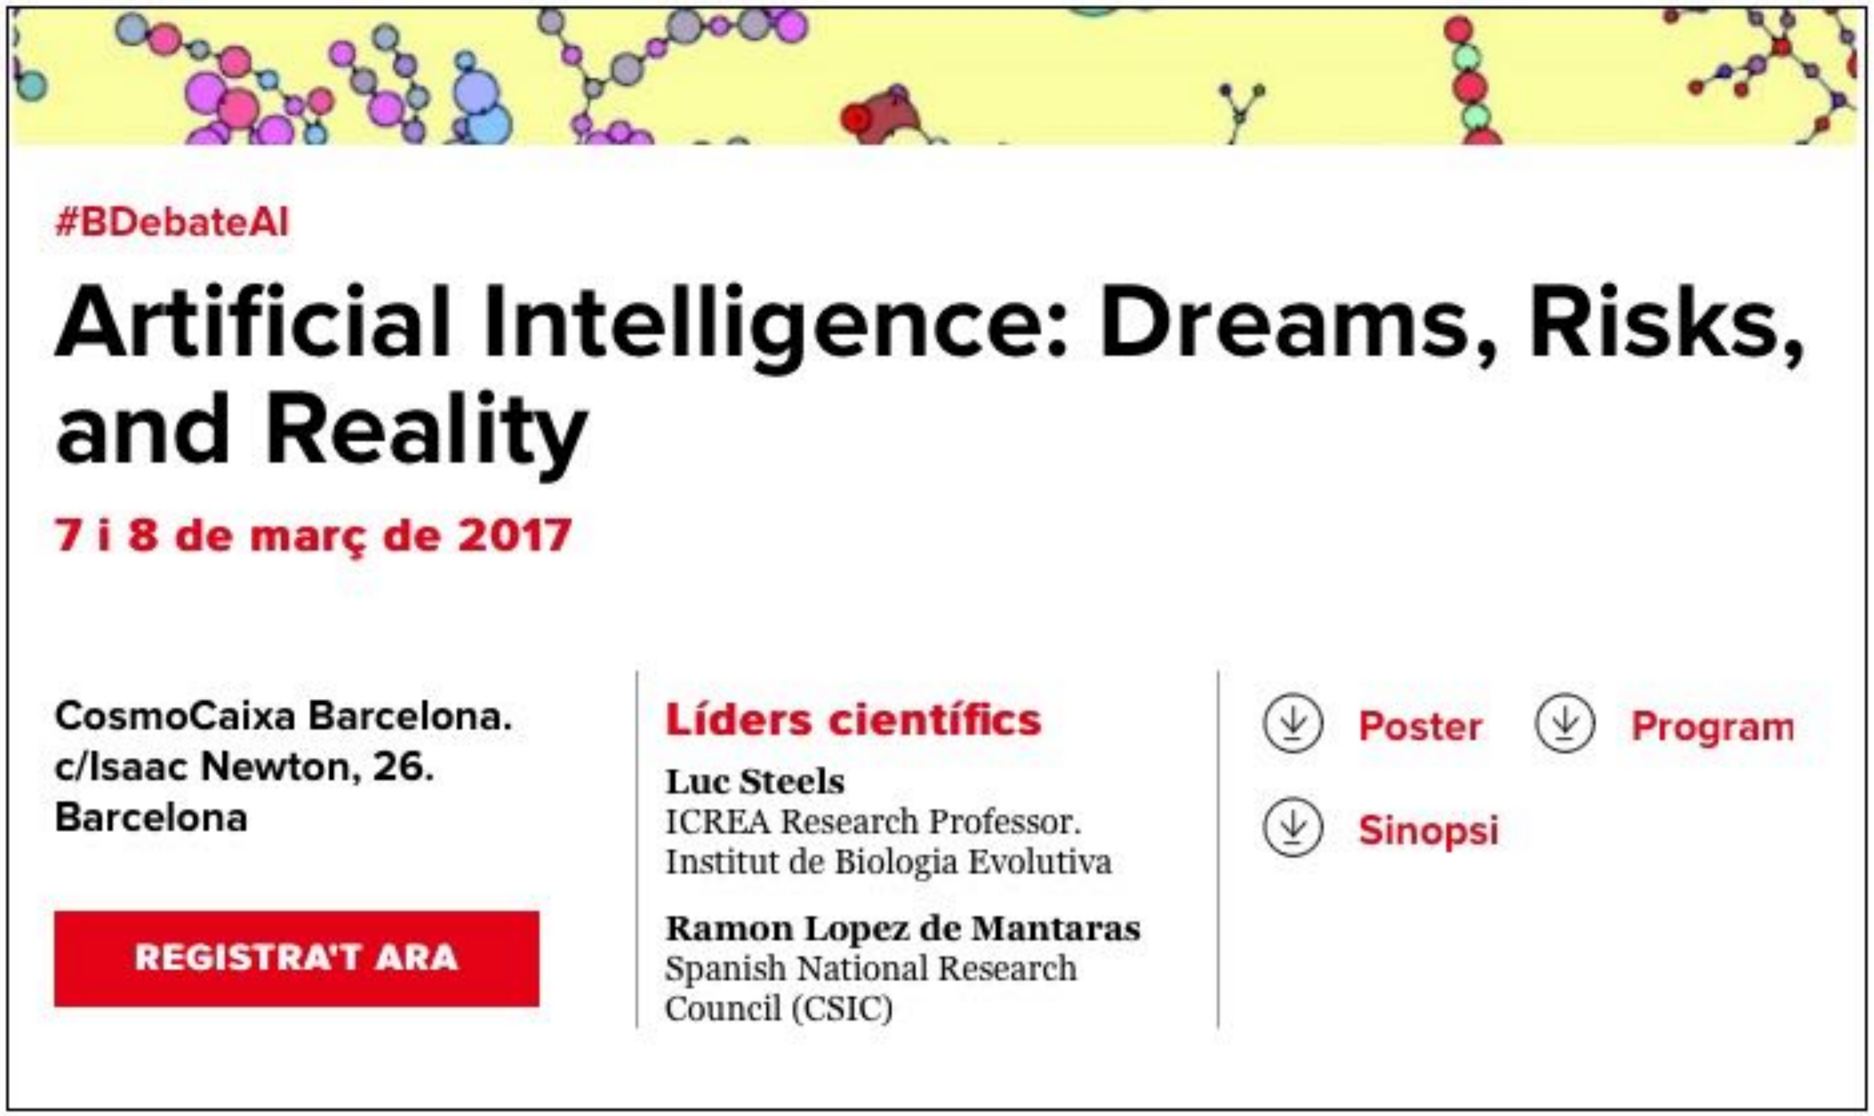

http://www.bdebate.org/ca/debat/artificial-intelligence-next-step-evolution

BARCELONA DECLARATION FOR THE PROPER DEVELOPMENT AND USAGE OF ARTIFICIAL INTELLIGENCE IN EUROPE

- Prudence
- Reliability
- Accountability
- Responsability
- Constrained autonomy
- Human role

BARCELONA DECLARATION FOR THE PROPER DEVELOPMENT AND USAGE OF ARTIFICIAL INTELLIGENCE IN EUROPE

Knowledge-based AI: top-down reasoning and problem solving strategies, language processing, and insight learning

Data-driven AI: bottom-up statistical machine learning algorithms to make predictions, complete partial data, or emulate behavior

AI Seasons

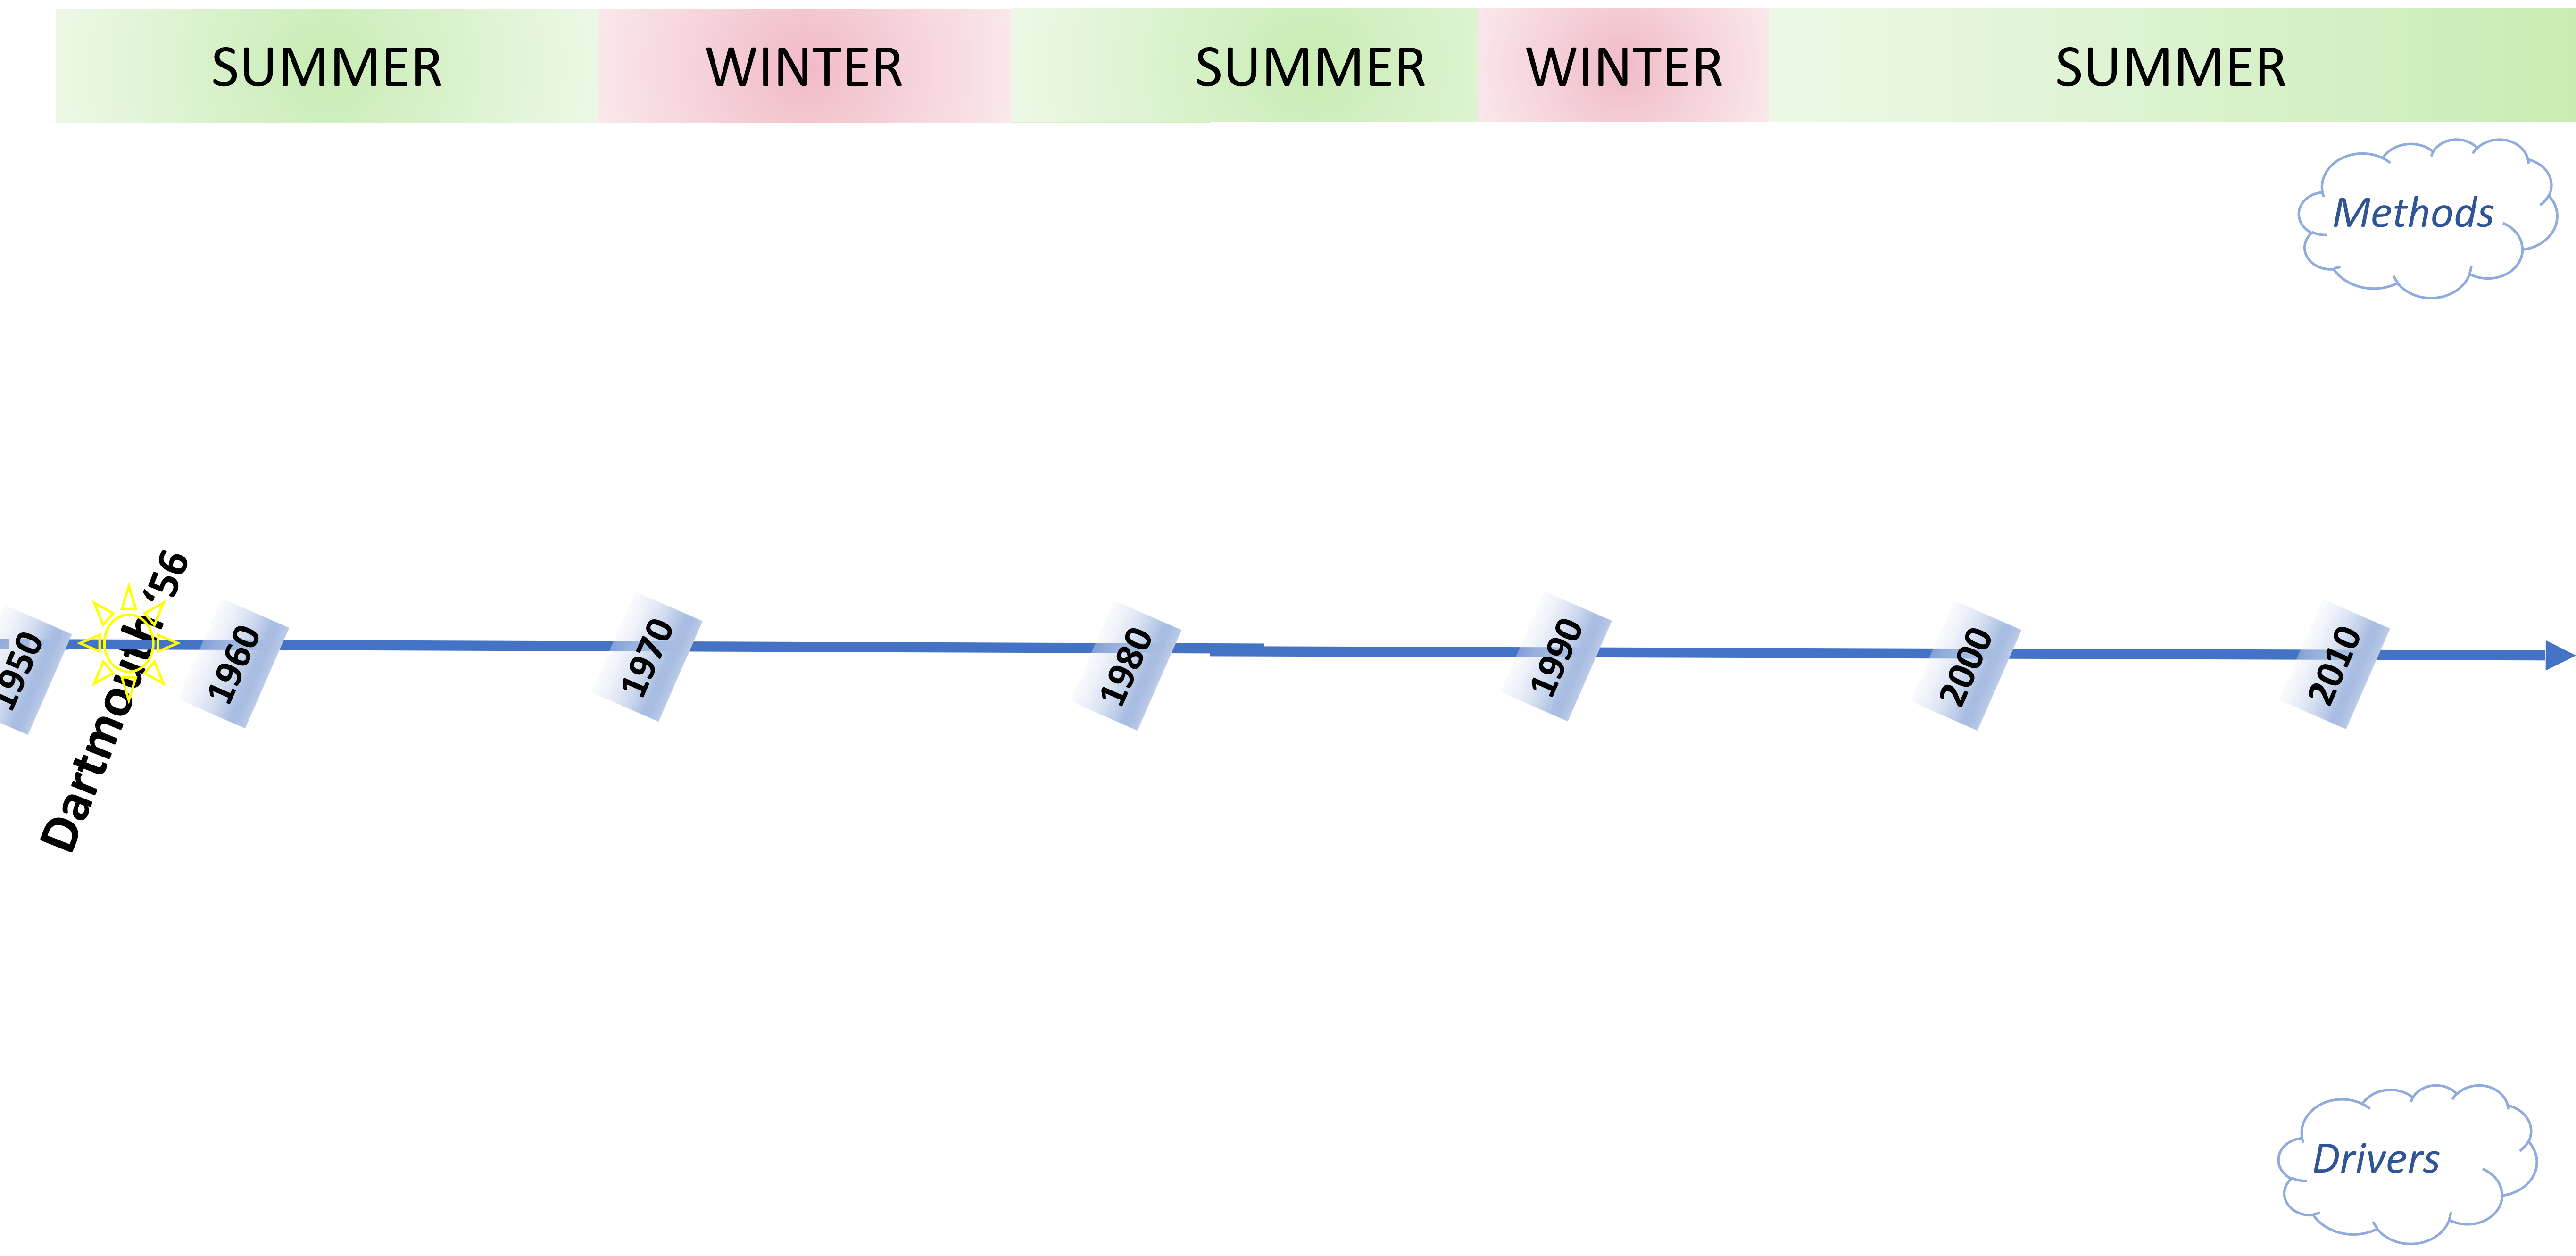
SUMMER
WINTER
SUMMER
WINTER
SUMMER
Methods
Dartmouth '56
1950
1960
1970
1980
1990
2000
2010
Drivers

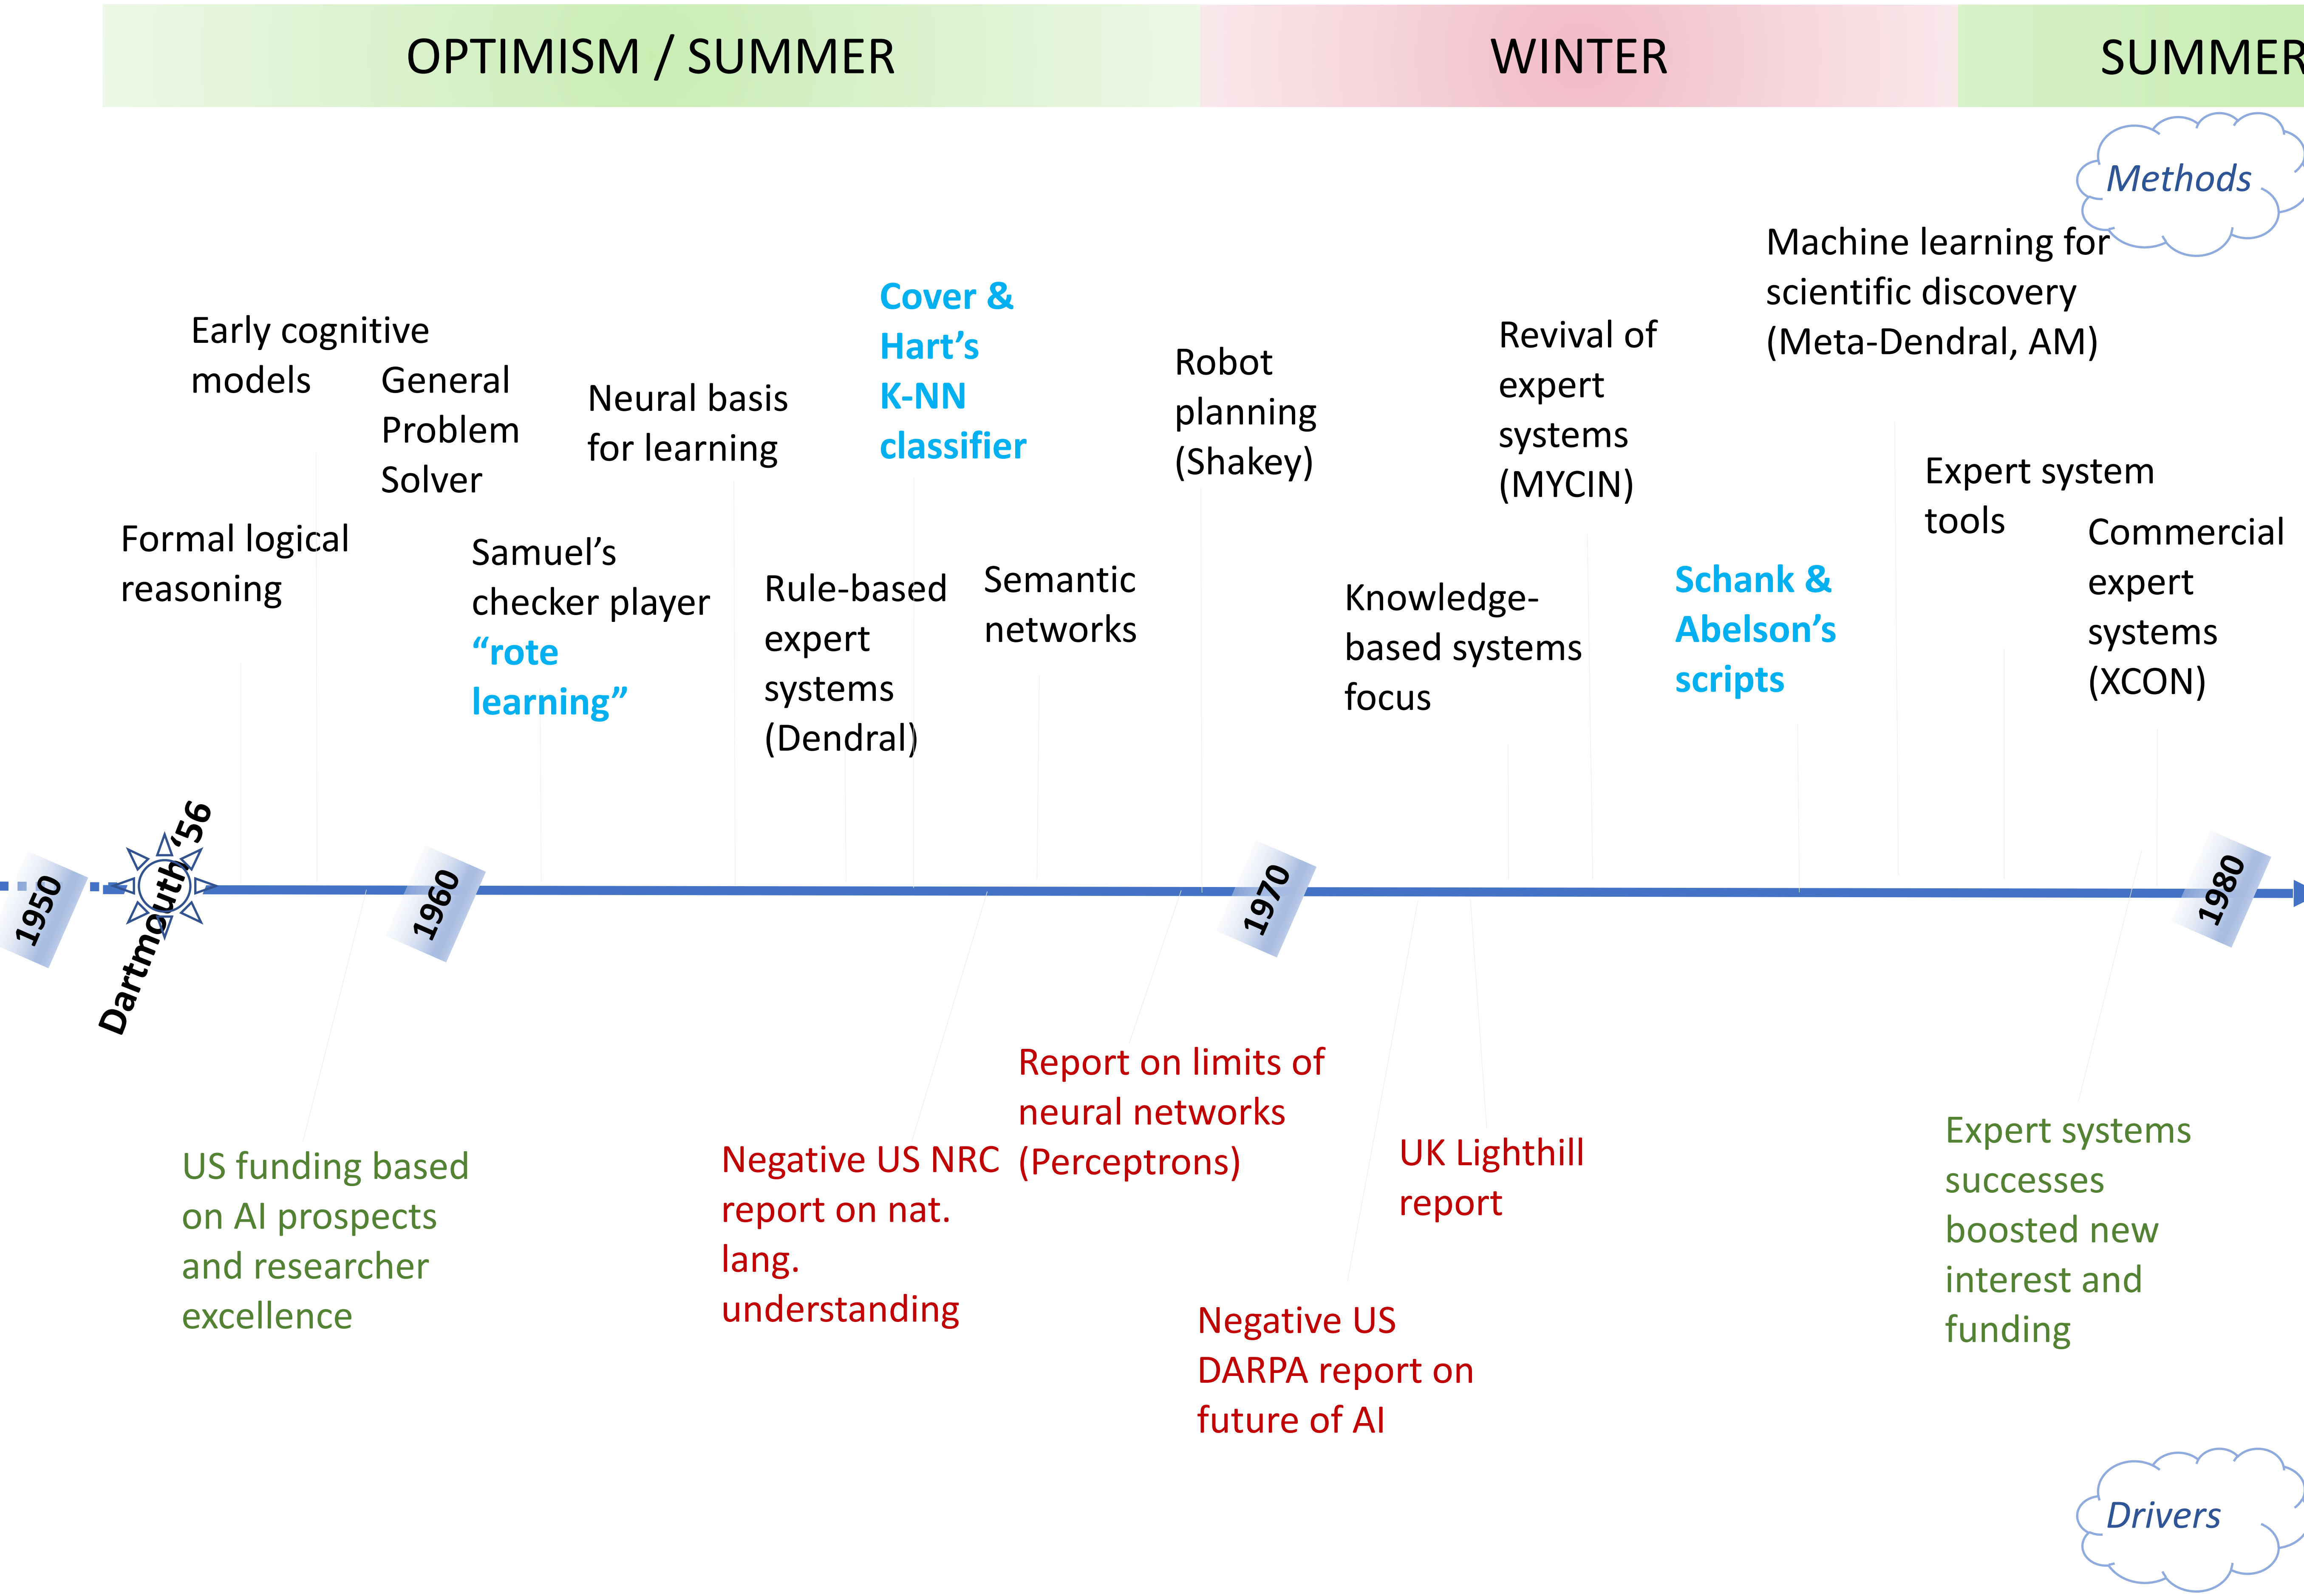
OPTIMISM / SUMMER
WINTER
SUMMER
Methods
Early cognitive models
General Problem Solver
Neural basis for learning
Cover & Hart's K-NN classifier
Robot planning (Shakey)
Revival of expert systems (MYCIN)
Machine learning for scientific discovery (Meta-Dendral, AM)
Formal logical reasoning
Samuel's checker player "rote learning"
Rule-based expert systems (Dendral)
Semantic networks
Knowledge-based systems focus
Schank & Abelson's scripts
Expert system tools
Commercial expert systems (XCON)
1950
Dartmouth '56
1960
1970
1980
US funding based on AI prospects and researcher excellence
Negative US NRC report on nat. lang. understanding
Report on limits of neural networks (Perceptrons)
UK Lighthill report
Negative US DARPA report on future of AI
Expert systems successes boosted new interest and funding
Drivers

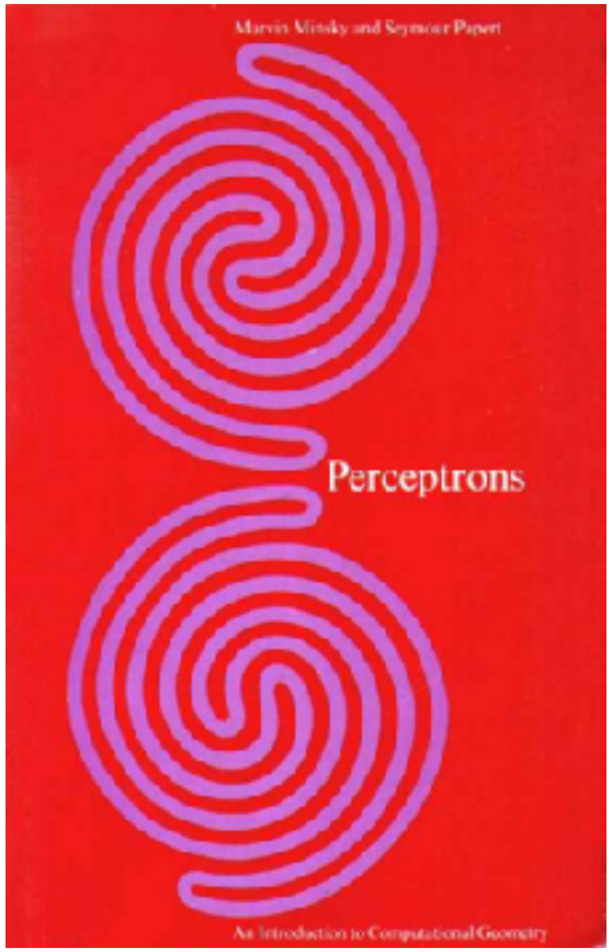

Minsky & Papert: Perceptrons (1969)

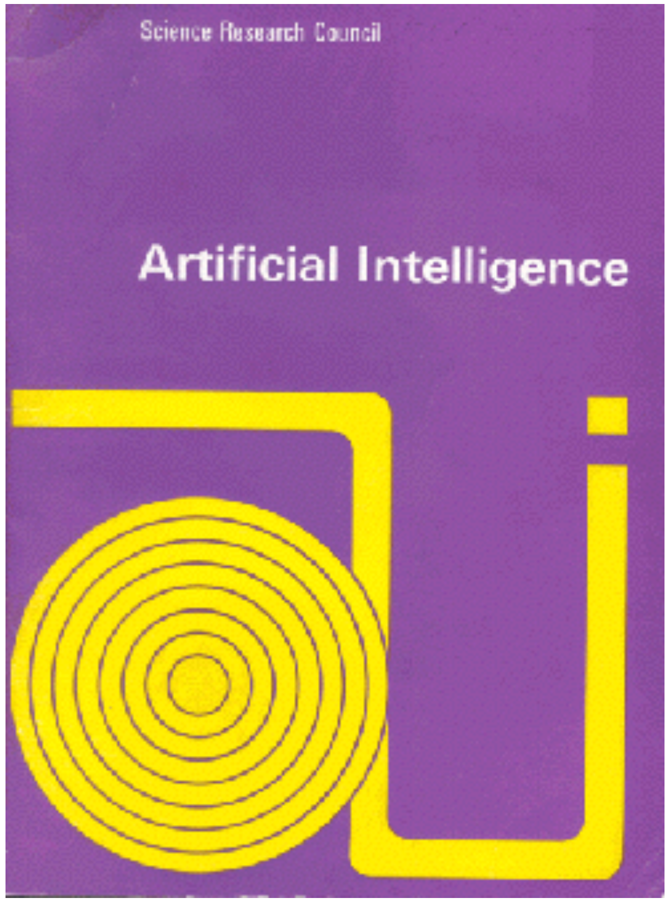

James Lighthill, 1973 (The Lighthill report)

DARPA:

Mansfield Amendment (1969): DARPA should fund "mission-oriented research, rather than basic undirected research"

American Study Group (1973): "AI research is unlikely to produce military applications in the foreseeable future."

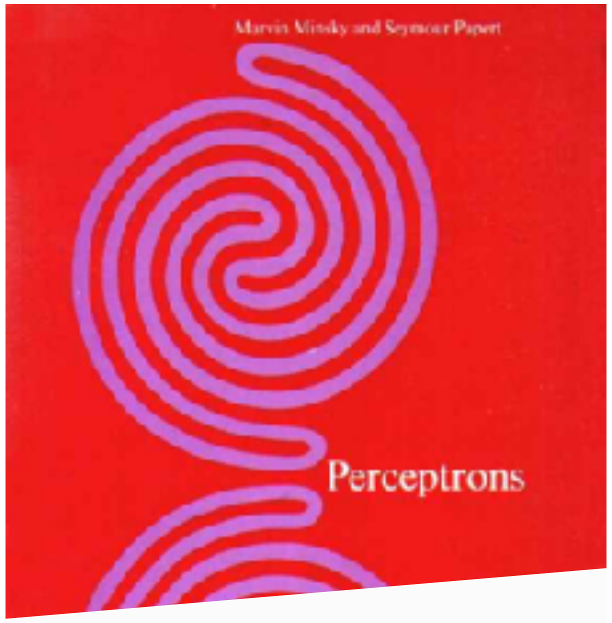

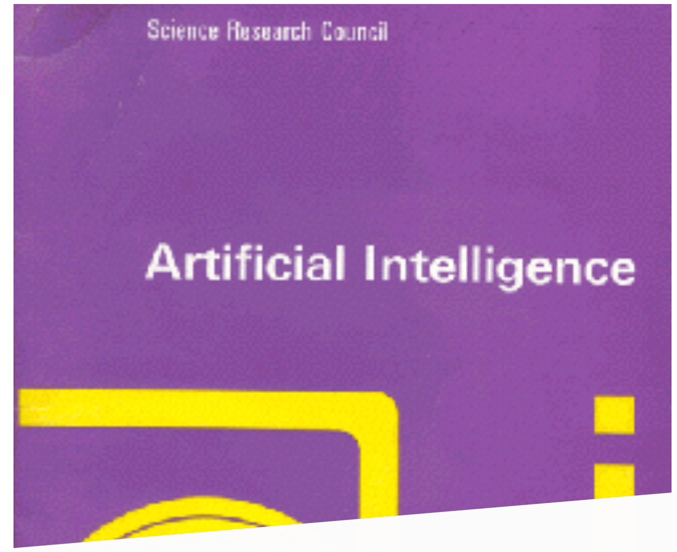

Configuration with R1/XCon (1978)

- Knowledge domain: Configuring VAX computers, to customers' specifications.
- Input: Required characteristics of the computer system.
- Output: Specification for the computer system.
- Inference engine: Forward chaining: the output specification was assembled in working memory.
- Knowledge representation: Production rules.
- DEC attempted to write a conventional program to do this task, with no success, then invited McDermott to write an AI system to do it. McDermott wrote R1/XCON. By 1986, it had processed 80,000 orders, and achieved 95-98% accuracy. It was reckoned to be saving DEC $25M a year.
- R1/XCON suffered from the shortcomings of simple production-rule-based systems.Expensive rewriting was needed to restore the operation of the system

American Study Group (1973): “AI research is unlikely to produce military applications in the foreseeable future.”

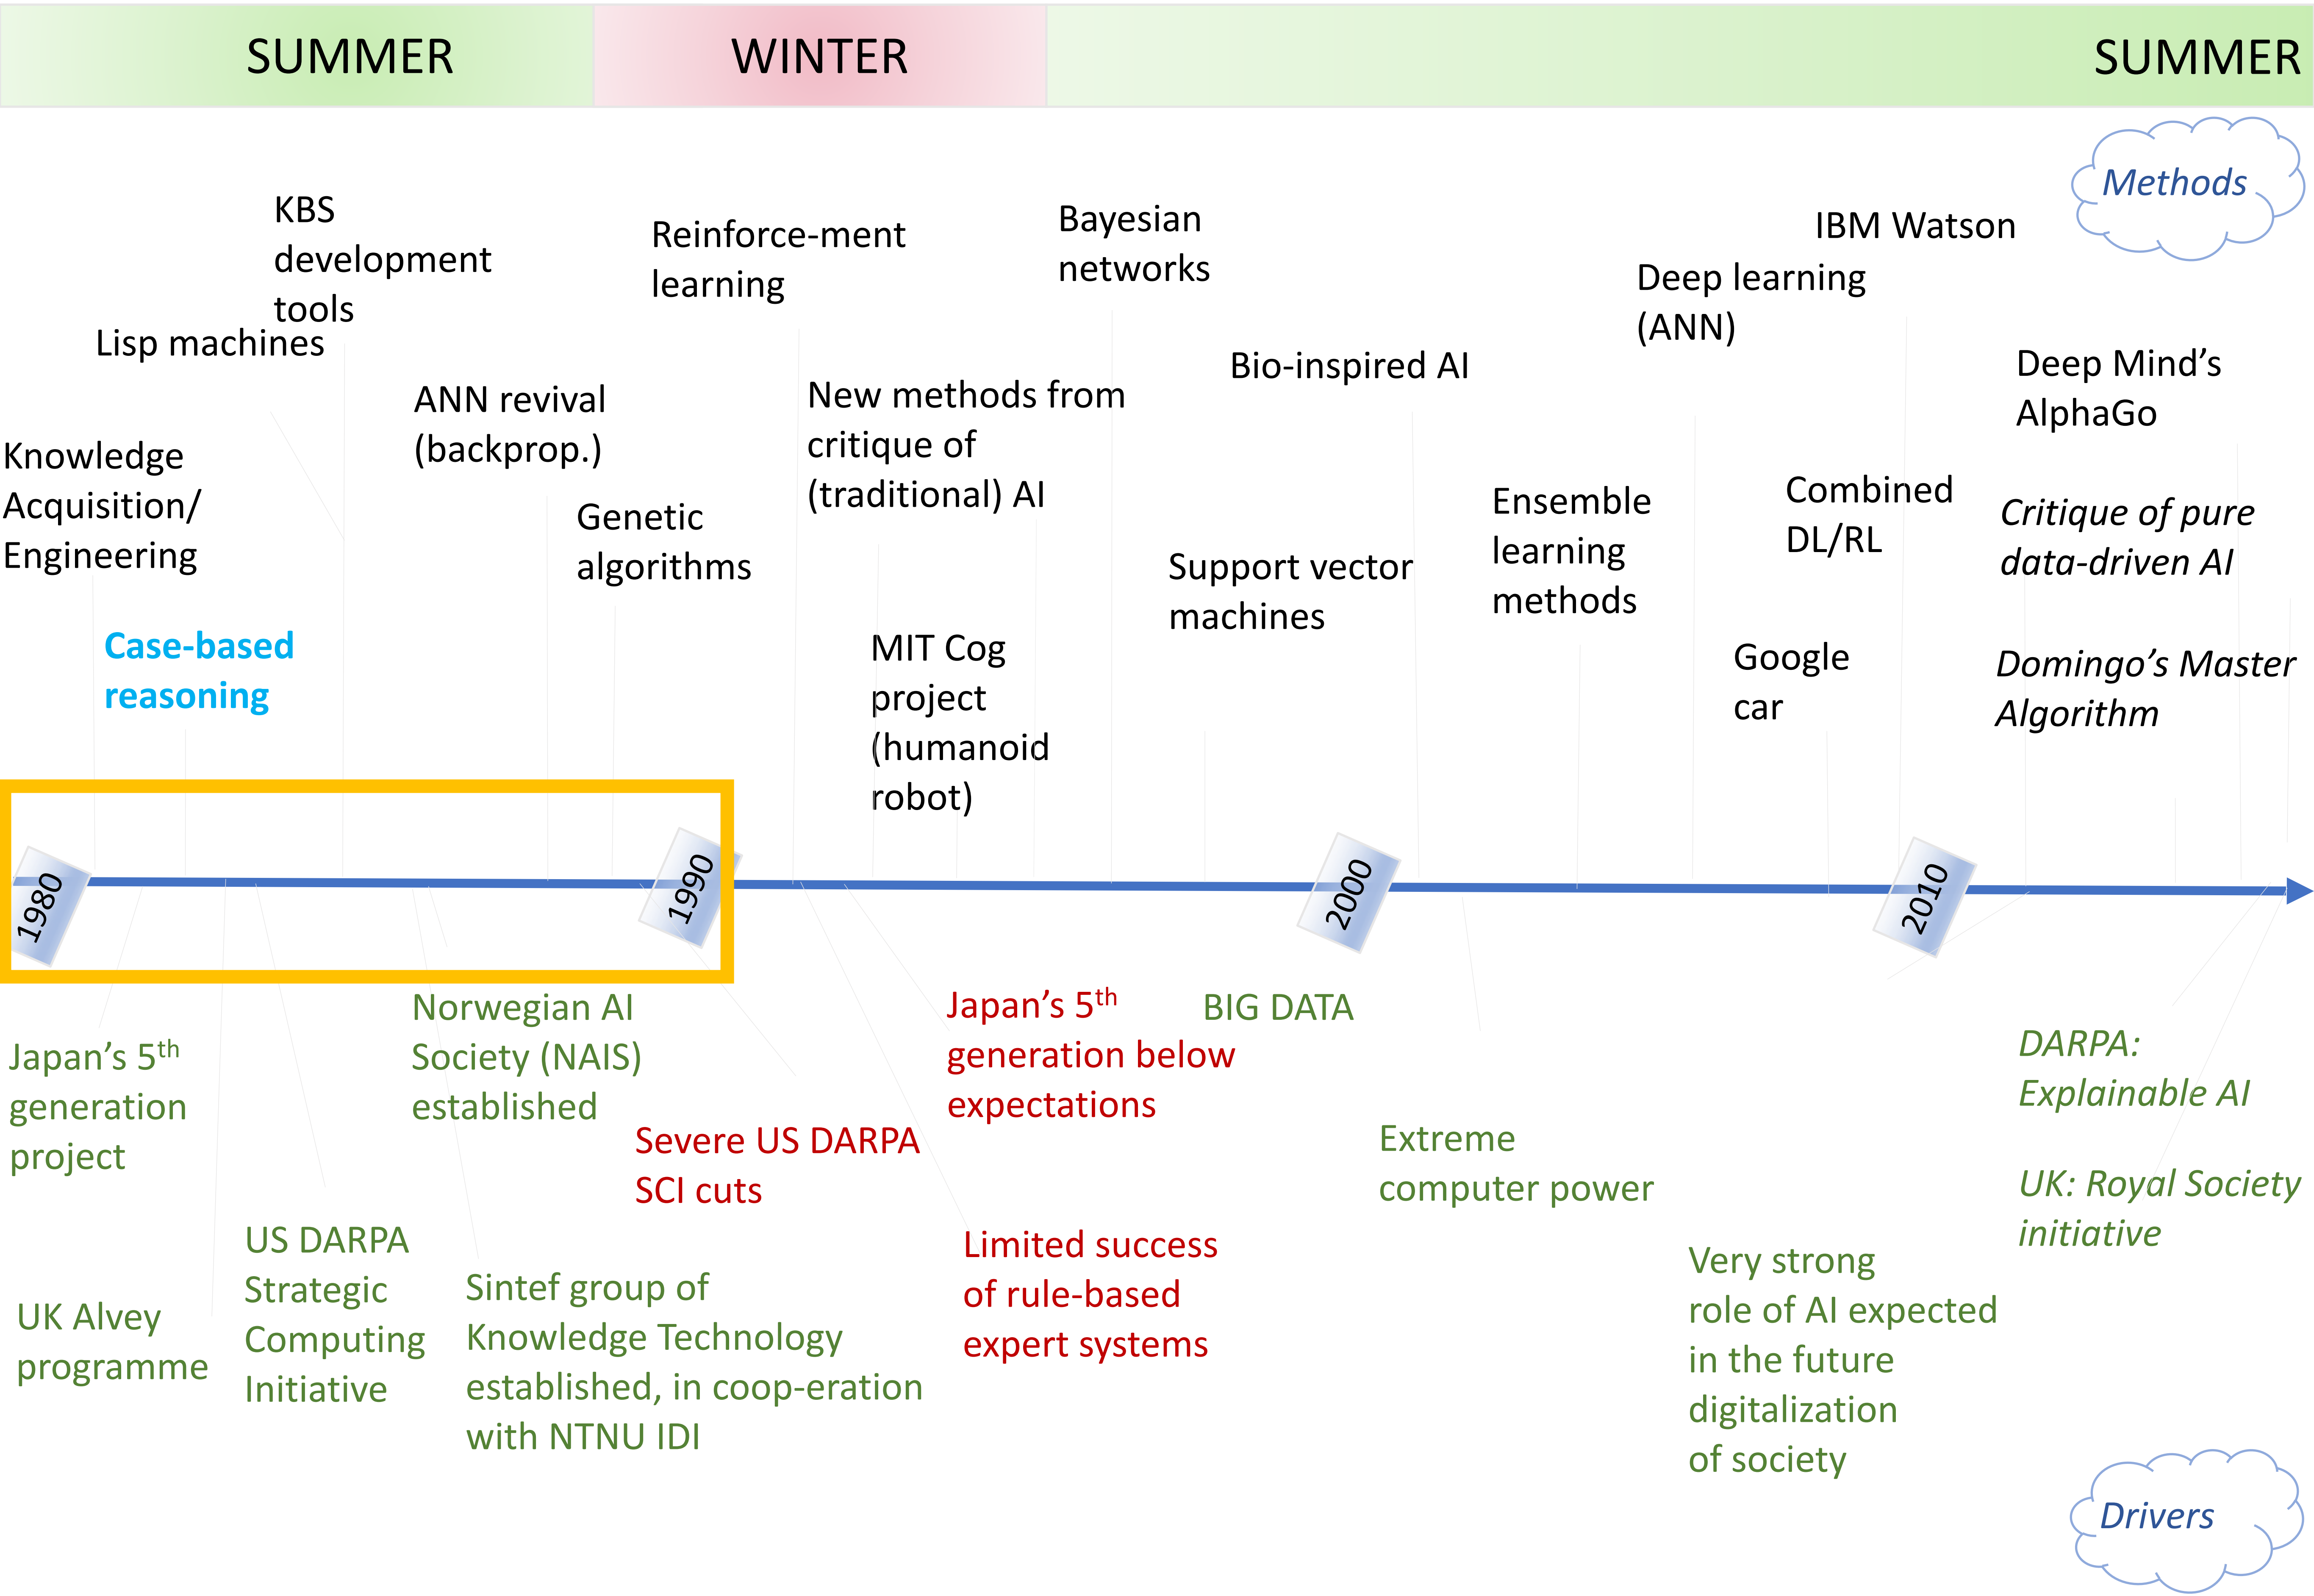
SUMMER
WINTER
SUMMER
Methods
KBS development tools
Lisp machines
Knowledge Acquisition/ Engineering
Case-based reasoning
ANN revival (backprop.)
Genetic algorithms
Reinforce-ment learning
New methods from critique of (traditional) AI
MIT Cog project (humanoid robot)
Bayesian networks
Bio-inspired AI
Support vector machines
Ensemble learning methods
Deep learning (ANN)
Google car
IBM Watson
Combined DL/RL
Deep Mind's AlphaGo
Critique of pure data-driven AI
Domingo's Master Algorithm
1980
1990
2000
2010
Japan's 5th generation project
UK Alvey programme
US DARPA Strategic Computing Initiative
Norwegian AI Society (NAIS) established
Sintef group of Knowledge Technology established, in coop-eration with NTNU IDI
Severe US DARPA SCI cuts
Japan's 5th generation below expectations
Limited success of rule-based expert systems
BIG DATA
Extreme computer power
Very strong role of AI expected in the future digitalization of society
DARPA: Explainable AI
UK: Royal Society initiative
Drivers

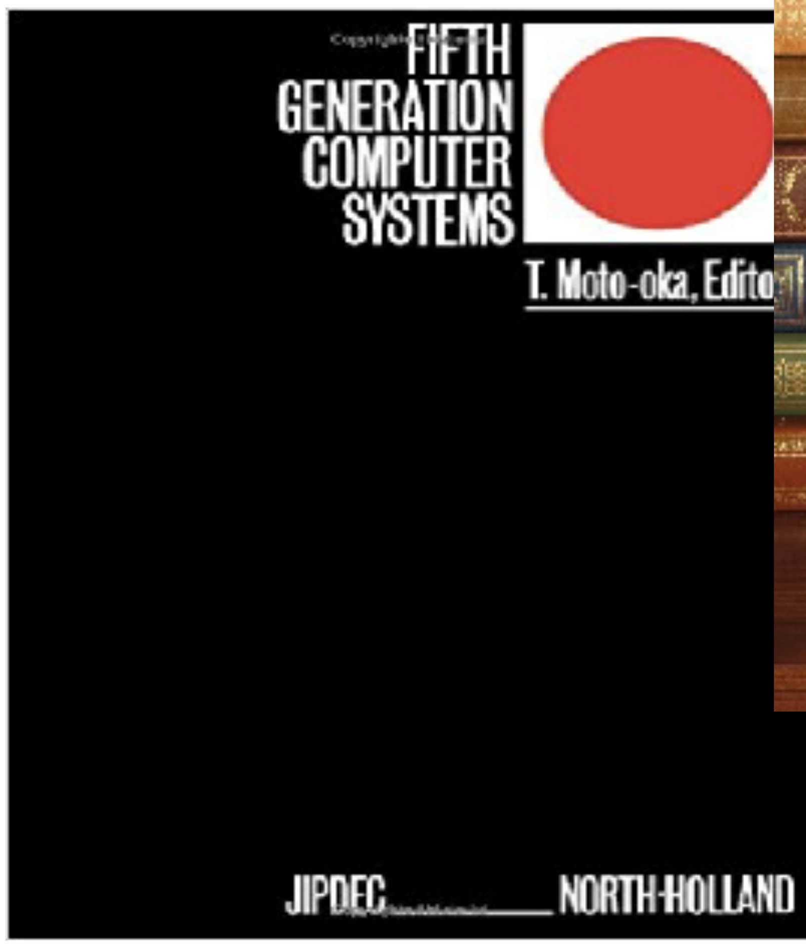

T. Moto-Oka, 1982

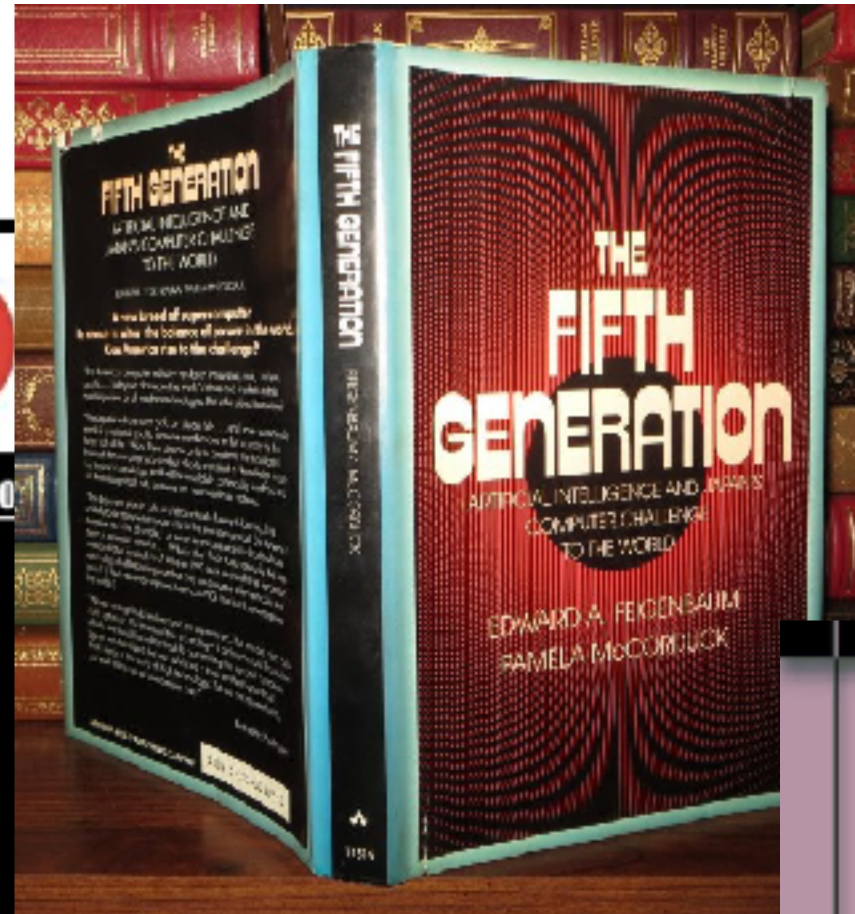

eigenbaum/McCorduck, 1983

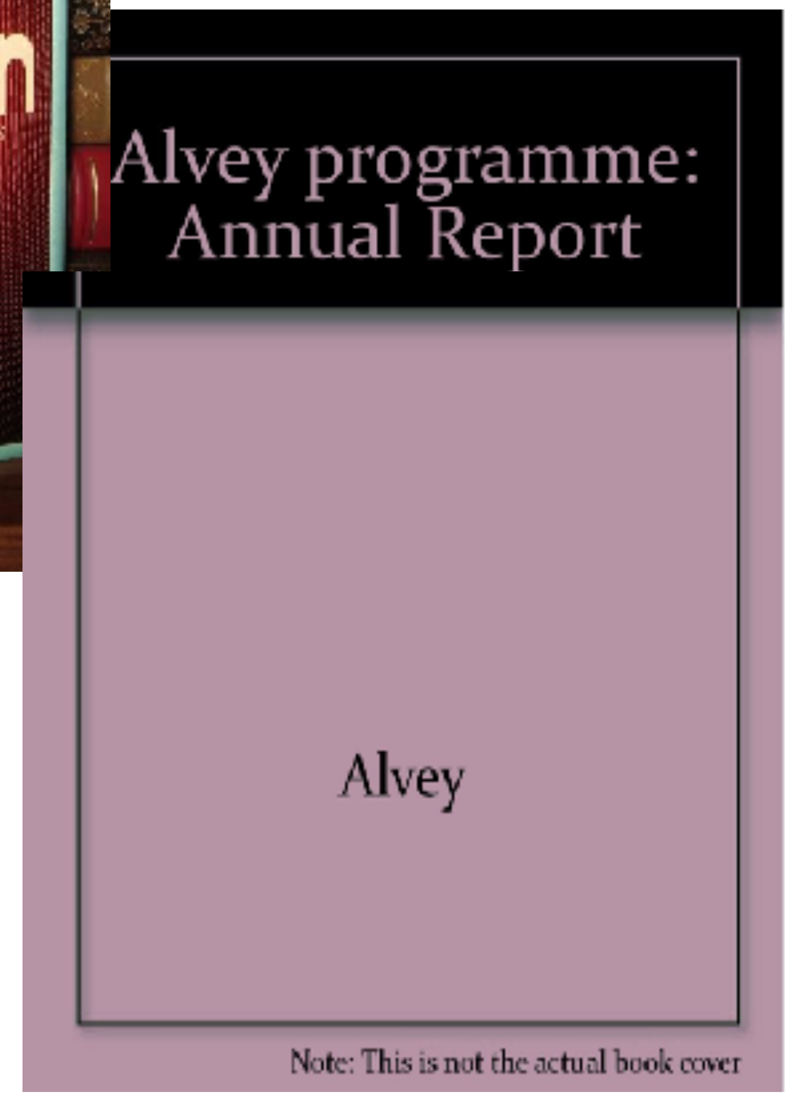

Some early CBR developments

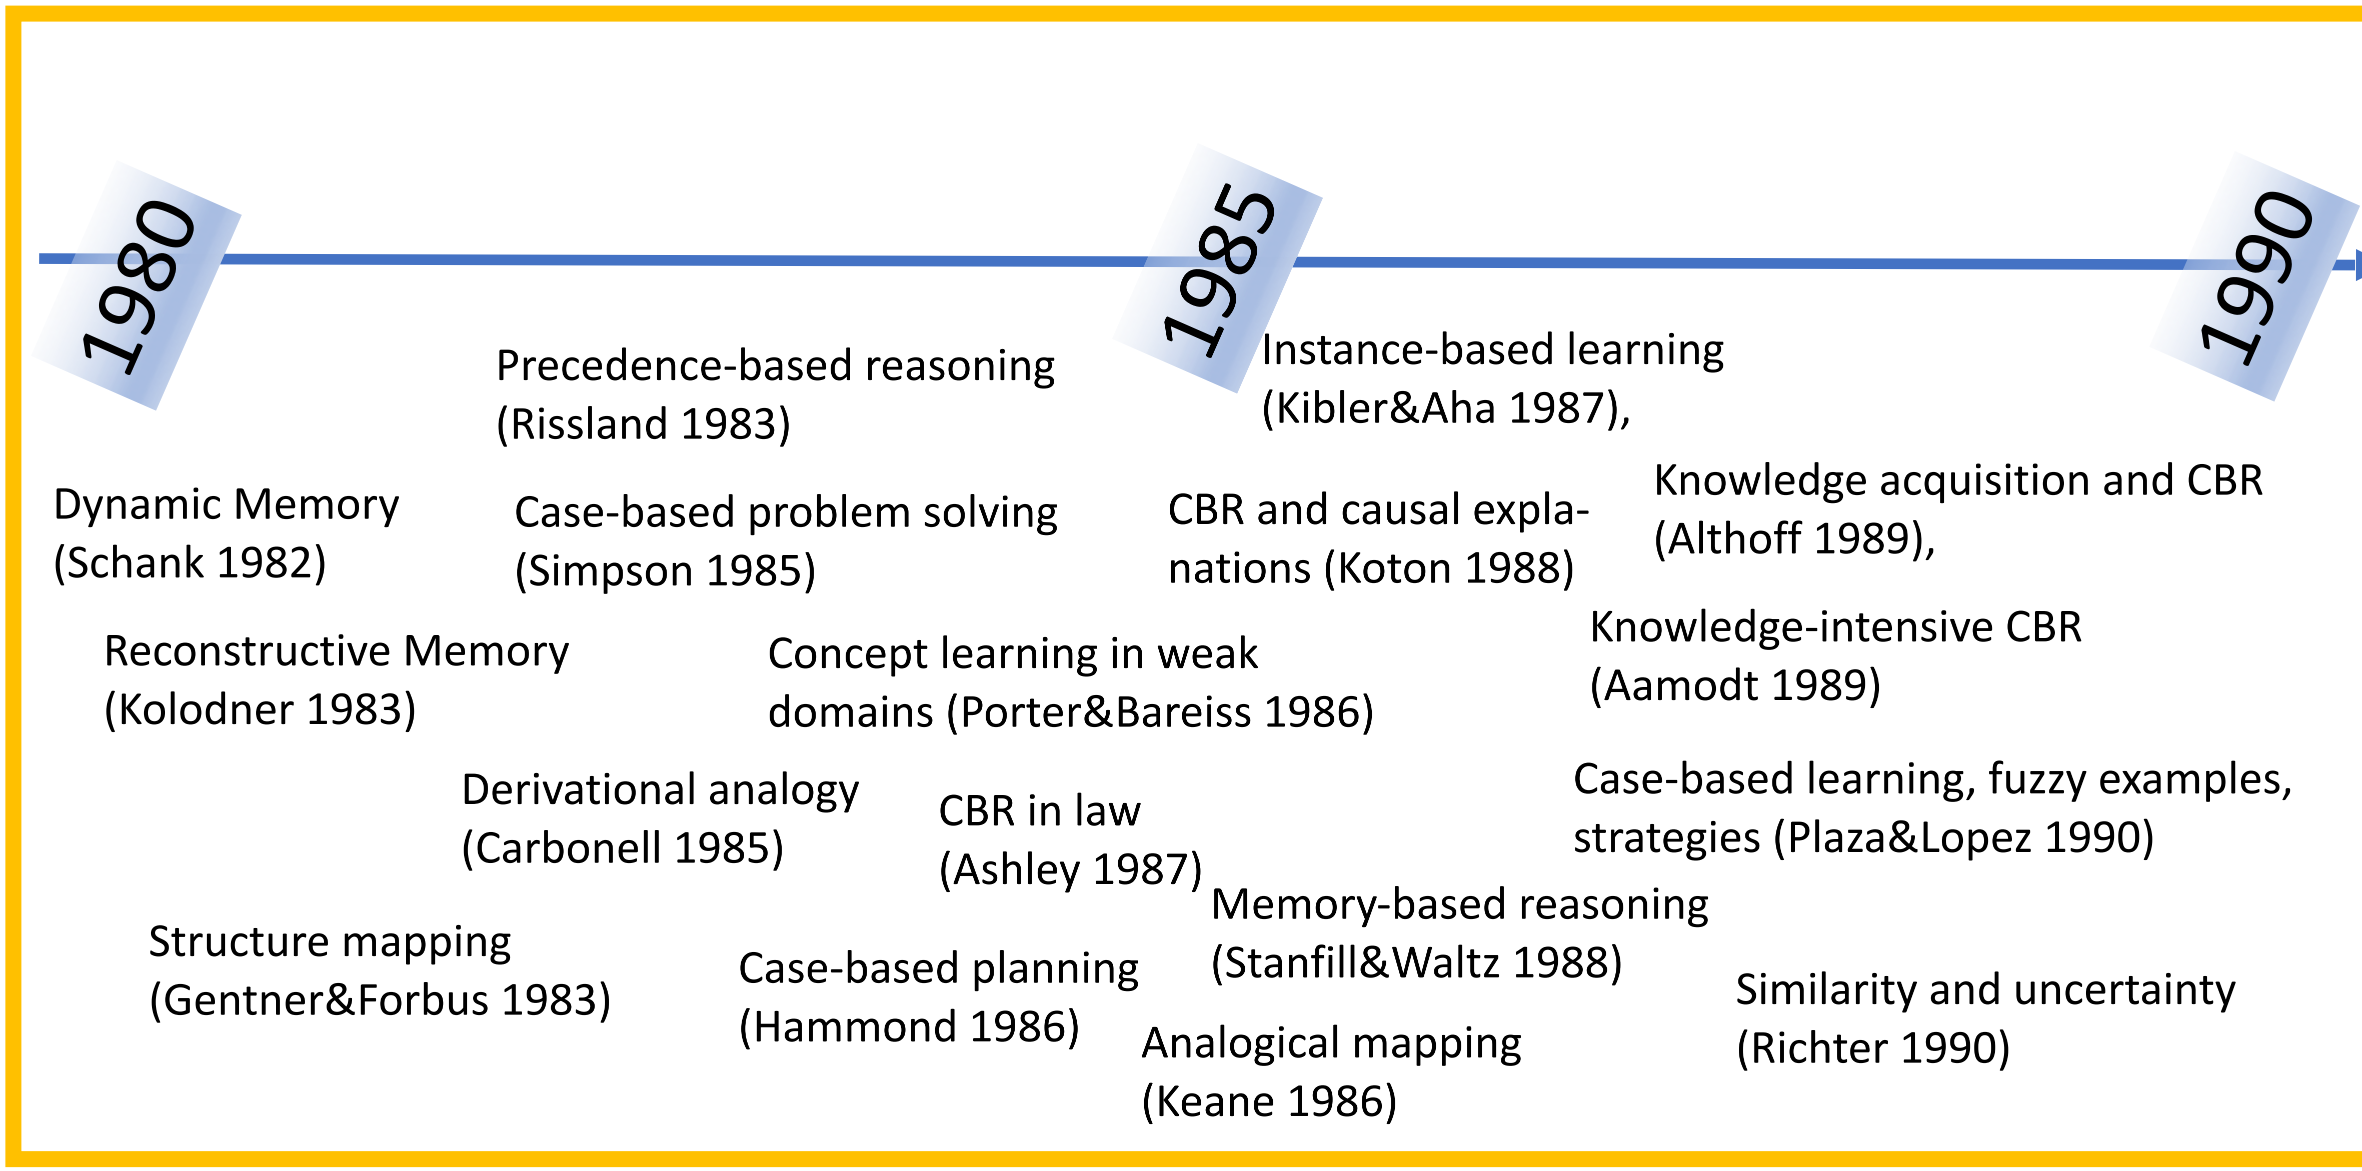

Main CBR foci:

- The Schankian school – theory of reminding and learning
- Instance-based learning
- Similarity
- Explanations
- Knowledge-intensive CBR
- Analogy reasoning


Case-Based Reasoning
Experiences, Lessons, & Future Directions
Case-Based Reasoning Workshop
May 1991
Sponsored by:
Defense Research Projects Agency
Information Science and Technology Office
Michael M. Richter
Rosina O. Weber
Case-Based Reasoning
A Textbook
Springer
Dynamic memory
A theory of reminding and learning in computers and people
ROGER C. SCHANK
LNAI State-of-the-Art Survey
Mario Lenz Brigitte Bartsc
Hans-Dieter Burkhard Ste
Springer
Copyrighted Material
CASE-BASED REASONING
Janet Kolodner
There is no proper reception or public entry to the judges' lobby
Separate circulation for the public, staff, and prisoners yields privacy and security
Proceedings
Techniques for Enterprise Systems

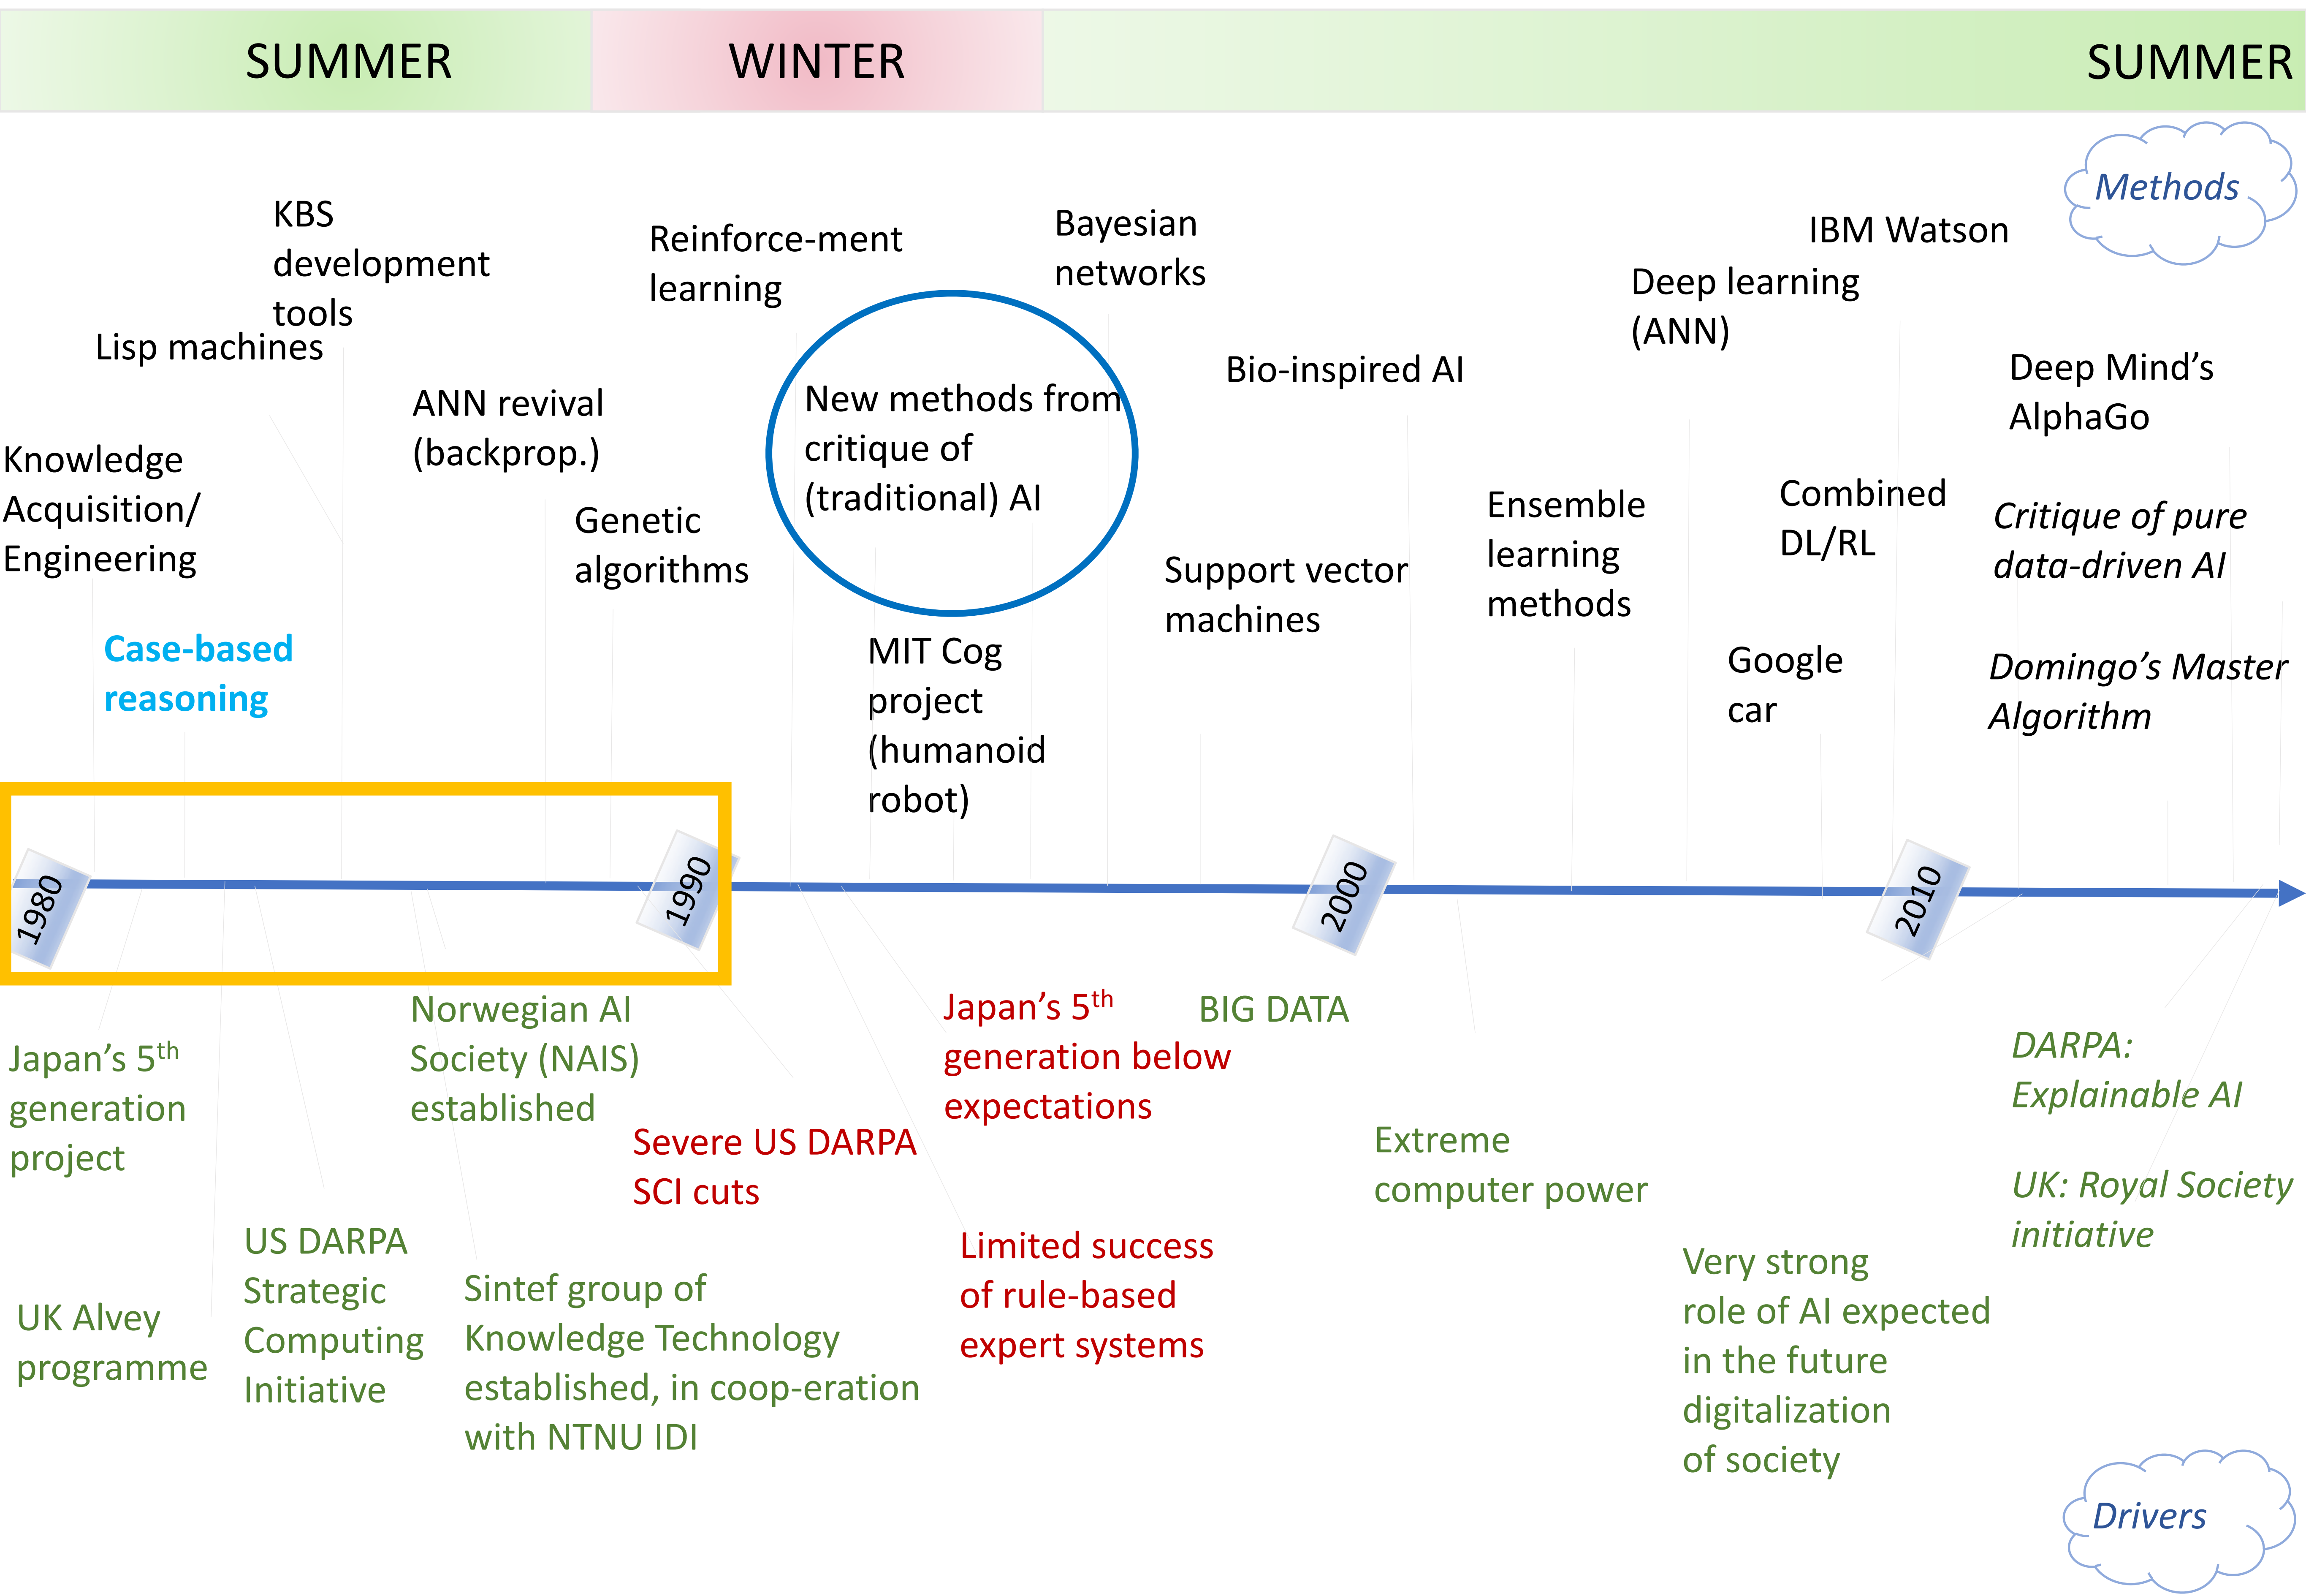
SUMMER
WINTER
SUMMER
Methods
KBS development tools
Reinforce-ment learning
Bayesian networks
IBM Watson
Deep learning (ANN)
Lisp machines
ANN revival (backprop.)
New methods from critique of (traditional) AI
Bio-inspired AI
Deep Mind's AlphaGo
Knowledge Acquisition/ Engineering
Genetic algorithms
Ensemble learning methods
Combined DL/RL
Critique of pure data-driven AI
Support vector machines
Case-based reasoning
MIT Cog project (humanoid robot)
Google car
Domingo's Master Algorithm
1980
1990
2000
2010
Norwegian AI Society (NAIS) established
Japan's 5th generation below expectations
BIG DATA
DARPA: Explainable AI
Japan's 5th generation project
Severe US DARPA SCI cuts
Extreme computer power
UK: Royal Society initiative
US DARPA Strategic Computing Initiative
Limited success of rule-based expert systems
Very strong role of AI expected in the future digitalization of society
UK Alvey programme
Sintef group of Knowledge Technology established, in coop-eration with NTNU IDI
Drivers

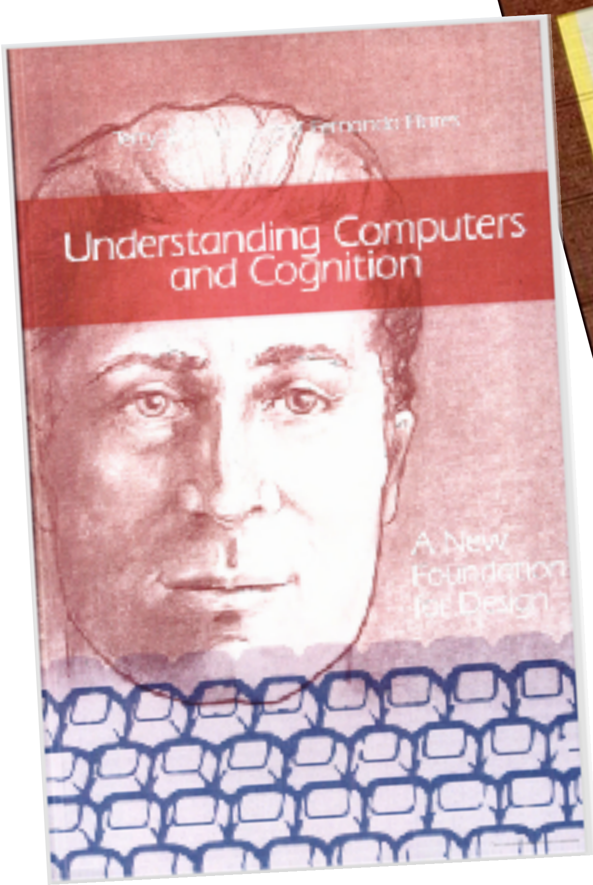

Winograd & Flores, 1987

1988

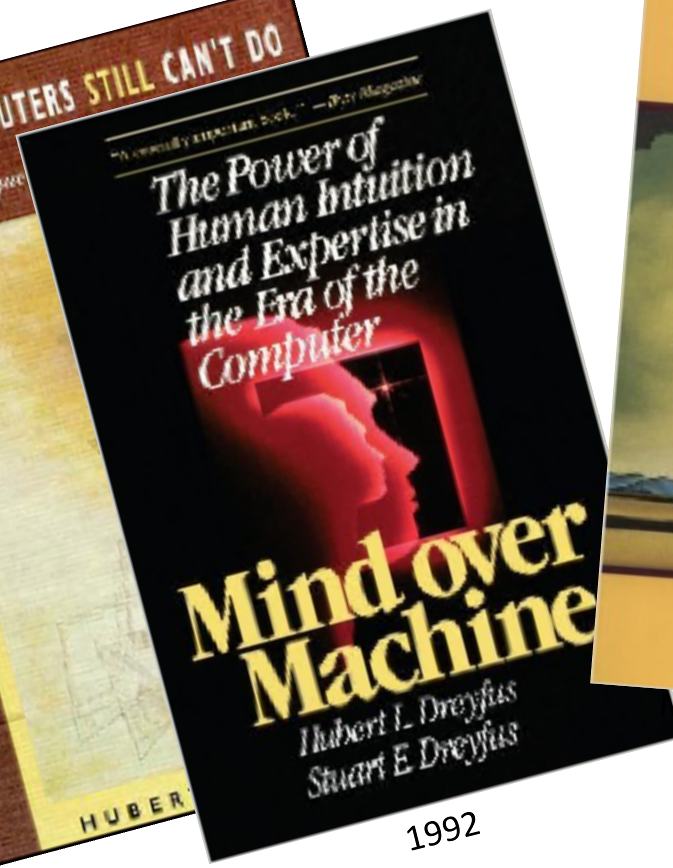

Dreyfus & Dreyfus

1992

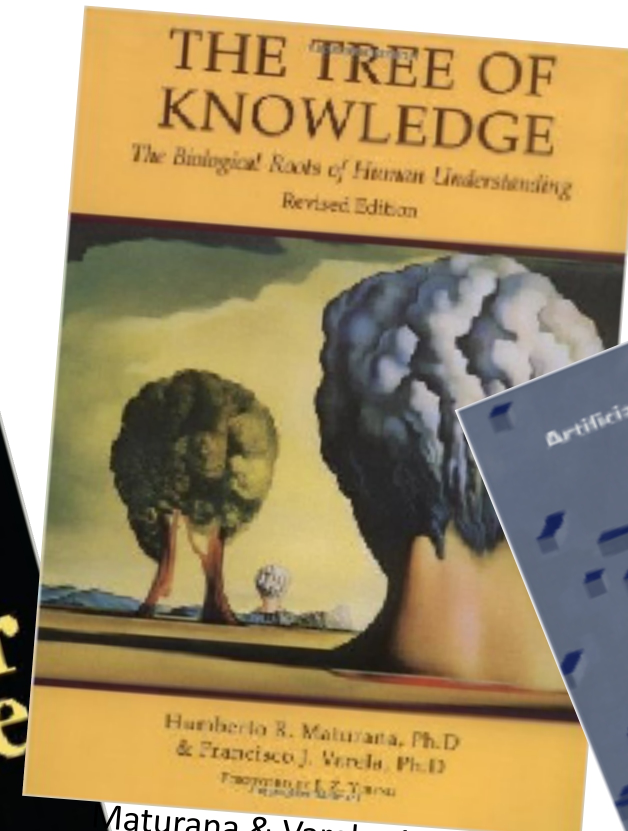

Maturana & Varela, 1992

Artificial Life Journal, 1993-

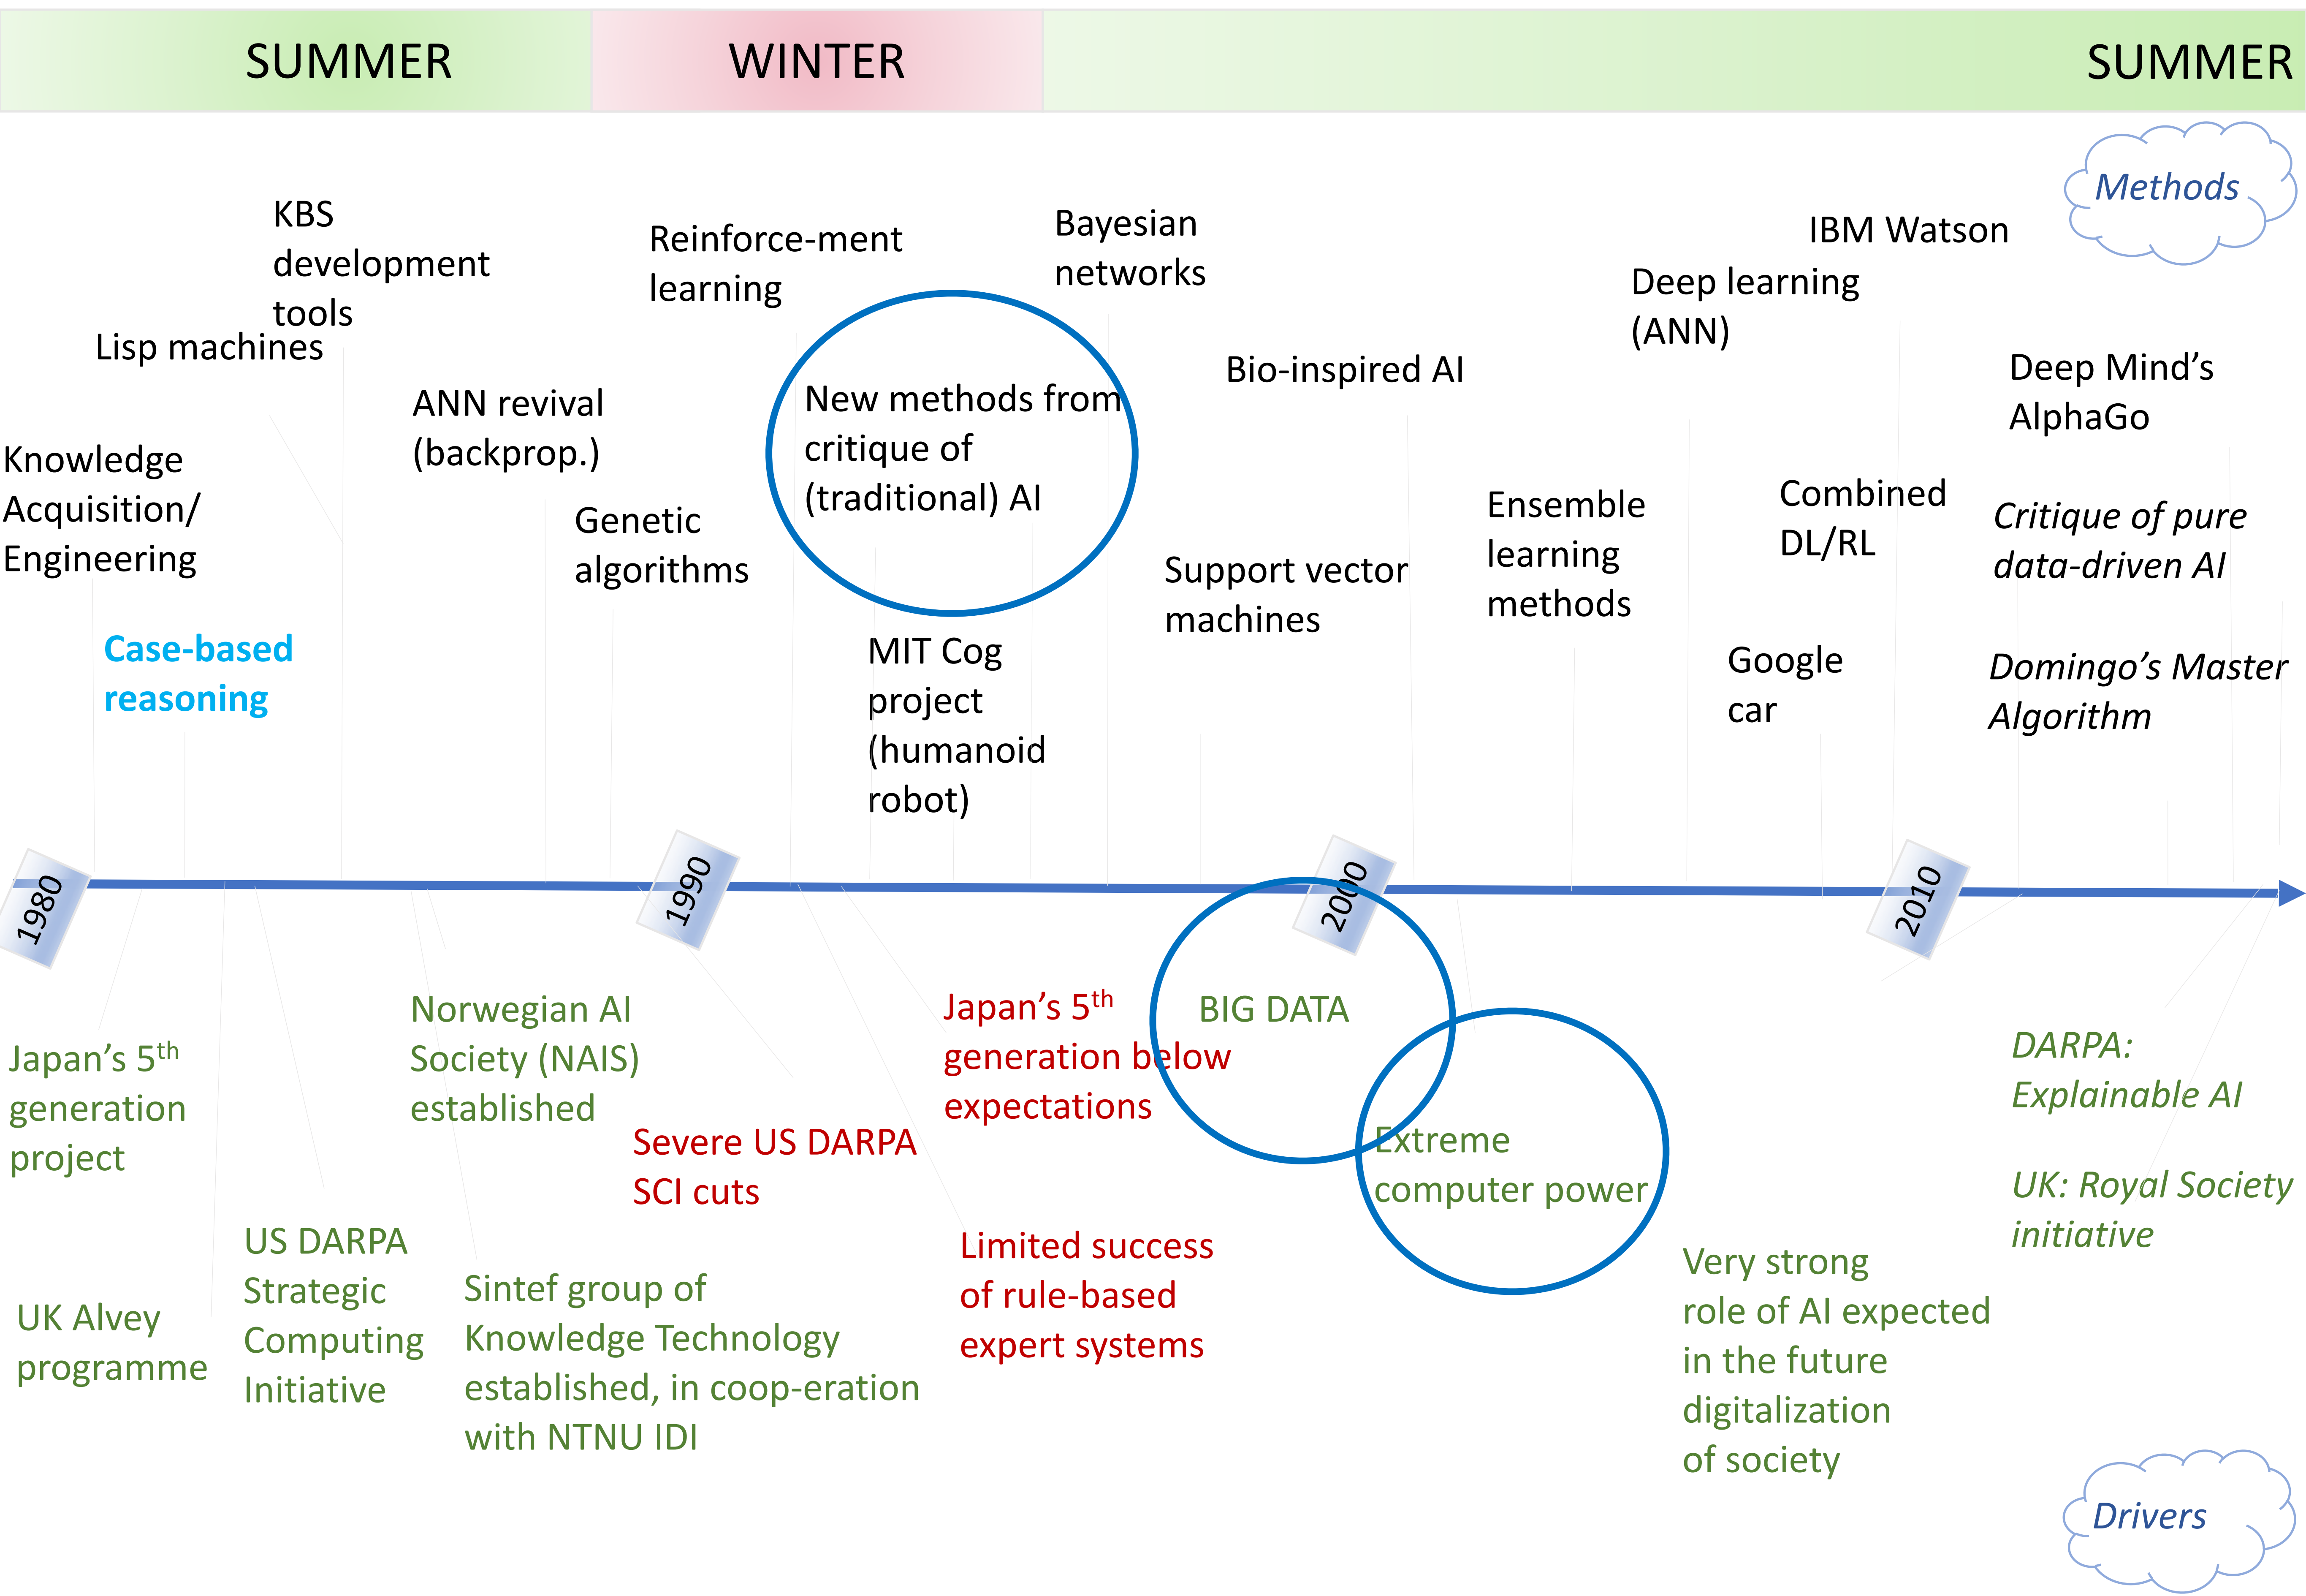

SUMMER
WINTER
SUMMER
Methods
KBS development tools
Lisp machines
Knowledge Acquisition/ Engineering
Case-based reasoning
ANN revival (backprop.)
Genetic algorithms
Reinforce-ment learning
New methods from critique of (traditional) AI
MIT Cog project (humanoid robot)
Bayesian networks
Bio-inspired AI
Support vector machines
Ensemble learning methods
Deep learning (ANN)
IBM Watson
Combined DL/RL
Google car
Deep Mind's AlphaGo
Critique of pure data-driven AI
Domingo's Master Algorithm
1980
1990
2000
2010
Japan's 5th generation project
UK Alvey programme
US DARPA Strategic Computing Initiative
Norwegian AI Society (NAIS) established
Sintef group of Knowledge Technology established, in coop-eration with NTNU IDI
Severe US DARPA SCI cuts
Japan's 5th generation below expectations
Limited success of rule-based expert systems
BIG DATA
Extreme computer power
Very strong role of AI expected in the future digitalization of society
DARPA: Explainable AI
UK: Royal Society initiative
Drivers

Knowledge-based vs. Data-driven CBR

Different approaches to how to capture experiences

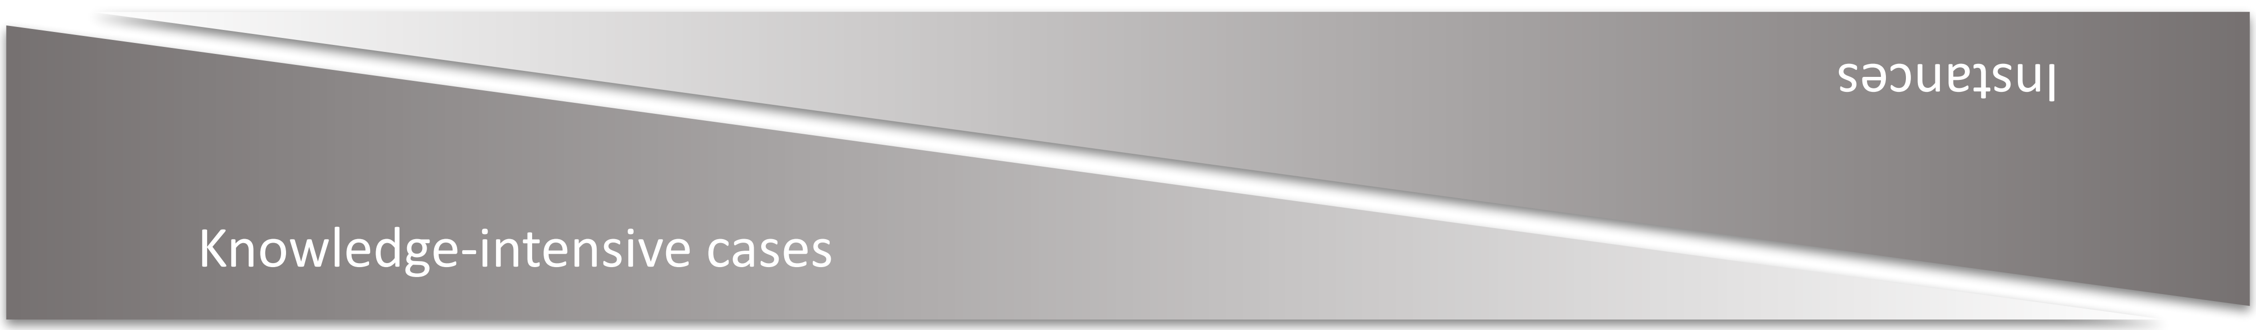

- Substantial generalized knowledge
- Few, comprehensive cases
- A case is a user experience

- Complex case structures
- Similarity assessment is an explanation
- Knowledge-based adaptation
- Knowledge-based learning

- No explicit generalized knowledge
- Many cases
- A case is a data record

- Simple case structures
- Global similarity metric
- No adaptation
- Learning by storing cases

Contentions

All currently succesful AI techniques are not **new**: basic ideas at least a decade old (with only small improvements)

Success of "killer apps" supports AI, but main differences are:

1) **lots** of data
2) high performance **computing**

“Silver bulletism”

“As a field I believe that we tend to suffer from what might be called *serial silver bullitism*, defined as follows:

the tendency to believe in a silver bullet for AI, coupled with the belief that previous beliefs about silver bullets were hopelessly naïve.”

(Hector Levesque, Research Excellence Lecture, IJCAI 2013)

No tricks!

- We should avoid being overly swayed by what appears to be the most promising approach of the day.
- We need to return to our roots in Knowledge Representation and Reasoning *for* language and *from* language.
- We should carefully study how simple knowledge bases might be used to make sense of the simple language needed to build slightly more complex knowledge bases, and so on.
- It is not enough to build knowledge bases without paying closer attention to the demands arising from their use.
- We should explore more thoroughly the space of computations between fact retrieval and full automated logical reasoning.

(Levesque, 2013)

The Cognitive Systems Movement

The Cognitive Systems Movement

Most of the original challenges still remain and provide many opportunities for research.

Because "AI" now has such limited connotations, we need a new label for research that:

- *De… …ies computational artifacts that … …n intelligence.*

…*gnitive systems*, a term promoted …)2).

Advances in Cognitive Systems 1 (2012) 1–2
Submitted 7/2012; published 7/2012

A New Journal for Cognitive Systems

Pat Langley PATRICK.W.LANGLEY@GMAIL.COM
Computing Science and Engineering, Arizona State University, Tempe, AZ 85287 USA
Computer Science Department, University of Auckland, Private Bag 92019, Auckland, New Zealand

This is the inaugural volume of a new journal, *Advances in Cognitive Systems*. The publication's main purpose is to communicate research progress in the field of cognitive systems, which aims to understand, in computational terms, and reproduce, in computational artifacts, the entire range of intelligent behavior observed in humans. As such, it provides a meeting place for those who remain committed to the original vision of artificial intelligence, from which it which inherits these audacious but exciting research goals.

New journals should be created rarely and only when three conditions hold. First, there should be a well-specified area of research that is broad enough to encompass many different approaches but still unified enough in aims and assumptions to support common ground among its contributors. Second, there should be a substantial body of ongoing work in this area, so that there will be enough material to publish on a sustainable basis. Finally, authors in the area should have encountered difficulty publishing their results in standard outlets, so that they need a new home for their research. I believe that all three conditions hold for the field of cognitive systems.

Pat Langley, Cognitive Systems Institute, Dec. 2015

Artificial General Intelligence

Artificial General Intelligence (AGI) is an emerging field aiming at the building of "thinking machines", that is, general-purpose systems with intelligence comparable to that of the human mind. While this was the original goal of Artificial Intelligence (AI), the mainstream of AI research has turned toward domain-dependent and problem-specific solutions; therefore it has become necessary to use a new name to indicate research that still pursues the "Grand AI Dream". Similar labels for this kind of research include "Strong AI", "Human- level AI", etc.

Journal of Artificial General Intelligence 5(1) 1-46, 2014
DOI: 10.2478/jagi-2014-0001
Submitted 2013-2-12
Accepted 2014-3-15

Artificial General Intelligence: Concept, State of the Art, and Future Prospects

BEN@GOERTZEL.ORG
Ben Goertzel
OpenCog Foundation
G/F, 51C Lung Mei Village
Tai Po, N.T., Hong Kong

Editor: Tsvi Achler

Abstract

In recent years broad community of researchers has emerged, focusing on the original ambitious goals of the AI field – the creation and study of software or hardware systems with general intelligence comparable to, and ultimately perhaps greater than, that of human beings. This paper surveys this diverse community and its progress. Approaches to defining the concept of Artificial General Intelligence (AGI) are reviewed including mathematical formalisms, engineering, and biology inspired perspectives. The spectrum of designs for AGI systems includes systems with symbolic, emergentist, hybrid and universalist characteristics. Metrics for general intelligence are evaluated, with a conclusion that, although metrics for assessing the achievement of human-level AGI may be relatively straightforward (e.g. the Turing Test, or a robot that can graduate from elementary school or university), metrics for assessing partial progress remain more controversial and problematic.

Keywords: AGI, general intelligence, cognitive science

1. Introduction

How can we best conceptualize and approach the original problem regarding which the AI field was founded: the creation of thinking machines with general intelligence comparable to, or greater than, that of human beings? The standard approach of the AI discipline (Russell and Norvig, 2010), as it has evolved in the 6 decades since the field's founding, views artificial intelligence largely in terms of the pursuit of discrete capabilities or specific practical tasks. But while this approach has yielded many interesting technologies and theoretical results, it has proved relatively unsuccessful in terms of the original central goals of the field.

Ray Kurzweil (Kurzweil, 2005) has us... "narrow AI" to refer to the creation of systems that carry out specific "intelligent" ...ntexts. For a narrow AI system,

We do not know how far we are from Artificial General Intelligence

Published on August 1, 2016

André Luis Paraense Follow
Postdoctoral Research Scientist @ UNICAMP / Mobile Full St...

The main reason for this lack of previsibility is simple: there are not enough people working on AGI, because most of the people are working on some kind of "narrow AI", aimed at resolving specific tasks.

KIMERA THE AGI COMPANY

What is artificial general intelligence? And does Kimera Systems' 'Nigel' qualify?

The creation of artificial general intelligence would be one of the biggest breakthr...

As an *artificial general intelligence* ... through an ever-growing neural netw... extend to any type of connected device and across ...

Google DeepMind publishes breakthrough Artificial General Intelligence architecture

INTELLIGENCE AND THE SENSES 15th March 2017 Matthew Griffin

…merge Modular **Deep Learning, Meta-Learning** and **Reinforcement Learning** into a single solution

…DeepMind's stab to become the first company to build the first, fabled, **Artificial General Intelligence** (AGI) solution.

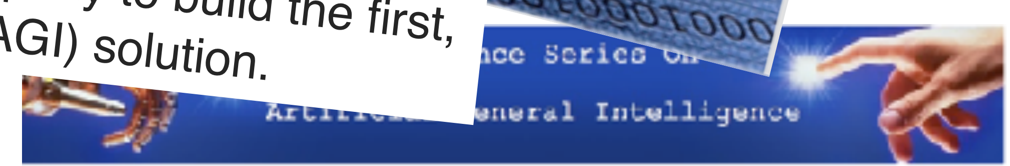

Mapping the Landscape of Human-Level Artificial General Intelligence

Sam S. Adams, Itamar Arel, Joscha Bach, Robert Coop, Rod Furlan, Ben Goertzel, J. Storrs Hall, Alexei Samsonovich, Matthias Scheutz, Matthew Schlesinger, Stuart C. Shapiro, John F. Sowa

AI MAGAZINE SPRING 2012

C1. The environment is complex, with diverse, interacting and richly structured objects.

C2. The environment is dynamic and open.

C3. Task-relevant regularities exist at multiple time scales.

C4. Other agents impact performance.

C5. Tasks can be complex, diverse and novel.

C6. Interactions between agent, environment and tasks are complex and limited.

C7. Computational resources of the agent are limited.

C8. Agent existence is long-term and continual.

Figure 1. Characteristics for AGI Environments, Tasks, and Agents.

R0. New tasks do not require re-programming of the agent

R1. Realize a symbol system

Represent and effectively use:

- R2. Modality-specific knowledge
- R3. Large bodies of diverse knowledge
- R4. Knowledge with different levels of generality
- R5. Diverse levels of knowledge
- R6. Beliefs independent of current perception
- R7. Rich, hierarchical control knowledge
- R8. Meta-cognitive knowledge

R9. Support a spectrum of bounded and unbounded deliberation

R10. Support diverse, comprehensive learning

R11 Support incremental, online learning

Figure 2. Cognitive Architecture Requirements for AGI.

Daniel Dennett:

Competence vs. Comprehension

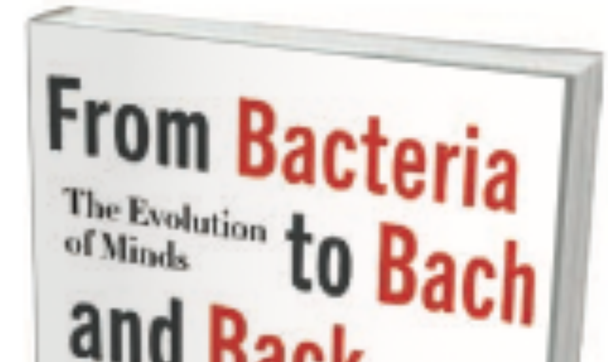

“competence without comprehension”

Just as computers can perform complex calculations without understanding arithmetic, so creatures can display finely tuned behaviour without understanding why they do so.

DANIEL C. DENNETT
W. W. Norton: 2017.

Does comprehension matter?

Do we want post-comprehension “science”?

Technological competence without comprehension?

Or DARPA’s “explainable AI” initiative . . . ?

Should we try to make persons out of them?

PRO: So they can explain their reasoning to us.

So they can develop their own imaginative curiosity, and epistemic goals.

CON: They will blur the lines of moral responsibility.

Dennett, Google Talk, Feb. 2017

Recent initiatives

The US

Broad Agency Announcement
Explainable Artificial Intelligence (XAI)
DARPA-BAA-16-53
August 10, 2016

DARPA

Defense Advanced Research Projects Agency
Information Innovation Office
675 North Randolph Street
Arlington, VA 22203-2114

The UK

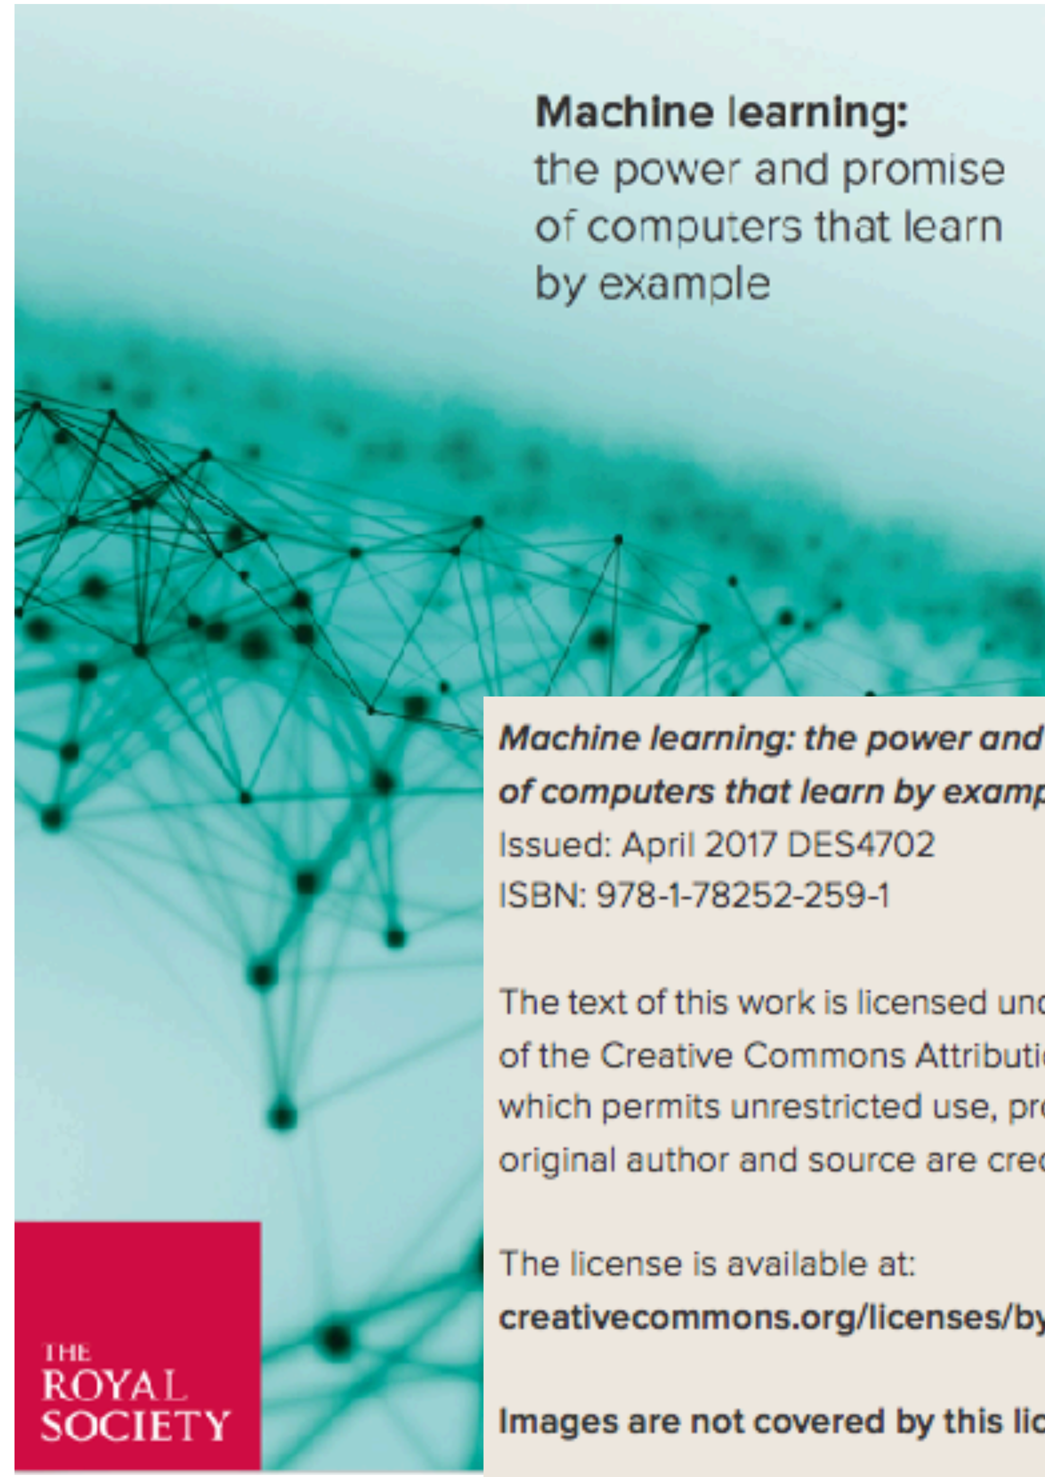

Machine learning: the power and promise of computers that learn by example

Issued: April 2017 DES4702
ISBN: 978-1-78252-259-1

This report can be viewed online at
royalsociety.org/machine-learning

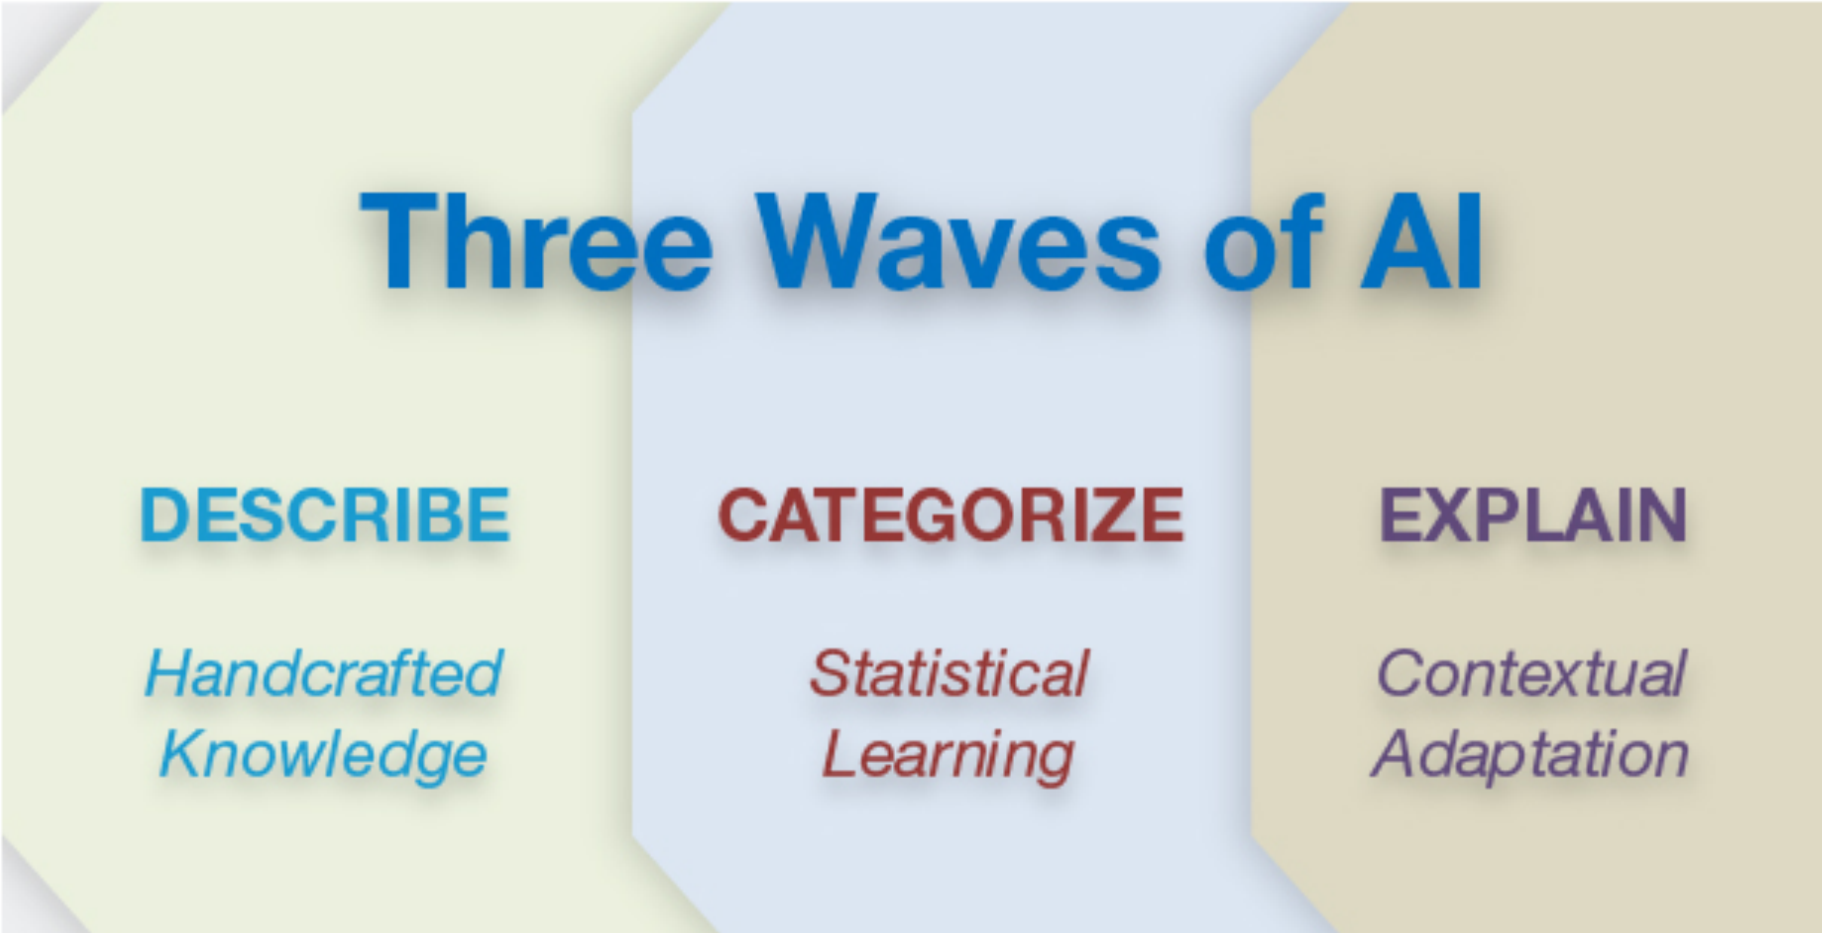

A DARPA Perspective on Artificial Intelligence (John Launchbury, Director I20, DARPA)

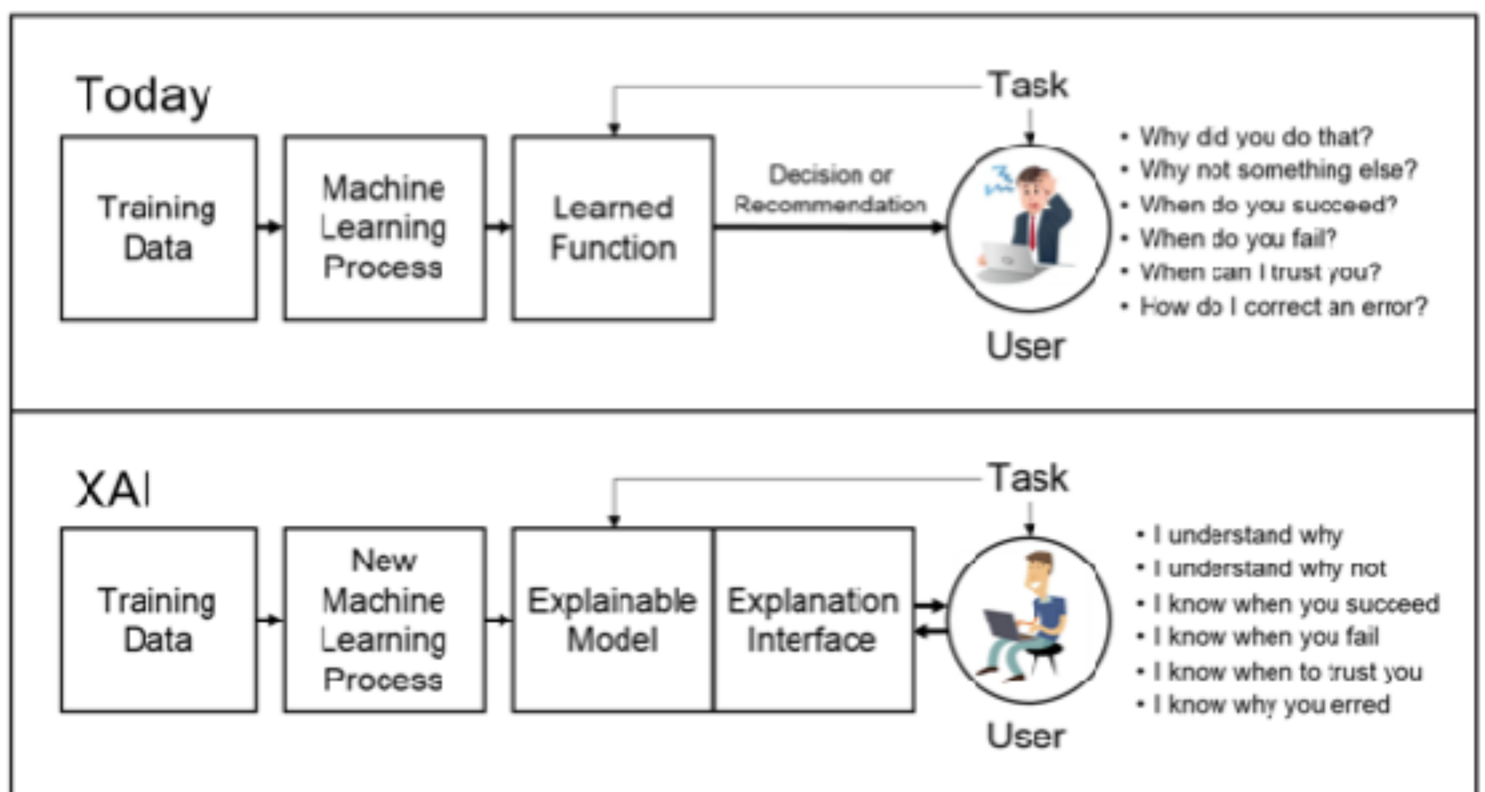

Master Algorithm

The 5 “tribes” of ML (P. Domingos)

Tribe	Origins	Problem	M.A.(solution)
Symbolists	Logic	Knowledge composition	Inverse deduction
Connectionist	Neuroscience	Credit assignment	Backpropagation
Evolutionaries	Evolutionary Biology	Structured discovery	Genetic Programing
Bayesians	Statistics	Uncertainty	Probabilistic Inference
Analogizers	Psychology	Similarity	Kernel machines

CBR? Clustering?

Master Algorithm

The 5 “tribes” of ML (P. Domingos)

Deep Learning (backpropagation)
is most successful now.
Why? What can we learn for its success?

“The analogizers are the least cohesive
of the five tribes”

Pedro Domingos. “The Master Algorithm”.

Master Algorithm

the 5 "tribes" of ML (P. Domingos)

"Perhaps in a future decade, machine learning will be dominated by deep analogy, combining in one algorithm the efficiency of nearest-neighbor, the mathematical sophistication of support vector machines, and the power and flexibility of analogical reasoning"

Pedro Domingos. "The Master Algorithm"

CBR strengths

Integrating learning and problem solving

Artificial Intelligence

Data-driven AI | Knowledge-intensive AI

CBR strengths

Artificial Intelligence

Data-driven AI　　　　Knowledge-intensive AI

Human cognition

System 1　　　　System 2

Thinking, Fast and Slow
Daniel Kahneman

Judgment under Uncertainty: Heuristics and Biases
Daniel Kahneman and Amos Tversky

CBR and cognitive models

Human cognition

System 1	System 2
Fast, automatic, frequent, emotional, stereotypic, subconscious	Slow, effortful, infrequent, logical, calculating, conscious

«Intuition is recognition»
(Herbert Simon)

CBR and cognitive models

Human cognition

System 1

Fast,
automatic,
frequent,
emotional,
stereotypic,
subconscious

Associative Memory (Daniel Kahneman) → Episodes, Stories, Experience

CBR and cognitive models

Challenge

An archive/base of cases $\longrightarrow$ Self-Organising Memory

data-driven similarity creation,
feature reduction & invention,
multi-level representation

Deep Learning Facts

(Joan Serrà, Telefonica R&D)
https://vimeo.com/album/4516480/video/211630902

- High Accuracy/performance
- **No Feature Engineering (ML development speedup)**
- Less “Complexity” and Specifications in Design
- Scale with Large Data (ML performance saturation)
- Transfer Learning (e.g. from numbers to letters)
- Flexibility (combining building blocks)
- “Unlikely” Learning (case-based learning of addition from images)
- Generative by Nature (generate new data similar to input date)

Supervised/Unsupervised Learning

- Unsupervised learning as dimensionality reduction
- Unsupervised learning as feature engineering
- the synergy of combining supervised /unsupervised learning
    - clustering + kNN
    - MF (Matrix factorization) à la PCA
        - Unsupervised = (1) Dimensionality reduction + (2) clustering
        - Supervised (labeled targets ~ regression)

Supervised/Unsupervised Learning

- Core "trick" in Deep Learning is how to combine Supervised/Unsupervised Leaning
  - E.g. Stacked Autoencoders
  - E.g. training convolutional nets

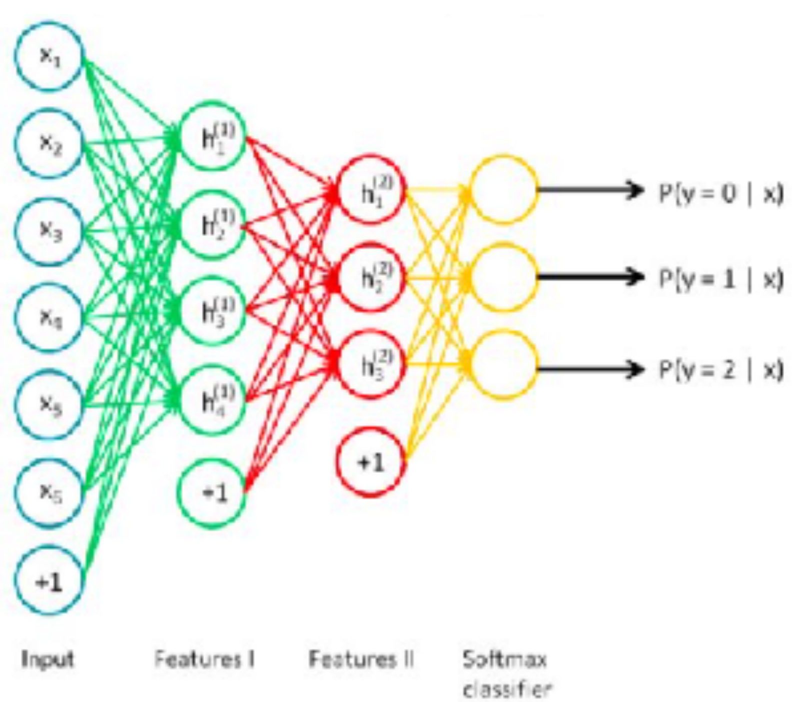

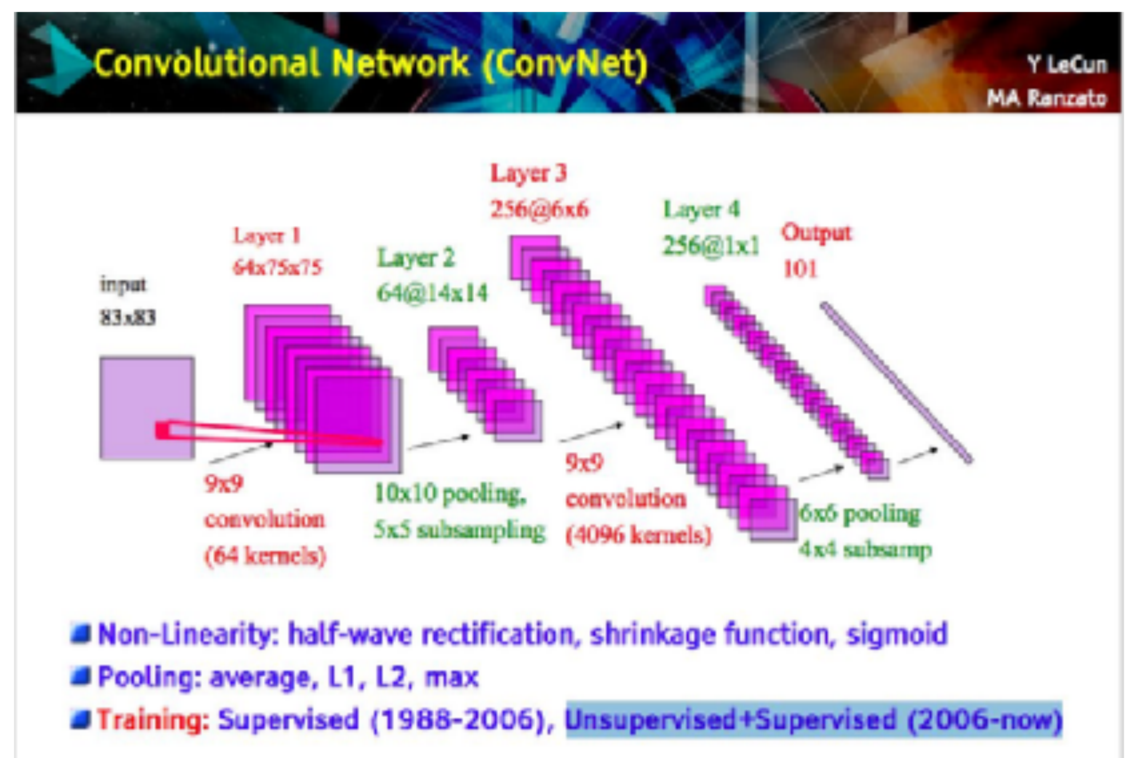

Xavier Amatriain (Quora)

10 More Lessons (Learned from building real-life Machine Learning Systems)

ConvNet

Training using Backpropagation

Process 1: convolution, ReLU (rectified linear unit) and pooling operations along with forward propagation in the Fully Connected layer

Process 2: measure total error and backpropagate (adjust weights)

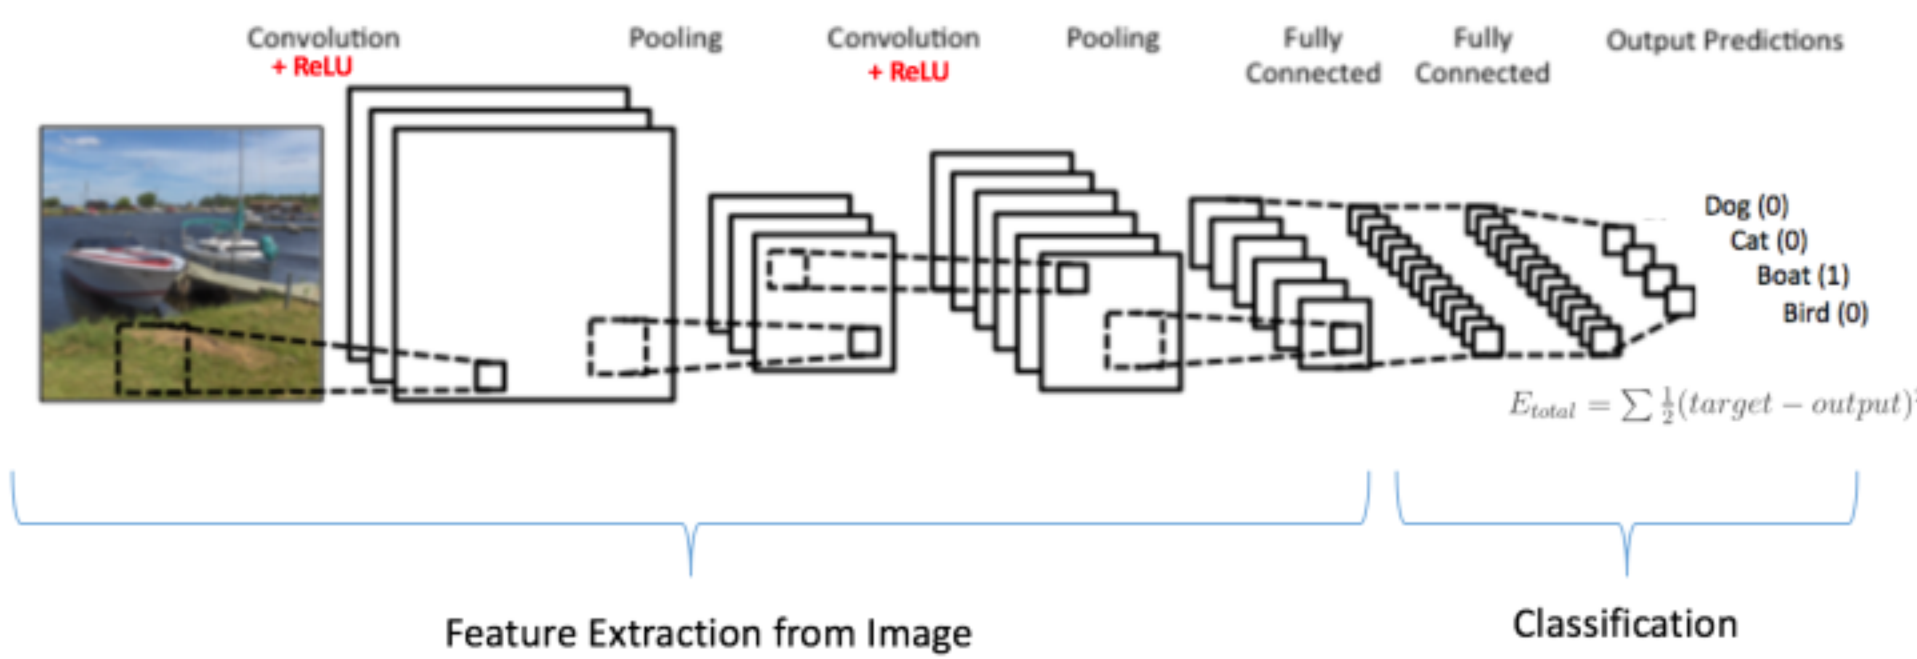

ConvNet

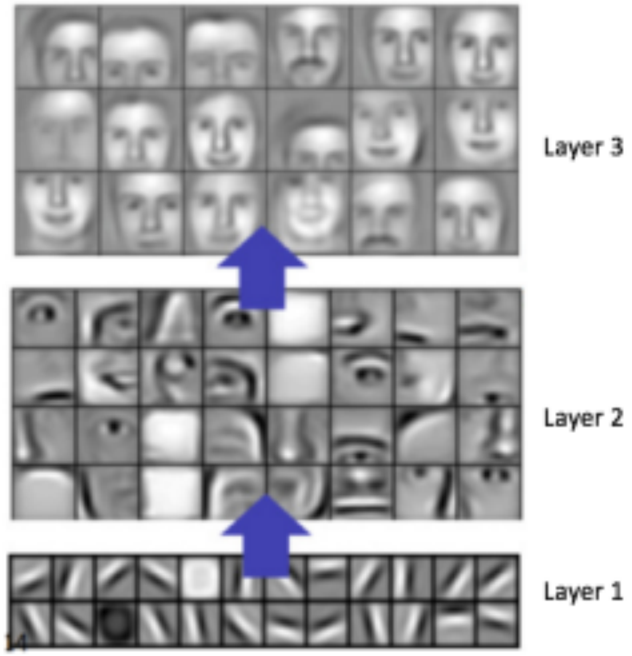
Layer 3
Layer 2
Layer 1

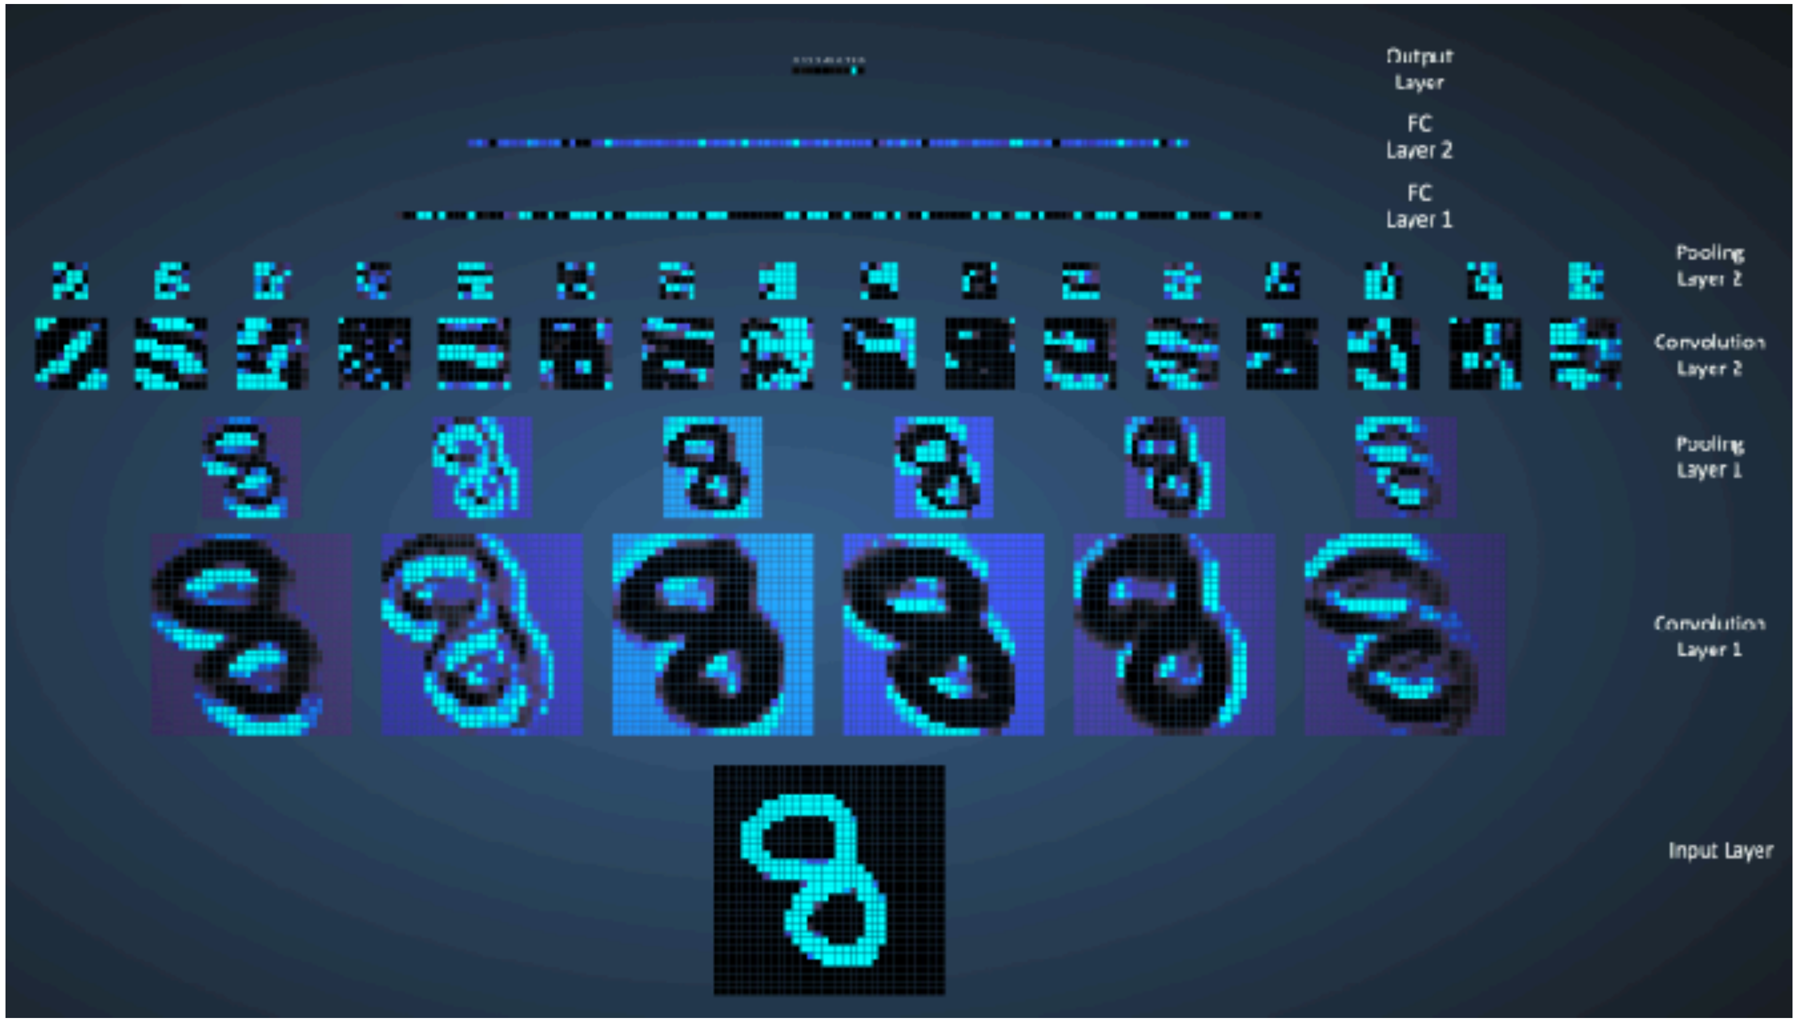
Output Layer
FC Layer 2
FC Layer 1
Pooling Layer 2
Convolution Layer 2
Pooling Layer 1
Convolution Layer 1
Input Layer

Challenge 1

- Given
  - a task/goal specifying output/solution
  - cases/instances/examples/(i,o) pairs

- Discover *automatically* from the given data
  - similarity measure
  - relevant features
  - multilevel representation

Self-Organising Memory

Challenge 1

CBR is viewed as integrating "learning" and "problem solving".
We should widen the scope to integrate
recognition tasks and deliberative tasks

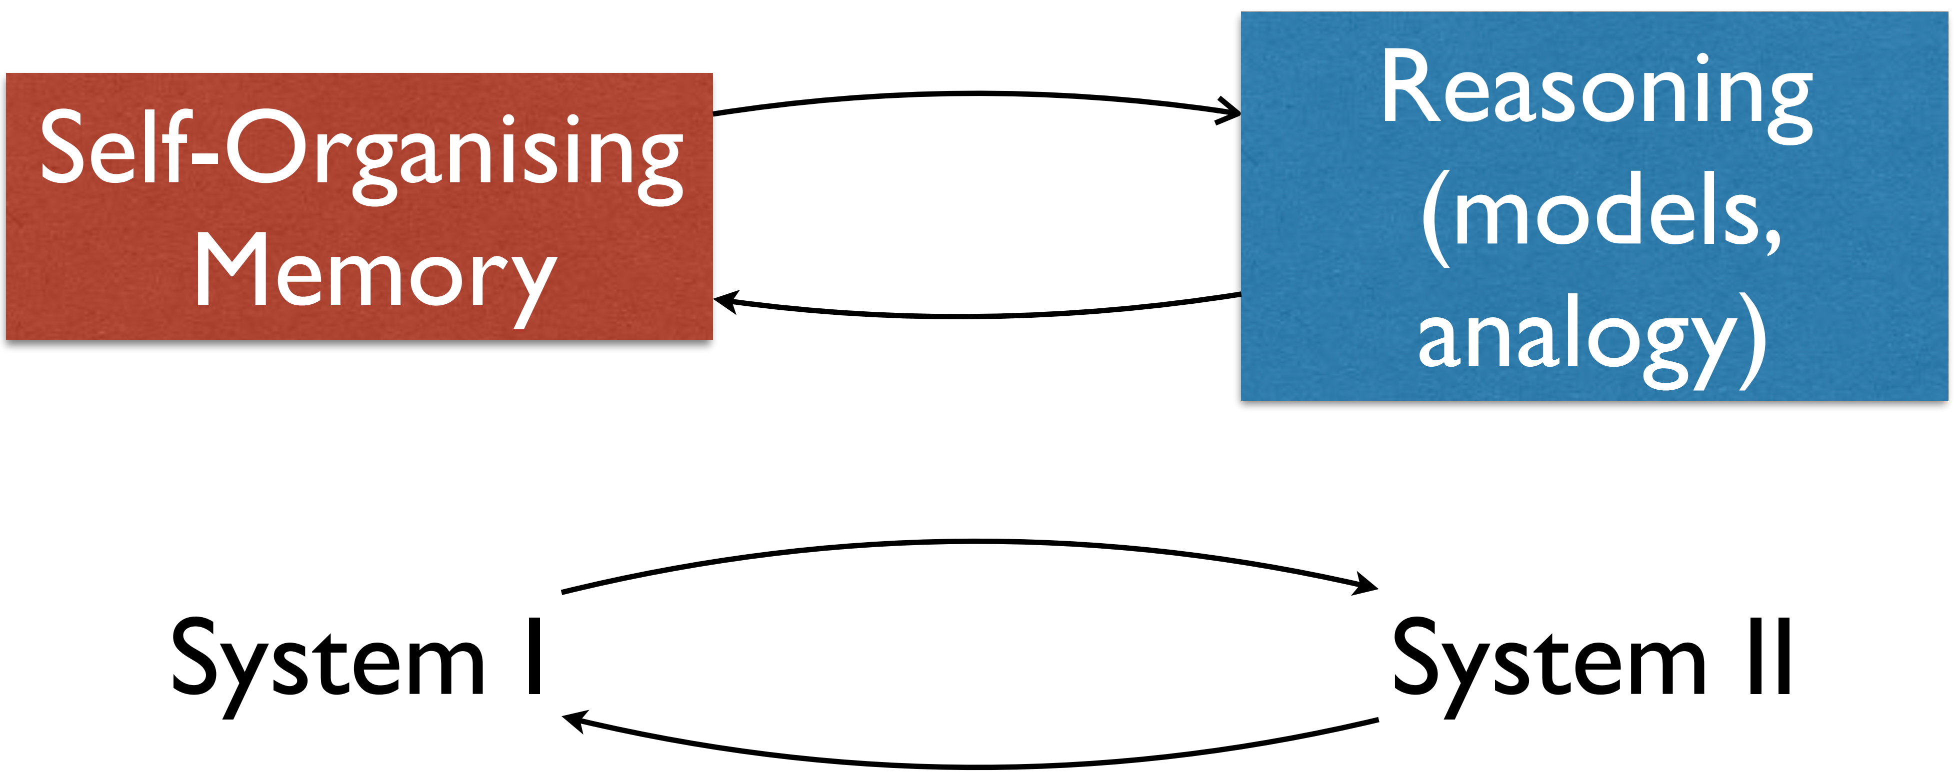

Analogy

"To Generalize is to be an Idiot;
To Particularize is the Alone Distinction of Merit"
William Blake
(comment to Joshua Reynolds' writings)

Analogy & Cognition

Analogy and categorization are the same:
"There is no fundamental difference between a <u>single memory</u> trace (instance/entity) and a <u>category</u> (concept)"
E.g. "The *moons* of Jupiter"

Analogy as the core of cognition
Douglas Hofstadter

Challenge 2

- Today’s AI need thousands of examples.

- How to learn from only one or two?
  - Better understanding the relations between CBR and
    - Analogy?
    - ‘XAI’?
    - One-shot learning?

Discussion

- How relevant is the new *cognitive turn* for CBR?
- What role can CBR play in the *current upswing* of AI?
- Can CBR offer a *new kind of synergy* of data-driven and knowledge-intensive approaches to AI?
- Other *new challenges* for CBR should we focus on?