Scaling up from one model to many: Individual Cognition and the Social Network

David Reitter
Carnegie Mellon University

Cognitive Teams

- Cognitive models are mostly observed in isolation.
- *Humans,* in the wild, ...

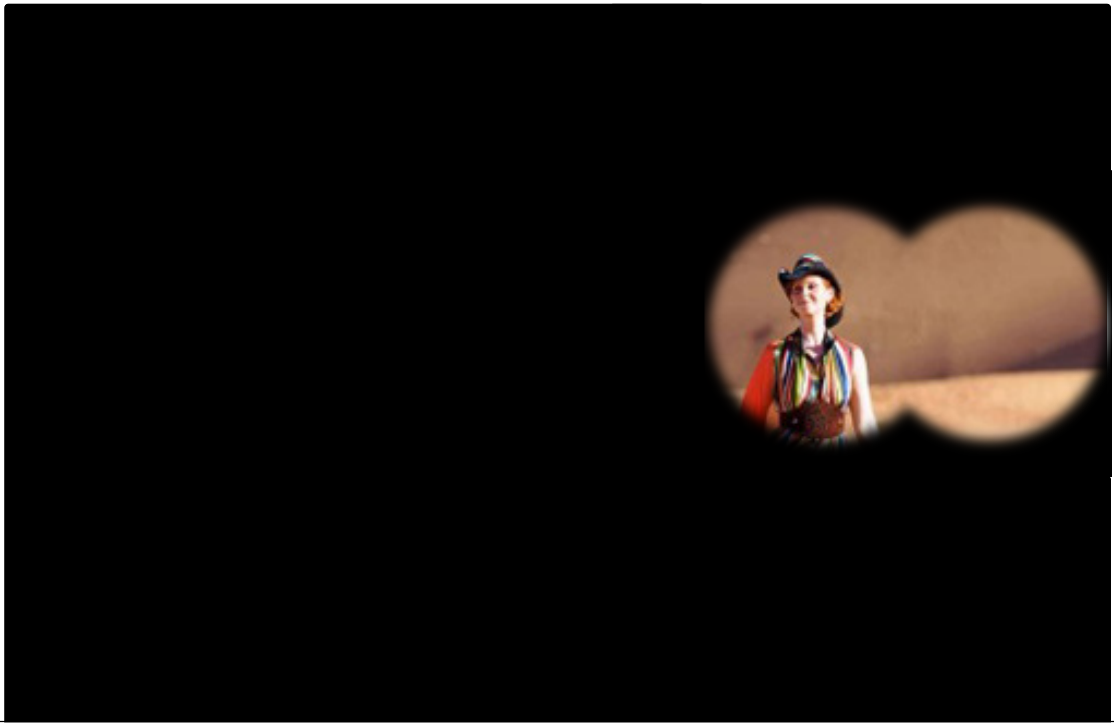

Cognitive Teams

- Cognitive models are mostly observed in isolation.
- *Humans,* in the wild, rarely occur in isolation.

David Reitter, Carnegie Mellon

3

Cognitive Teams

- Cognitive models are mostly observed in isolation.
- *Humans,* in the wild, rarely occur in isolation.
  - Innate social behavior has evolved over many species, as has cognition.
  - Is there co-adaptation?

David Reitter, Carnegie Mellon

Cognitive Teams

- Cognitive models are mostly observed in isolation.
- *Humans,* in the wild, rarely occur in isolation.
- Need for team/group modeling in training applications, basic questions of cultural evolution, behavior of organizations, etc.

Individual and Team Small Arms Training Systems
Instructional Strategies and Team Modeling
Team Training and Performance Measurement
Training Technology for Distributed and Joint Systems
Medical Team Performance and Simulation Training
NAWC Training Systems Division BAA Jan 2011

David Reitter, Carnegie Mellon

Some questions

- Cognition and social structure: co-evolution?
- Communication
  - How do communication policies interact with network structure and individual memory? BRIMS 2011
  - How do individuals use metacognitive awareness to regulate communication (vs. task execution)?
    - Robots aren't very good at it
  - How do individuals in dyads/groups/communities agree on common interaction and communication schemes? Pickering&Garrod 2004
- External/Internal memory
  - How does individual memory differ from memory externalized and distributed into transient and permanent networks of people, agents, or models? Merlin 1991

Social structure and its effects

- What is the role of network structure for a team, and position within the network for the individual?
- In a network, only network neighbors play the naming game
  - Small-World network
    (low path length, high clustering coefficient, assortatively mixed by degree)
  - Grid (torus)
  - Random Graph
  - Trees

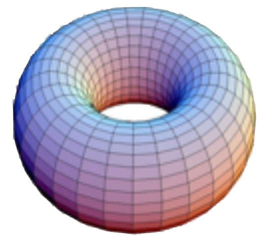

- ACT-R simulation
  - Controlled: mean degree (except trees), number of nodes
  - Here: 85 nodes, mean deg. 5., 50 rep. per condition

Pictionary Game

Reitter&Lebiere 2011, J-CSR

- Two-player *naming game*; one player is to convey a given concept by way of a drawing
- Over time, participants standardize drawings *Domain Terminology*
- Player pairings may change according to a network graph

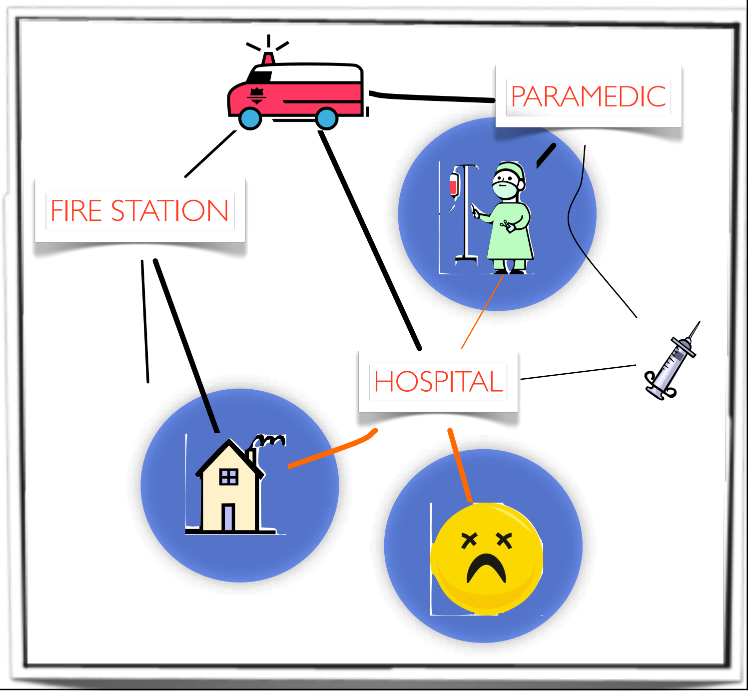

David Reitter, Carnegie Mellon

Pictionary Accuracy

Reitter&Lebiere 2010, CMCL

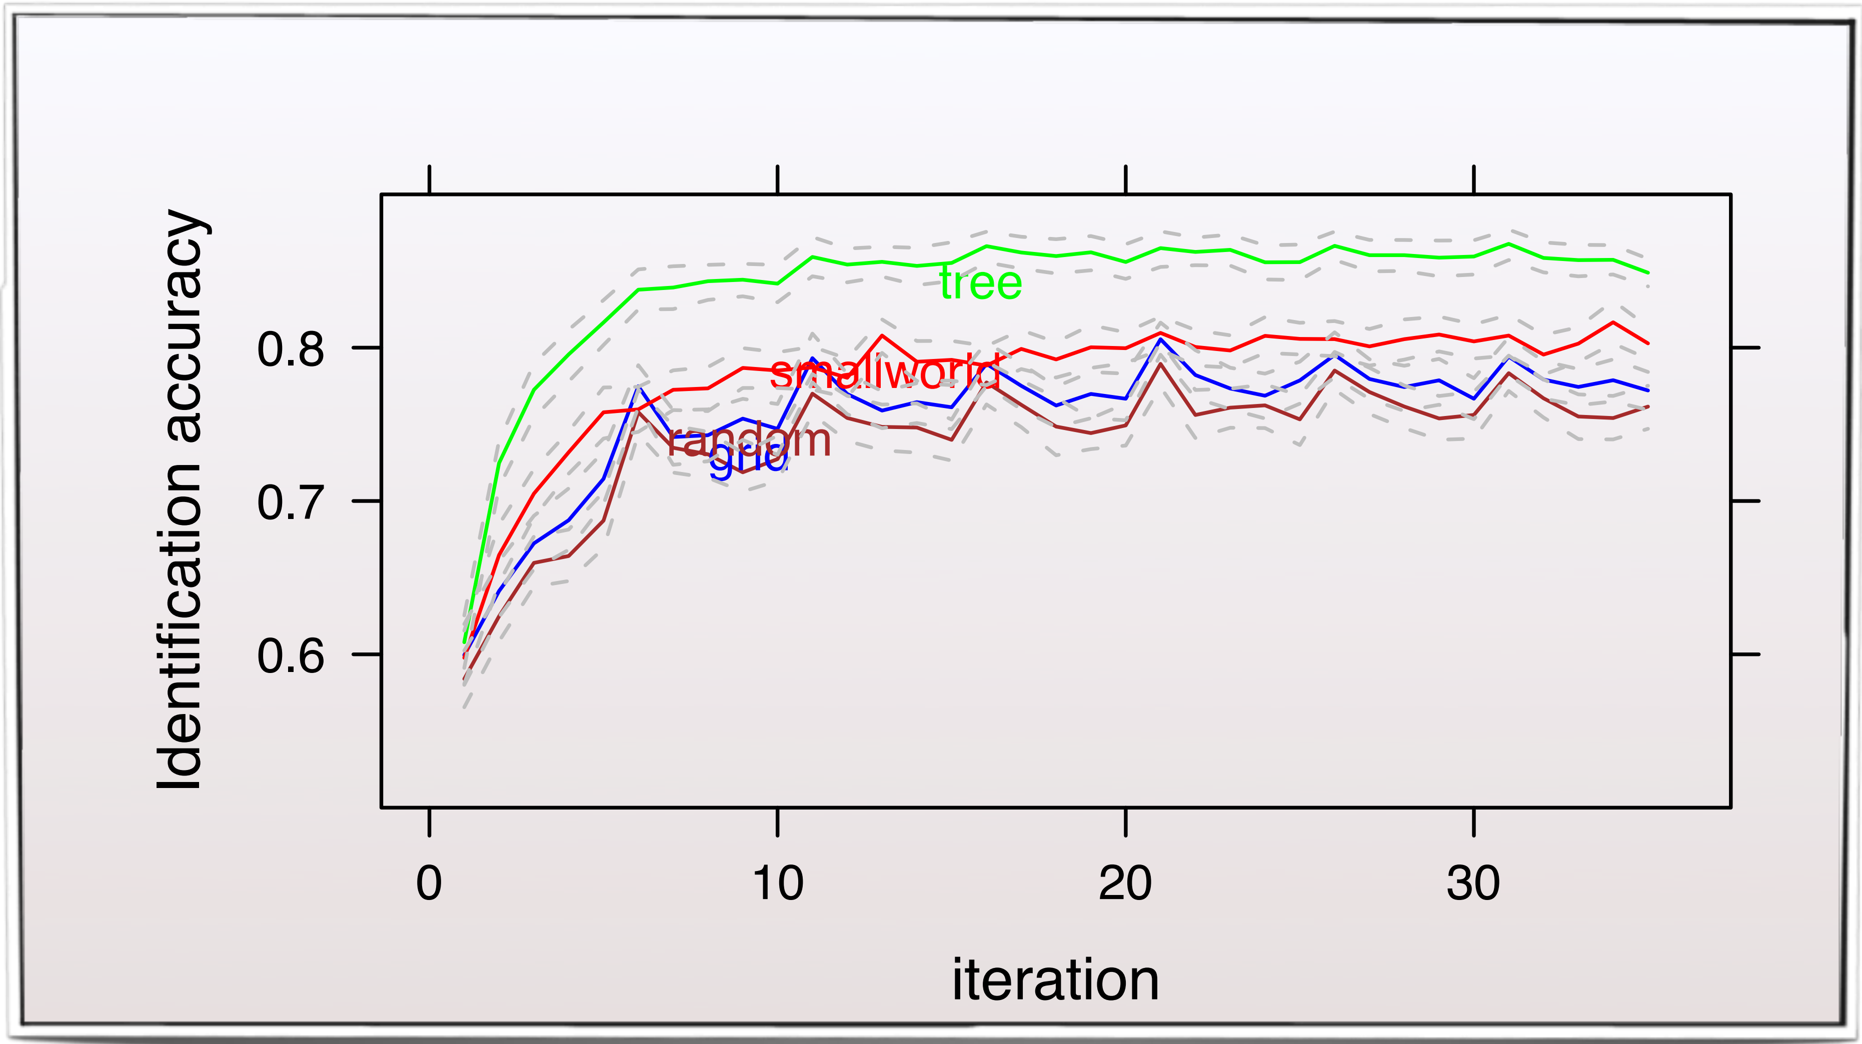

David Reitter, Carnegie Mellon

9

Pictionary Accuracy

Reitter&Lebiere 2010, CMCL

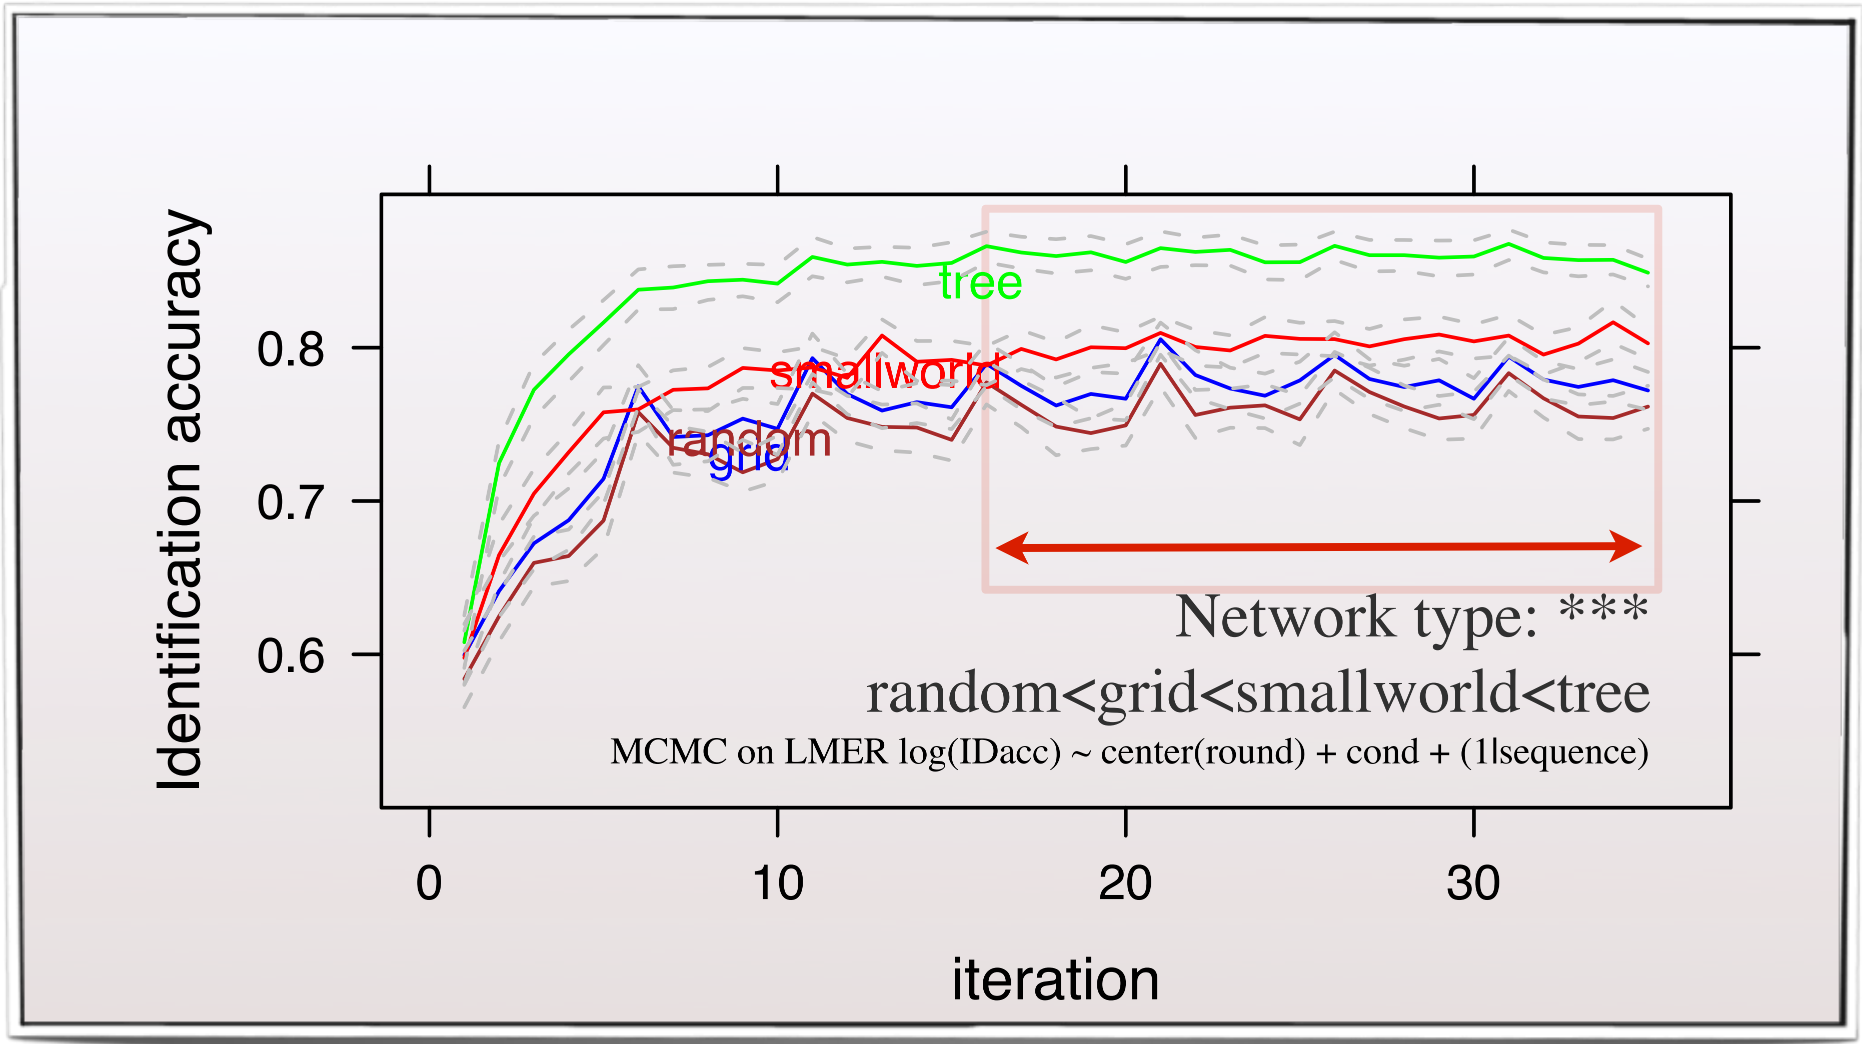

David Reitter, Carnegie Mellon

Overview

- The need for models of interaction
- How do we get there?
  - Scalability in cognitive modeling
  - Scalable experimental paradigms

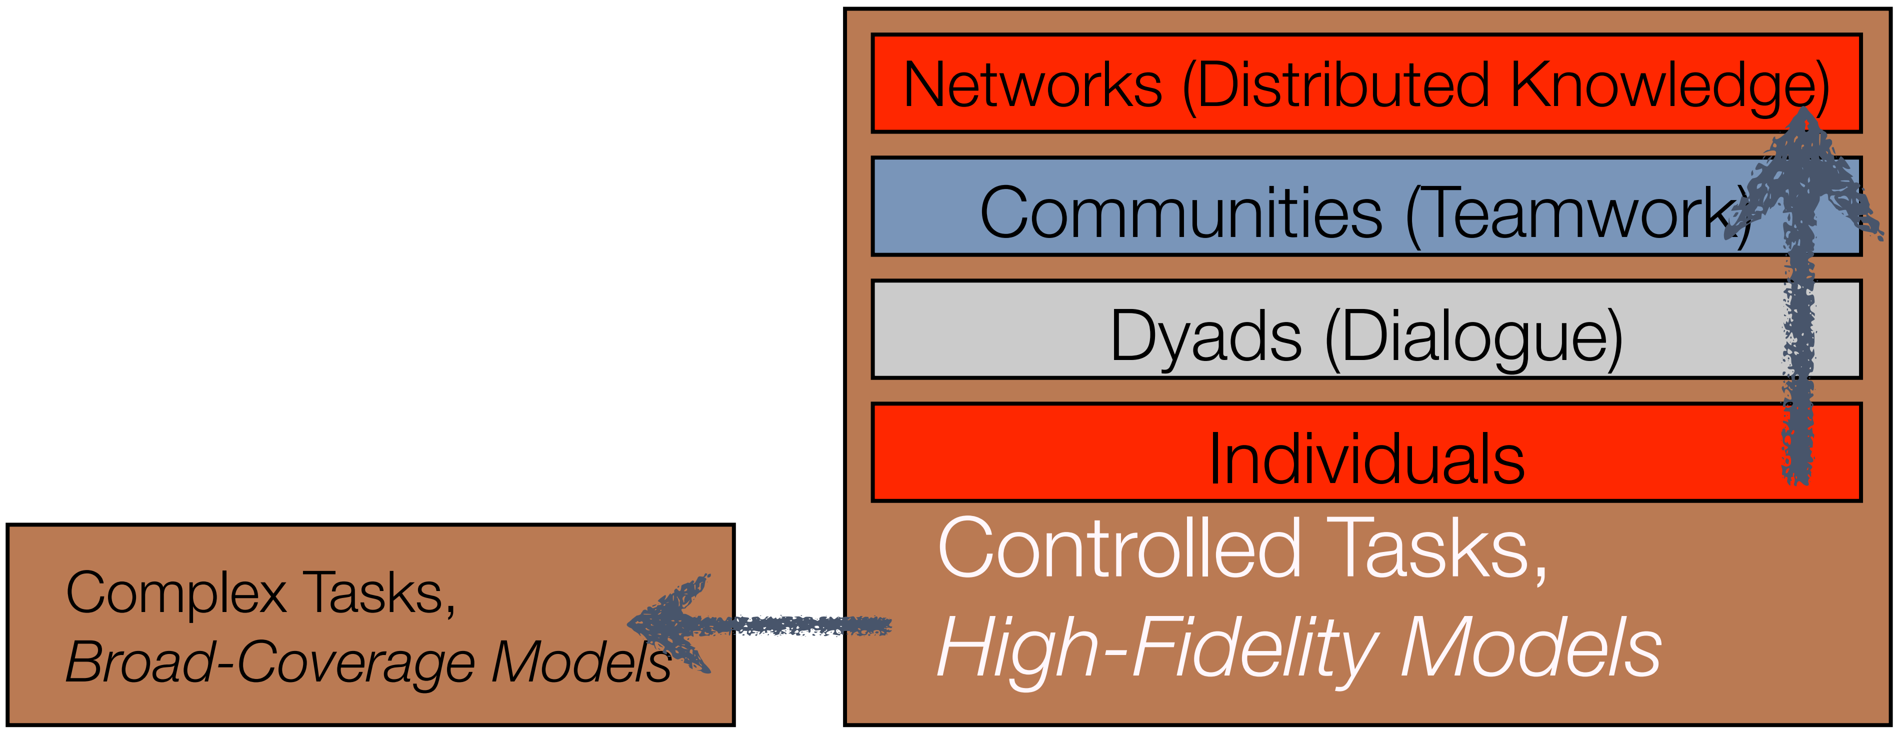
Networks (Distributed Knowledge)
Communities (Teamwork)
Dyads (Dialogue)
Individuals
Controlled Tasks,
High-Fidelity Models
Complex Tasks,
Broad-Coverage Models

From experiments to models

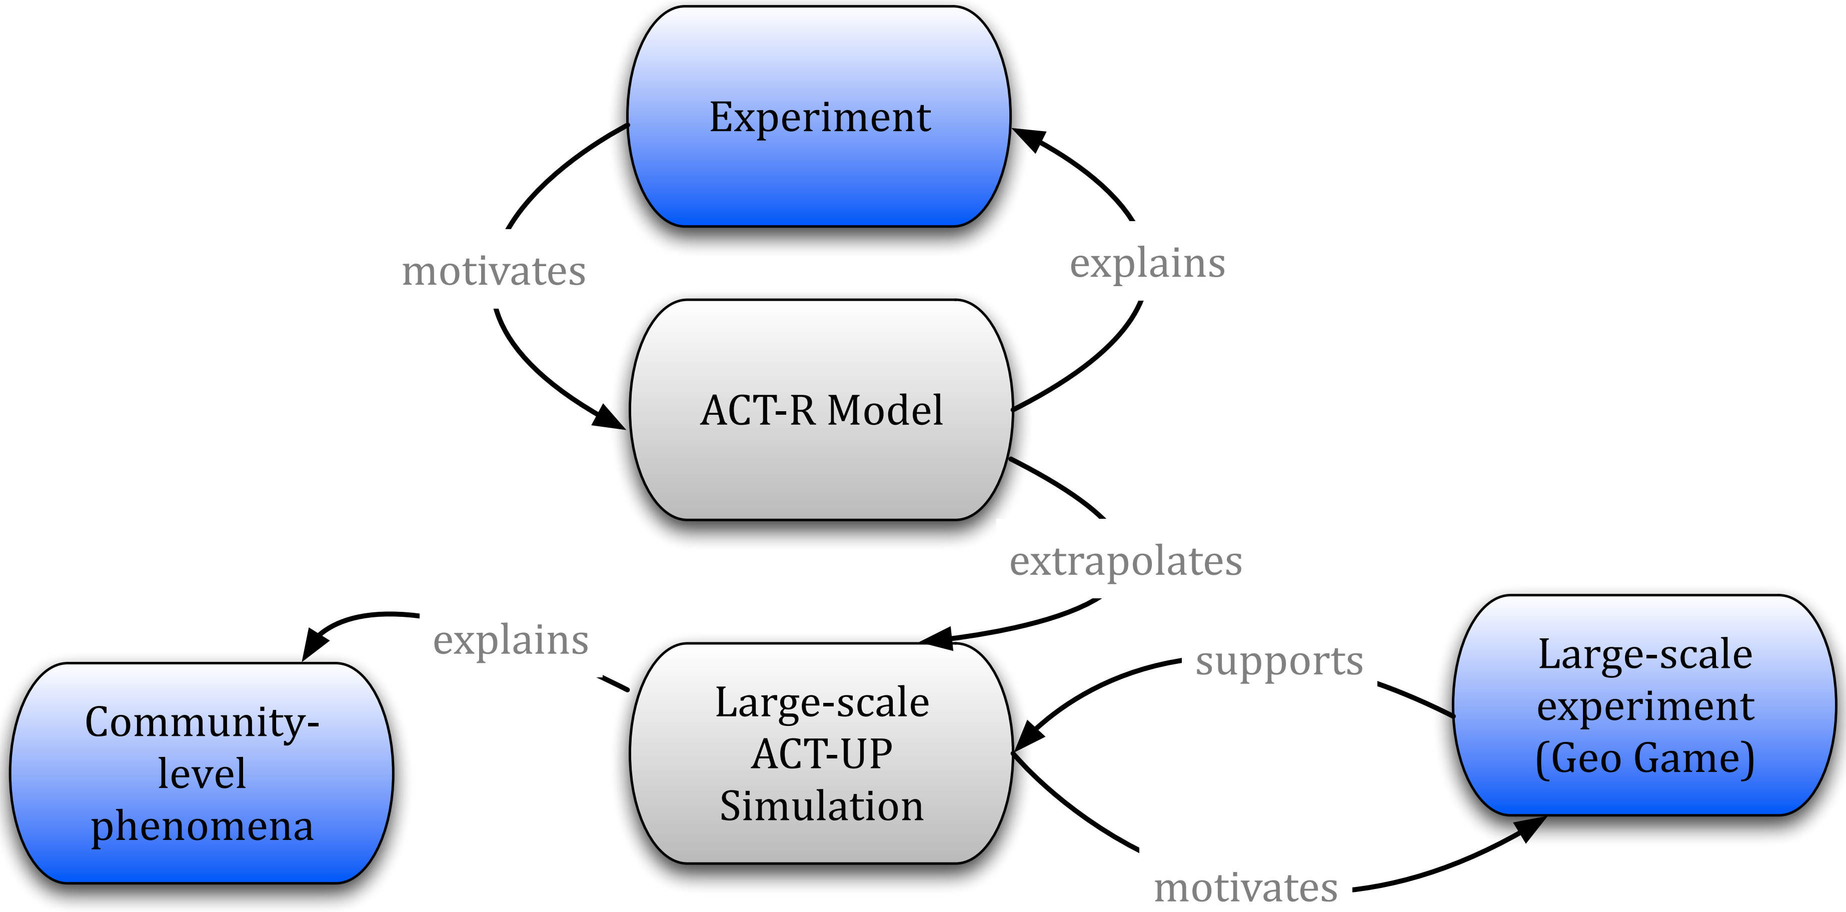

David Reitter, Carnegie Mellon

How do we achieve scalability?

- Higher-level representations
  - HLSR - “An abstract language for cognitive modeling” (Jones et al)
    - compiled to ACT-R, Soar
    - Relations and Transforms

```
transform MoveDiskToPeg(d isa Disk,p isa Peg) (
  # Consider if a goal to put disk d on peg p
  consider-if ( goal<DiskOnPeg>(d, p) )
  body ( DiskClearToMoveToPeg(d, p)
         DiskIsOnPeg(d, other-peg)
         consider-instead(
           DiskIsOnPeg(d, other-peg),
           new<DiskIsOnPeg>(d, p)))
)

# If DiskCleartoMoveToPeg or DiskIsOnPeg
#  fails, an impasse is generated.  A query
#  can retrieve a goal to resolve this impasse,
#  where "trans" binds to the transform instance
impasse<MoveDiskToPeg>(trans)
```

Figure 2: An example of a transform in HLSR

David Reitter, Carnegie Mellon

How do we achieve scalability?

- Higher-level representations
  - HLSR - “An abstract language for cognitive modeling” (Jones et al)
  - Herbal “ A high-level language and development environment for cognitive models” (Cohen et al)
    - Ontology, states, operators
    - Compiles to Soar and different expertise levels of ACT-R

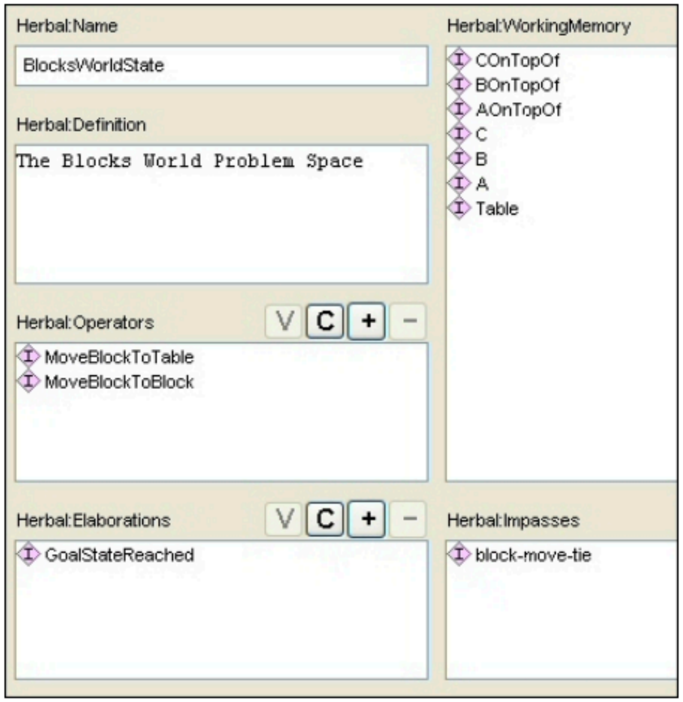

Figure 5.2.1 Top State for the Blocks World Model.

David Reitter, Carnegie Mellon

How do we achieve scalability?

- Higher-level representations
  - HLSR - “An abstract language for cognitive modeling” (Jones et al)
  - Herbal “ A high-level language and development environment for cognitive models” (Cohen et al)
    - Ontology, states, operators
    - Compiles to Soar and different expertise levels of ACT-R
- Re-usability
- Efficiency and scalability (DM)
  - cf., Douglass et al. (2009, 2010) RML1

David Reitter, Carnegie Mellon

ACT-R

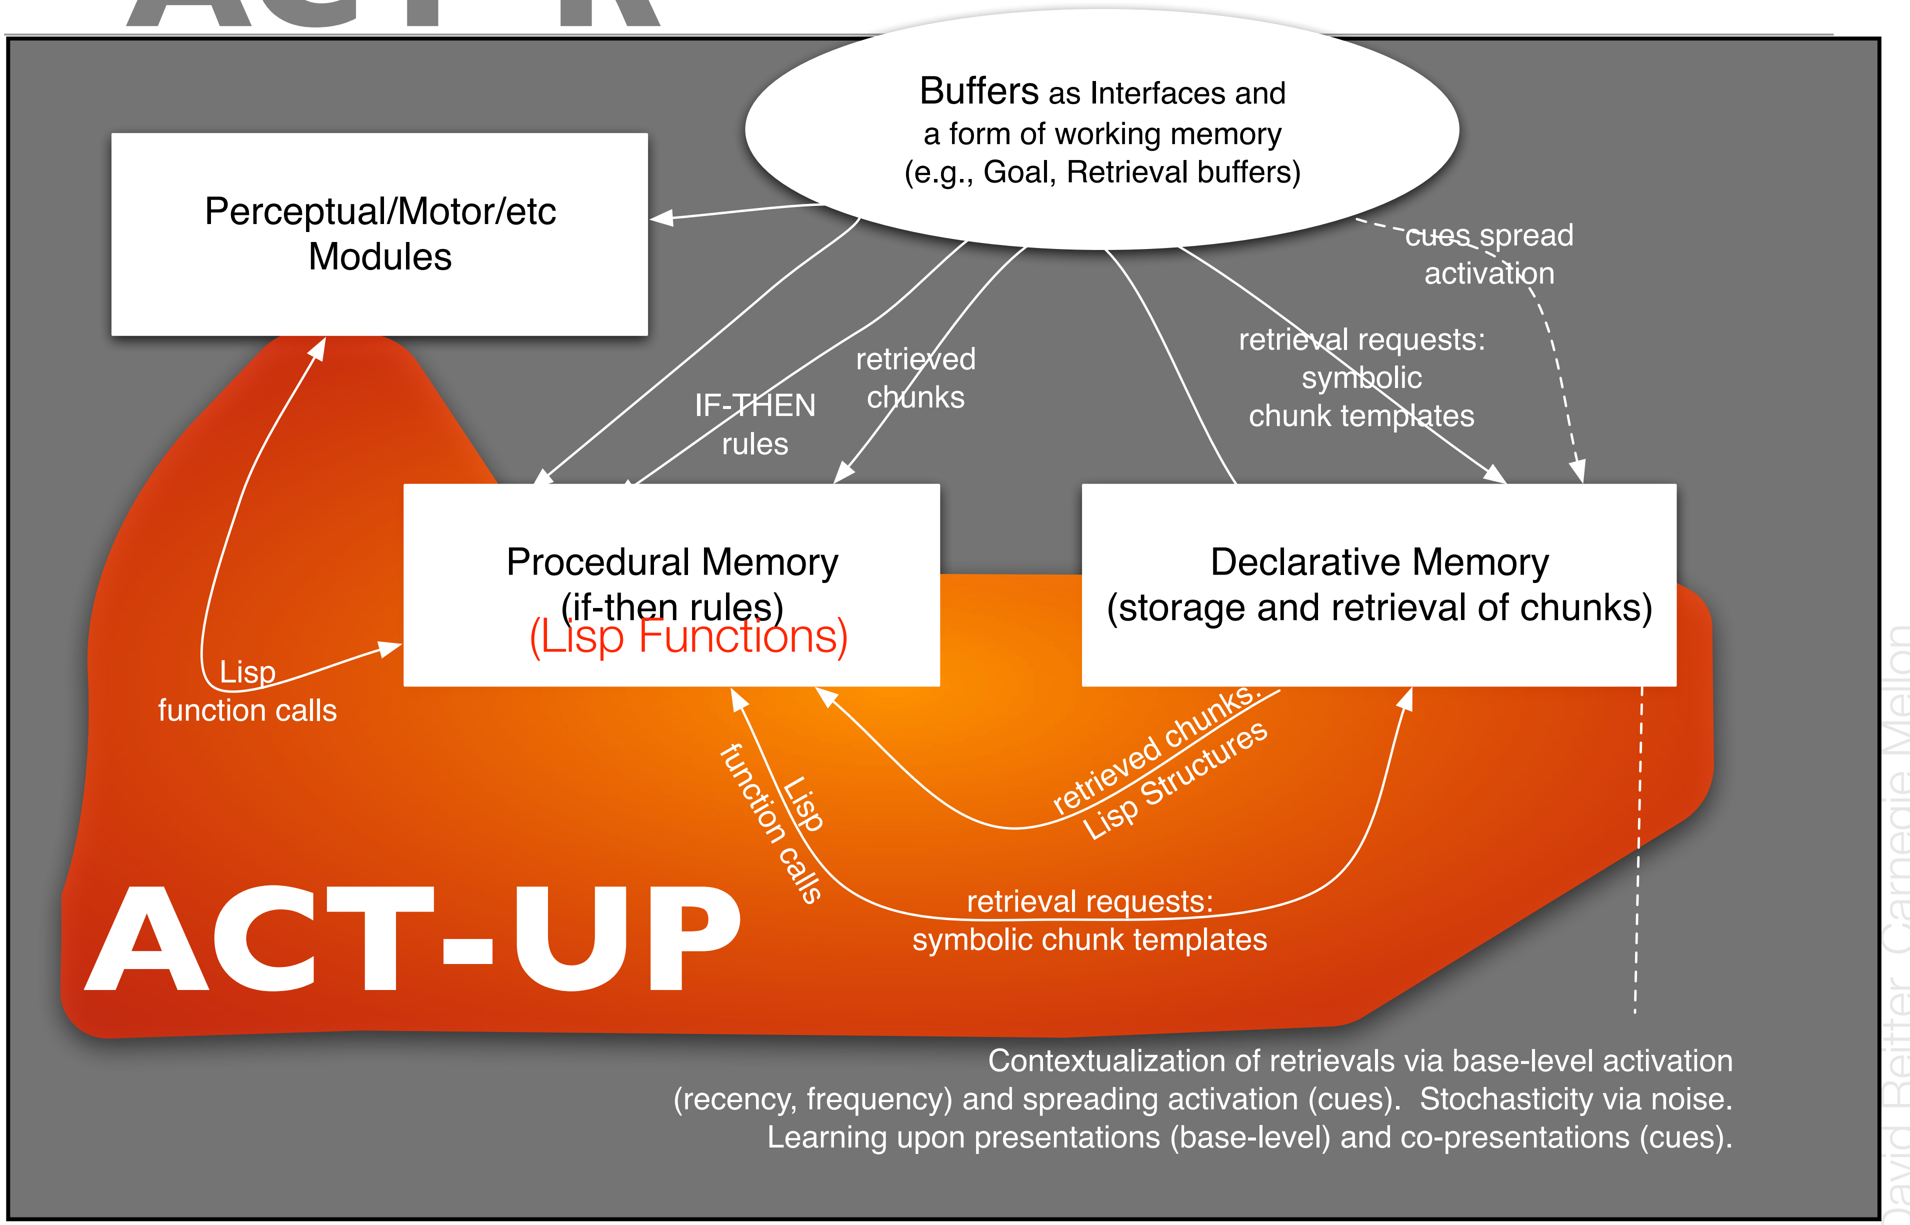
Buffers as Interfaces and
a form of working memory
(e.g., Goal, Retrieval buffers)
Perceptual/Motor/etc
Modules
cues spread
activation
retrieved
chunks
IF-THEN
rules
retrieval requests:
symbolic
chunk templates
Procedural Memory
(if-then rules)
(Lisp Functions)
Declarative Memory
(storage and retrieval of chunks)
Lisp
function calls
Lisp
function calls
retrieved chunks:
Lisp Structures
retrieval requests:
symbolic chunk templates
ACT-UP
Contextualization of retrievals via base-level activation
(recency, frequency) and spreading activation (cues). Stochasticity via noise.
Learning upon presentations (base-level) and co-presentations (cues).
David Reitter, Carnegie Mellon

ACT-UP is not ACT-R 6...

- ACT-UP Interface is synchronous
  - Serial execution
  - Deterministic strategies defined as programs
  - Compare: ACT-Simple (Salvucci&Lee 2003)
- Parallelism (e.g., perceptual/motor modules) possible
  - Threaded Cognition
- Non-deterministic rule choice is possible
  - Reinforcement-learning as in ACT-R 6
- Optional typing (with inheritance)

Two “stories”

- ACT-UP is an implementation of ACT-R (with some fundamental changes)
  - 1. Make hard claims in model cores, underspecify and estimate the rest
  - 2. Rapidly prototype, reuse and scale up models

David Reitter, Carnegie Mellon

ACT-UP: Bias-free estimates

- ACT-R models are end-to-end models: some are overly specific
- ACT-UP treats models as hard claims
  - Evaluate each specified component against data
  - Underspecify the rest and fit free parameters (e.g., timing)

```
;; The Model
;;;; define a production
(defproc check-factual-sentence (person location target term)
  (let ((cfd
         (retrieve-chunk
          ;; hard constraints:
          (append (list :chunk-type 'comprehendfact)
                  (if (eq term 'person)
                      (list :arg1 person)
                      (list :arg2 location)))
          :cues (list person location))))
    (pass-time *model-time-parameter*)
    (if (and cfd ; if not retrieved, answer would be NO
             (equal person (comprehendfact-arg1 cfd))
             (equal location (comprehendfact-arg2 cfd)))
        ;; answer "K" (yes)
        (if target t nil) ; YES
        ;; answer "d" (no)
        (if target nil t))))
```

Efficiency

- Sentence production (syntactic priming) model
  - 30 productions in ACT-R, 720 lines of code
  - 82 lines of code in ACT-UP (3 work-days)
  - ACT-R 6: 14 sentences/second
    - "tree" optimization for production selection enabled
  - ACT-UP: 380 sentences/second

Scalability

- Language evolution model
  - Simulates domain vocabulary emergence (ICCM 2009, JCSR 1010)
  - 40 production rules in ACT-R (could not prototype)
  - 8 participants interacting in communities
- In larger community networks: 1000 agents, 84M interactions (about 1 minute sim. time each), 37 CPU hours

David Reitter, Carnegie Mellon

ACT-UP availability

- http://act-up.psy.cmu.edu/
  - SBCL, OpenMCL/CCL, LispWorks, (Allegro)

22

Experimental and modeling paradigms

- Iterated learning
  - Kirby et al., Griffiths
  - Individual re-acquires its own output (e.g., a word) or an isomorph
  - Applies individual biases and learns again

- Group-level phenomena indicating individual cognitive biases?
  - Griffiths&Tenenbaum 2006: Bayesian priors for everyday distributions found in groups
- One-on-one simple games
  - Fay, Garrod, et al.: Maze Game, Pictionary Task
  - Networks: Kearns, Judd, et al.
    - e.g., Coloring Task

David Reitter, Carnegie Mellon

The Geo Game

- Foraging task
  - Player is to find target items (selected by system)
  - Items are spread out and hidden on a map

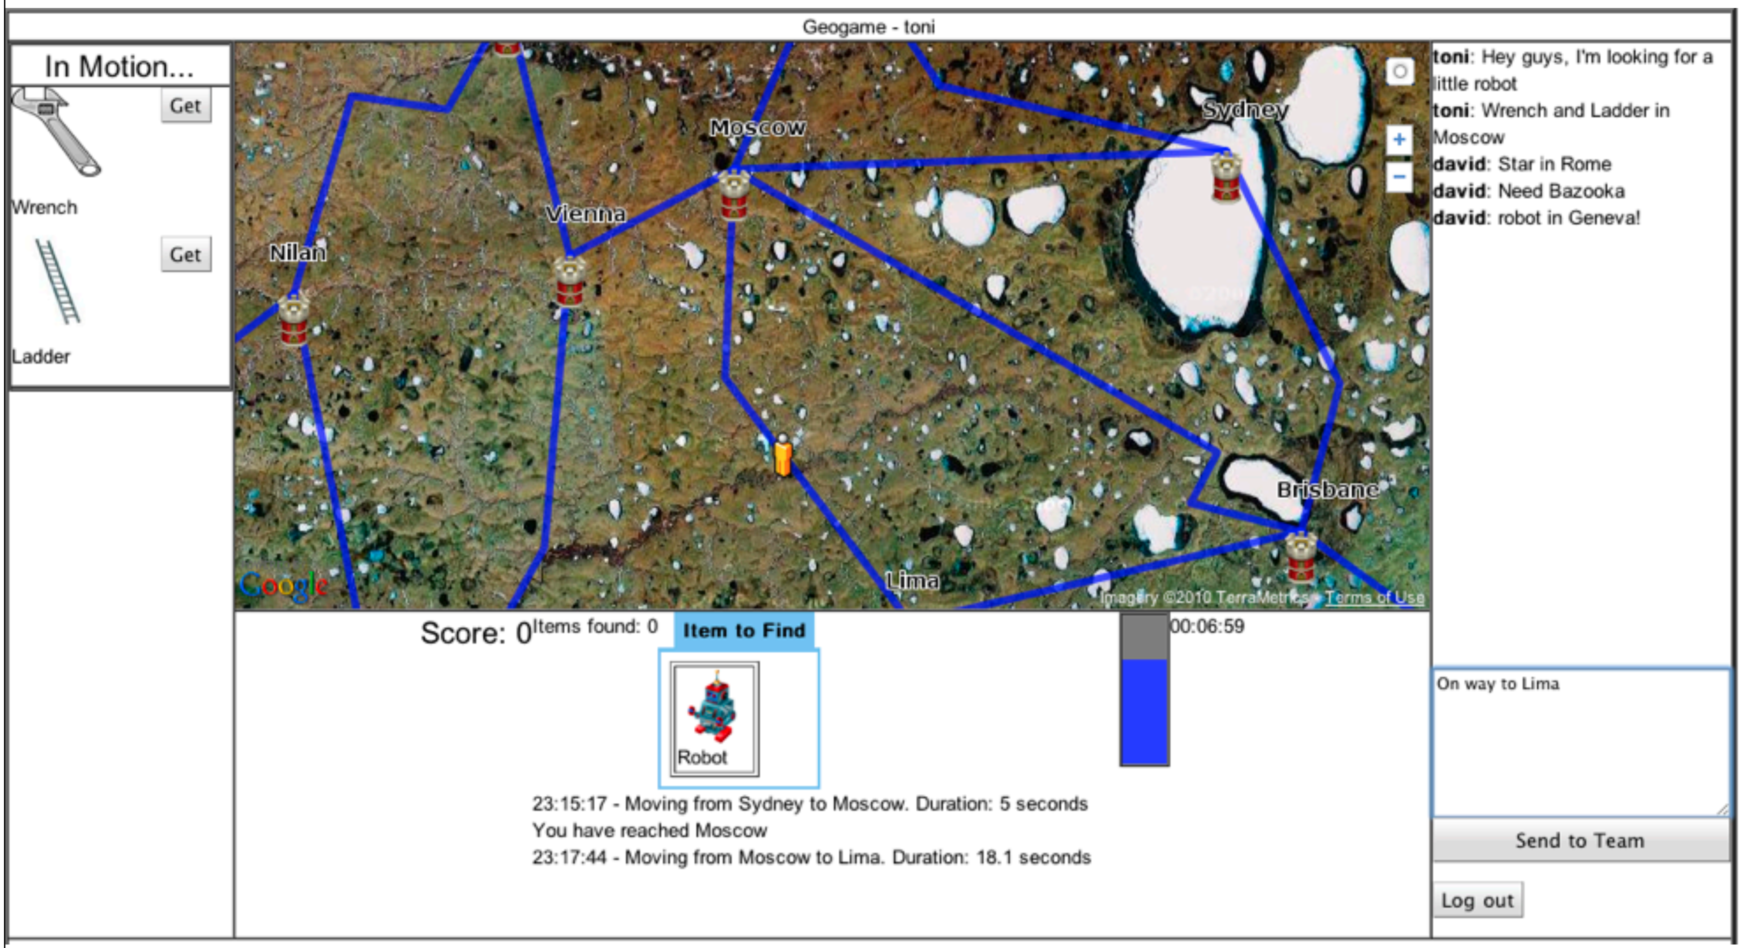

24

The Geo Game

- Foraging task
  - Player is to find target items (selected by system)
  - Items are spread out and hidden on a map

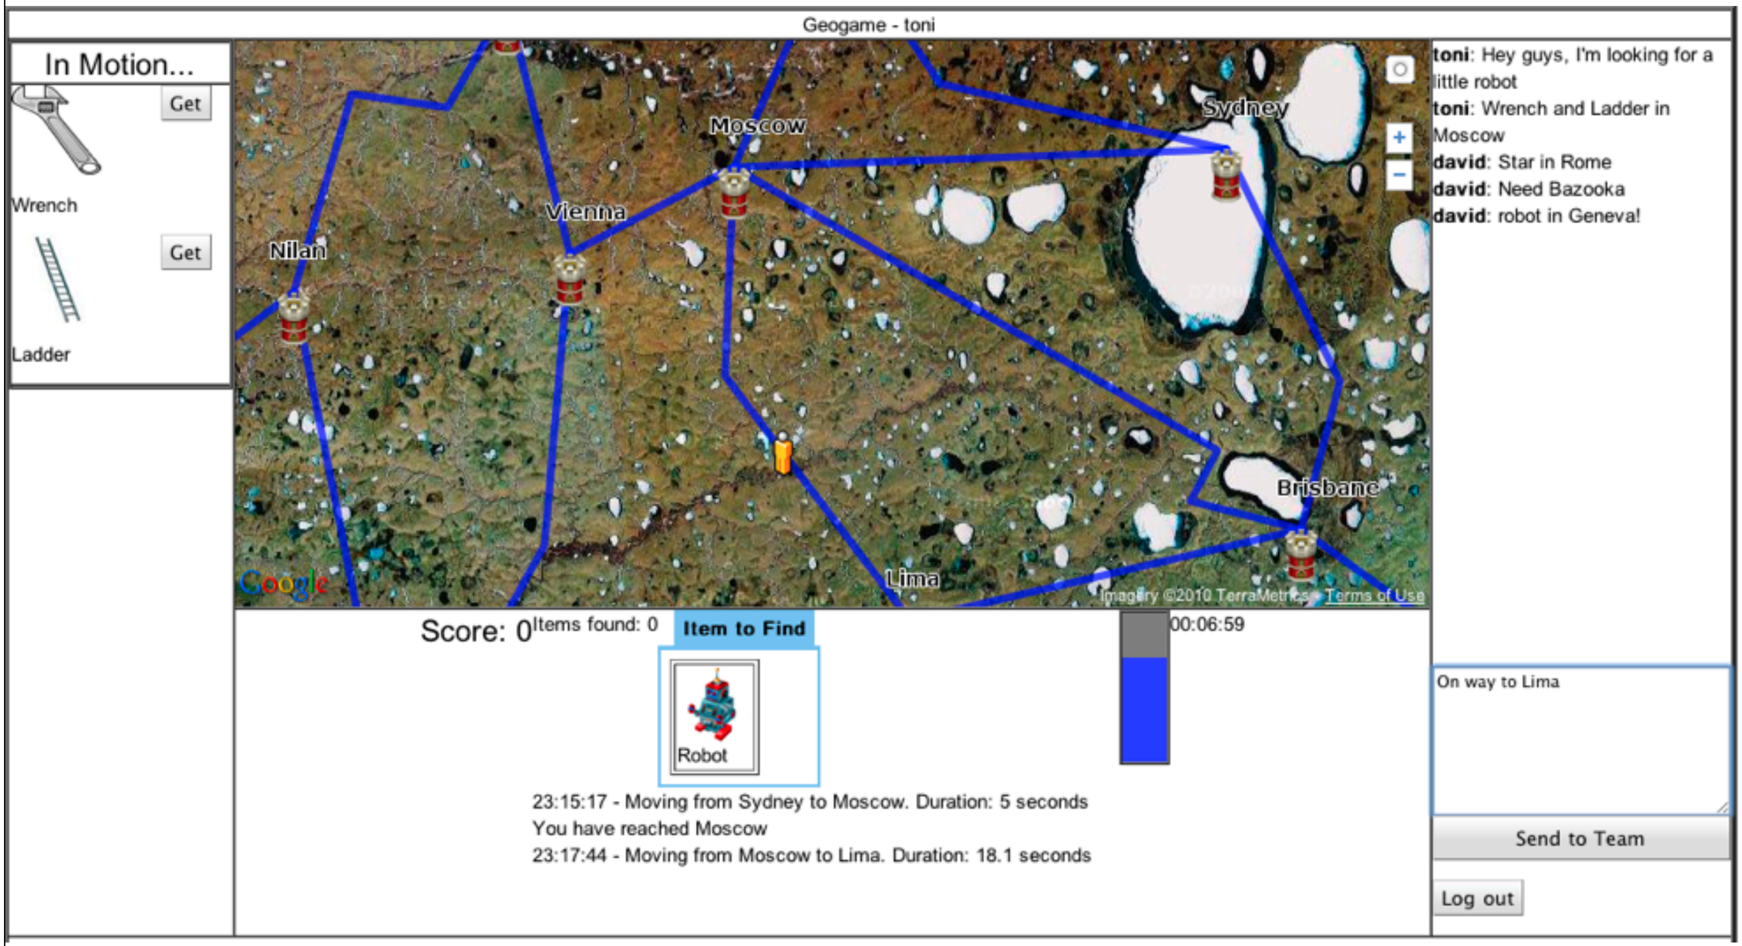

The Geo Game

with: C. Lebiere, J. Vinokurov, K. Sycara, A. Juarez, et al.

- Foraging task
  - Player is to find target items (selected by system)
  - Items are spread out and hidden on a map

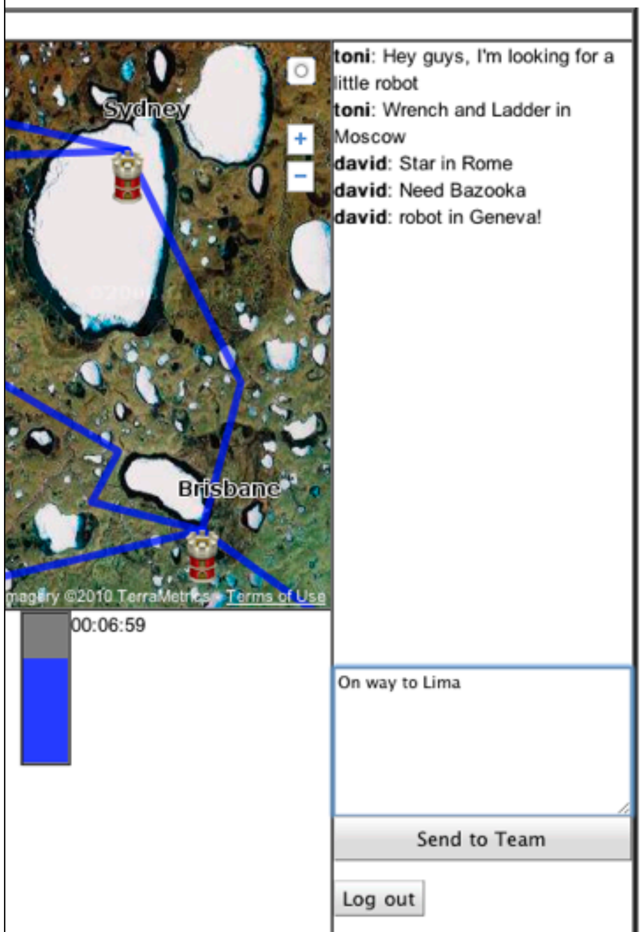

- Communication
  - Players are to communicate with others
    - Exchange of item-location information
  - Communication graph
    - Every player has fixed set of network neighbors to whom they broadcast (*friends*)

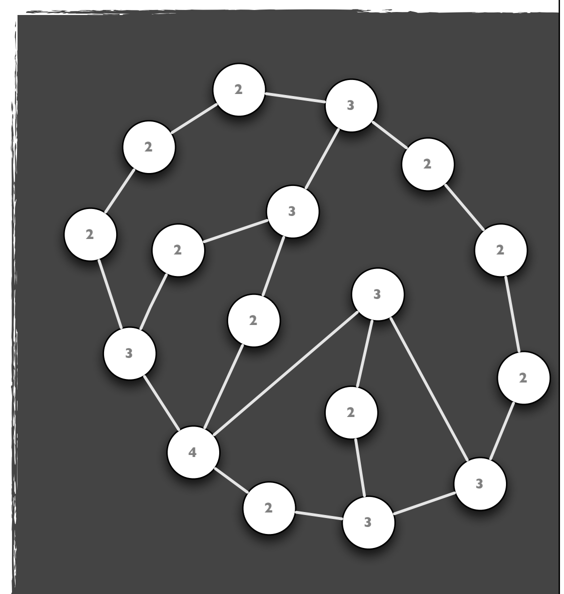

The Geo Game

with: C. Lebiere, J. Vinokurov, K. Sycara, A. Juarez, et al.

- Foraging task
  - Player is to find target items (selected by system)
  - Items are spread out and hidden on a map

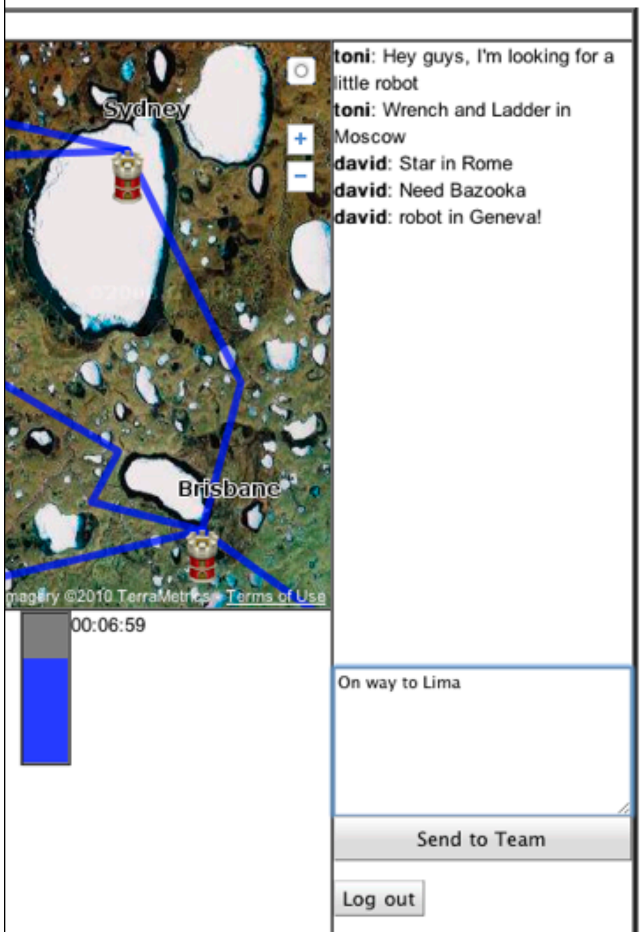

- Communication
  - Players are to communicate with others
    - Exchange of item-location information
  - Communication graph
    - Every player has fixed set of network neighbors to whom they broadcast (*friends*)

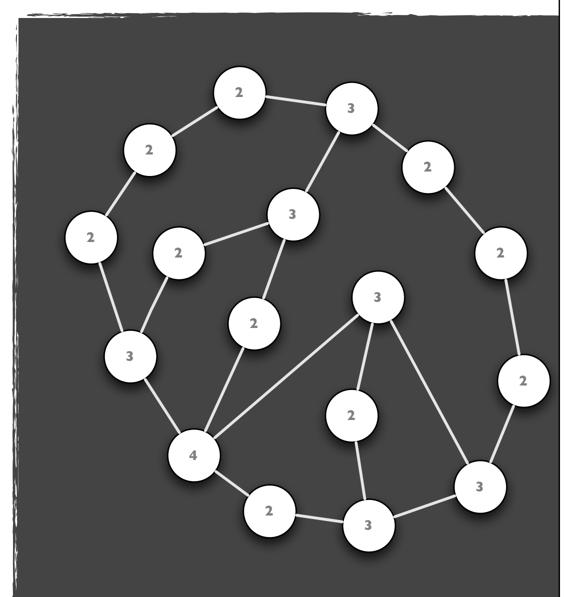

Cognitive Model

Reitter&Lebiere 2011, CogSci

- Where is the Cake?
  - Either needed by node itself
  - Or requested to answer another node's request
- Communicate or Work?
  - round-robin message reading and navigating
  - stop at city attracts attention
- Declarative memory: <cake, Vienna>
- ACT-R: frequency, recency

$$B_i = \log \sum_{j=1}^{n} t_j^{-d} + \beta_i$$

Anderson

- *Used to predict effect of randomization in task and in individuals: what to control in the experiment?*
  - *Distributions of parameters?*

"The dirty little secret"

David Reitter, Carnegie Mellon

Experiment

- Communication policies
  - dump: pass on all available information
    - item locations in current city
    - item locations reported by others
  - target: request information and only pass on useful facts

Result

- If human networks serve as useful information filters, then a targeted communication policy should
  - increase the overall performance (points)
  - increase efficiency (points per message)

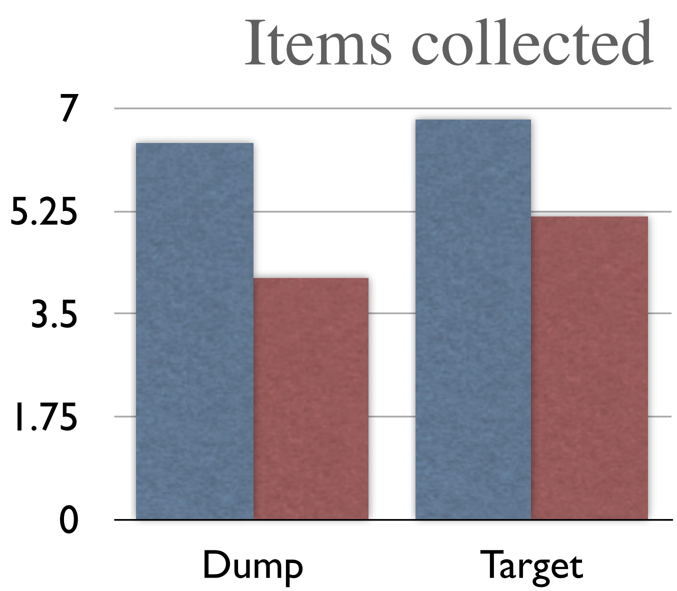

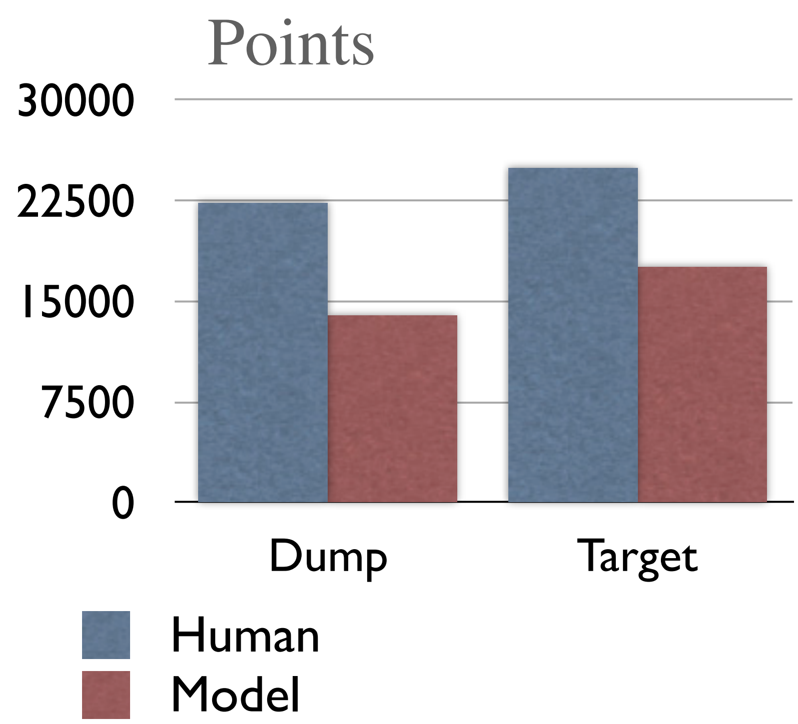

What if...

Reitter&Lebiere 2011, CogSci

- With the "social task", does memory decay help?
- Decay
  - bll=0.0 (Maximal memory, no recency effect)
  - bll=0.5 (realistic level, ACT-R)
  - bll=3.0 (Minimal memory, strong recency effect)
- Points: mean of 17 players (30 minutes)

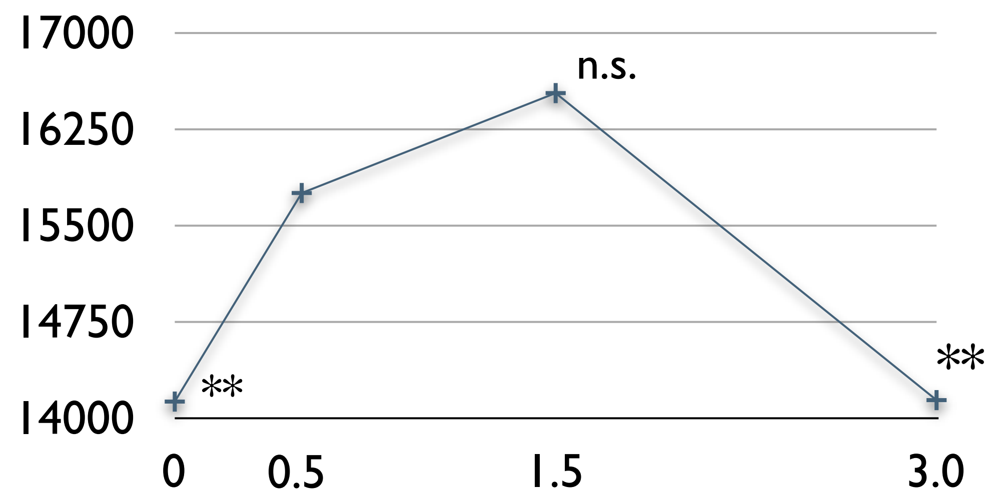

ACT-R Base-Level Learning Parameter
Default: 0.5
Significance tests: compared to default

- see also: Parameter exploration in models (Gluck et al., 2010) / Rick Moore

David Reitter, Carnegie Mellon

Conclusions

- Models of teams and larger communities
  - practical needs in e.g., training systems
  - basic questions of socio-cognitive interaction
- How to do it?
  - Cognitive modeling paradigms
    - HSLR, Herbal
    - ACT-UP
    - Geo Game model
  - Empirical designs
    - iterated learning
    - simple and complex games

David Reitter, Carnegie Mellon