Introduction to Cognitive Robotics

Module 7: Cognitive Architectures

Lecture 2: Example cognitive architectures: Soar, ACT-R, CLARION, BBD

David Vernon
Carnegie Mellon University Africa

www.vernon.eu

Example Cognitive Architectures

Surveys:

Biologically Inspired Cognitive Architectures Society, Comparative Repository of Cognitive Architectures, http://bicasociety.org/cogarch/architectures.htm (25 cognitive architectures)

A Survey of Cognitive and Agent Architectures, University of Michigan, http://ai.eecs.umich.edu/cogarch0/ (12 cognitive architectures)

W. Duch, R. J. Oentaryo, and M. Pasquier. "Cognitive Architectures: Where do we go from here?", Proc. Conf. Artificial General Intelligence, 122-136, 2008. (17 cognitive architectures)

D. Vernon, G. Metta, and G. Sandini, "A Survey of Artificial Cognitive Systems: Implications for the Autonomous Development of Mental Capabilities in Computational Agents", IEEE Transactions on Evolutionary Computation, Vol. 11, No. 2, pp. 151-180, 2007. (14 cognitive architectures)

D. Vernon, C. von Hofsten, and L. Fadiga. "A Roadmap for Cognitive Development in Humanoid Robots", Cognitive Systems Monographs (COSMOS), Vol. 11, Springer, 2010. Chapter 5 and Appendix I (20 cognitive architectures)

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. (84 cognitive architectures)

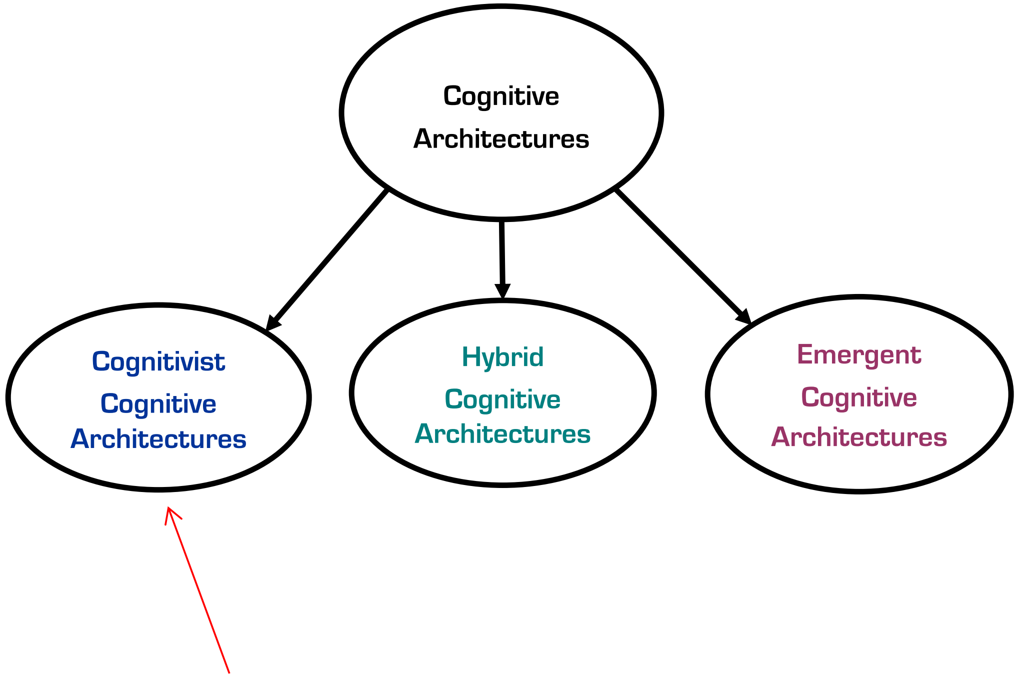

Kotseruba and Tsotsos refer to these as Symbolic Architectures

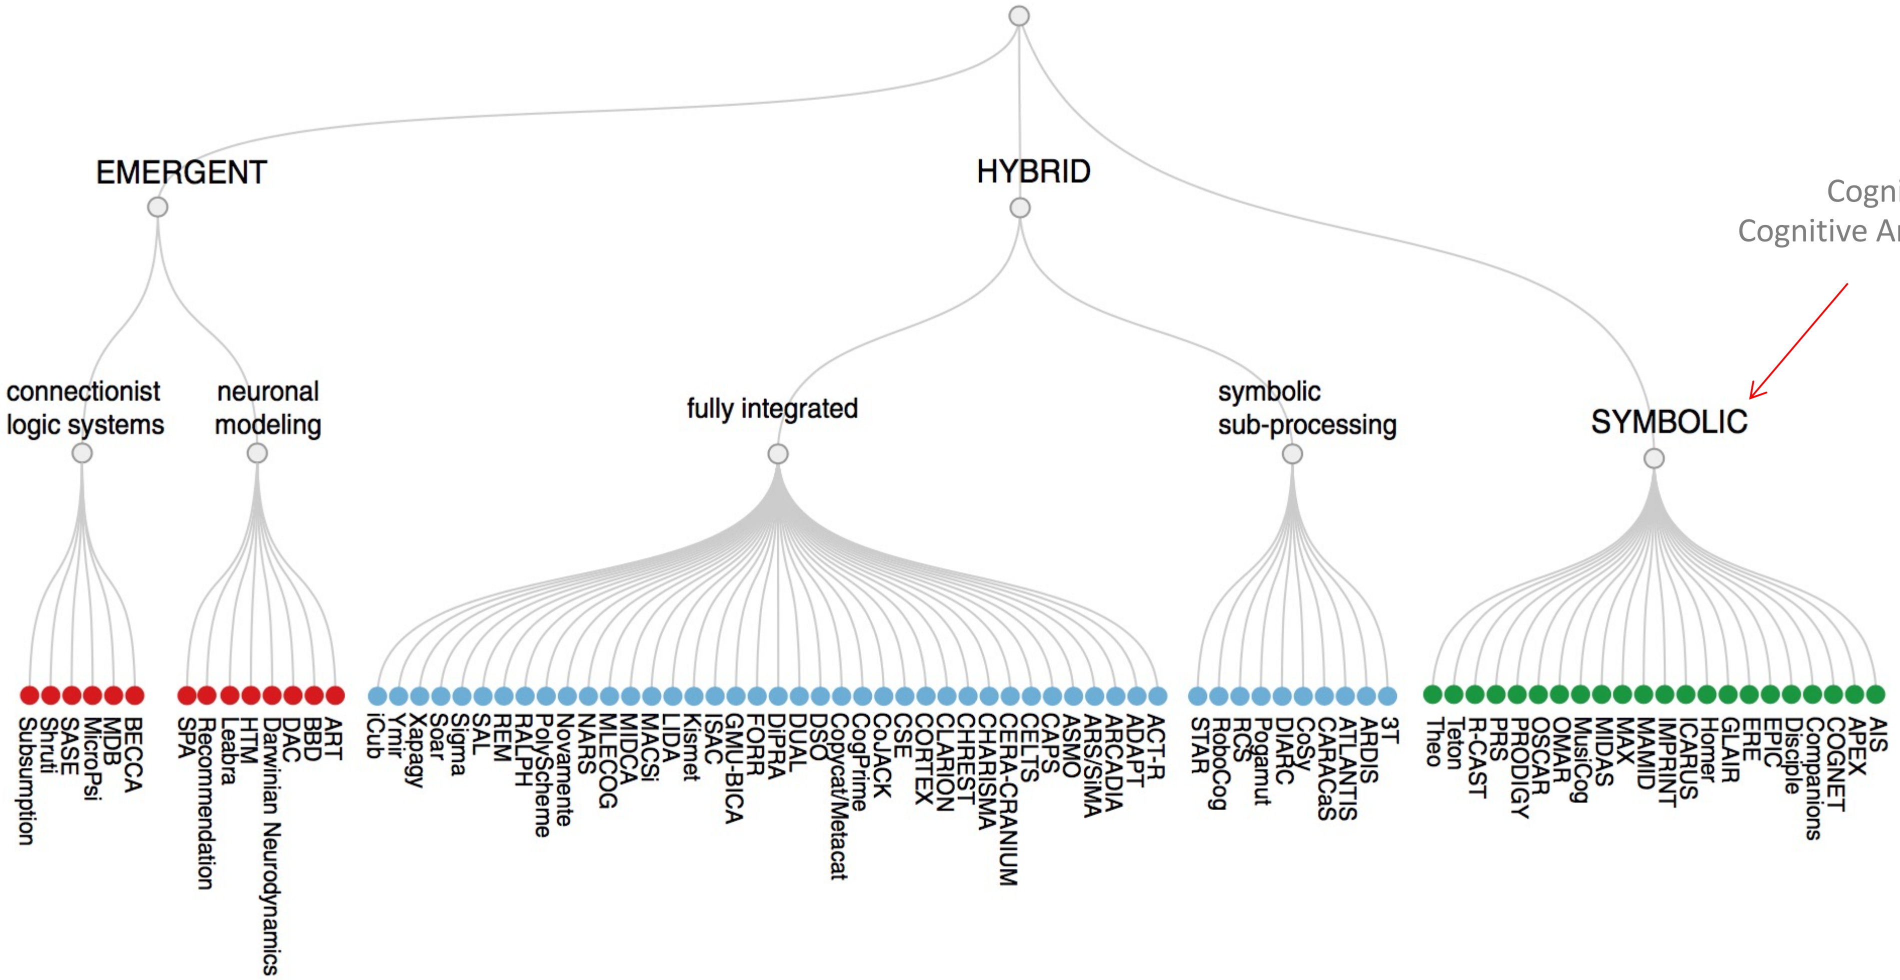

Cognitivist
Cognitive Architectures

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.

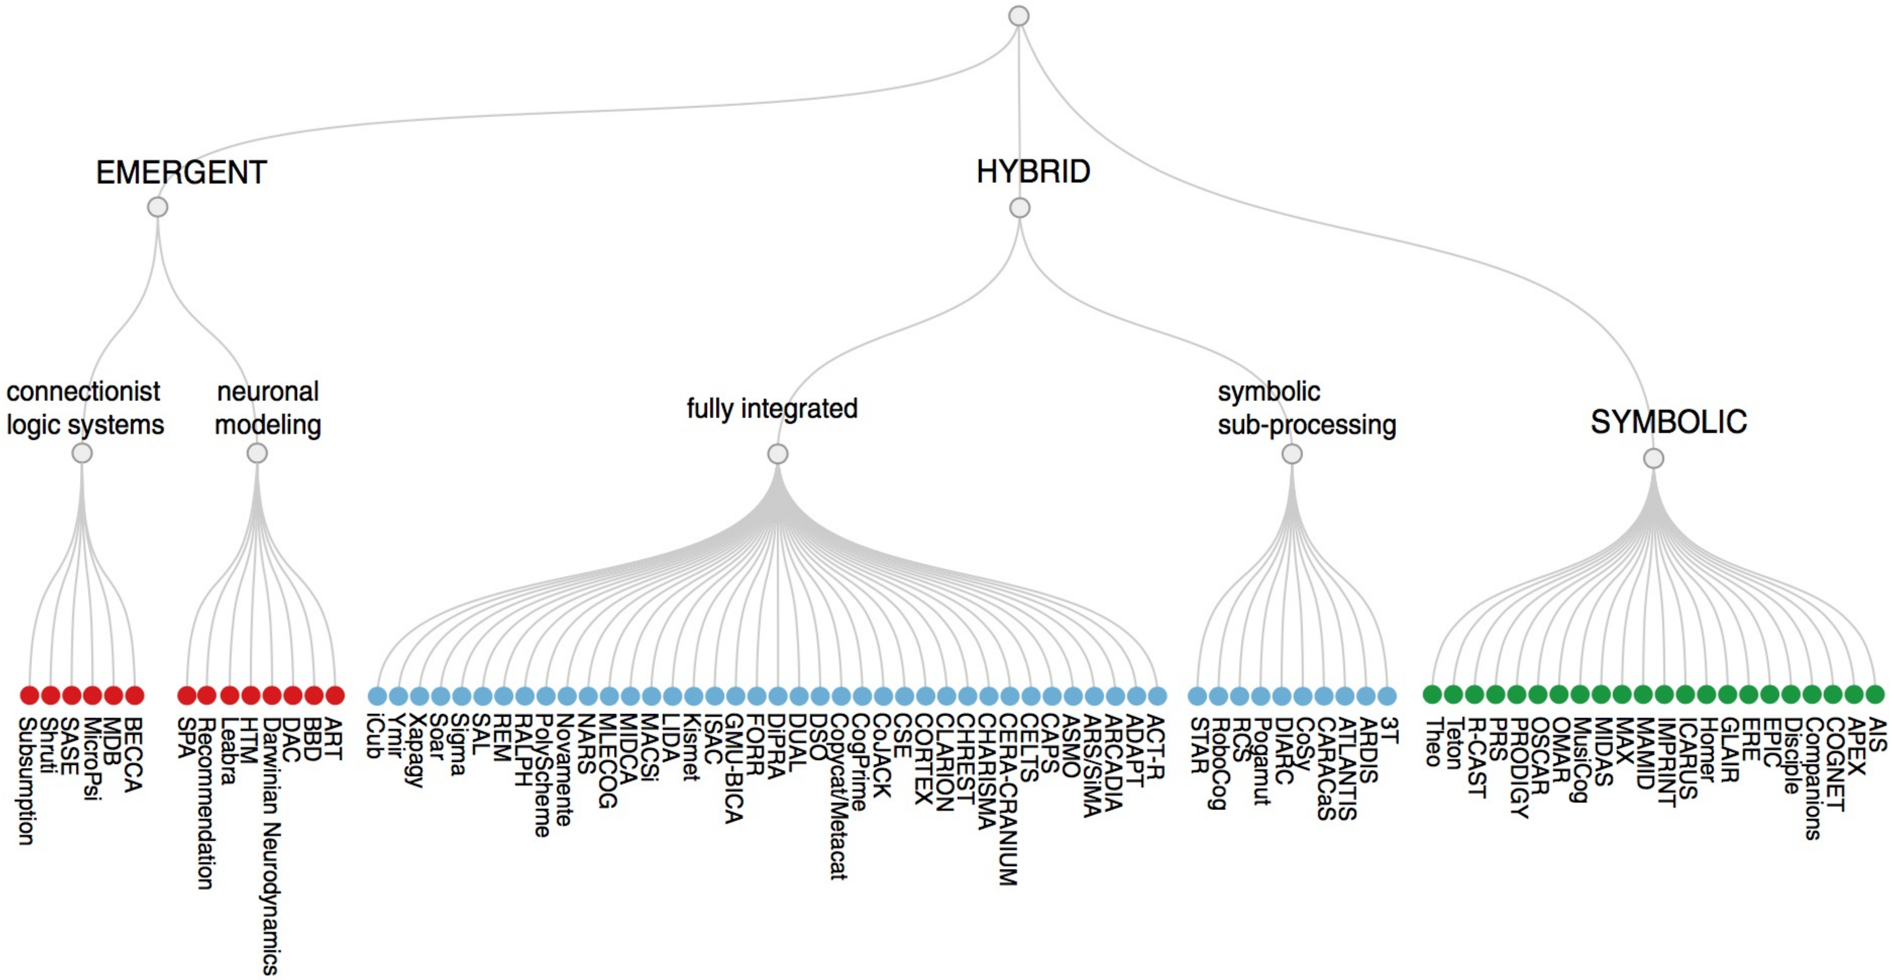

We will briefly sample a few of the more well-known cognitive architectures

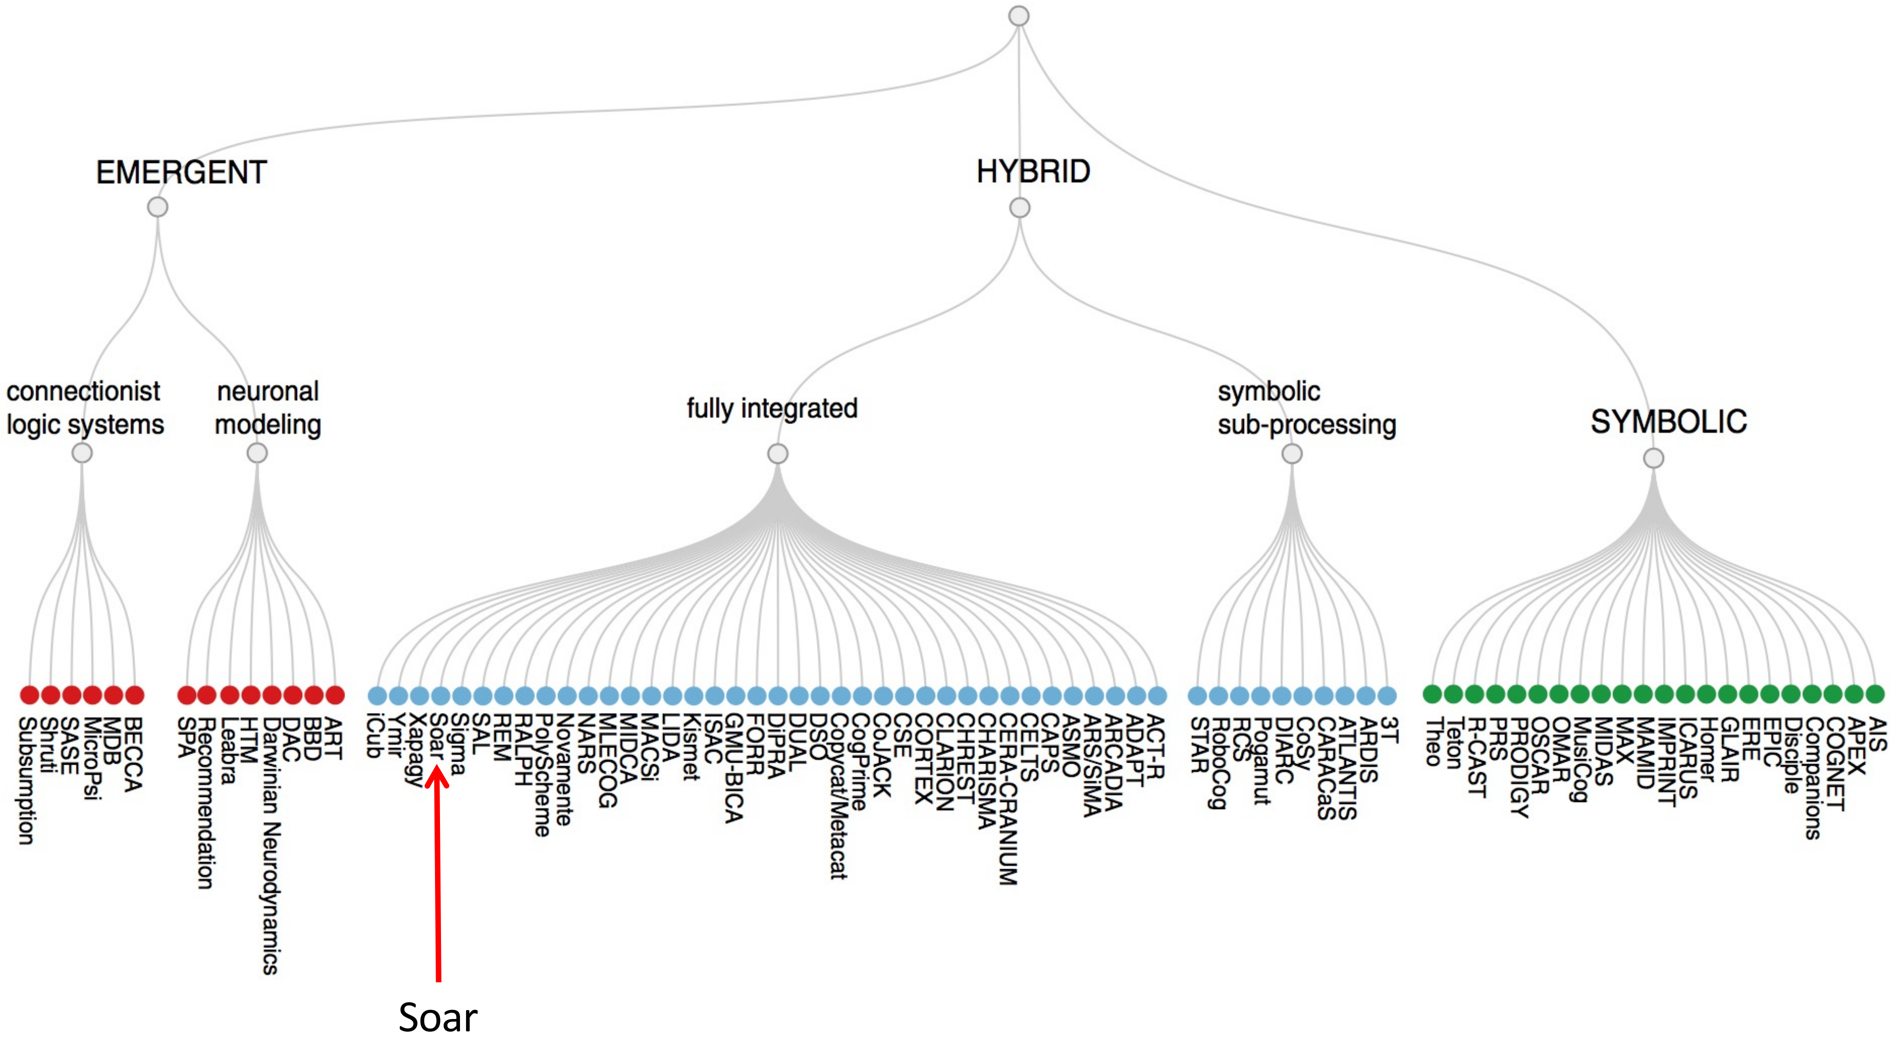
EMERGENT
HYBRID
connectionist logic systems
neuronal modeling
fully integrated
symbolic sub-processing
SYMBOLIC
Subsumption
Shruti
SASE
MicroPsi
MDB
BECCA
SPA
Recommendation
Leabra
HTM
Darwinian Neurodynamics
DAC
BBD
ART
iCub
Ymir
Xapagy
Soar
Sigma
SAL
REM
RALPH
PolyScheme
Novamente
NARS
MLECOG
MIDCA
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
Soar

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

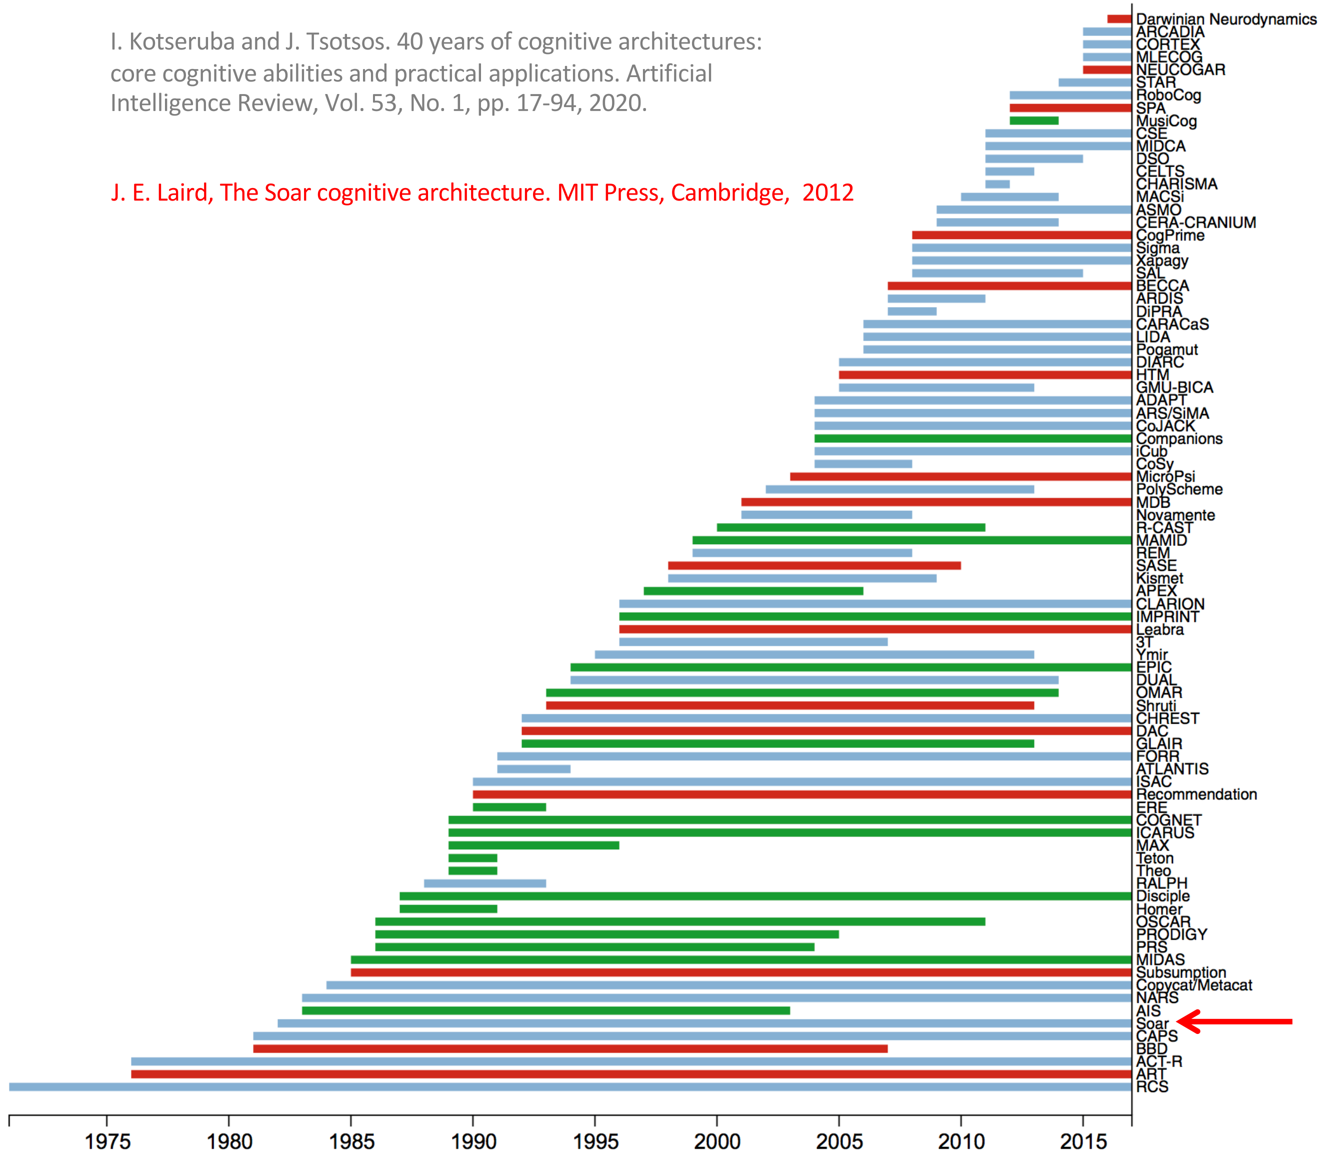

Soar

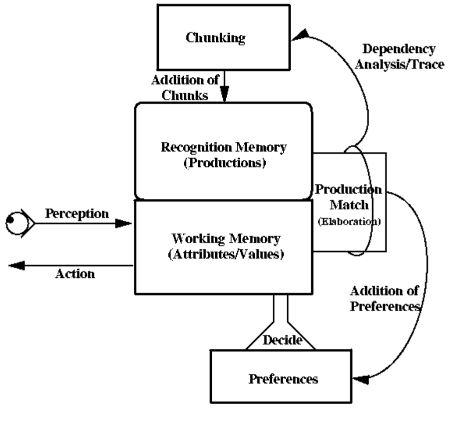

(Based on Figure 3.1, pg 20, The Soar's User Manual, Version 6)

- A. Newell's candidate for a Unified Theory of Cognition
- 1983 – 2021; now version 9.6
- Production (rule-based) system
- Cyclic operation
  1. Production cycle: fire all rules that match information in the symbolic working memory; update memory, fire all rules …
  2. Decision cycle: select an action
- Universal sub-goaling: create a new goal and expose more knowledge when an impasse is encountered
- Learns a new rule when an impasse is resolved

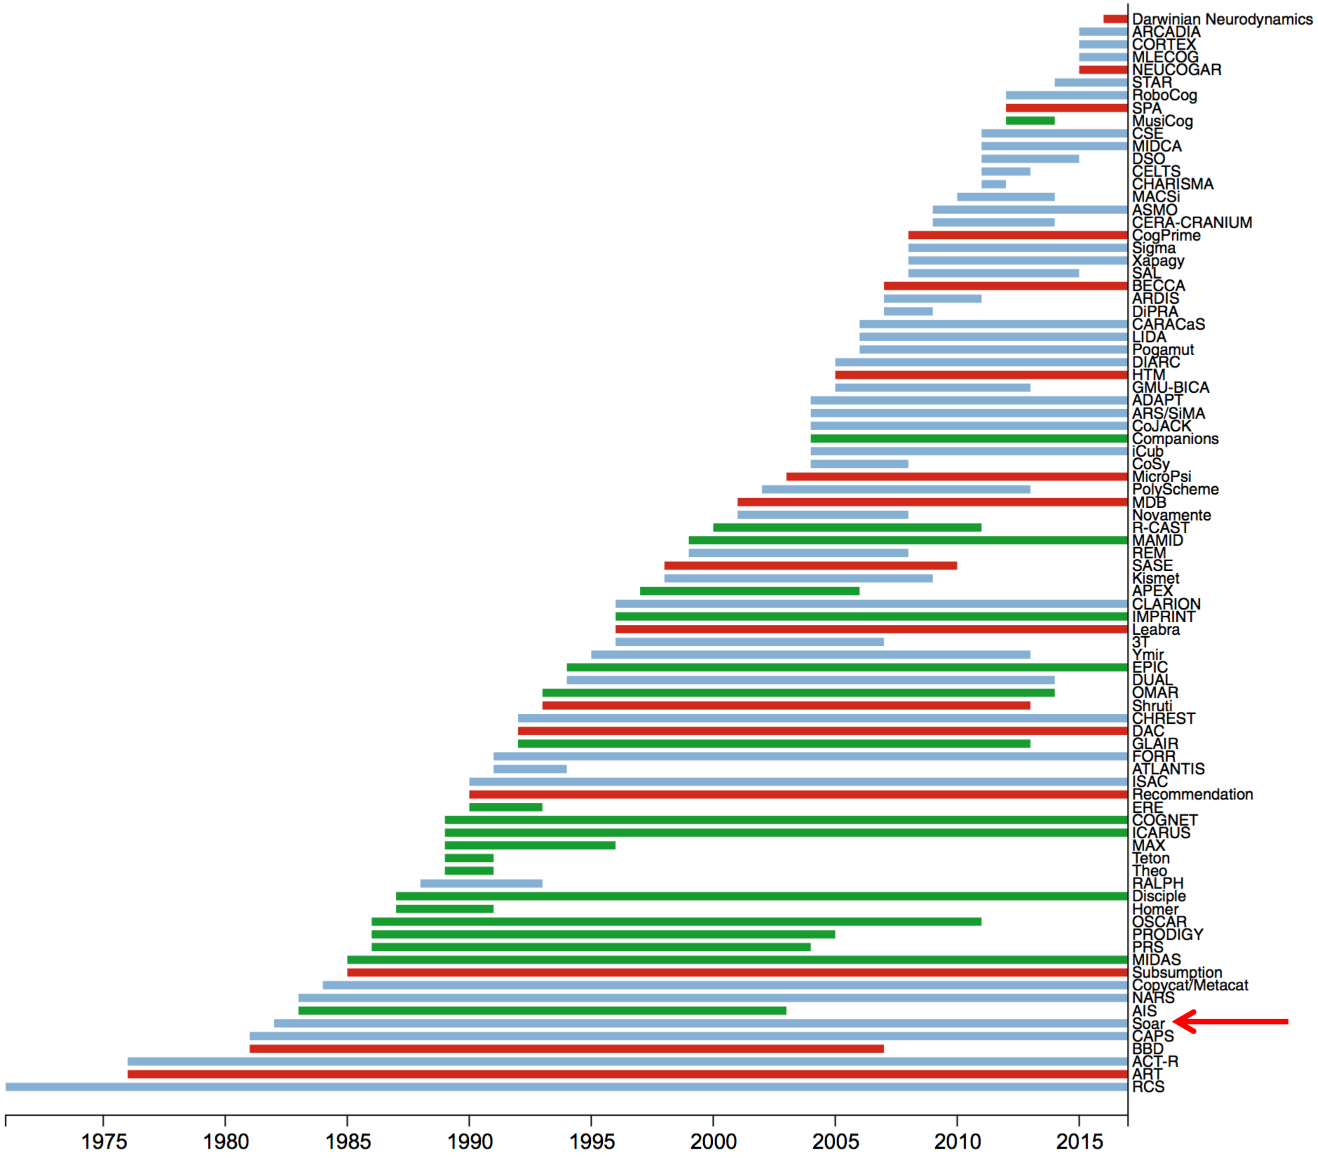

A.1 Cognitivist Cognitive Architectures

A.1.1 The Soar Cognitive Architecture

The Soar system [211, 326, 220, 222] is Newell's candidate for a Unified Theory of Cognition [271] and, as such, it is an architypal cognitivist cognitive architecture (as well as being an iconic one).It is a production (or rule-based) system[1] that operates in a cyclic manner, with a production cycle and a decision cycle. It operates as follows. First, all productions that match the contents of declarative (working) memory fire. A production that fires may alter the state of declarative memory and cause other productions to fire. This continues until no more productions fire. At this point, the decision cycle begins in which a single action from several possible actions is selected. The selection is based on stored action preferences. Thus, for each decision cycle there may have been many production cycles. Productions in Soar are low-level; that is to say, knowledge is encapsulated at a very small grain size.

One important aspect of the decision process concerns a process known as *universal sub-goaling*. Since there is no guarantee that the action preferences will be unambiguous or that they will lead to a unique action or indeed any action, the decision cycle may lead to an 'impasse'. If this happens, Soar sets up an new state in a new problem space — sub-goaling — with the goal of resolving the impasse. Resolving one impasse may cause others and the sub-goaling process continues. It is assumed that degenerate cases can be dealt with (*e.g.* if all else fails, choose randomly between two actions). Whenever an impasse is resolved, Soar creates a new production rule which summarizes the processing that occurred in the sub-state in solving the sub-goal. Thus, resolving an impasse alters the system super-state, *i.e.* the state in which the impasse originally occurred. This change is called a result and becomes the outcome of the production rule. The condition for the production rule to fire is derived from a dependency analysis: finding what declarative memory items matched in the course of determining the result. This change in state is a form of learning and it is the only form that occurs in Soar, *i.e.* Soar only learns new production rules. Since impasses occur often in Soar, learning is pervasive in Soar's operation.

[1] A production is effectively an IF-THEN condition-action pair. A production system is a set of production rules and a computational engine for interpreting or executing productions.

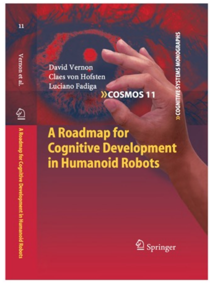

D. Vernon, C. von Hofsten, and L. Fadiga. A Roadmap for Cognitive Development in Humanoid Robots, Cognitive Systems Monographs (COSMOS), Vol. 11, Springer, 2010.

Soar

(Laird et al. 2012)

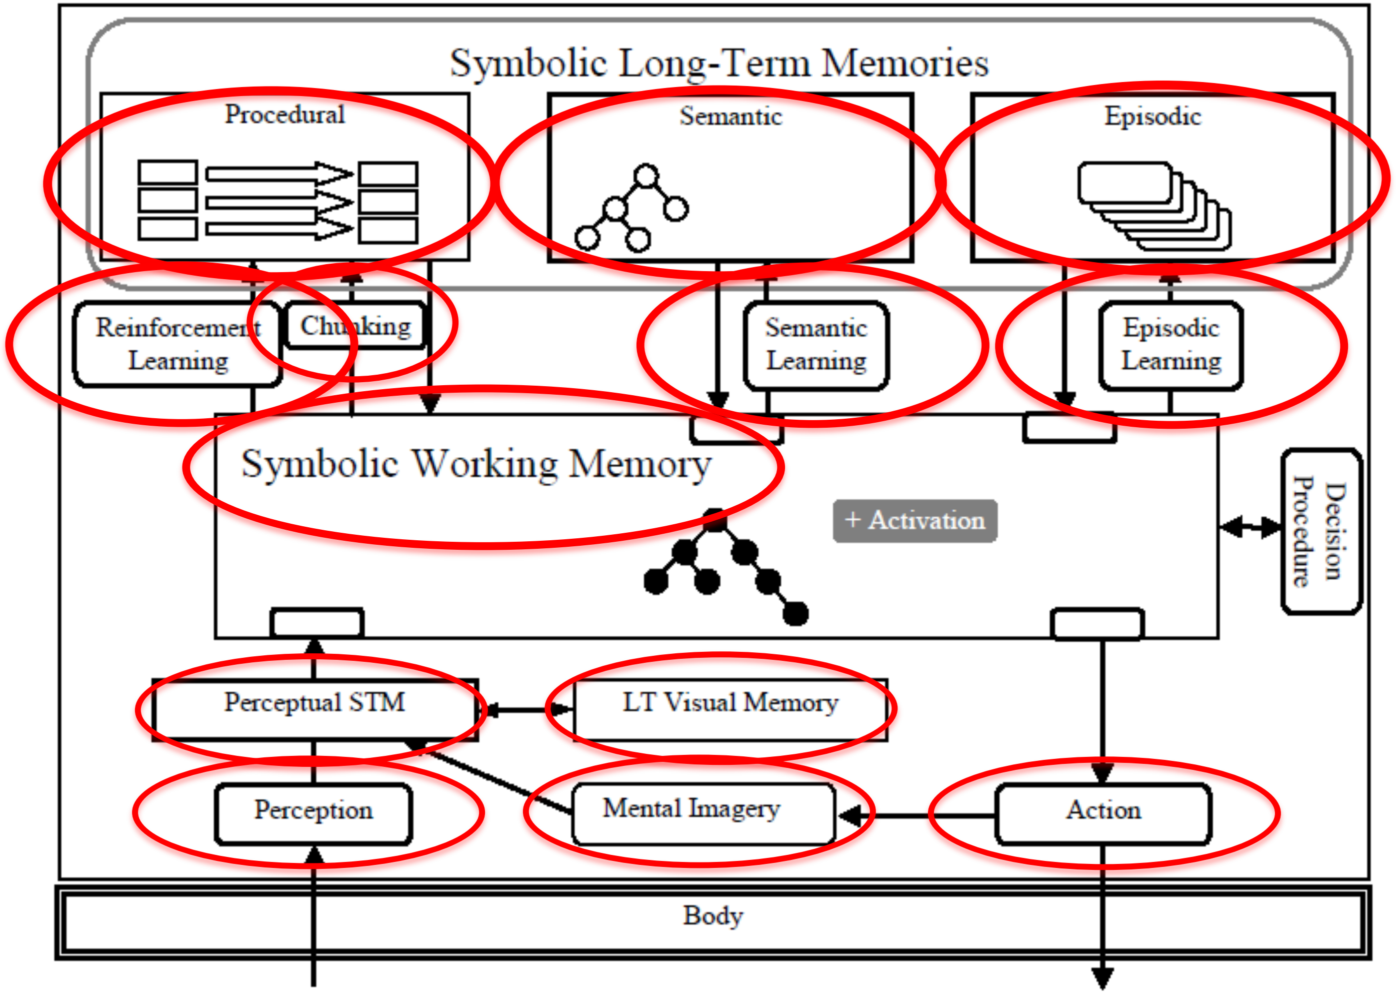

TransAIR Conference "Demoratize AI"
https://transair-bridge.org/conference-2/

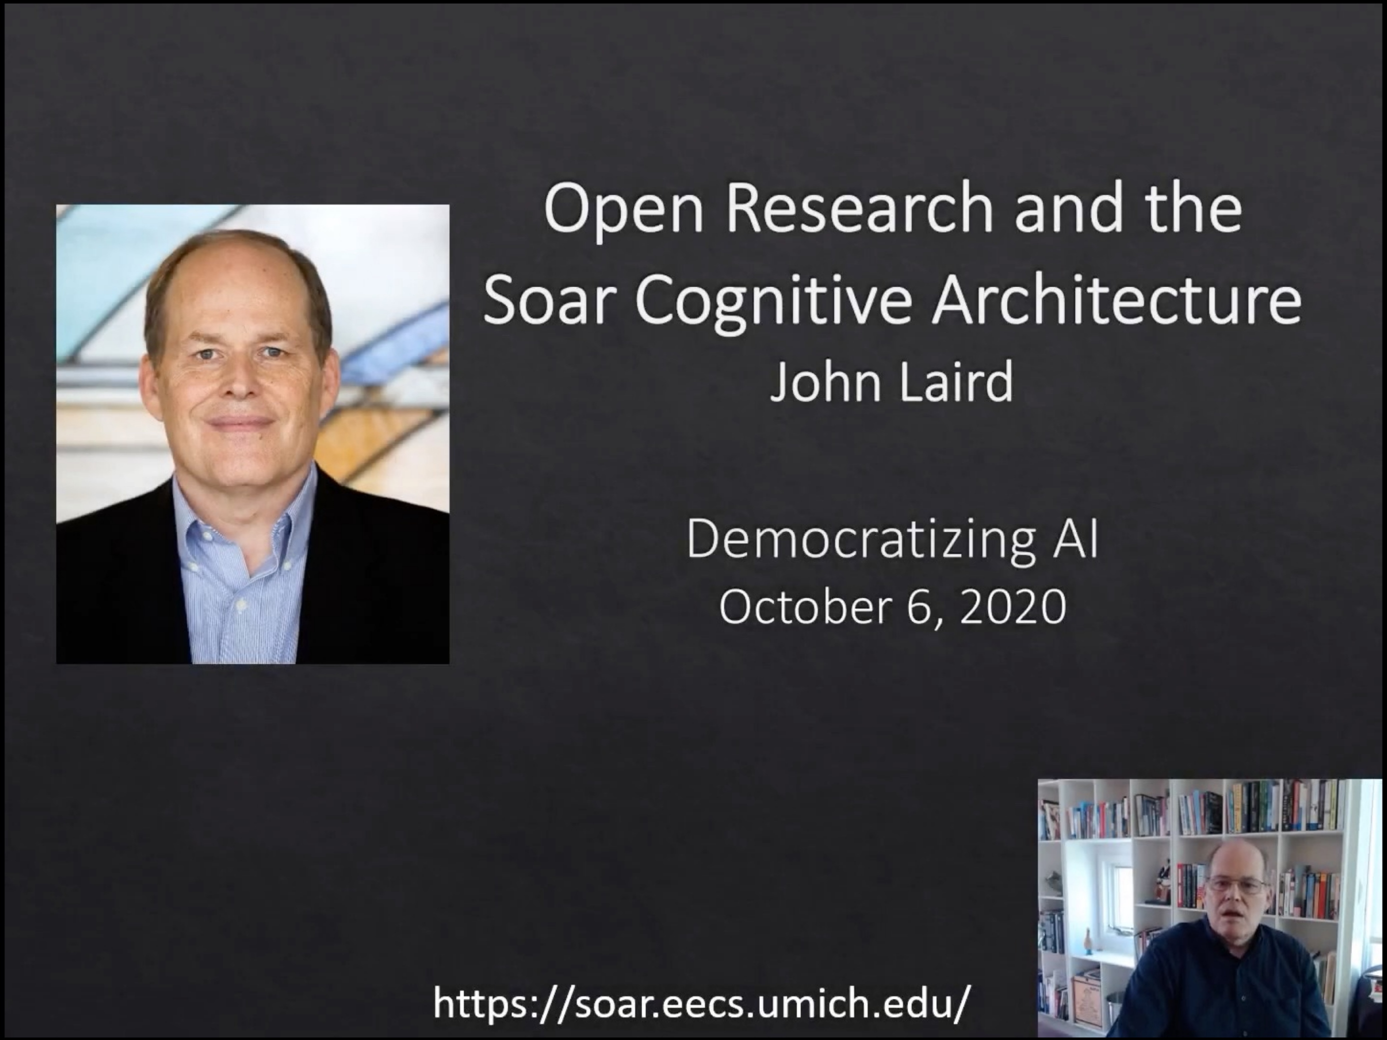

The following is a short excerpt
For the full talk see: https://www.youtube.com/watch?v=2pNsfBj7XSA&feature=youtu.be

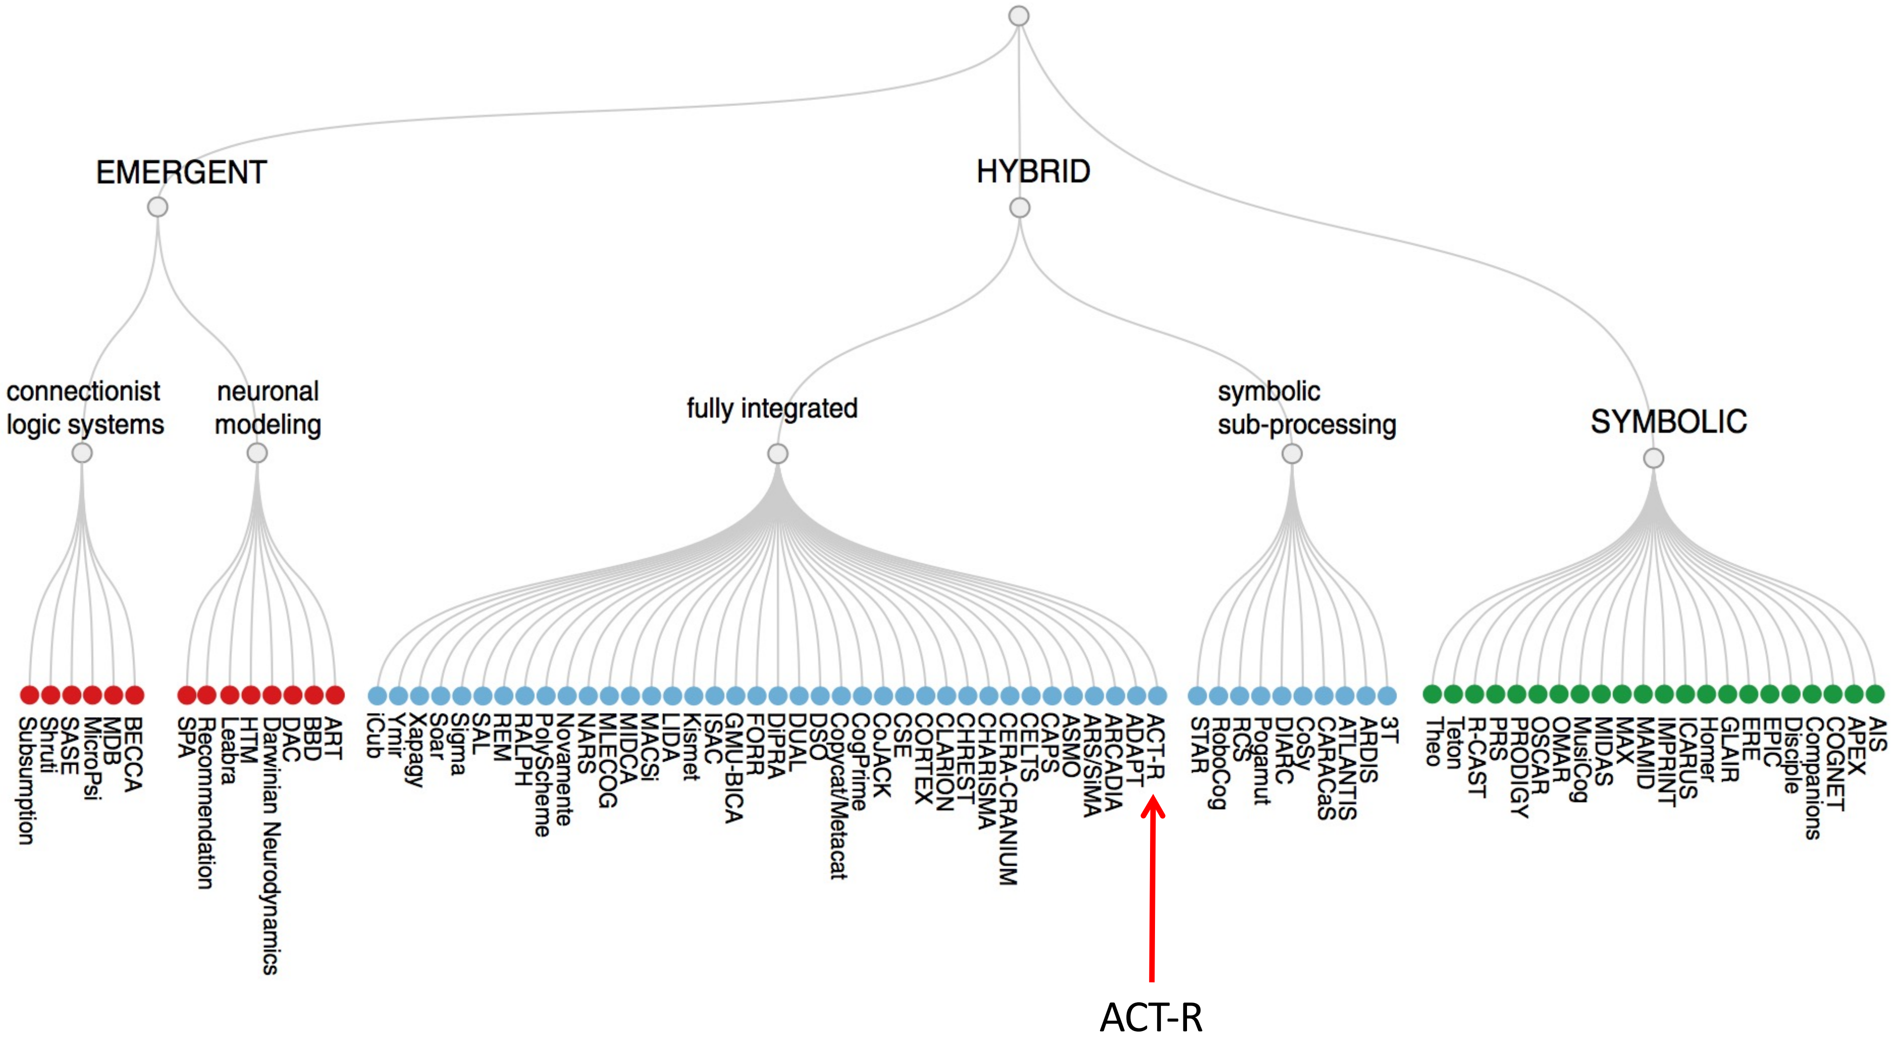
EMERGENT
HYBRID
SYMBOLIC
connectionist logic systems
neuronal modeling
fully integrated
symbolic sub-processing
Subsumption
Shruti
SASE
MicroPsi
MDB
BECCA
SPA
Recommendation
Leabra
HTM
Darwinian Neurodynamics
DAC
BBD
ART
iCub
Ymir
Xapagy
Soar
Sigma
SAL
REM
RALPH
PolyScheme
Novamente
NARS
MIDCA
MLECOG
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.

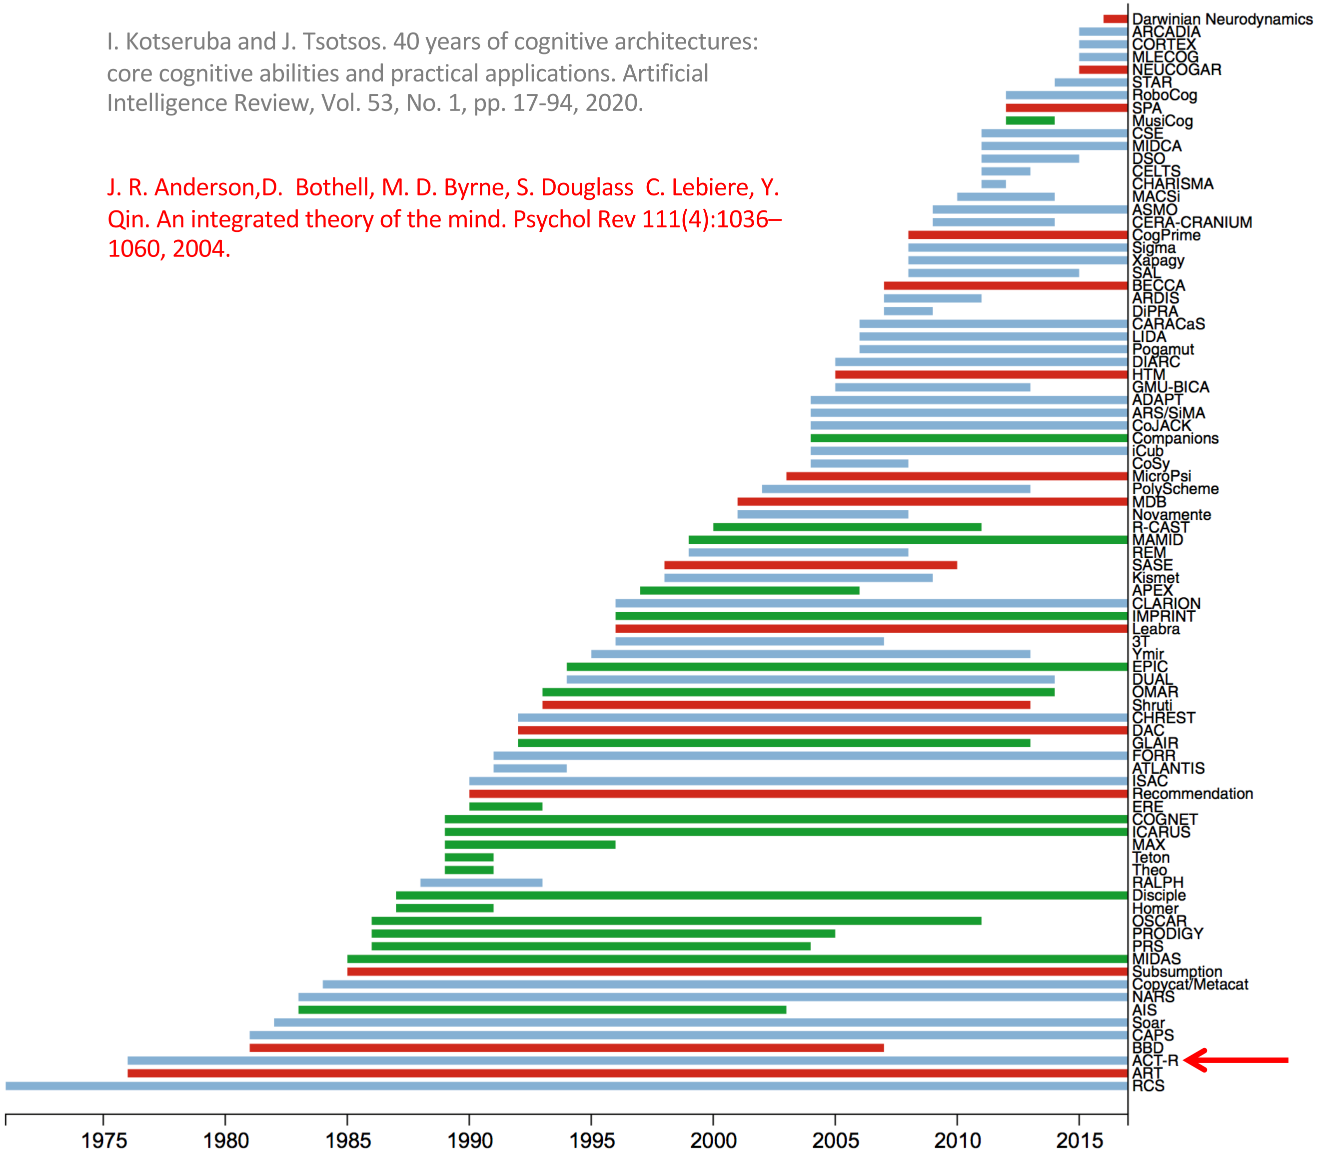

ACT-R

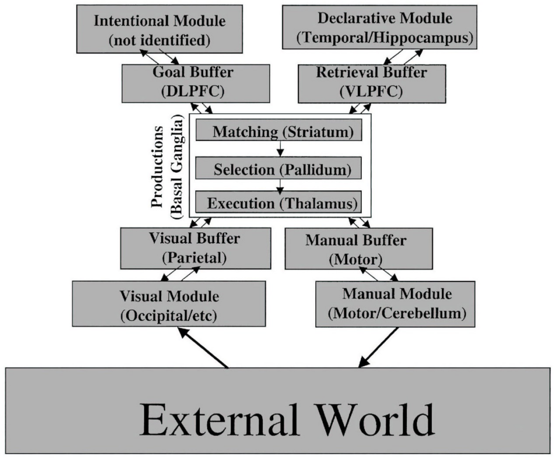

- J. Anderson's candidate for a Unified Theory of Cognition
- 1996, 2004; now version 7
- Production system with five modules: Intentional, Declarative, Visual, Manual, Production
- Cyclic operation: executes one production per cycle
  - Pattern of information in the buffers is recognized
  - A single production is selected and fires
  - The buffers are updated
- Each cycle takes approximately 50 ms.

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.

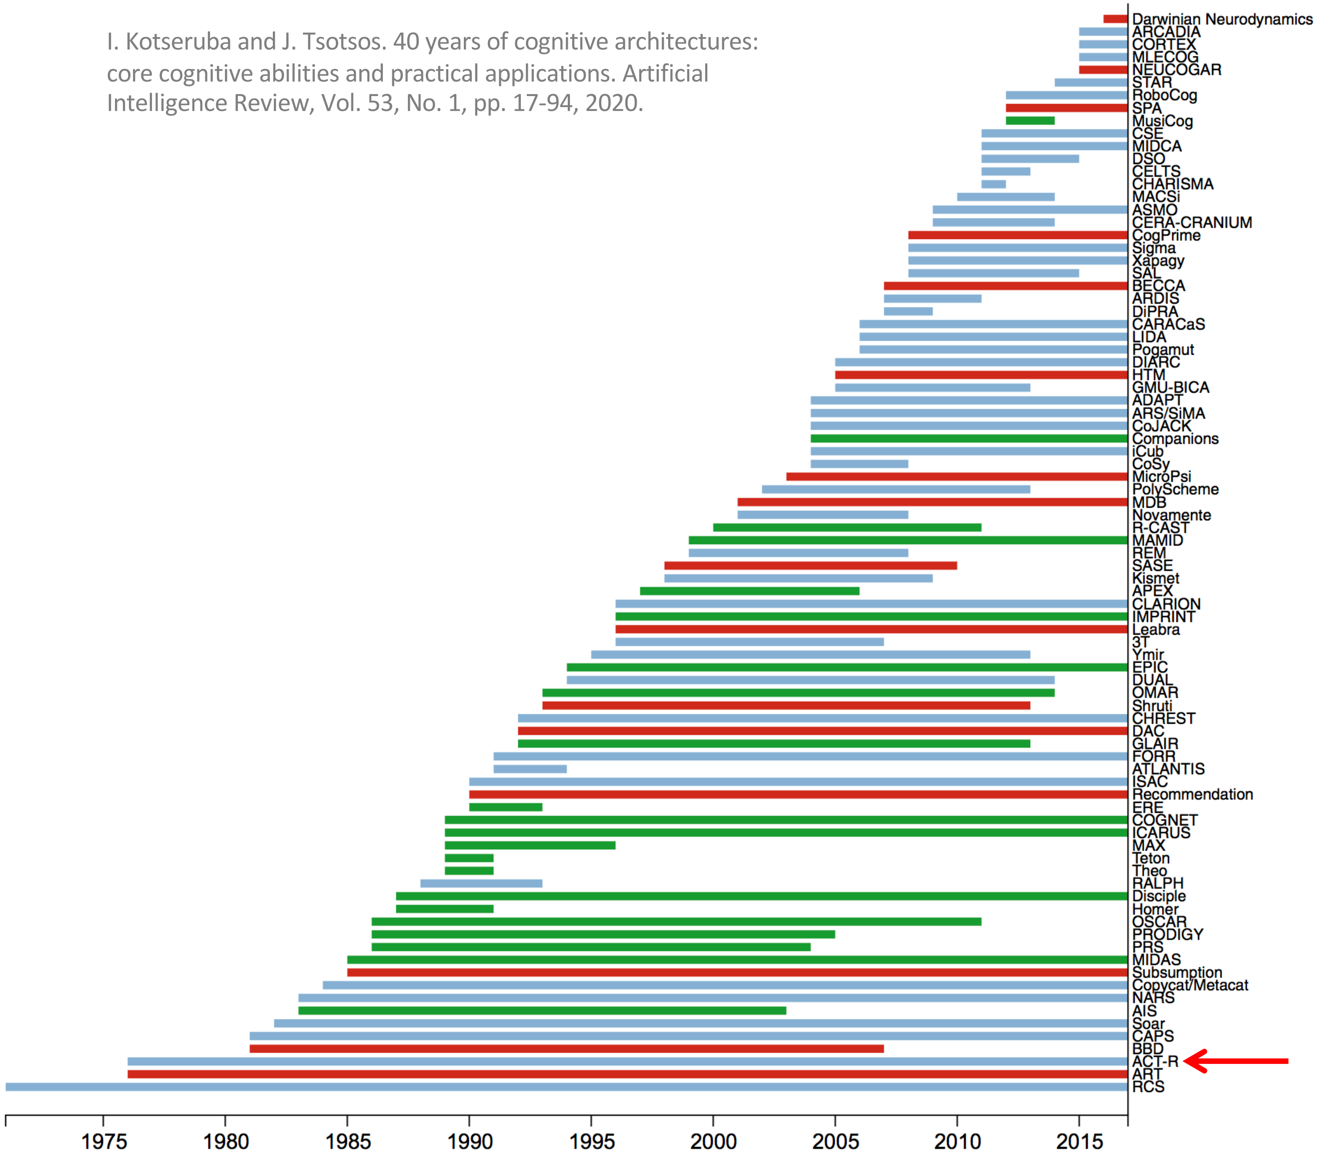

162 A Catalogue of Cognitive Architectures

A.1.3 ACT-R — Adaptive Control of Thought - Rational

The ACT-R [6, 7] cognitive architecture is a widely-regarded candidate for a unified theory of cognition. It focusses on modular decomposition and offers a theory of how these modules are integrated to produce coherent cognition. The architecture comprises five specialized modules, each devoted to processing a different kind of information (see Figure A.1). There is a vision module for determining the identity and position of objects in the visual field, a manual module for controlling hands, a declarative module for retrieving information from long-term information, and a goal module for keeping track of the internal state when solving a problem. Finally, it also has a production system that coordinates the operation of the other four modules. It does this indirectly via four buffers into which each module places a limited amount of information.

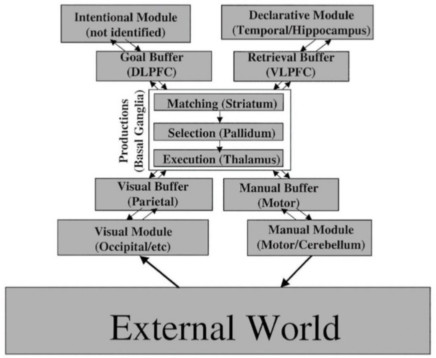

Fig. A.1 The ACT-R Cognitive Architecture (from [7])

ACT-R operates in a cyclic manner in which the patterns of information held in the buffers (and determined by external world and internal modules) are recognized, a single production fires, and the buffers are updated. It is assumed that this cycle takes approximately 50 ms.

There are two serial bottle-necks in ACT-R. One is that the content of any buffer is limited to a single declarative unit of knowledge, called a 'chunk'. This implies that only one memory can be retrieved at a time and indeed that a single object can be encoded in the visual field at any one time. The second bottle-neck is that only one production is selected to fire in any one cycle. This contrasts with both Soar and

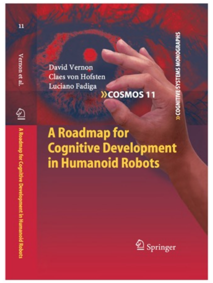

D. Vernon, C. von Hofsten, and L. Fadiga. A Roadmap for Cognitive Development in Humanoid Robots, Cognitive Systems Monographs (COSMOS), Vol. 11, Springer, 2010.

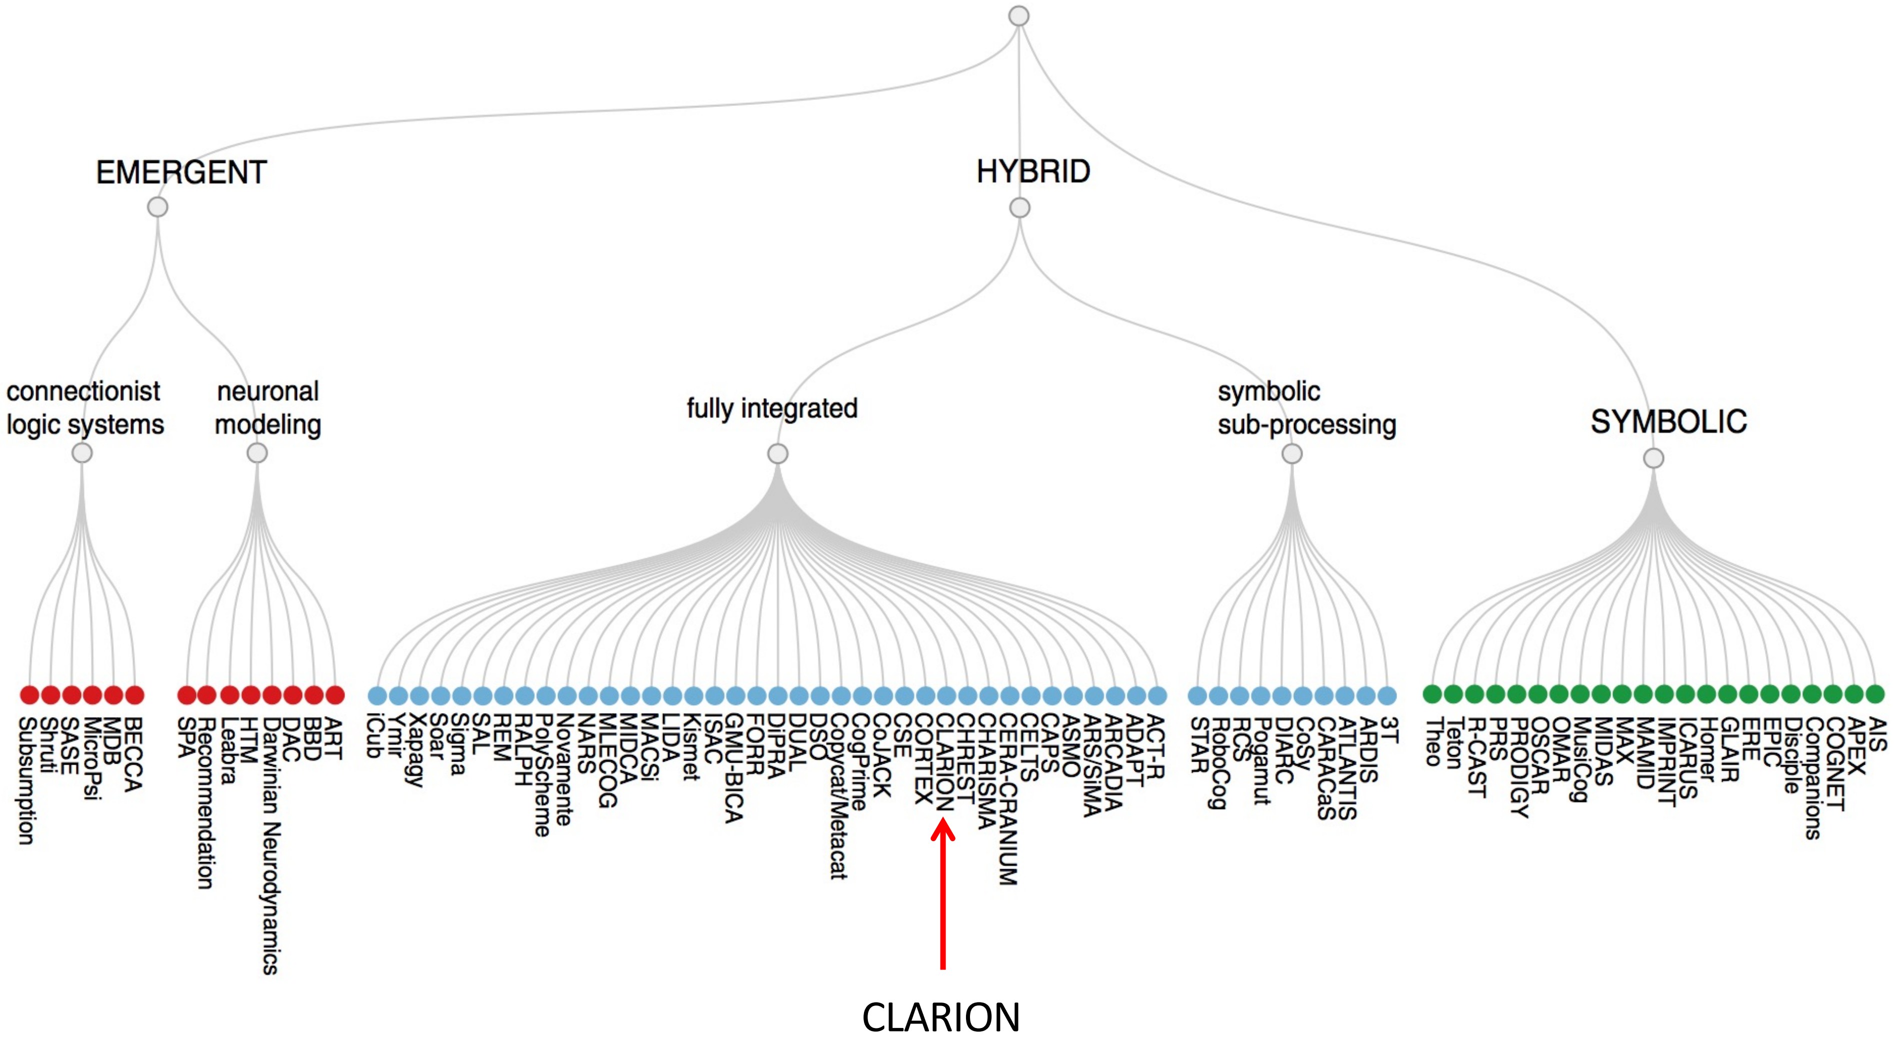
EMERGENT
HYBRID
connectionist logic systems
neuronal modeling
fully integrated
symbolic sub-processing
SYMBOLIC
Subsumption
Shruti
SASE
MicroPsi
MDB
BECCA
SPA
Recommendation
Leabra
HTM
Darwinian Neurodynamics
DAC
BBD
ART
iCub
Ymir
Xapagy
Soar
Sigma
SAL
REM
RALPH
PolyScheme
Novamente
NARS
MLECOG
MIDCA
MACSi
LIDA
Kismet
ISAC
GMU-BICA
FORR
DiPRA
DUAL
DSO
CogPrime
Copycat/Metacat
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
CLARION

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.

R. Sun. Anatomy of the mind: exploring psychological mechanisms and processes with the CLARION cognitive architecture. Oxford University Press, Oxford. 2016

R. Sun. The CLARION Cognitive Architecture: Toward a Comprehensive Theory of the Mind. The Oxford Handbook of Cognitive Science, S. Chipman (Ed.), 2017.

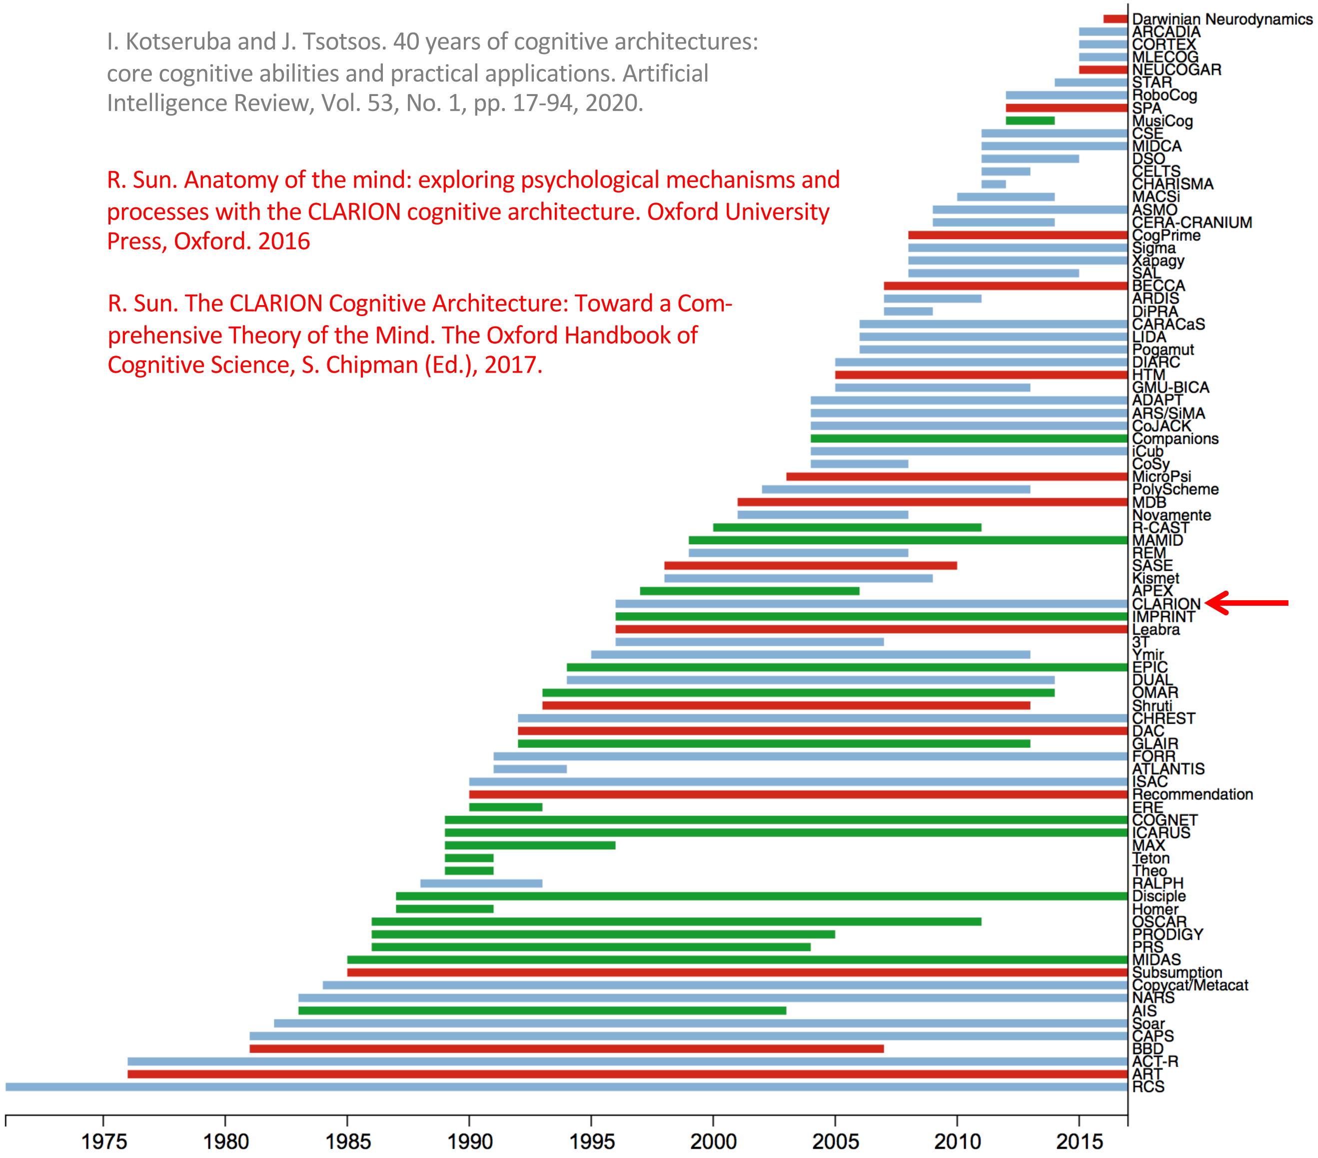

CLARION

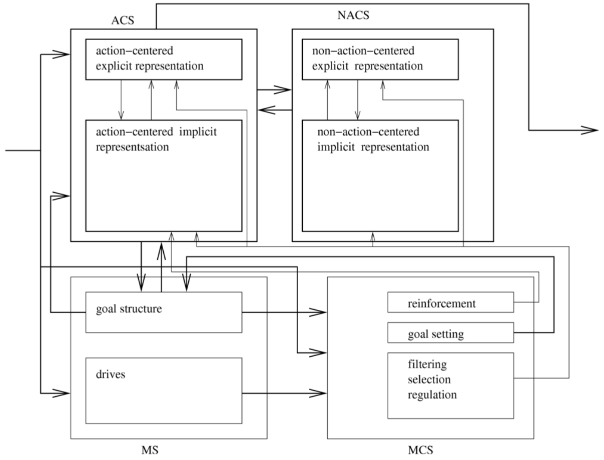

Four sub-systems

1. ACS – Action-centred subsystem
2. NACS – Non-action-centred subsystem
3. MS – Motivational subsystem
4. MCS – meta-cognitive subsystem

All four subsystems have two levels of knowledge representation

- Implicit connectionist bottom level
- Explicit symbolic top level

Implicit and explicit levels interact and cooperate both in action selection and in learning

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.

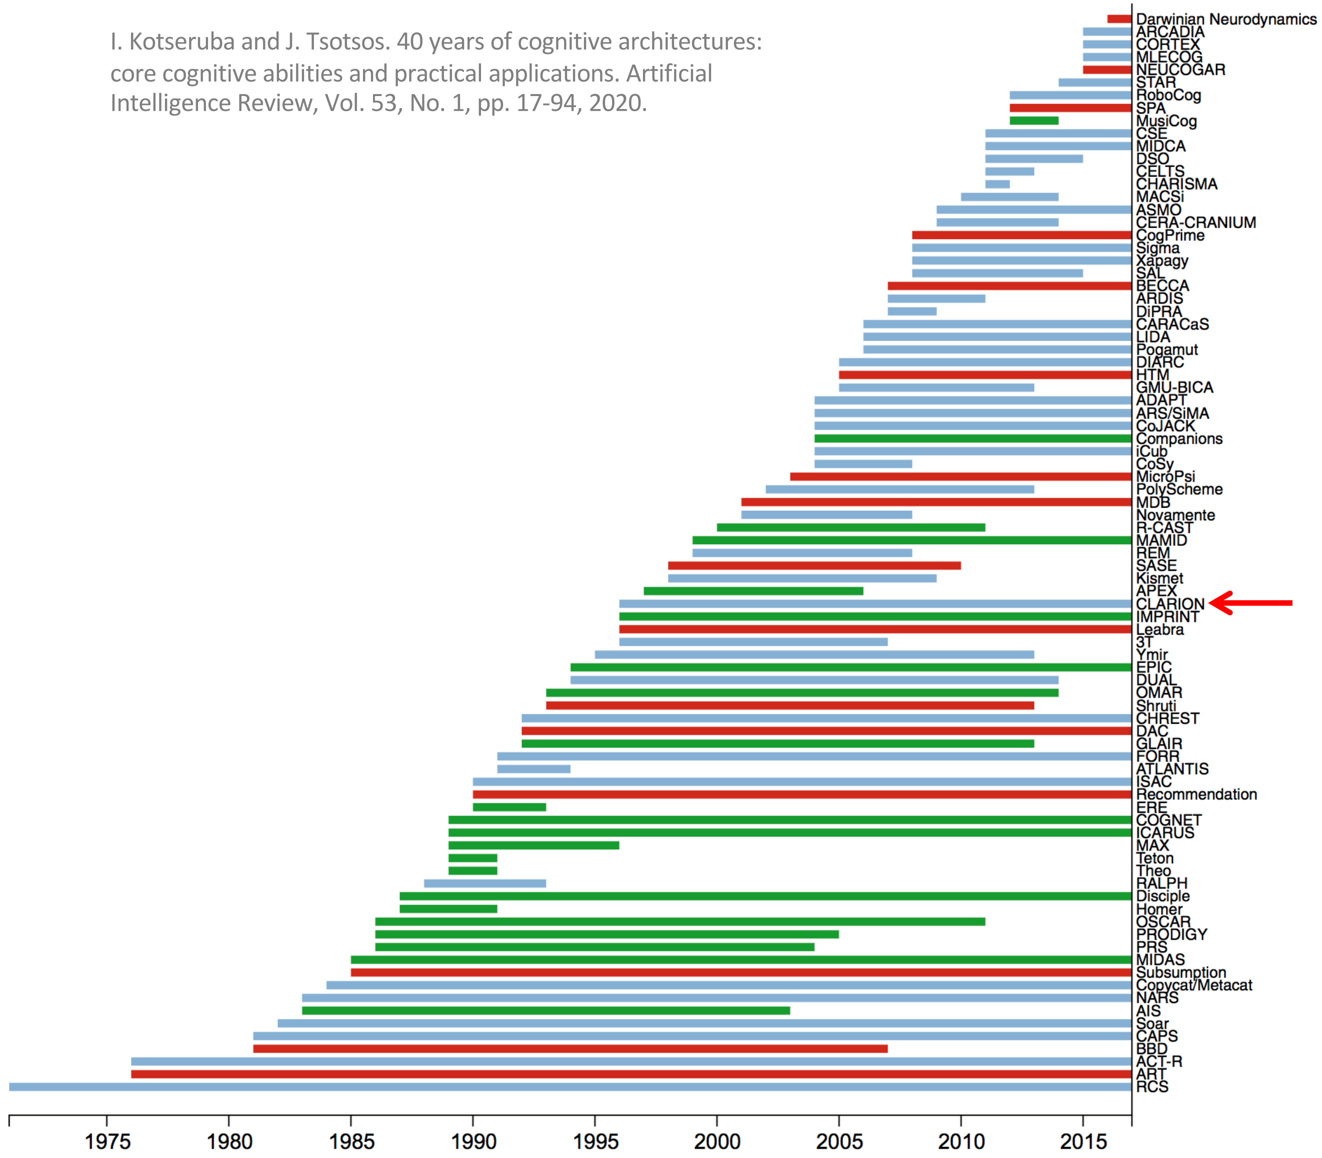

194 A Catalogue of Cognitive Architectures

A.3.6 *The CLARION Cognitive Architecture*

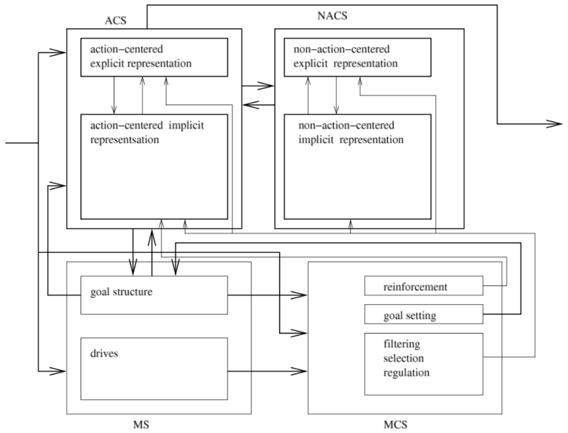

Fig. A.10 The CLARION hybrid cognitive architecture (from [364]). ACS stand for the action-centered subsystem, NACS for the non-action-centred subsystem, MS for the motivational subsystem, and MCS for the meta-cognitive subsystem. All four subsystems have two types of representation: implicit (connectionist) and explicit (symbolic).

CLARION [362, 363, 364] is an architypal hybrid cognitive architecture, deploying both connectionist and symbolic representations. It comprises four subsystems:

1. An action-centred subsystem (ACS);
2. A non-action-centred subsystem (NACS);
3. A motivational subsystem (MS);
4. A meta-cognitive subsystem (MCS).

All four subsystems have two levels of knowledge representation: an implicit connectionist bottom level and an explicit symbolic top level. The implicit and explicit levels interact and cooperate both in action selection and in learning.

The action-centred subsystem controls both external physical movements and internal "mental" operations. Given some observational state, i.e. a set of sensory features, the bottom level evaluates the desirability of all possible actions. The desirability is learned by reinforcement learning using the Q-Learning algorithm [392]. At the same time, the top level identifies possible actions from a rule network, again based on the observed sensory features. The bottom-level and top-level action are compared and the most appropriate top-level action is selected and executed. The

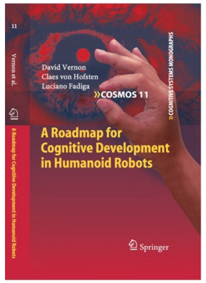

D. Vernon, C. von Hofsten, and L. Fadiga. A Roadmap for Cognitive Development in Humanoid Robots, Cognitive Systems Monographs (COSMOS), Vol. 11, Springer, 2010.

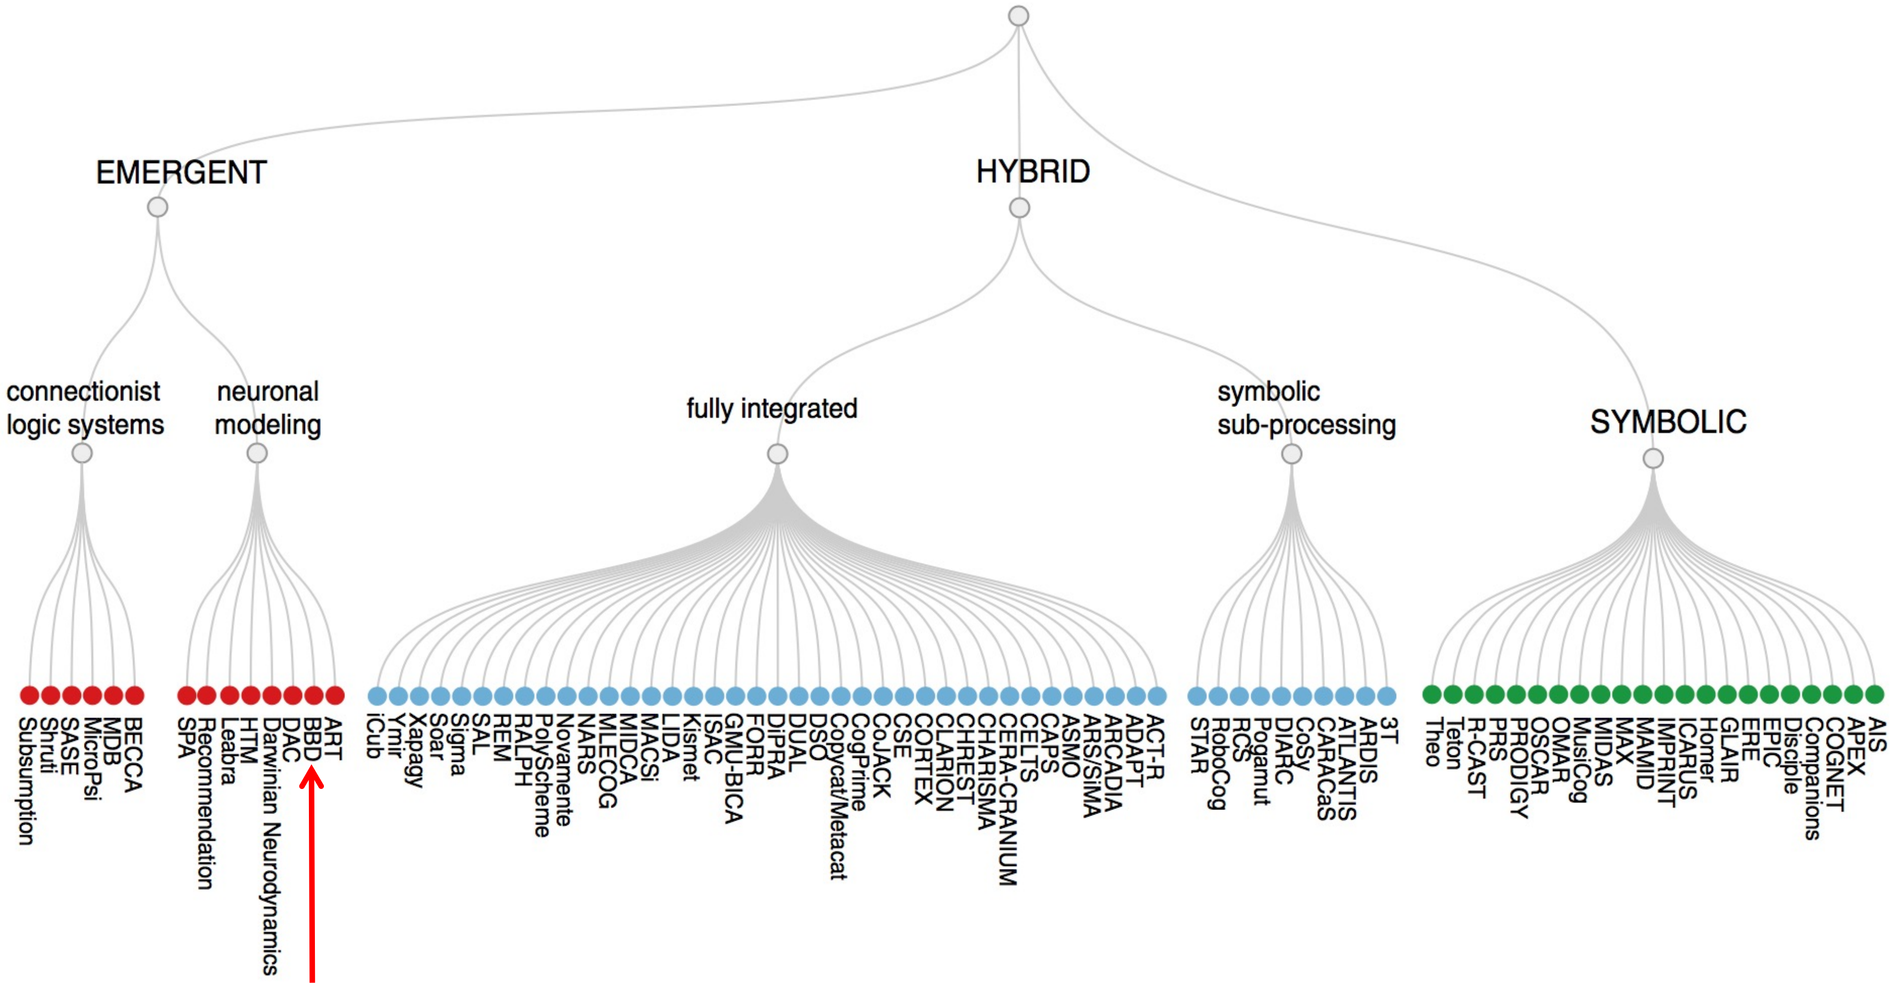

Brain-Based Devices (BBD)

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. L. Krichmar and G. M. Edelman. Brain-based devices for the study of nervous systems and the developmentof intelligent machines. Artificial Life 11(1–2):63–77, 2005.

J.G. Fleischer and G.M. Edelman. Brain based Devices. IEEE Robotics and Automation Magazine, 2009.

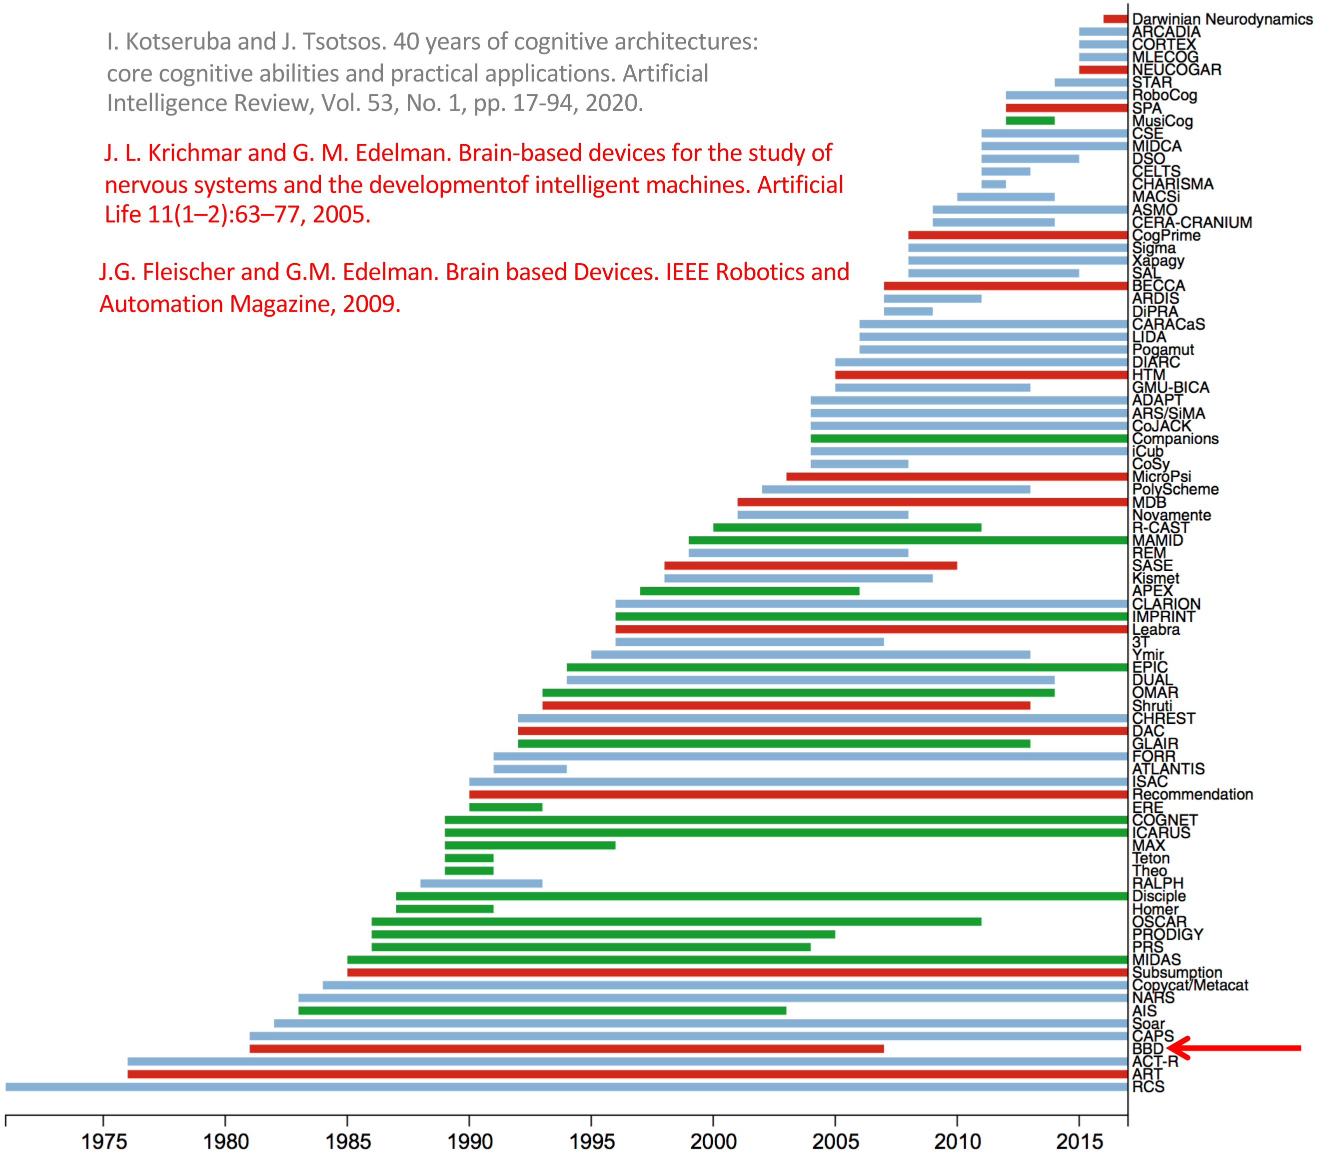

BBD

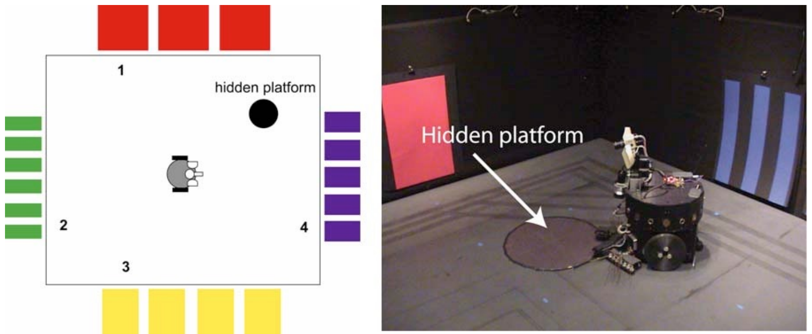

- Series of robot platforms focussed on developmental cognition
- Brain-based Devices (BBD)
  - Simulated nervous system
  - Develop spatial and episodic memory
  - Recognition capabilities
  - Autonomous experiential learning
- Neuromimetic: mimic the neural structure of the brain
- Differ from connectionist approaches: focus on
  - Nervous system as a whole
  - Constituent parts
  - Their interaction

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. L. Krichmar and G. M. Edelman. Brain-based devices for the study of nervous systems and the developmentof intelligent machines. Artificial Life 11(1–2):63–77, 2005.

J.G. Fleischer and G.M. Edelman. Brain based Devices. IEEE Robotics and Automation Magazine, 2009.

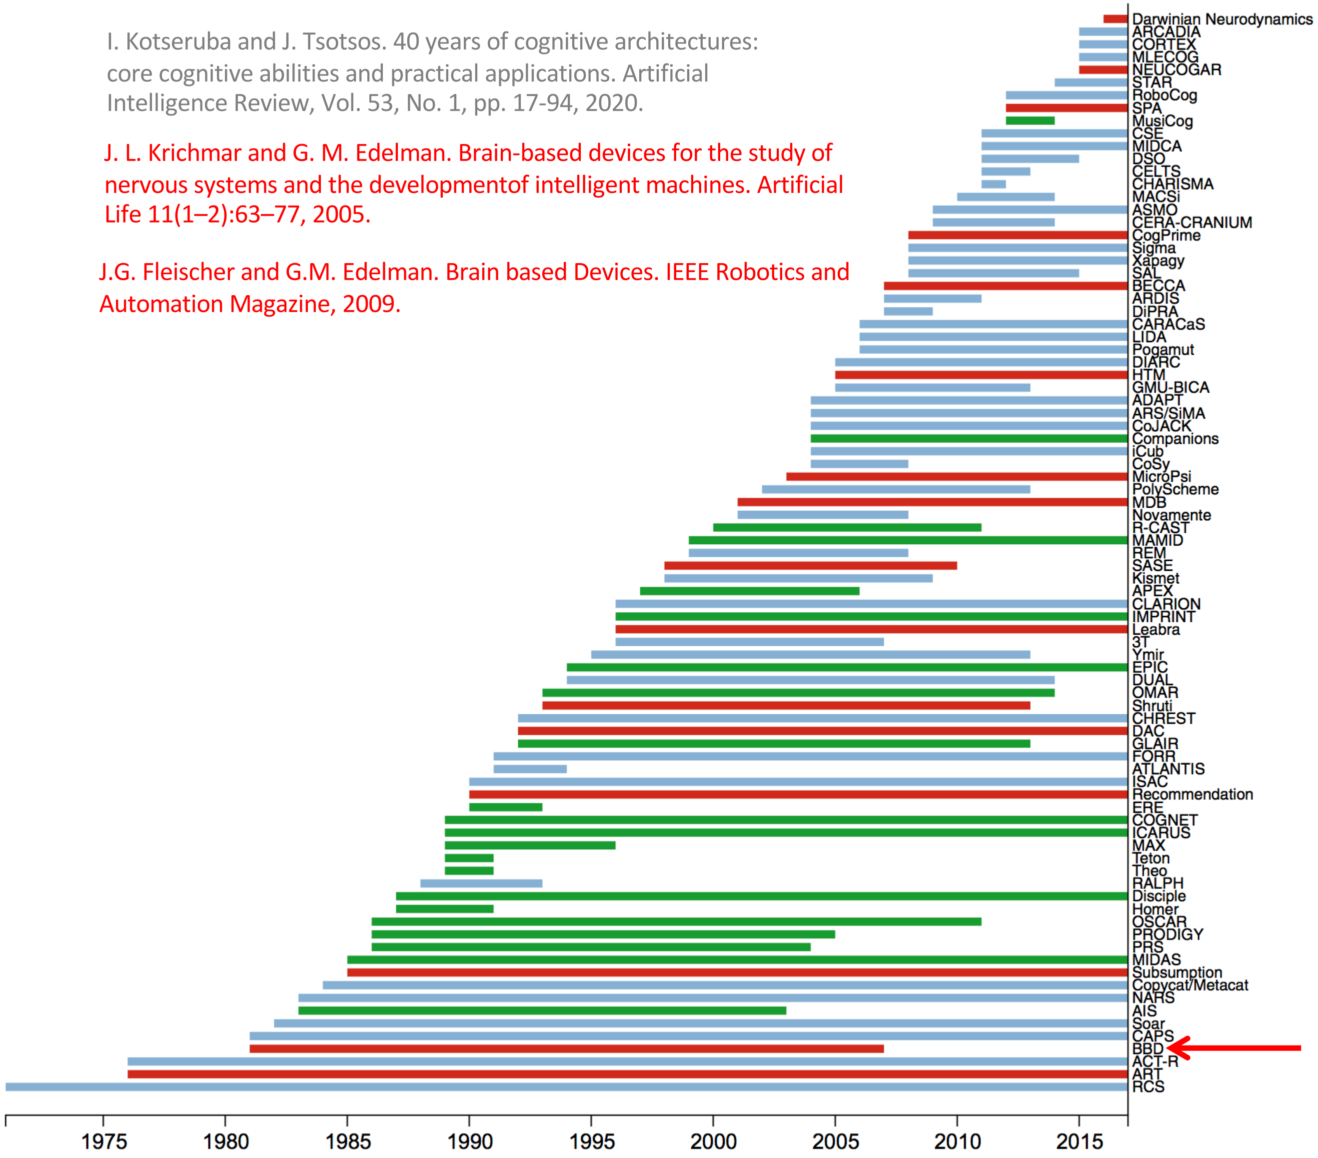

A.2 Emergent Cognitive Architectures 181

A.2.5 *Darwin: Neuromimetic Robotic Brain-Based Devices*

Kirchmar *et al.* [203, 204, 205, 206, 207, 334] have developed a series of robot platforms called Darwin to experiment with developmental agents. These systems are 'brain-based devices' (BBDs) which exploit a simulated nervous system that can develop spatial and episodic memory as well as recognition capabilities through autonomous experiential learning. As such, BDDs are a neuromimetic approach in the emergent paradigm that is most closely aligned with the enactive and the connectionist models. It differs from most connectionist approaches in that the architecture is much more strongly modelled on the structure and organization of the brain than are conventional artificial neural networks, *i.e.* they focus on the nervous system as a whole, its constituent parts, and their interaction, rather than on a neural implementation of some individual memory, control, or recognition function.

The principal neural mechanisms of the BDD approach are synaptic plasticity, a reward (or value) system, reentrant connectivity, dynamic synchronization of neuronal activity, and neuronal units with spatiotemporal response properties. Adaptive behaviour is achieved by the interaction of these neural mechanisms with sensorimotor correlations (or contingencies) which have been learned autonomously by active sensing and self-motion.

Darwin VIII is capable of discriminating reasonably simple visual targets (coloured geometric shapes) by associating it with an innately preferred auditory cue. Its simulated nervous system contains 28 neural areas, approximately 54,000 neuronal units, and approximately 1.7 million synaptic connections. The architecture comprises regions for vision (V1, V2, V4, IT), tracking (C), value or saliency (S), and audition (A). Gabor filtered images, with vertical, horizontal, and diagonal selectivity, and red-green colour filters with on-centre off-surround and off-centre on-surround receptive fields, are fed to V1. Sub-regions of V1 project topographically to V2 which in turn projects to V4. Both V2 and V4 have excitatory and inhibitory reentrant connections. V4 also has a non-topographical projection back to V2 as well as a non-topographical projection to IT, which itself has reentrant adaptive connections. IT also projects non-topographically back to V4. The tracking area (C) determines the gaze direction of Darwin VIII's camera based on excitatory projections from the auditory region A. This causes Darwin to orient toward a sound source. V4 also projects topographically to C causing Darwin VIII to centre its gaze on a visual object. Both IT and the value system S have adaptive connections to C which facilitates the learned target selection. Adaptation is effected using the Hebbian-like Bienenstock-Cooper-Munroe (BCM) rule [41]. From a behavioural perspective, Darwin VIII is conditioned to prefer one target over others by associating it with the innately preferred auditory cue and to demonstrate this preference by orienting towards the target.

Darwin IX can navigate and categorize textures using artificial whiskers based on a simulated neuroanatomy of the rat somatosensory system, comprising 17 areas, 1101 neuronal units, and approximately 8400 synaptic connections.

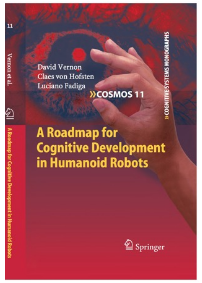

D. Vernon, C. von Hofsten, and L. Fadiga. A Roadmap for Cognitive Development in Humanoid Robots, Cognitive Systems Monographs (COSMOS), Vol. 11, Springer, 2010.

Recommended Reading

D. Vernon, Artificial Cognitive Systems – A Primer, MIT Press, 2014; Chapter 3.

D. Vernon, Cognitive Architectures, in Cognitive Robotics, A. Cangelosi and M. Asada, MIT Press, 2022.

A. Lieto, M. Bhatt, A. Oltramari, and D. Vernon, "The Role of Cognitive Architectures in General Artificial Intelligence", editorial for a special issue on "Cognitive Architectures for Artificial Minds", Cognitive Systems Research, Vol. 48, pp. 1-3, 2017.

I. Kotseruba and J. Tsotsos. 40 years of cognitive architectures: core cognitive abilities and practical applications. Artificial Intelligence Review, 2020

J. E. Laird, C. Lebiere, and P. S. Rosenbloom, “A standard model of the mind: toward a common computational framework across artificial intelligence, cognitive science, neuroscience, and robotics”, AI Magazine, Vol. 38, pp. 13-26.

D. Vernon, C. von Hofsten, and L. Fadiga. "A Roadmap for Cognitive Development in Humanoid Robots", Cognitive Systems Monographs (COSMOS), Vol. 11, Springer, 2010; Chapter 5 and Appendix A.

References

Krichmar, J. L. (2012). Design principles for biologically inspired cognitive architectures. Biologically Inspired Cognitive Architectures 1, 73–81.

Langley, P. (2005). An adaptive architecture for physical agents. In IEEE/WIC/ACM International Conference on Intelligent Agent Technology, Compiegne, France, pp. 18–25. IEEE Computer Society Press.

Langley, P. (2006). Cognitive architectures and general intelligent systems. AI Magazine 27(2), 33–44.

Langley, P., J. E. Laird, and S. Rogers (2009). Cognitive architectures: Research issues and challenges. Cognitive Systems Research 10 (2), 141– 160.

Lehman, J. F., J. E. Laird, and P. S. Rosenbloom (1998). A gentle intro- duction to soar, an architecture for human cognition. In S. Sternberg and D. Scarborough (Eds.), Invitation to Cognitive Science, Volume 4: Methods, Models, and Conceptual Issues. Cambridge, MA: MIT Press.

Newell, A. (1990). Unified Theories of Cognition. Cambridge MA: Harvard University Press.

References

Ritter, F. E. and R. M. Young (2001). Introduction to this special issue on using cognitive models to improve interface design. International Journal of Human-Computer Studies 55, 1–14.

Samsonovich, A. (2010). Toward a unified catalog of implemented cog- nitive architectures. In Proc. the Conference on Biologically Inspired Cognitive Architectures, pp. 195–244.

Sun, R. (2004). Desiderata for cognitive architectures. Philosophical Psychology 17(3), 341–373.

Sun, R. (2007). The importance of cognitive architectures: an analysis based on CLARION. Journal of Experimental & Theoretical Artificial Intelligence 19(2), 159–193.

Vernon, D., C. von Hofsten, and L. Fadiga (2016). Desiderata for developmental cognitive architectures. Biologically Inspired Cognitive Architectures 18, 116–127.

Ziemke, T. (2016). The body of knowledge: On the role of the living body in grounding embodied cognition. BioSystems 148, 4–11.

Recommended Videos

Daniel Wolpert, Columbia University: The Real Reason for Brains

https://www.youtube.com/watch?v=7s0CpRfyYp8

John E. Laird, University of Michigan: Open Research and the Soar Cognitive Architecture

https://www.youtube.com/watch?v=2pNsfBj7XSA&feature=youtu.be

John E. Laird, University of Michigan: The Soar Cognitive Architecture: Current and Future Capabilities

https://www.youtube.com/watch?v=BUiWk-DqLaA

Kazuhiko Kawamura, Vanderbilt University: Cognitive Robotics and Control:

https://www.youtube.com/watch?v=7i_l80w2mtg

Jeffrey Krichmar, University of California: Neurorobotics: Connecting the Brain, Body and Environment

https://www.youtube.com/watch?v=rb2OQH7ghW8

Ron Sun, Rensselaer Polytechnic Institute: Clarion: A comprehensive, Integrative Cognitive Architecture

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

Recommended Videos

These and other short videos on cognitive architectures can be found at the 2021 TransAIR Workshop on Cognitive Architectures for Robot Agents

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