Towards a New Cognitive Hourglass

Uniform Implementation of Cognitive Architecture via Factor Graphs

Paul S. Rosenbloom
Department of Computer Science & Institute for Creative Technologies
University of Southern California
July 24, 2009

Cognitive Architecture

- Hypothesis about fixed structure underlying cognition
  - Defines core memories, reasoning processes, learning mechanisms, external interfaces, etc.
- Yields intelligent behavior when combined with knowledge in memories
  - Including more advanced reasoning, learning, etc.
- May model human cognition, strive for human-like intelligence, or be purely artificial

Example: Soar (3-8)

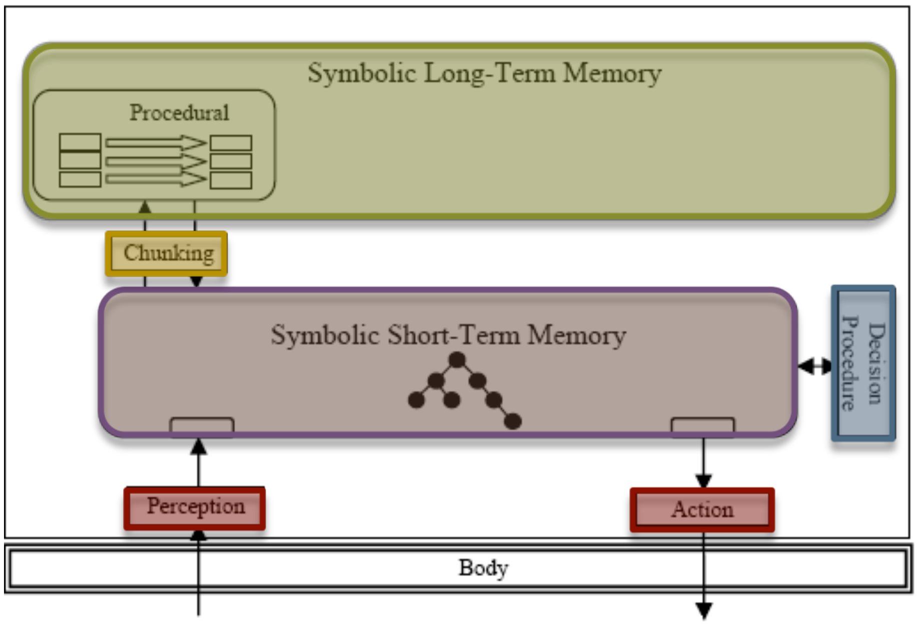

Soar 3-8

- Symbolic working memory
- Long-term memory of rules
- Decide what to do based on preferences retrieved into working memory
- Reflect when can't decide
- Learn results of reflection
- Interact with world

Systems Levels in Cognition

- In the large, cognitive architecture is a theory about one or more *systems levels* in an intelligent being
  - Usually part of Cognitive Band
- At each level, a combination of structures and processes implements basic elements at the next higher level

TIME SCALE OF HUMAN ACTION

Scale (sec)	Time Units	System	World (theory)
10^{7}	months		SOCIAL BAND
10^{6}	weeks		
10^{5}	days		
10^{4}	hours	Task	RATIONAL BAND
10^{3}	10 min	Task	
10^{2}	minutes	Task	
10^{1}	10 sec	Unit task	COGNITIVE BAND
10^{0}	1 sec	Operations	
10^{-1}	100 ms	Deliberate act	
10^{-2}	10 ms	Neural circuit	BIOLOGICAL BAND
10^{-3}	1 ms	Neuron	
10^{-4}	100 µs	Organelle	

(Newell, 1990)

Hierarchical View of Soar (3-8)

Scale	Functionality	Mechanism	Details
1 sec	Reflective	Problem Space Search	Impasse/Subgoal if can't Decide
100 ms	Deliberative	Decision Cycle	Preference-based Decisions upon Quiescence
10 ms	Reactive	Elaboration Cycle	Parallel Rule Match & Firing

Learning by Chunking

Architecture/Level Girth

- An architecture/level's *girth* is its range of structures & processes
  - Uniformity versus diversity
- Uniformity: Minimal mechanisms combining in general ways
  - Appeals to simplicity and elegance
  - The "physicist's approach"
  - Challenge: Achieving full range of required functionality/coverage
- Diversity: Large variety of specialized mechanisms
  - Appeals to functionality and optimization
  - The "biologist's approach"
  - Challenge: Achieving integrability, extensibility and maintainability
- Across a hierarchy, level girth may stay comparable or vary
  - Physicists and biologists likely assume uniform
  - Network researchers assume hourglass

The Internet hourglass

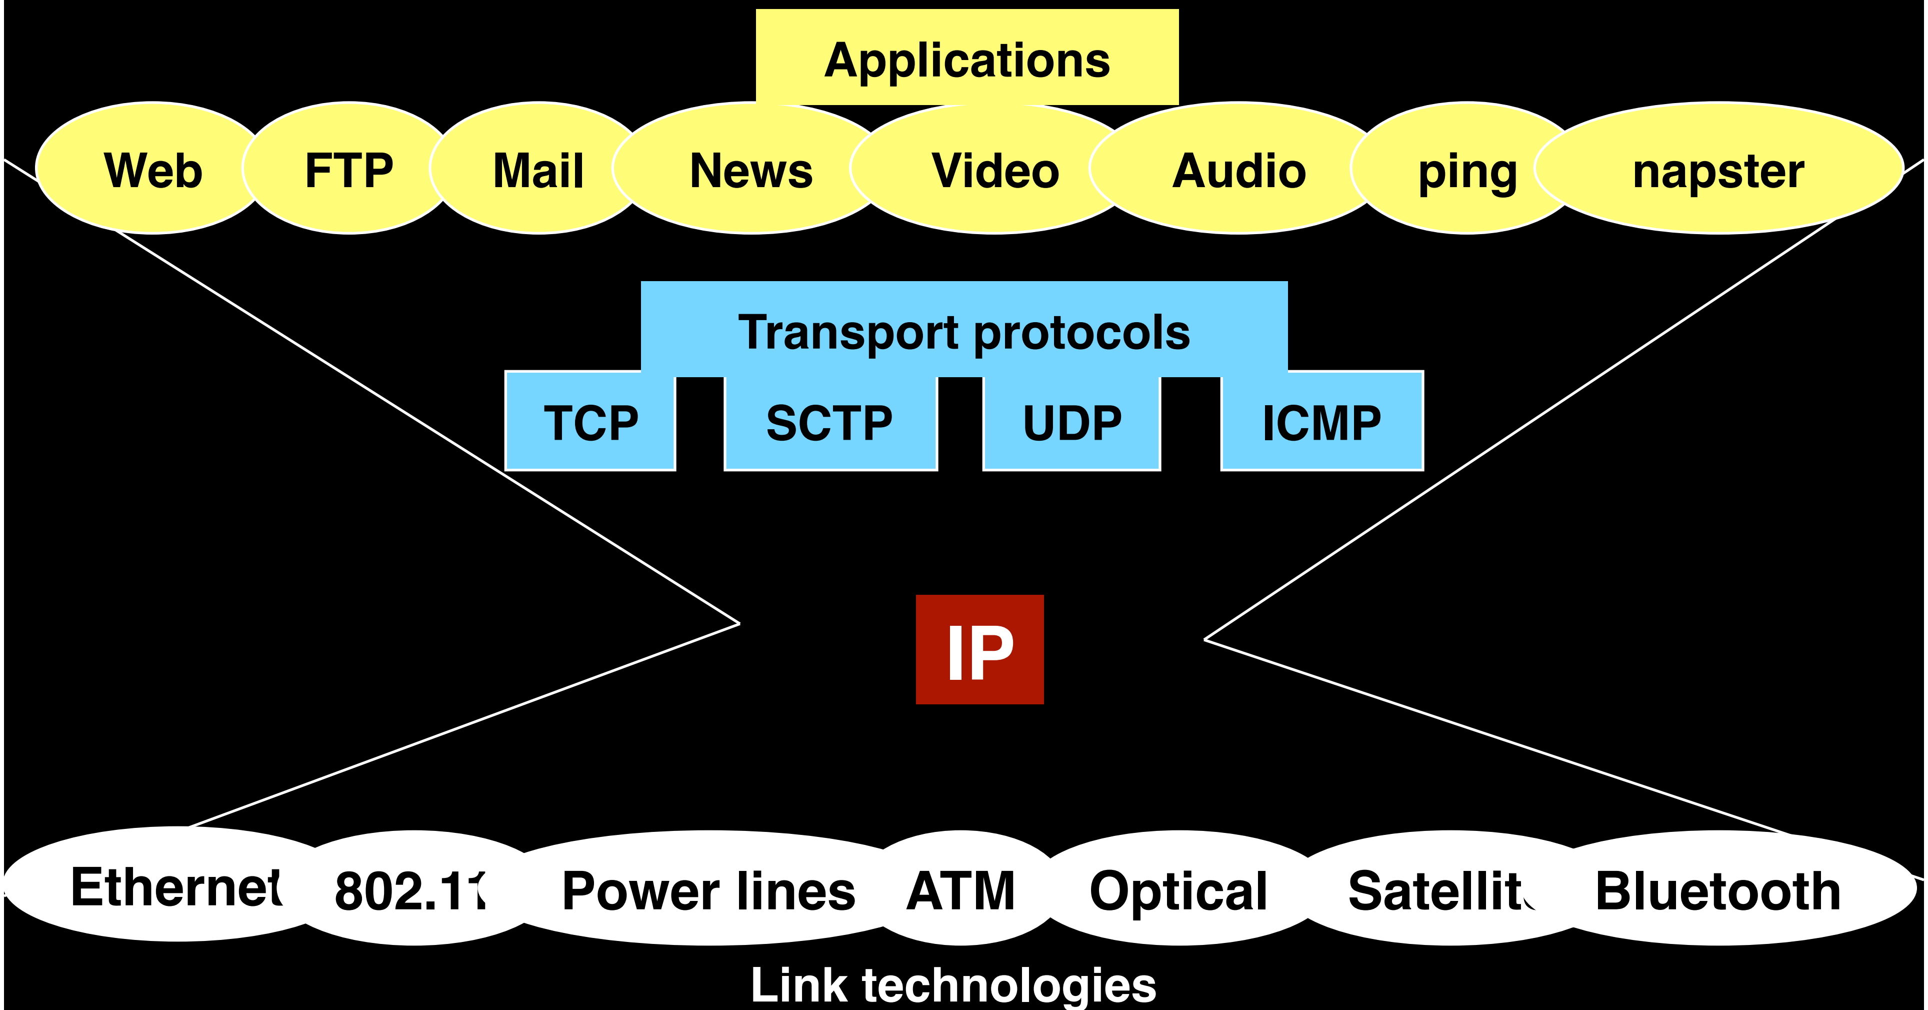

From Hari Balakrishnan

The Internet hourglass

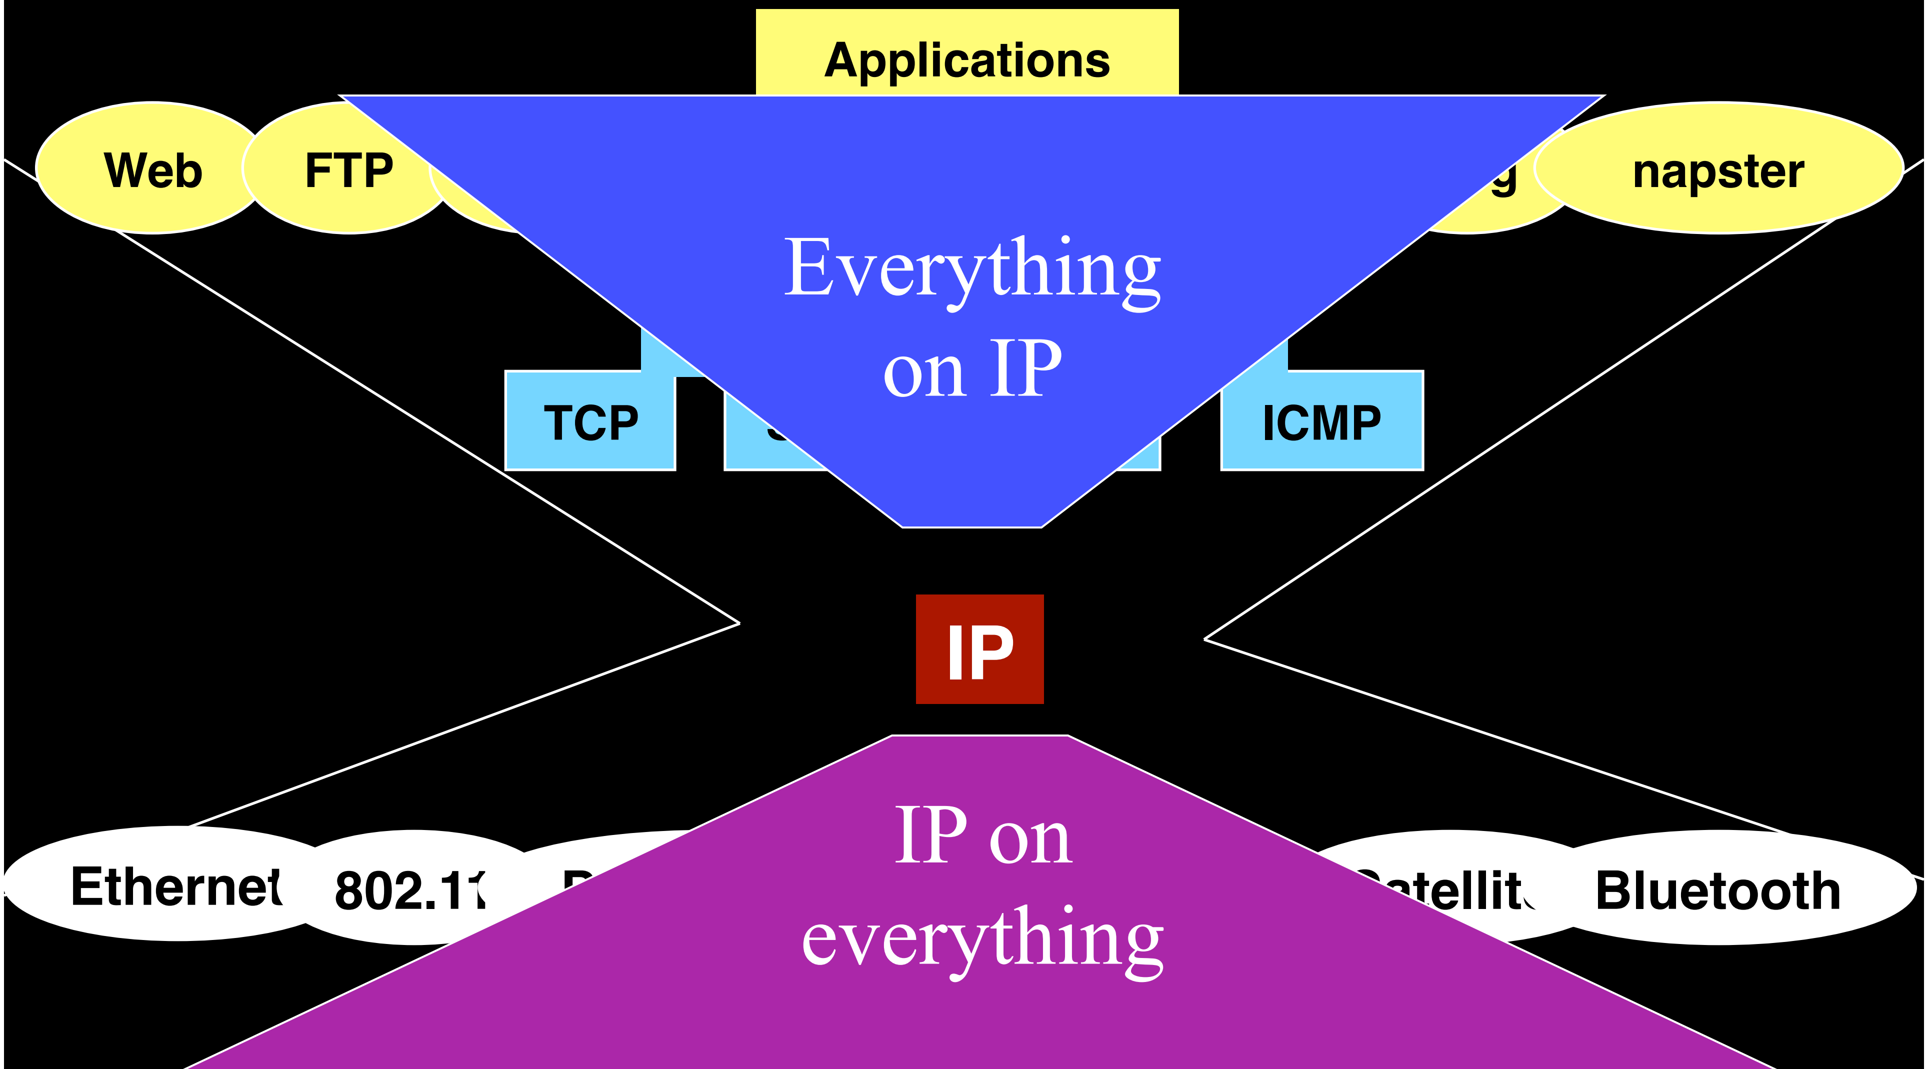

From Hari Balakrishnan

What About Cognition?

- Top (applications) is clearly diverse
  - Key part of what architectures try to explain
- Bottom is likely diverse as well
  - Physicalism: Grounded in diversity of biology
  - Strong AI: Also groundable in other technologies
- Is the waist uniform or diverse?
  - Hourglass or rectangle
  - Traditionally question about the architecture

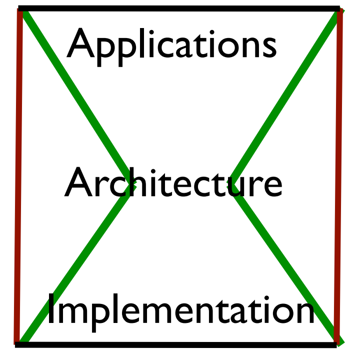

Examples

- Soar (3-8) is a traditional uniform architecture
- ACT-R is a traditional diverse architecture

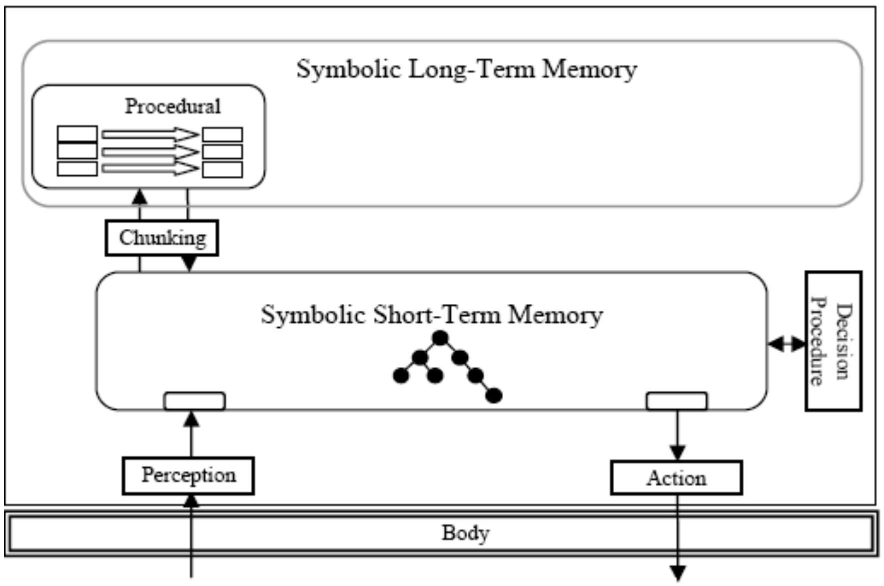

Soar 3-8

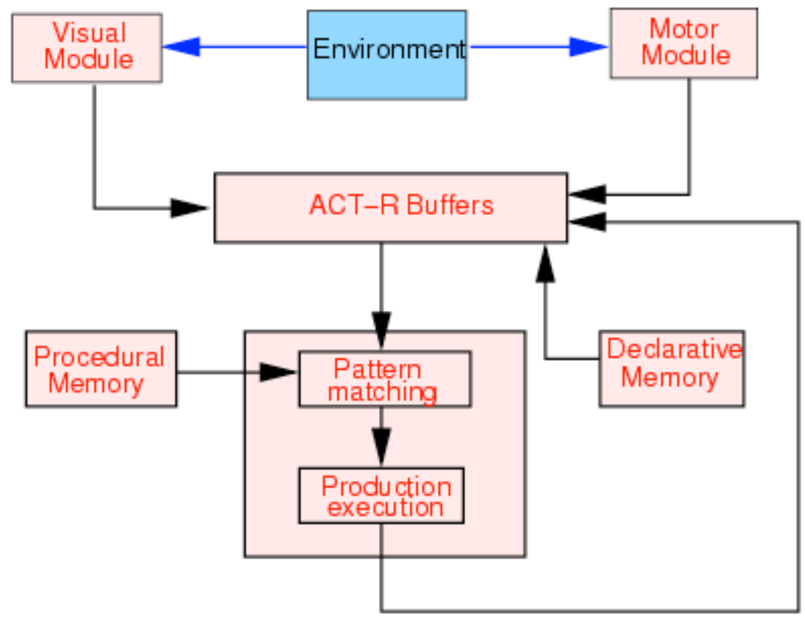

ACT-R

Examples

- Soar is a traditional uniform architecture
- ACT-R is a traditional diverse architecture
- Recently Soar 9 has become highly diverse

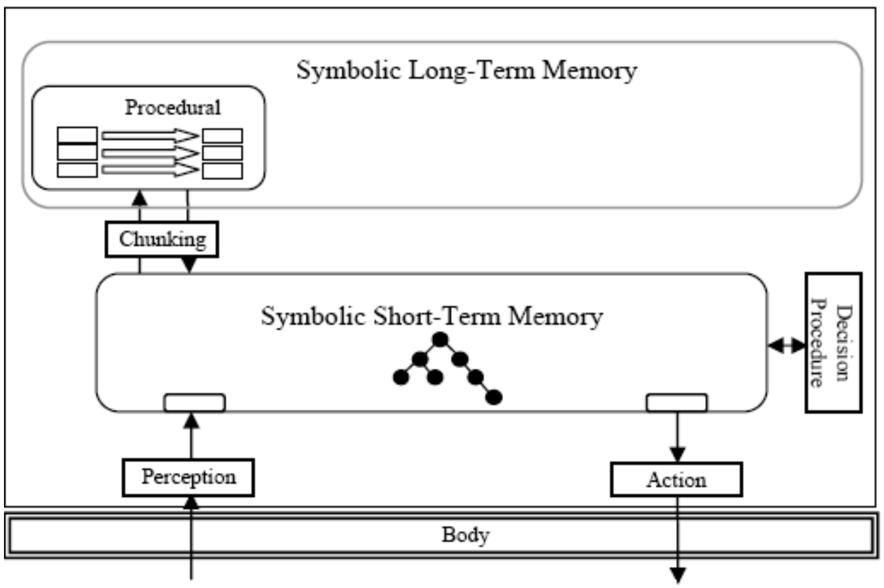

Soar 3-8

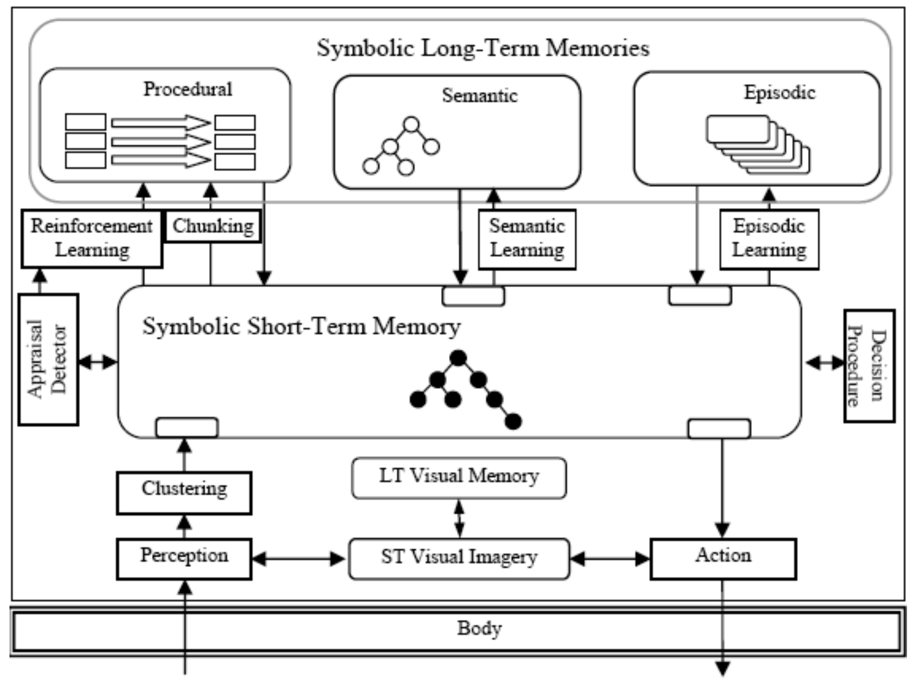

Soar 9

Towards Reconciling Uniformity and Diversity

- Accept diversity at the architectural level
- Move search for uniformity down to *implementation level*
  - *Biological Band* in humans
    - Locus of neural modeling
  - *Computational Band* in AI
    - Normally just Lisp, C, Java, etc.
      - Impacts efficiency and robustness but usually not part of theory
- Base on *graphical models*
  - Uniform approach to symbol, probability & signal processing
  - Related to neural networks but broader

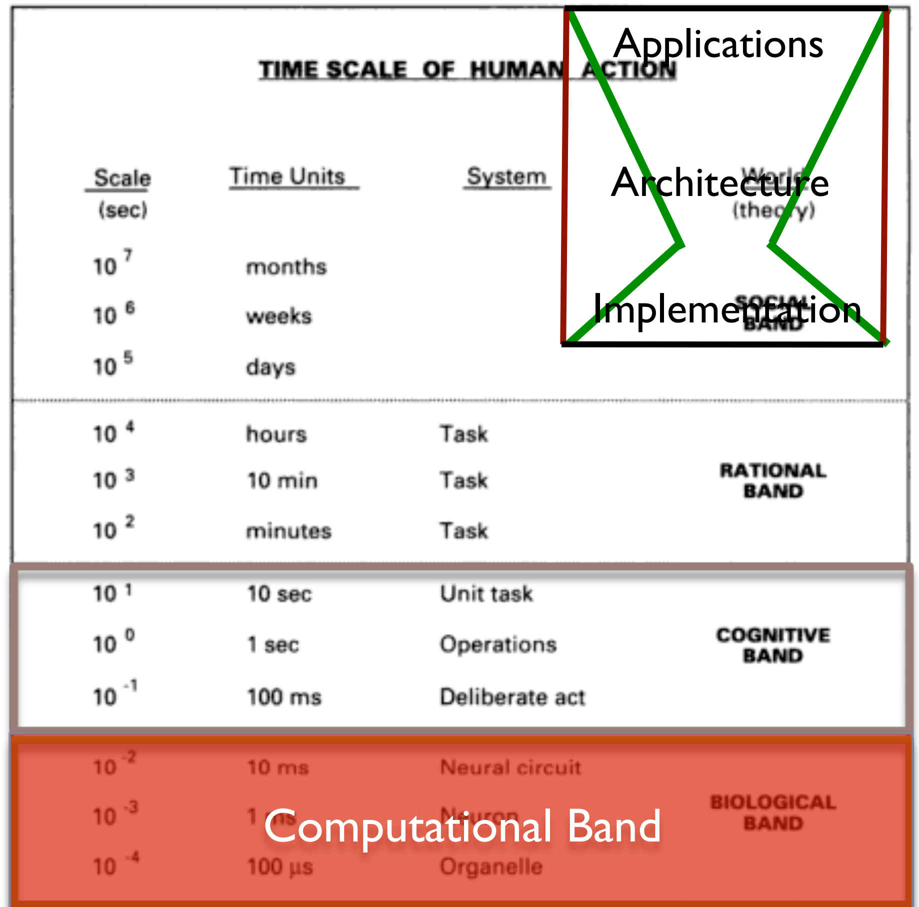

(Newell, 1990)

Graphical Models

- Efficient computation with multivariate functions: $f(u,w,x,y,z)$
  - By decomposition over partial independencies among arguments
  - For constraints, probabilities, speech, etc.
- Come in a variety of related flavors
  - **Bayesian networks**: Directed, variable nodes
    - E.g, $p(u,w,x,y,z) = p(u)p(w)p(x|u,w)p(y|x)p(z|x)$
  - **Markov networks**: Und., variable nodes & clique potentials
    - Basis for *Markov logic* and *Alchemy*
  - **Factor graphs**: Und., variable & factor nodes
    - E.g., $f(u,w,x,y,z) = f_1(u,w,x)f_2(x,y,z)f_3(z)$
- Compute marginals via variants of
  - Sum-product (message passing)
  - Monte Carlo (sampling)

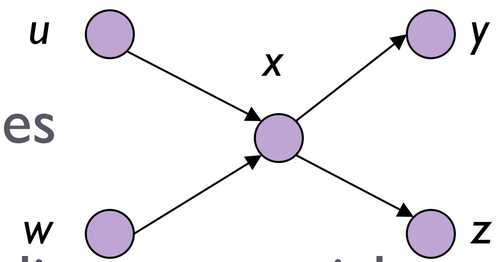

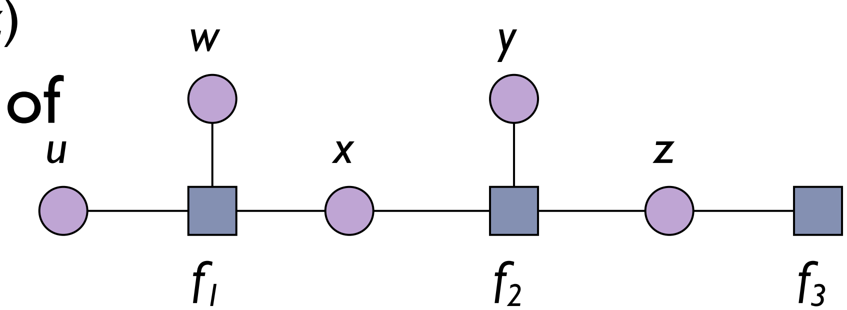

Potential for the Implementation Level

- State-of-the-art algorithms for *symbol*, *probability* and *signal* processing all derivable from the *sum-product algorithm*
  - Belief propagation in Bayesian networks
  - Forward-backward in hidden Markov models
  - Kalman filters, Viterbi algorithm, FFT, turbo decoding
  - Arc-consistency in constraint diagrams
- Potential to go beyond existing architectures to yield an effective and uniform basis for:
  - Fusing symbolic and probabilistic reasoning (mixed)
  - Unifying cognition with perception and motor control (hybrid)
  - Bridging from symbolic to neural processing
- Raises hope of a uniform implementation level that integrates broad functionality at the architecture level

Scope of Sum-Product Algorithm

Message/Variable Domain

Message/Variable Range

	Discrete	Continuous
Boolean	Symbols	
Numeric	Probability (Distribution)	Signal & Probability (Density)

- *Mixed* models combine Boolean and numeric ranges
- *Hybrid* models combine discrete and continuous domains
- *Hybrid mixed* models combine all possibilities
- *Dynamic hybrid mixed* models add a temporal dimension

Research Strategy

- Goals
  - Evaluate extent to which graphical models can provide a uniform implementation layer for existing architectures
  - Develop novel, more functional architectures
    - Enhancing and/or hybridizing existing architectures
    - Starting from scratch leveraging strengths of graphical models
    - *Improving elegance, functionality, extensibility, integrability and maintainability*
- Initial approach
  - Reimplement and enhance the Soar architecture
    - One of the longest standing and most broadly applied architectures
    - Exists in both uniform (Soar ≤8) and diverse (Soar 9) forms
  - Start from the bottom up, implementing uniform version while looking for opportunities to more uniformly incorporate Soar 9's diversity plus critical capabilities beyond all versions of Soar

Progress to Date

- *Elaboration cycle* implementation via factor graphs
  - **Production match** and firing
- *Decision cycle* implementation via Alchemy (Markov logic)
  - Elaboration phase and decision procedure
- With both also went beyond existing capability
  - *Lower complexity bound* for production match
  - *Mixed* elaboration phase with simple *semantic memory* and *trellises*
- Still preliminary, partial implementations
  - Sufficient to demonstrate initial feasibility
  - Insufficient for full evaluation of uniformity and functionality

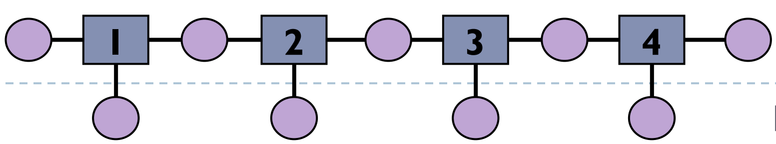

Simple Mapping of Production Match onto Factor Graphs

P1: Inherit Color

C1: (<v0> ^type <v1>)

C2: (<v1> ^color <v2>)

-->

A1: (<v0> ^color <v2>)

Model as a Boolean function:

$P_1(v_0,v_1,v_2) = C_1(v_0,v_1)C_2(v_1,v_2)A_1(v_0,v_2)$

WM is 3D Boolean array (*obj* x *att* x *val*)
- 1 when triple in WM
- 0 otherwise

Messages are Boolean vectors
- 1 when variable value possible
- 0 when variable value ruled out

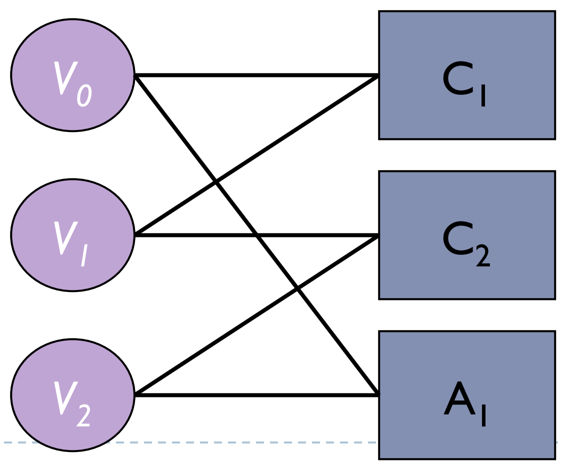

Issues:
1. Hiding WM and constant tests
2. Confusing binding combinations
3. Ignoring constant conditions
4. Array sizes and combinatorics

Handling Array Size and Combinatorics

- Solution to other problems yields a new match algorithm that reduces cost to exponential in the graph's *treewidth* rather than number of conditions
  - But arrays can still be very large (in both dimension and rank)
- Exploit uniformity across arrays (e.g., WM almost all 0)
  - Use an N dimensional generalization of quad/octrees (*exptrees*)
    - If entire space has one value, assign it as the value
    - Otherwise partition into 2^N regions at next level, and recur
- Views WM (and messages) as piece-wise constant functions
  - Currently extending to piecewise linear functions
    - Natural compact representation for probabilities, signals, images, etc.
- Could also consider more adaptive partitioning/clustering

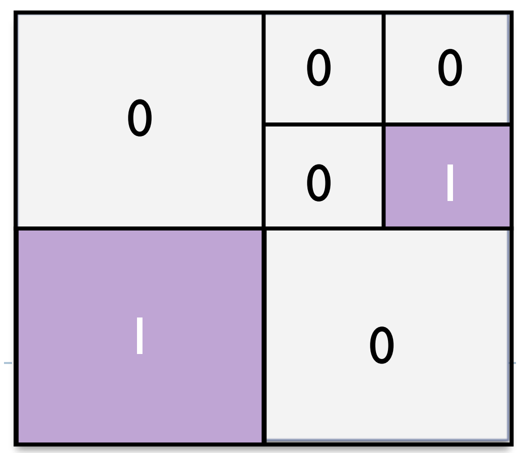

Significant efficiency gains on test cases

1. Exceeded Heap Space ➔ *132 sec*
2. *1.7 sec* ➔ *.25 sec* (factor of 7)

Conclusion

- Despite increasing diversity within cognitive architectures, there is still hope for uniformity in the cognitive hierarchy
  - The narrow waist in the cognitive hourglass
- Uniformity at the implementation level should facilitate:
  - Exploring, understanding and improving existing architectures
  - Developing novel architectures with increased elegance and functionality
- Factor graphs (and graphical models in general) promise to yield a wide diversity of capabilities in a uniform manner
  - Mixed, hybrid and neural processing
  - Procedural and declarative (semantic/episodic) memories
- Initial progress made, but much more still to be done
  - Complete reimplementation and extension of Soar
  - Reexamine other cognitive architectures and hybrids among them
  - Experiment with radically new architectures enabled by graphical models