Certificate I: Understanding AI and Machine Learning in Africa

Course AIML01: Artificial Intelligence – Past, Present, and Future

Module 2: The Nature of AI

Lecture 1: Symbolic AI and GOFAI

Carnegie Mellon University
Africa

Symbolic AI

- Symbolic AI is often referred to as good old-fashioned artificial intelligence: GOFAI

- Symbolic AI is one of the key historical, methodological, and epistemological approaches to AI

  – Historical, in the sense that if forms one of the cornerstones of AI

  – Methodological, in the sense that it is still an effective approach to the implementation of AI

  – Epistemological, in that it addresses what kinds of facts (or knowledge) about the world are available to an AI agent, how these facts can be represented, and what rules allow valid conclusions to be drawn from these facts

Symbolic AI

- Symbolic AI has its origins in the 1950s, in particular at the 1956 Dartmouth Workshop which, as we've seen, is considered by many as the inception of the discipline of artificial intelligence

- It constituted the primary, classical approach in the first 30 years of AI research, before the second AI Winter and the advent of connectionist AI and machine learning

For more detail, see Boden M. (2014). GOFAI. In The Cambridge Handbook of Artificial Intelligence Frankish, K and Ramsey, W, Editors.

Symbolic AI

The term "symbolic" refers to the fact that AI algorithms and programs are based on a set of

- symbols
- symbol manipulation processes

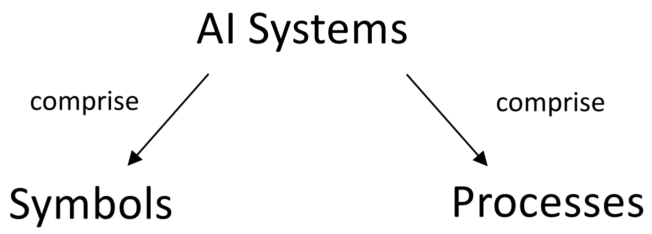

Physical Symbol Systems

Allen Newell and Herbert Simon proposed the concept of a Physical Symbol System

Computer Science as Empirical Inquiry: Symbols and Search

Allen Newell and Herbert A. Simon

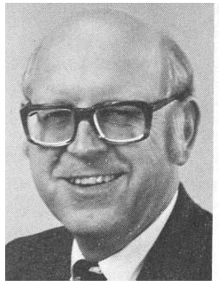


Computer science is the study of the phenomena surrounding computers. The founders of this society understood this very well when they called themselves the Association for Computing Machinery. The machine—not just the hardware, but the programmed, living machine—is the organism we study.

This is the tenth Turing Lecture. The nine persons who preceded us on this platform have presented nine different views of computer science. For our organism, the machine, can be studied at many levels and from many sides. We are deeply honored to appear here today and to present yet another view, the one that has permeated the scientific work for which we have been

Key Words and Phrases: symbols, search, science, computer science, empirical, Turing, artificial intelligence, intelligence, list processing, cognition, heuristics, problem solving.

CR Categories: 1.0, 2.1, 3.3, 3.6, 5.7.

Copyright © 1976, Association for Computing Machinery, Inc. General permission to republish, but not for profit, all or part of this material is granted provided that ACM's copyright notice is given and that reference is made to the publication, to its date of issue, and to the fact that reprinting privileges were granted by permission of the Association for Computing Machinery.

The authors' research over the years has been supported in part by the Advanced Research Projects Agency of the Department of Defense (monitored by the Air Force Office of Scientific Research) and in part by the National Institutes of Mental Health.

Authors' address: Carnegie-Mellon University, Pittsburgh.

Physical Symbol Systems

Allen Newell and Herbert Simon proposed the concept of a Physical Symbol System

> "a set of entities, called symbols, which are physical patterns that can occur as component of another type of entity called an expression (or symbol structure)"
>
> (Newell and Simon, 1976: 116)

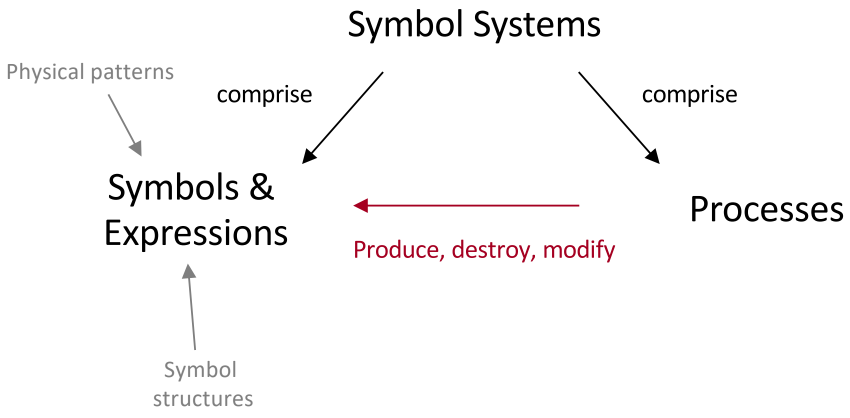

Physical Symbol Systems

Allen Newell and Herbert Simon proposed the concept of a Physical Symbol System

- These symbols are purely formal and meaningless entities
- In practice, they are normally interpreted by the programmer to have a particular semantic content such as words, numbers, images, and actions

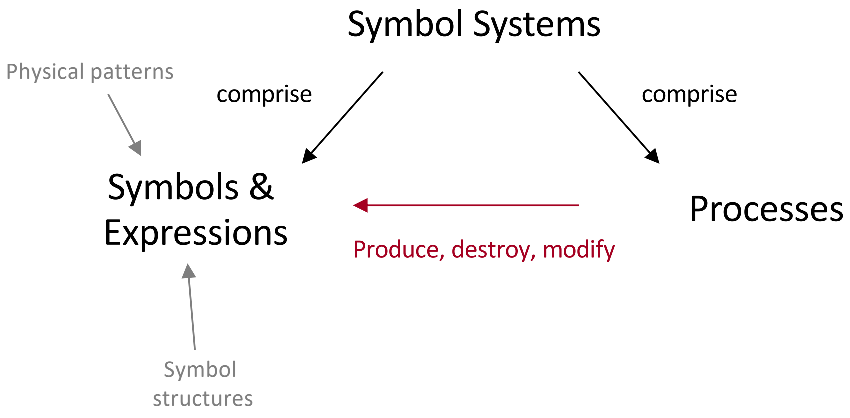

Physical Symbol Systems

Allen Newell and Herbert Simon proposed the concept of a Physical Symbol System

- These symbols are purely formal and meaningless entities
- The problem of attaching semantic meaning to syntactic symbols by linking them with entities in the environment is known as the symbol grounding problem

Symbols and symbol systems are purely syntactic entities

(just like letters & words with grammar that stipulates how thy can be combined in valid sentences)

But they have no semantic content

(just like the characters 象征 written in a foreign language, e.g. Chinese, that you don't speak or understand)

Often, this grounding is accomplished by the programmer who interprets the symbols to have a particular semantic content, such as words, numbers, images, and actions

Logic Formalisms

The symbol expressions are created using logic formalisms, such as propositional logic

- A sentence in propositional logic is either true or false and comprises
  - Propositional symbols (which can be true or false)
  - Boolean connectives (AND, OR, NOT, …)
- Propositional logic describes specific instances of things
  - e.g., whether a specific object is both red and round

`Red AND Round`
(also written `Red ∧ Round`)

This expression is true if the value of the symbol `Red` is true and the value of the symbol `Round` is true

`AND` (equivalently ∧) is the Boolean conjunction connective

Logic Formalisms

The symbol expressions are created using logic formalisms, such as first-order logic that

- Deals with objects in general and the relations among them
- Can express facts about some or all of the objects and can represent general laws or rules
- Also referred to has first-order predicate calculus (or just predicate calculus)

The predicate `Apple(Red, Round)` defines an `Apple` relation between `Red` and `Round`

We can define facts

We can define rules

We can form queries

We can make inferences to draw conclusions from facts and rules

Logic Formalisms

The symbol expressions can also be arranged in IF-THEN production rules

```
IF apple, THEN eat
```

Logic Formalisms

Semantic network: a symbol system where

- Each node has a symbol

  e.g., `Red`, `Apple`, `Fruit`

- Links have a label for the semantic relationship between nodes

  e.g., `IS A` or `HAS`

- Hierarchical relationship between nodes

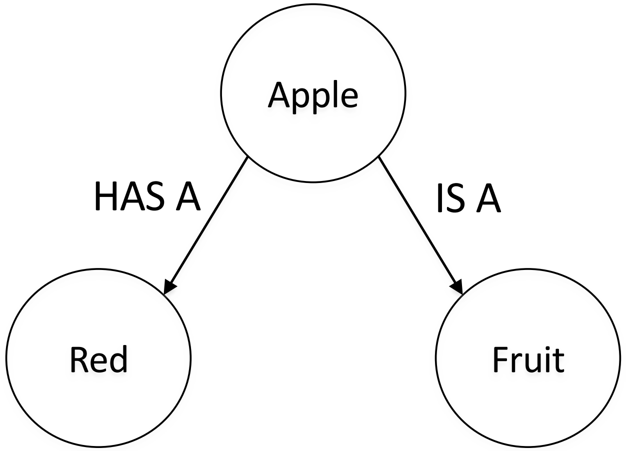

Logic Formalisms

A collection of symbol structures for a specific domain

- constitutes a knowledge base that is used by the system
- to reason about the problem

To use inference to draw conclusion about how to solve the problem

Heuristic Search

- Symbol systems solve problems by using the processes of heuristic search
  (Newell and Simon,1976: 116)

- The search for the optimal link between the problem definition and its solution must be guided by heuristics

Rules of thumb that are helpful in guiding the search toward the solution in an optimal way (or, at least, an efficient way)

Computer Science as Empirical Inquiry: Symbols and Search

Allen Newell and Herbert A. Simon

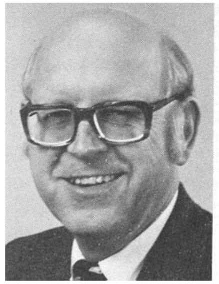

Computer science is the study of the phenomena surrounding computers. The founders of this society understood this very well when they called themselves the Association for Computing Machinery. The machine—not just the hardware, but the programmed, living machine—is the organism we study.

This is the tenth Turing Lecture. The nine persons who preceded us on this platform have presented nine different views of computer science. For our organism, the machine, can be studied at many levels and from many sides. We are deeply honored to appear here today and to present yet another view, the one that has permeated the scientific work for which we have been

Key Words and Phrases: symbols, search, science, computer science, empirical, Turing, artificial intelligence, intelligence, list processing, cognition, heuristics, problem solving.
CR Categories: 1.0, 2.1, 3.3, 3.6, 5.7.
Copyright © 1976, Association for Computing Machinery, Inc. General permission to republish, but not for profit, all or part of this material is granted provided that ACM's copyright notice is given and that reference is made to the publication, to its date of issue, and to the fact that reprinting privileges were granted by permission of the Association for Computing Machinery.
The authors' research over the years has been supported in part by the Advanced Research Projects Agency of the Department of Defense (monitored by the Air Force Office of Scientific Research) and in part by the National Institutes of Mental Health.
Authors' address: Carnegie-Mellon University, Pittsburgh.

Heuristic Search

AI heuristic search and planning algorithms are widely used today for

- Scheduling and logistics
- Data mining
- Games
- Searching the web
- Planning in robotics

The Physical Symbol Systems Hypothesis

An important aspect of the GOFAI approach is that symbol systems can model human intelligence:

The Physical Symbol Systems Hypothesis

"A physical symbol system has the necessary and sufficient means for general intelligent action."

(Newell and Simon 1976:116).

Computer Science as Empirical Inquiry: Symbols and Search

Allen Newell and Herbert A. Simon

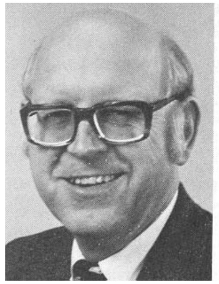

Computer science is the study of the phenomena surrounding computers. The founders of this society understood this very well when they called themselves the Association for Computing Machinery. The machine—not just the hardware, but the programmed, living machine—is the organism we study.

This is the tenth Turing Lecture. The nine persons who preceded us on this platform have presented nine different views of computer science. For our organism, the machine, can be studied at many levels and from many sides. We are deeply honored to appear here today and to present yet another view, the one that has permeated the scientific work for which we have been

Key Words and Phrases: symbols, search, science, computer science, empirical, Turing, artificial intelligence, intelligence, list processing, cognition, heuristics, problem solving.
CR Categories: 1.0, 2.1, 3.3, 3.6, 5.7.
Copyright © 1976, Association for Computing Machinery, Inc. General permission to republish, but not for profit, all or part of this material is granted provided that ACM's copyright notice is given and that reference is made to the publication, to its date of issue, and to the fact that reprinting privileges were granted by permission of the Association for Computing Machinery.
The authors' research over the years has been supported in part by the Advanced Research Projects Agency of the Department of Defense (monitored by the Air Force Office of Scientific Research) and in part by the National Institutes of Mental Health.
Authors' address: Carnegie-Mellon University, Pittsburgh.

The Physical Symbol Systems Hypothesis

In the same seminal paper, Newell and Simon introduce a second hypothesis:

The Heuristic Search Hypothesis

“The task of intelligence is to avert the ever-present threat of the exponential explosion of search."

(Newell and Simon 1976:116).

The solutions to problems are represented as symbol structures

A physical symbol system exercises its intelligence in problem-solving by effective and efficient search

Generating and progressively modifying symbol structures until it produces a solution structure

GOFAI Systems

Today, we refer to them as knowledge representation and reasoning systems.

- A classic example of a GOFAI system is an expert system
  - a program that represents the knowledge of the human expert in a specific domain
  - using a a set of IF-THEN production rules
  - which can be used to offer advice to non-experts or provide solutions to experts

- MYCIN was one of the first expert systems
  - To support medical doctors
  - in the diagnosis and treatment of infectious diseases

GOFAI Systems

- Today, expert systems have been developed in a wide range of domains & applications

  - Commerce
  - Education
  - Medicine
  - Military

- Some are capable of highly complex planning on the order of tens of thousands of search steps (Franklin 2014)

GOFAI Systems

Major strengths of GOFAI

- Ability to model hierarchical and sequential tasks
- Language processing
- Problem solving
- Games
- Represent knowledge bases
- Reason using logical inference

A knowledge base – also known as an ontology – representing the objects that robots manipulate in everyday activities

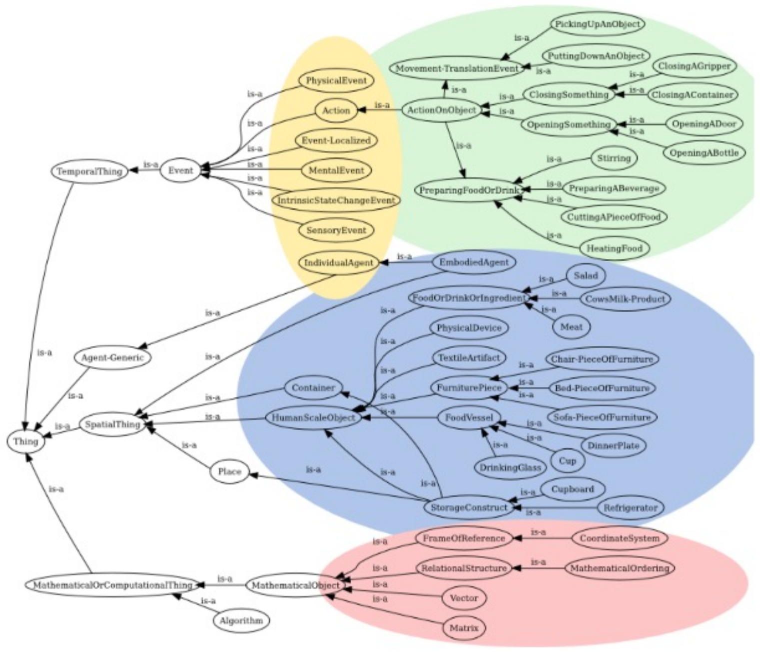

M. Beetz, M. Tenorth, and J. Winkler, “Open-EASE – a knowledge processing service for robots and robotics/AI researchers,” in IEEE International Conference on Robotics and Automation (ICRA), Seattle, Washington, USA, 2015.

GOFAI Systems

Some limitations of GOFAI

- These AI programs can be brittle
  - they can produce incorrect decisions or inferences when there is missing or contradictory data

- They are subject to the frame problem
  - How do you represent the effects of actions without having to represent explicitly a large number of non-effects

GOFAI Systems

Some limitations of GOFAI

- The symbol grounding problem
  - The problem of attaching semantic meaning to syntactic symbols by linking them with entities in the environment

- They cannot learn new knowledge

GOFAI Systems

These limitations led to

Not helped by the Initial strong claims about the power of symbolic AI to deal with general intelligence and any problem domain

- An AI Winter in the 1980s
- A switch in focus to connectionist and machine learning approaches

GOFAI Systems

Even so, GOFAI had some significant achievements

Aka knowledge representation and reasoning systems.

- Their widespread use of commercial expert systems

- Their essential role in games industry
  - To control the intelligent behaviour of the virtual agents
  - The historical victory of the IBM Deep Blue system in 1997, beating the chess world champion Garry Kasparov

- IBM Watson's victory in 2011 over two human champions in the TV game *Jeopardy!*

Cognitive Architectures

GOFAI also made a major contribution to the field of cognitive science in the guise of Unified Theories of Cognition (UTC)

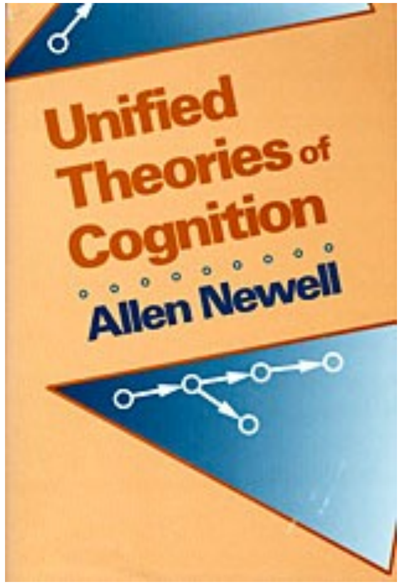

https://www.hup.harvard.edu/catalog.php?isbn=9780674921016

Cognitive Architectures

UTCs cover a broad range of issues in modelling human cognition and intelligence

- Attention
- Memory
- Problem solving
- Decision making
- Learning

from several aspects

- Psychology
- Neuroscience
- Computer Science

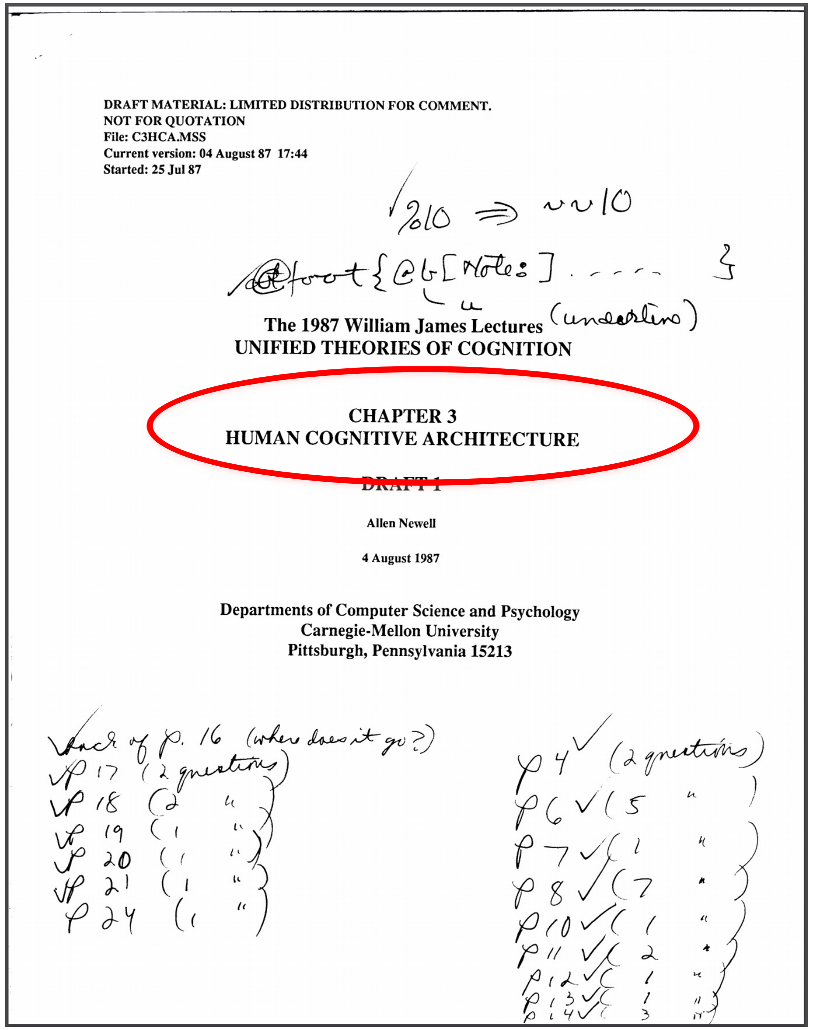

DRAFT MATERIAL: LIMITED DISTRIBUTION FOR COMMENT.
NOT FOR QUOTATION
File: C3HCA.MSS
Current version: 04 August 87 17:44
Started: 25 Jul 87

The 1987 William James Lectures
UNIFIED THEORIES OF COGNITION

CHAPTER 3
HUMAN COGNITIVE ARCHITECTURE

DRAFT 1

Allen Newell

4 August 1987

Departments of Computer Science and Psychology
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

http://digitalcollections.library.cmu.edu/awweb/awarchive?type=file&item=352120

Cognitive Architectures

A cognitive architecture orchestrates the core cognitive abilities

Allowing the agent to exhibit flexible context-sensitive behaviour,
prospectively selecting and controlling the actions
that are required to achieve given goals

Perception
Attention
Action selection
Memory
Learning
Reasoning
Meta-reasoning
Prospection

Cognitive Architectures

Overall structure and organization of a cognitive system

- Essential modules
- Essential relations between these modules
- Essential algorithmic and representational details in each module

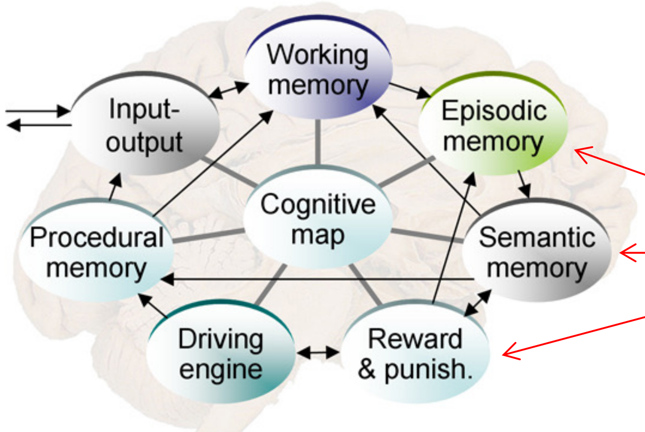

Commitment to formalisms for representation and processes

(GMU-BICA Architecture: Samsonovich 2005)

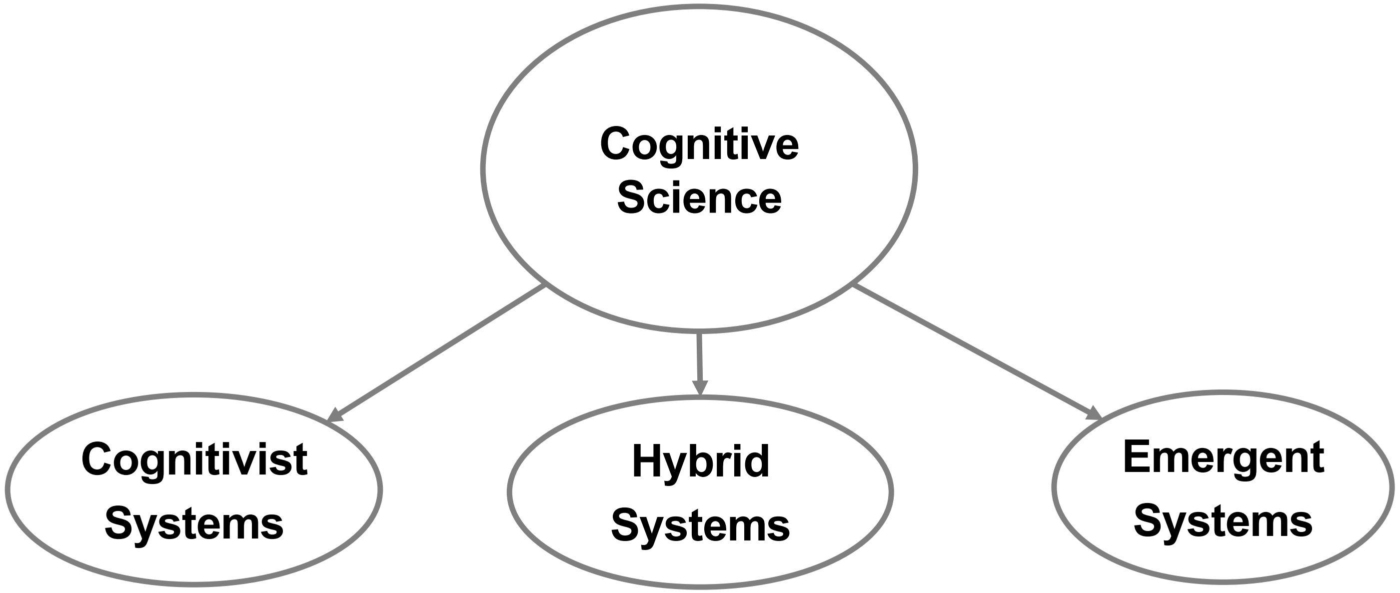

There are three paradigms of cognitive science

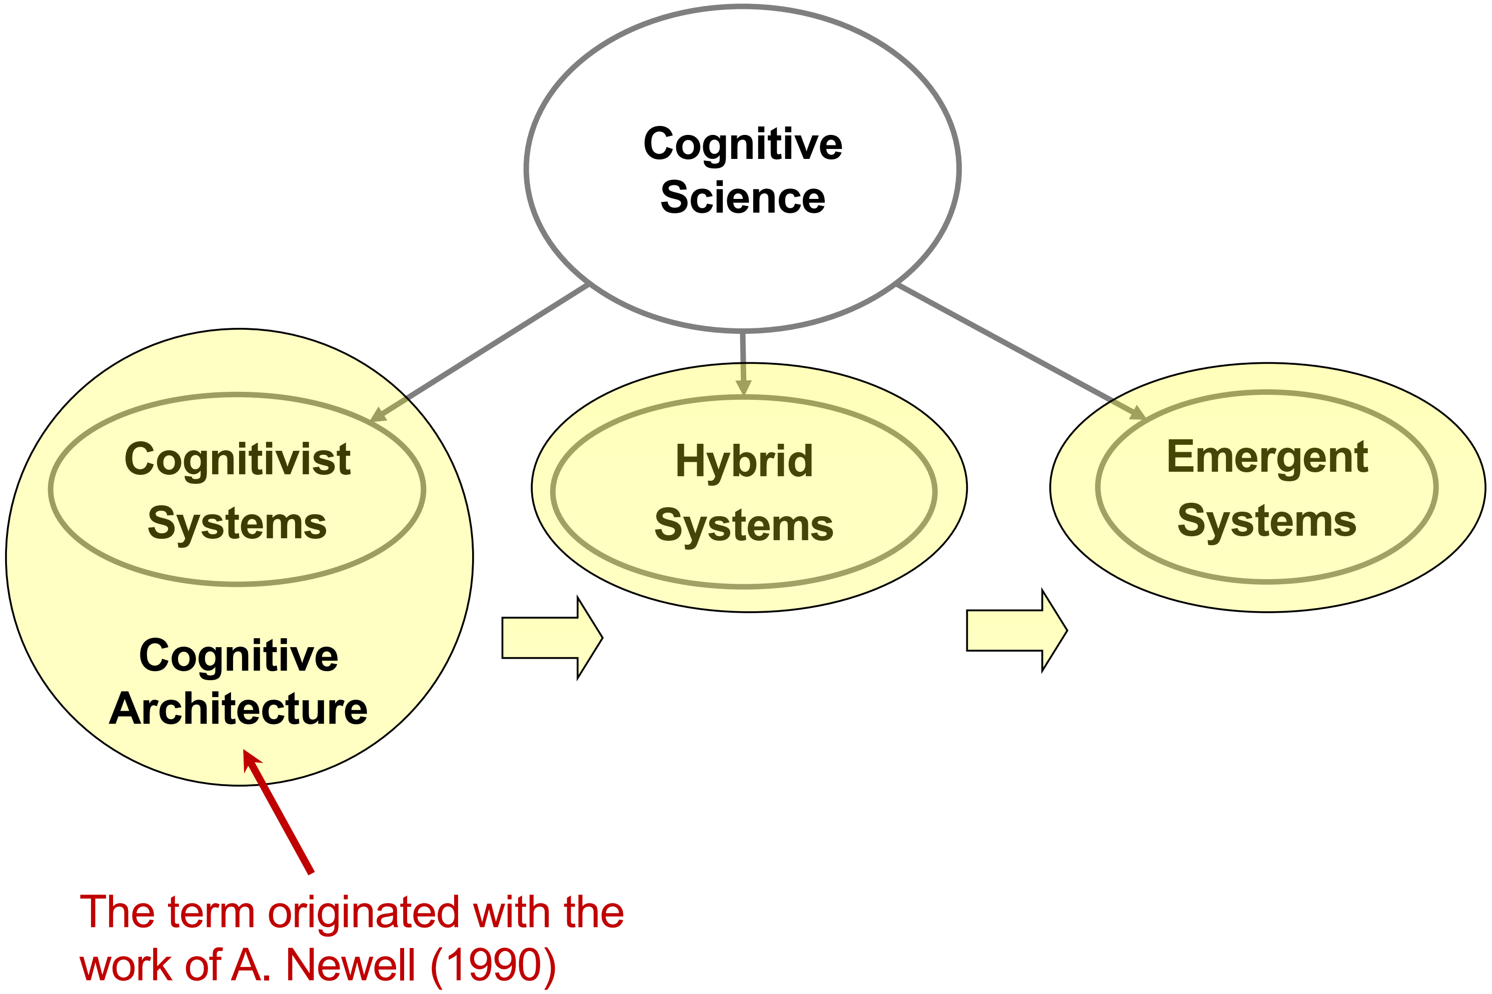
Cognitive Science
Cognitivist Systems
Cognitive Architecture
Hybrid Systems
Emergent Systems
The term originated with the work of A. Newell (1990)

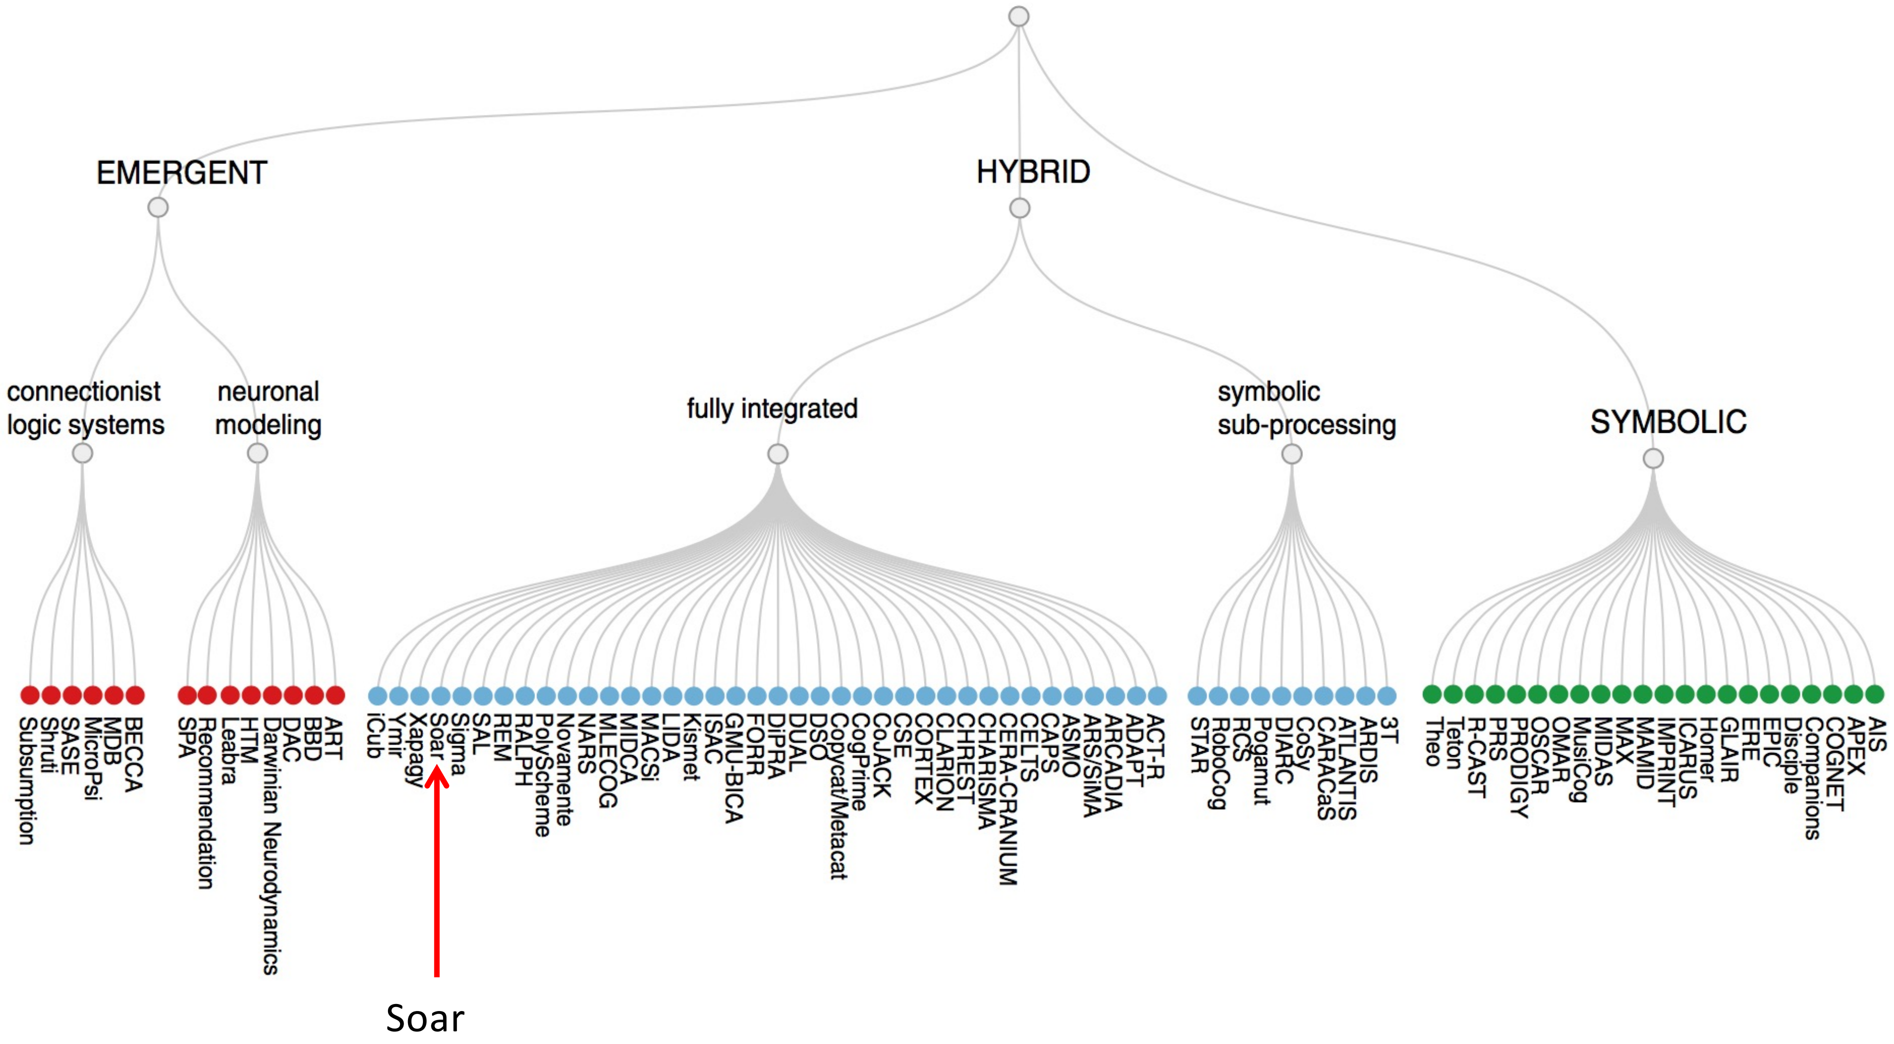

I. Kotseruba and J. Tsotsos. 40 years of cognitive architectures: core cognitive abilities and practical applications. Artificial Intelligence Review, Vol. 53, No. 1, pp. 17-94, 2020.

I. Kotseruba and J. Tsotsos. 40 years of cognitive architectures: core cognitive abilities and practical applications. Artificial Intelligence Review, Vol. 53, No. 1, pp. 17-94, 2020.

J. E. Laird, The Soar cognitive architecture. MIT Press, Cambridge, 2012

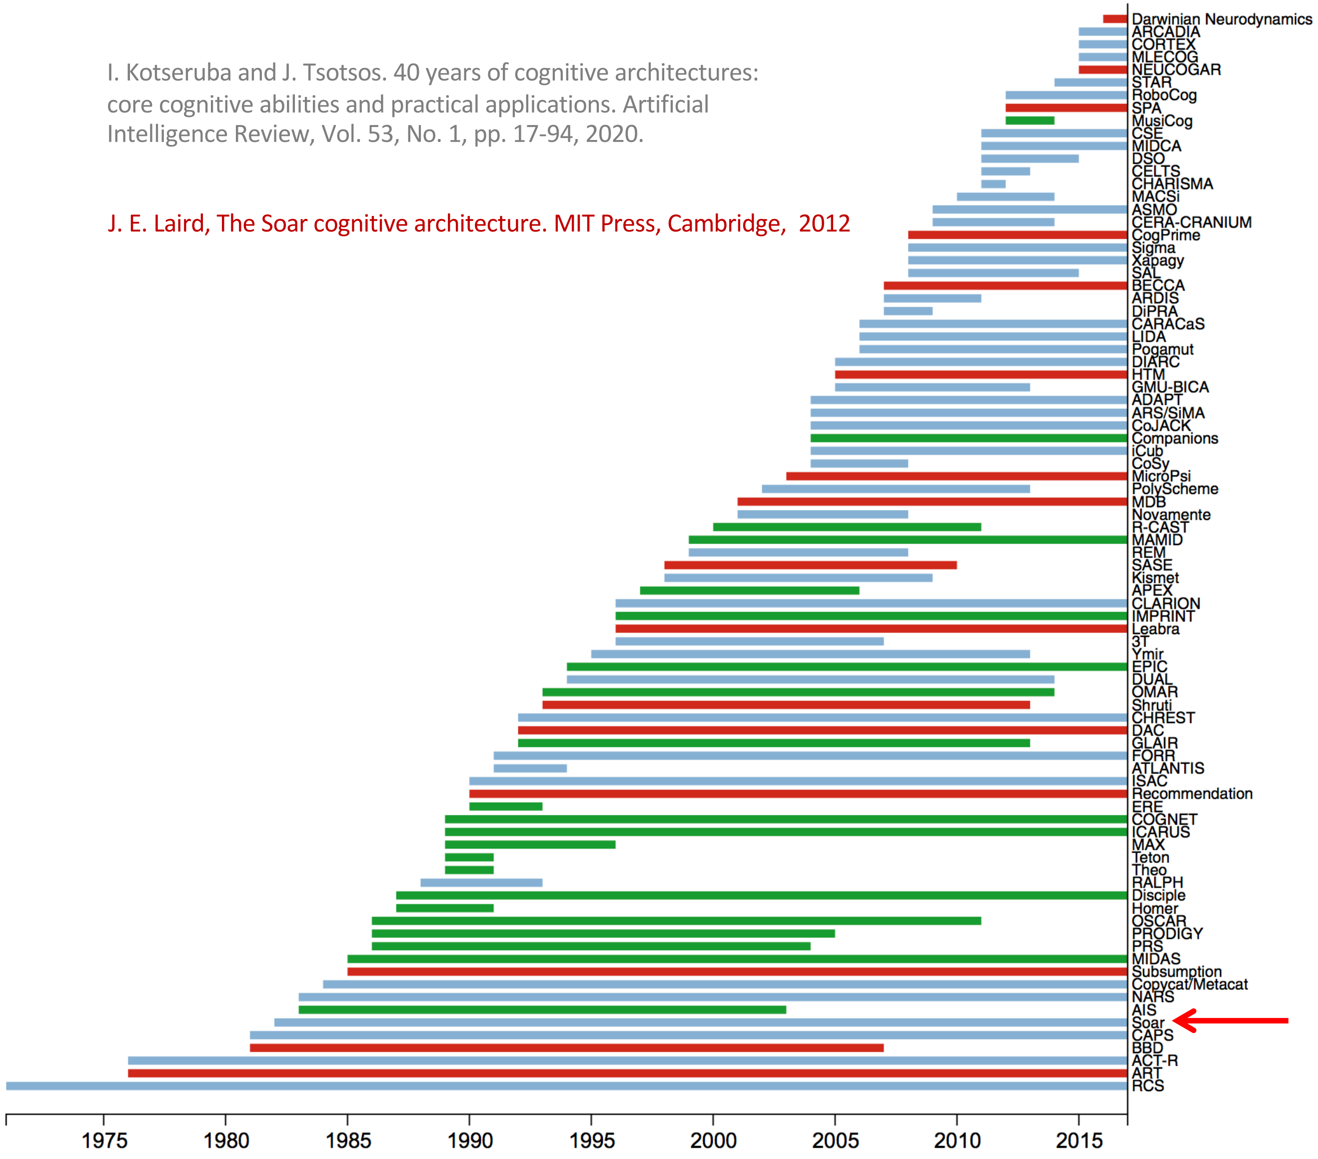

Soar

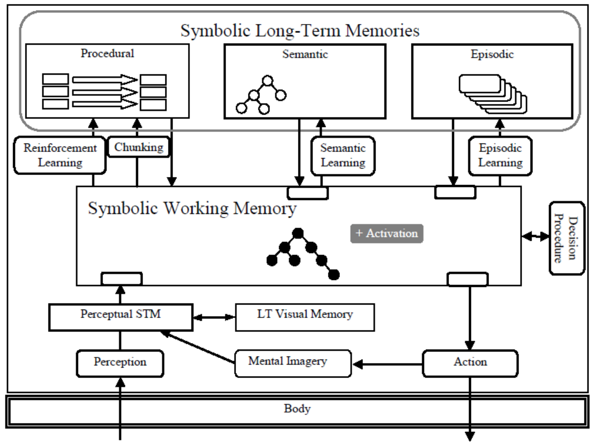

- A. Newell's candidate for a Unified Theory of Cognition
- 1983 – 2022; now version 9.6
- Production (rule-based) system

John Laird

2021 TransAIR Workshop on Cognitive Architectures for Robot Agents

https://transair-bridge.org/workshop-2021/

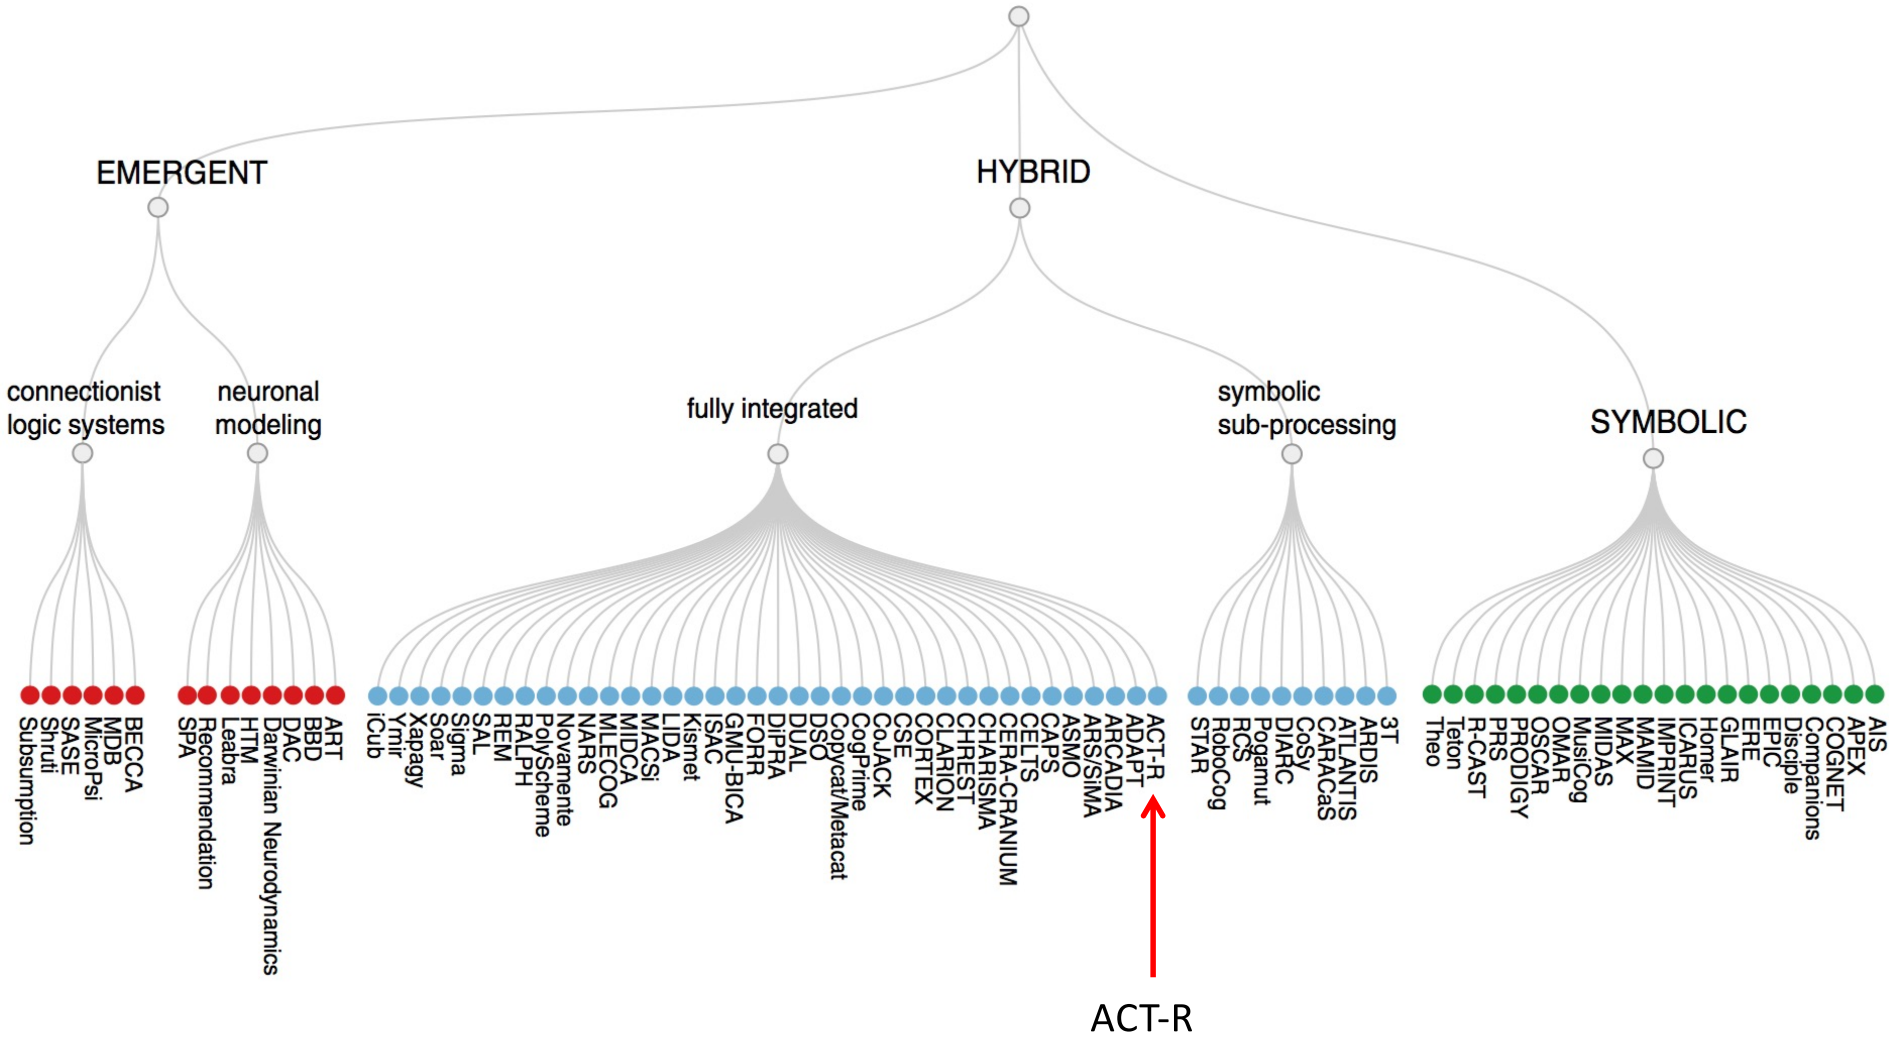
EMERGENT
HYBRID
SYMBOLIC
connectionist logic systems
neuronal modeling
fully integrated
symbolic sub-processing
Subsumption
Shruti
SASE
MicroPsi
MDB
BECCA
SPA
Recommendation
Leabra
HTM
Darwinian Neurodynamics
DAC
BBD
ART
iCub
Ymir
Xapagy
Soar
Sigma
SAL
REM
RALPH
PolyScheme
Novamente
NARS
MLECOG
MIDCA
MACSi
LIDA
Kismet
ISAC
GMU-BICA
FORR
DiPRA
DUAL
DSO
Copycat/Metacat
CogPrime
CoJACK
CSE
CORTEX
CLARION
CHREST
CHARISMA
CERA-CRANIUM
CELTS
CAPS
ASMO
ARS/SiMA
ARCADIA
ADAPT
ACT-R
STAR
RoboCog
RCS
Pogamut
DIARC
CoSy
CARACaS
ATLANTIS
ARDIS
3T
Theo
Teton
R-CAST
PRS
PRODIGY
OSCAR
OMAR
MusiCog
MIDAS
MAX
MAMID
IMPRINT
ICARUS
Homer
GLAIR
ERE
EPIC
Disciple
Companions
COGNET
APEX
AIS
ACT-R

I. Kotseruba and J. Tsotsos. 40 years of cognitive architectures: core cognitive abilities and practical applications. Artificial Intelligence Review, Vol. 53, No. 1, pp. 17-94, 2020.

J. R. Anderson,D. Bothell, M. D. Byrne, S. Douglass C. Lebiere, Y. Qin. An integrated theory of the mind. Psychol Rev 111(4):1036–1060, 2004.

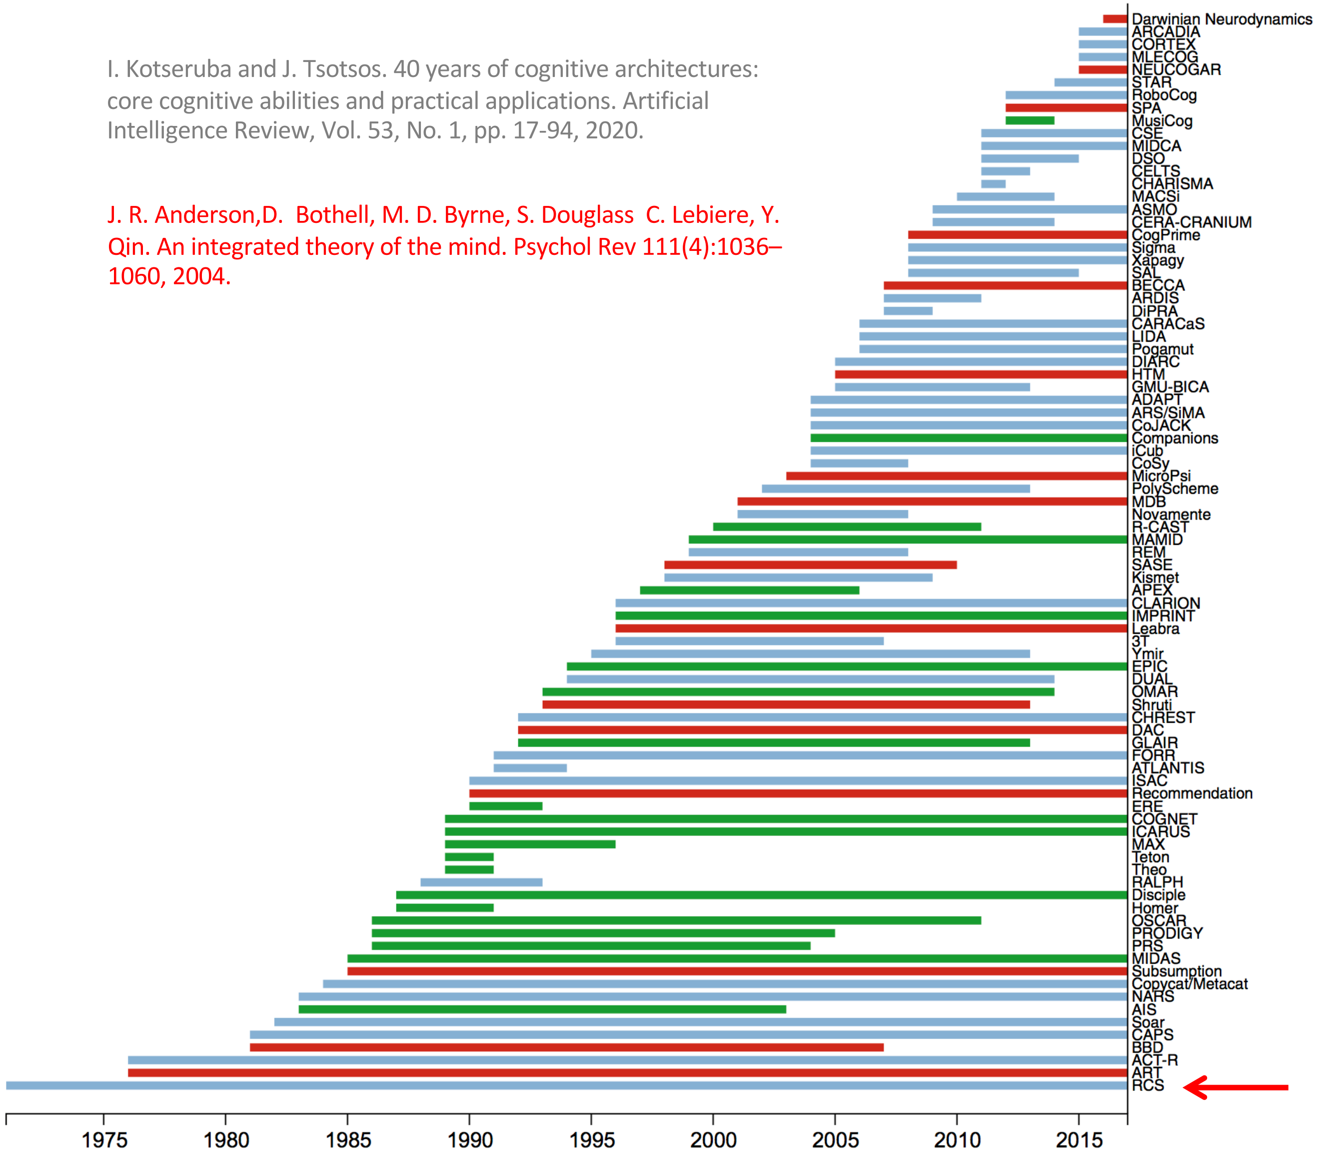

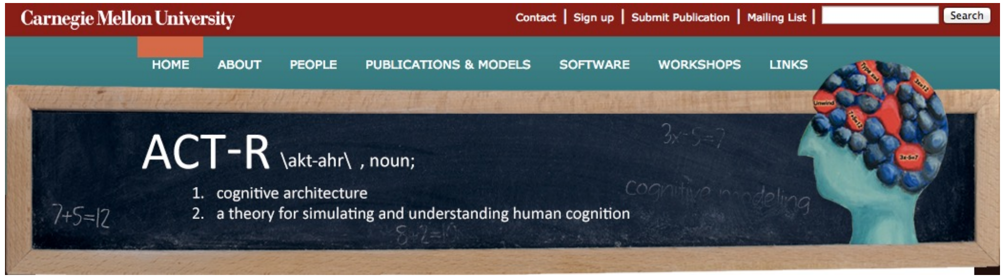

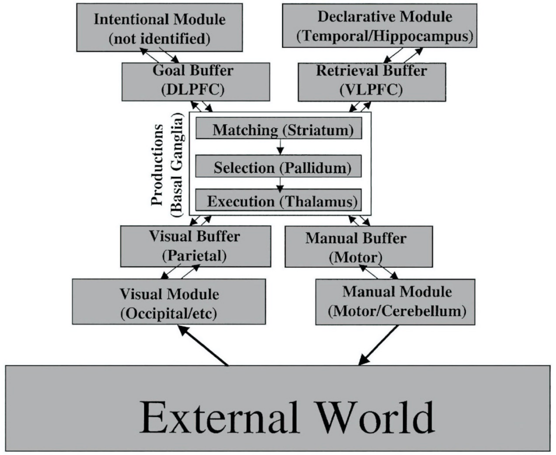

- J. Anderson's candidate for a Unified Theory of Cognition
- 1996, 2004; now version 7
- Production system with five modules: Intentional, Declarative, Visual, Manual, Production