Stable Configurations for Cooperation in Complex Agents with Minimal Communication

or

Where Should Complexity Go?

Joanna J. Bryson

Harvard Primate Cognitive Neuroscience Laboratory

Franklin W. Olin College of Engineering

jbryson@wjh.harvard.edu

http://www.ai.mit.edu/∼joanna

Synopsis

- General Concept: People worry too much about communication in MAS and not enough about the agents themselves.
- Evidence: Animals are complex but their communication is simple.
- Special Concept: There is a trade-off between the amount of social structure required in a (primate) society and the capacity for conflict resolution (Flack and de Waal).

Outline

- Communication vs. Behavior Complexity
- How to Build Complex Agents
- Example: Socializing Without Communicating
- A Little More on Social Organization

Why Should Communication be Limited?

(in Nature)

- In a complex dynamic environment, an agent has best access to its own opportunities and dangers.
- Signals must persist for multi-tasking agents to reliably observe them $\Rightarrow$ no rapid or detailed instructions.
- Perception is unreliable $\Rightarrow$ signals must be distinct $\Rightarrow$ not many different ones.
- Evolution favors behavior uniform within species *and* advantageous to conspecifics $\Rightarrow$ high-level reference sufficient.

Do These Criteria Apply to Artificial Agents?

- Yes if communication is limited (e.g. by distance) or highly time-critical (e.g. in Robocup).
- If they don't, maybe we are discussing distributed intelligence in one agent, rather than a real multi-agent system (MAS).

Outline

- Communication vs. Behavior Complexity
- How to Build Complex Agents
- Example: Socializing Without Communicating
- A Little More on Social Organization

Start with Modular Intelligence

- Modules are relatively simple to design (or evolve).
- Provide the control for agent's actions (expressed and/or internal),
- the perception required for that control, and
- the memory required for perception or control.

This extends Behavior-Based AI (Brooks 1986, 1991) to be more like Object-Oriented Design. A little like Fodor (1983) but no central processor, less encapsulation.

Add Arbitration for Resource Conflicts

- Hierarchical Reactive Plans
  - help limit search, and
  - provide state on decisions (to defeat dithering).
- Support three types of action-selection problems:
  - Some things need to be checked at all times: drives.
  - Some only need considering in particular contexts: competences.
  - Some things reliably follow from others: action patterns.

Arbitration is not Fully Centralized Control

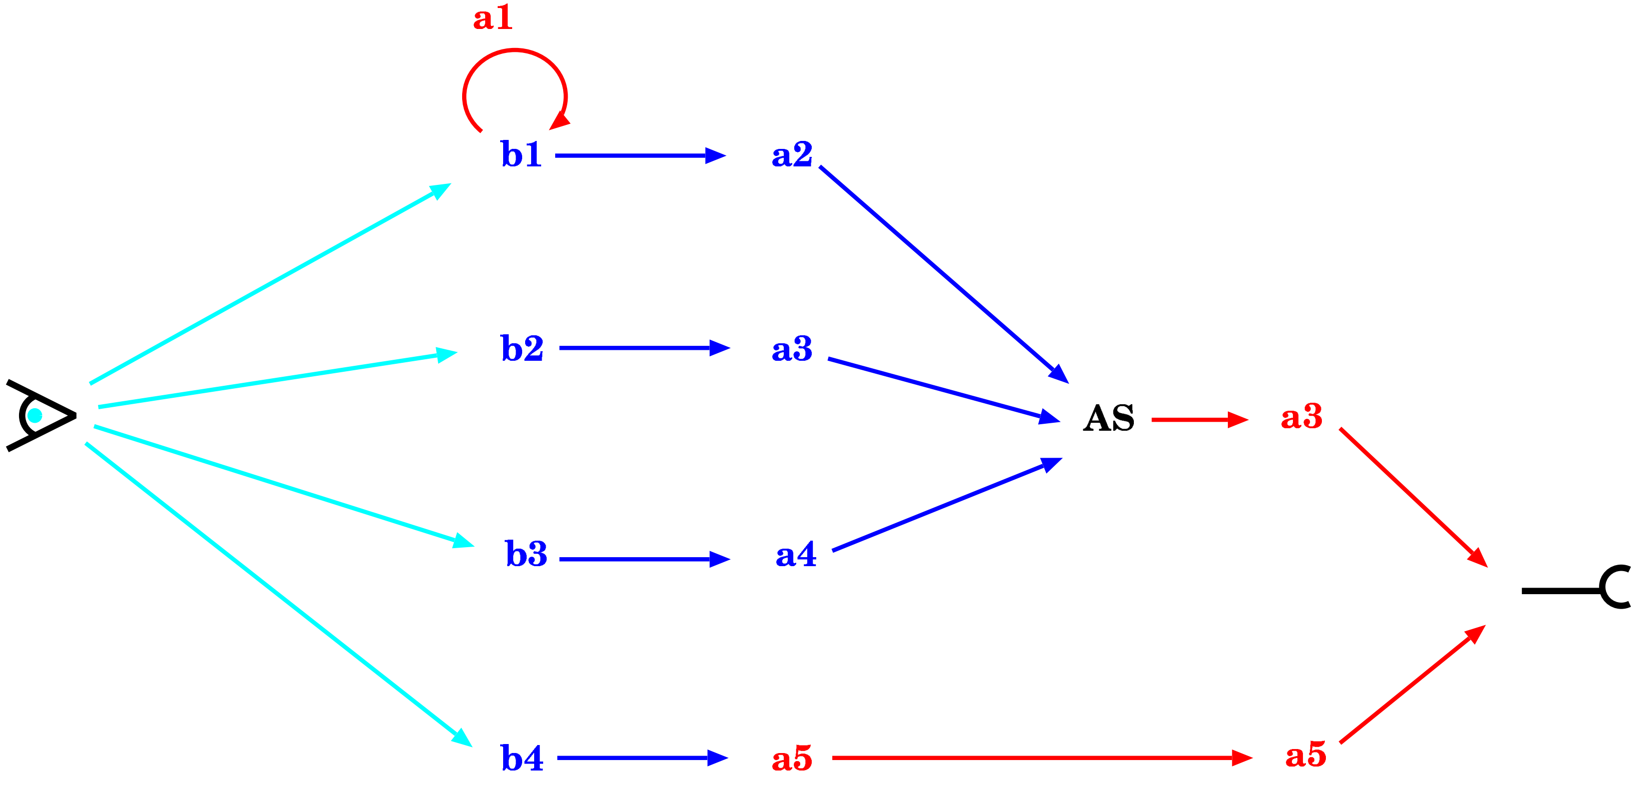

Behavior-Oriented Design (BOD) Methodology

- Initial Decomposition
- Cyclic Development

(Bryson & Stein IJCAI 2001, Bryson MIT AI PhD 2001)

Initial Decomposition

1. Specify (high-level) what the agent will do.
2. Describe activities as sequences of actions. reactive plans
3. Identify sensory and action primitives from these sequences.
4. Identify the state necessary to enable the primitives, cluster primitives by shared state. behaviors
5. Identify and prioritize goals or drives. drive collection
6. Select a first behavior to implement.

Cyclic Development

- Scale the system.
  - Code behaviors and / or plans.
  - Test and debug code.
- Simplify the design.
  - Revise the specifications.

Outline

- Communication vs. Behavior Complexity
- How to Build Complex Agents
- Example: Socializing Without Communicating
- A Little More on Social Organization

Basic Social Behaviors

	Navigate	**Groom**	**Explore**
state	x, y, size, name focus-of-attn	drive-level partner	drive-level direction-of-interest
actions	approach unentangle do-it	groom choose-partner engage	choose-new-location wait

- *life* (D)
  - untangle (tangled?) — $\boxed{\text{untangle}}$
  - groom (C) (want-to-groom?)
    - (partner-chosen?) (aligned?) — $\boxed{\text{groom}}$
    - (partner-chosen?) (touching?) — $\boxed{\text{align}}$
    - (partner-chosen?) — $\boxed{\text{approach}}$
    - ($\top$) — $\boxed{\text{choose-partner}}$
  - explore (C) (want-novel-loc?)
    - (place-chosen?) (there-yet?) — $\boxed{\text{lose-target}}$
    - (place-chosen?) — $\boxed{\text{explore-that-a-way}}$
    - ($\top$) — $\boxed{\text{choose-explore-target}}$
  - wait ($\top$) — $\boxed{\text{wait}}$

- *life* (D)
  - untangle (tangled?) → untangle
  - groom (C) (want-to-groom?)
    - (partner-chosen?) (aligned?) → notify groom
    - (being-groomed?) → choose-groomer-as-partner
    - (partner-chosen?) (touching?) → align
    - (partner-chosen?) → approach
    - (⊤) → choose-partner
  - receive (being-groomed?) → tolerate-grooming
  - explore (C) (want-novel-loc?)
    - (place-chosen?) (there-yet?) → lose-target
    - (place-chosen?) → explore-that-a-way
    - (⊤) → choose-explore-target
  - wait (⊤) → wait

Basic Social Behaviors

	Navigate	**Groom**	**Explore**
state	x, y, size name focus-of-attn	drive-level, partner groomed-by groomed-at	drive-level direction-of-interest
actions	approach unentangle do-it	groom, engage choose-partner tolerate, notify	choose-new-location wait

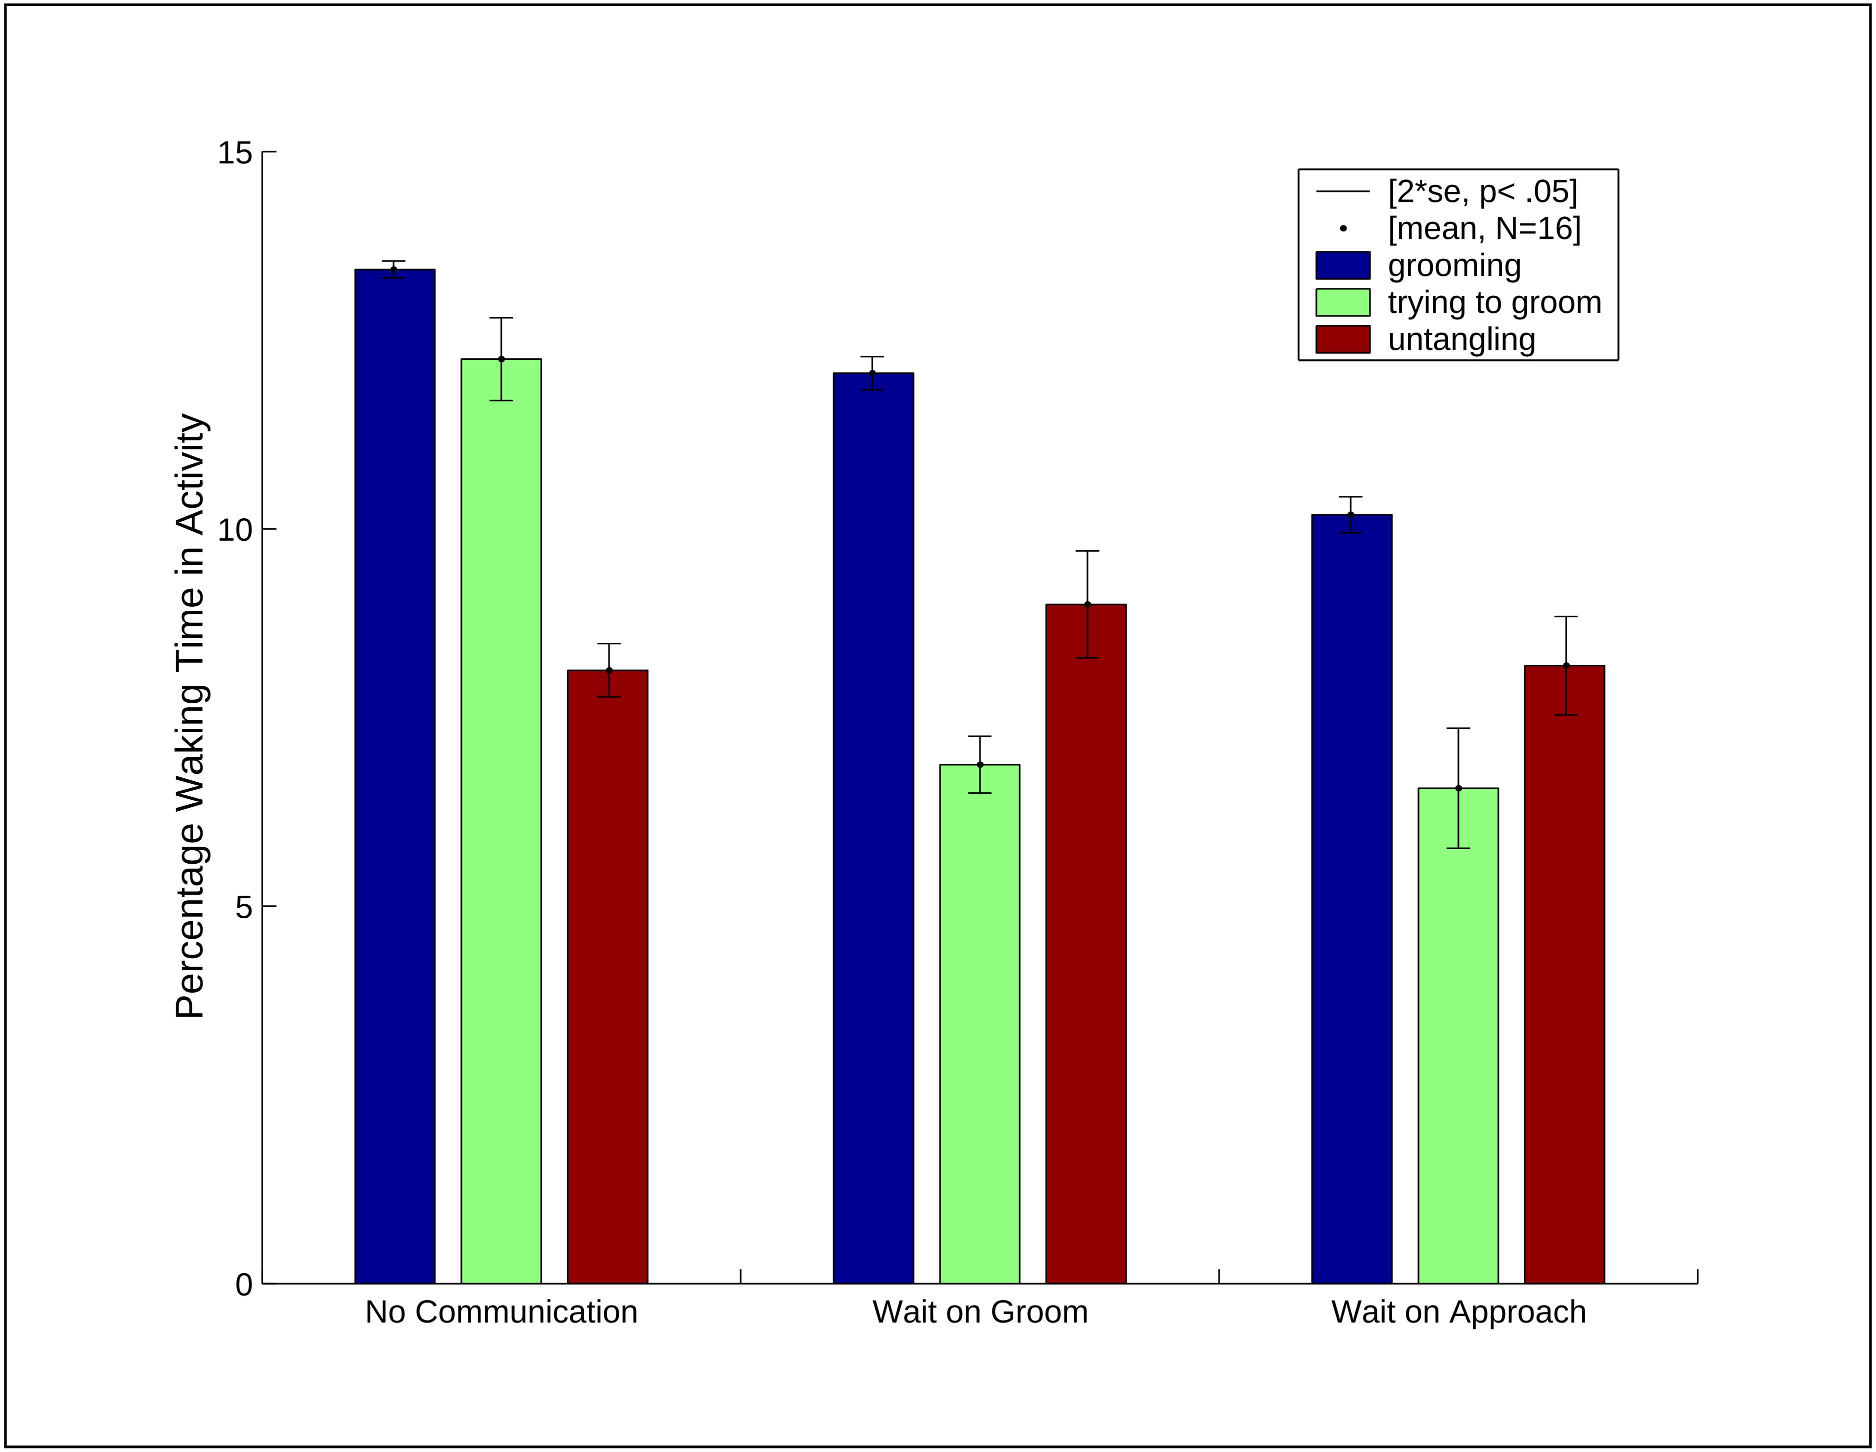

15
10
5
0
Percentage Waking Time in Activity
[2*se, p< .05]
[mean, N=16]
grooming
trying to groom
untangling
No Communication
Wait on Groom
Wait on Approach

Outline

- Communication vs. Behavior Complexity
- How to Build Complex Agents
- Example: Socializing Without Communicating
- A Little More on Social Organization

A Little More on Social Organization

- Related primate species have differing degrees of hierarchical vs. egalitarian social orders.
- The ones that are more egalitarian have more reconciliation strategies and reconcile more frequently.
- Weird analology: Kind of like token ring vs. ethernet.
- Egalitarians fight more, but with less severity.
- Some strategies can be learned, other seems innate.
- Goal of modeling effort.

Summary

- General Radical Concept: Maybe if things communicate much they aren't really agents in a MAS, but are modules in an agent.
- Less Radical Concept: Maybe we should worry more explicitly about software engineering.
- Special Radical Concept: Maybe studying primate social organization will be useful to MAS research.

Acknowledgements

- Much of this work was done on DAML funding under the supervision of Lynn Andrea Stein, of Olin College of Engineering.
- Primate society conceptualizing done in collaboration with Jessica Flack of Emory University.

[Talk Boundary]

Is BOD Biologically Plausible?

- Elements in BOD
  - Behaviors: Neocortex
  - Action Selection: Basal Ganglia, Periaqueductal Gray Matter
  - Drive Collection: Amygdalic System
- Other Possible Additions
  - Smoothing and Interpolation (Cerebellum?)
  - Episodic Memory (Hippocampus?)

Revising the Specification: State

- Prefer the simplest.
  1. Control State
  2. Deictic State
  3. Specialized State (learning)
  4. Meta-State (learning to learn)
- Exceptions:
  - Eliminate Plan Redundancy
  - Reduce Plan Complexity

Revising the Specification: Control

- Prefer the simplest.
  - Single Primitive $>$ Sequence
  - Sequence $>$ BRP
  - Control State $>$ Variable State
- Exceptions:
  - Want part of primitive $\Rightarrow$ sequence.
  - Sequence elements repeated, skipped $\Rightarrow$ BRP.
  - Use variable state to:
    - Replace lots of triggers.
    - Generalize control state.

Example

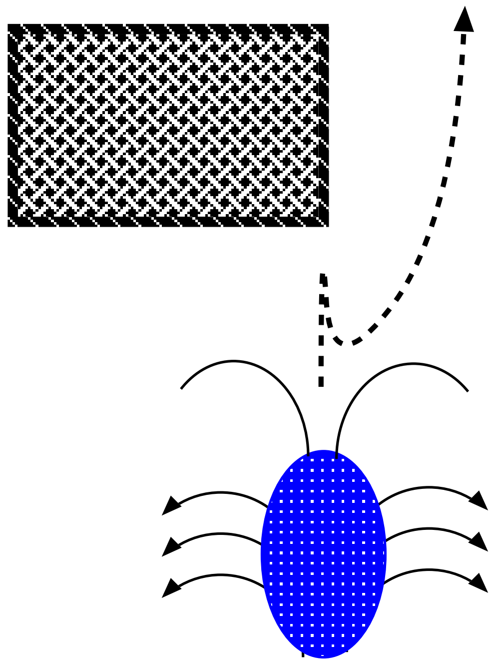

Control State Only

$$\textbf{walk} \Rightarrow \Big\uparrow \left\langle \begin{array}{rl} (\textsf{left-feeler-hit}) \Rightarrow & \textbf{avoid-obstacle-left} \\ (\textsf{right-feeler-hit}) \Rightarrow & \textbf{avoid-obstacle-right} \\ \Rightarrow & \textsf{walk-straight} \end{array} \right\rangle$$

$$\textbf{avoid-obstacle-left} \Rightarrow \langle \textsf{walk backwards} \rightarrow \textsf{walk right} \rightarrow \textsf{walk left} \rangle$$

$$\textbf{avoid-obstacle-right} \Rightarrow \langle \textsf{walk backwards} \rightarrow \textsf{walk left} \rightarrow \textsf{walk right} \rangle$$

Deictic State as Well

$$\xleftarrow[\textbf{compensate-avoid}]{\textbf{avoid-hit, feeler-hit,}} \boxed{\begin{array}{c}\text{deictic-avoid}\\ \hline \text{hit-left}\end{array}} \xleftarrow[\text{feeler info}]{}$$

$$\textbf{walk} \Rightarrow \Big\uparrow \left\langle \begin{array}{r} (\text{feeler-hit}) \Rightarrow \textbf{avoid-obstacle} \\ \Rightarrow \text{walk-straight} \end{array} \right\rangle$$

$$\textbf{avoid-obstacle} \Rightarrow \langle \text{walk backwards} \rightarrow \text{avoid hit} \rightarrow \text{compensate avoid} \rangle$$

Specialized State (rather than Deictic)

$$\xleftarrow[\textbf{store-obstacle}]{\textbf{back-up, find-way}} \boxed{\begin{array}{c}\text{specialized-avoid}\\ \hline \text{local-map}\end{array}} \xleftarrow[\text{feeler info}]{}$$

$$\textbf{walk} \Rightarrow \Big\uparrow \left\langle \begin{array}{rl} (\text{feeler-hit}) & \Rightarrow \text{store-obstacle back-up} \\ & \Rightarrow \text{find-way} \end{array} \right\rangle$$

Learning

- Learning in Behaviors
- Learning of Plans
  - Search
  - Evolution
  - Imitation
- Learning Behaviors
  - Existing work is one behavior in BOD.
  - Learning dynamical models (e.g. Hogg, Brand)

A Simple Behavior

$\boxed{\underline{\text{screeching}}}$

A Behavior with State

screeching
screeching-now?
pulse-duration

Behaviors with Perception

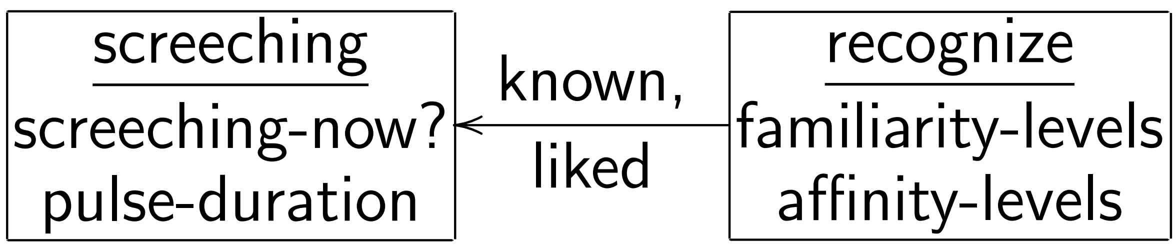

Behaviors that Aren't Objects

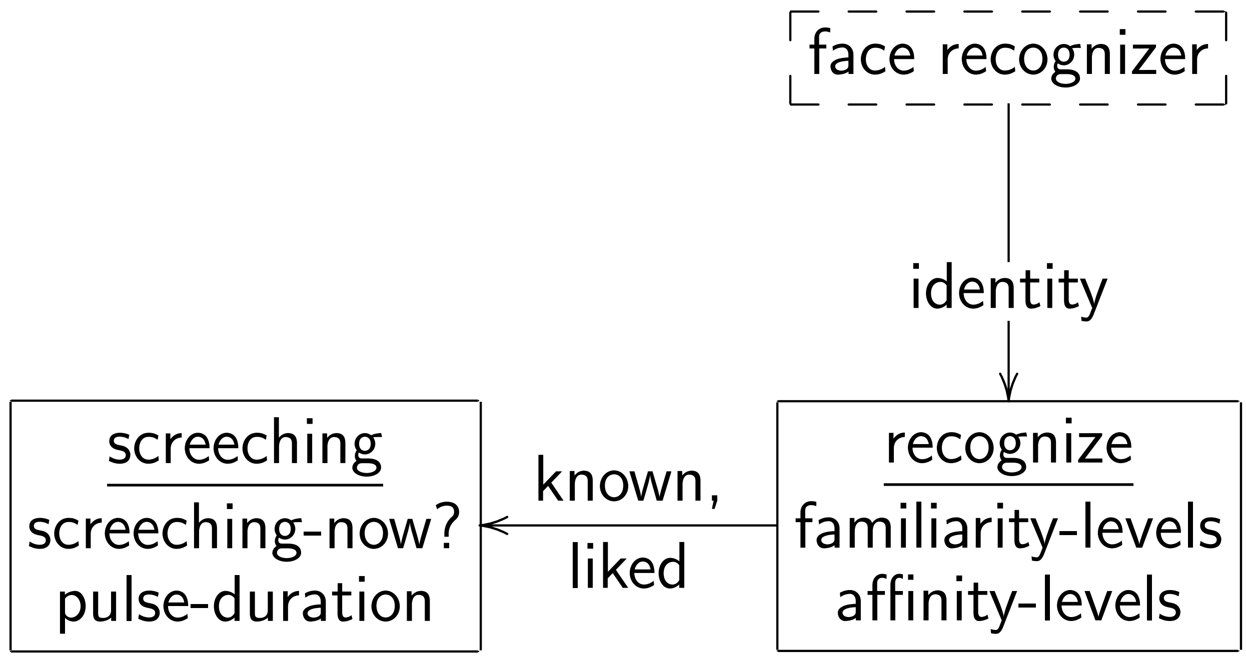

Behaviors with Processes and/or Triggers

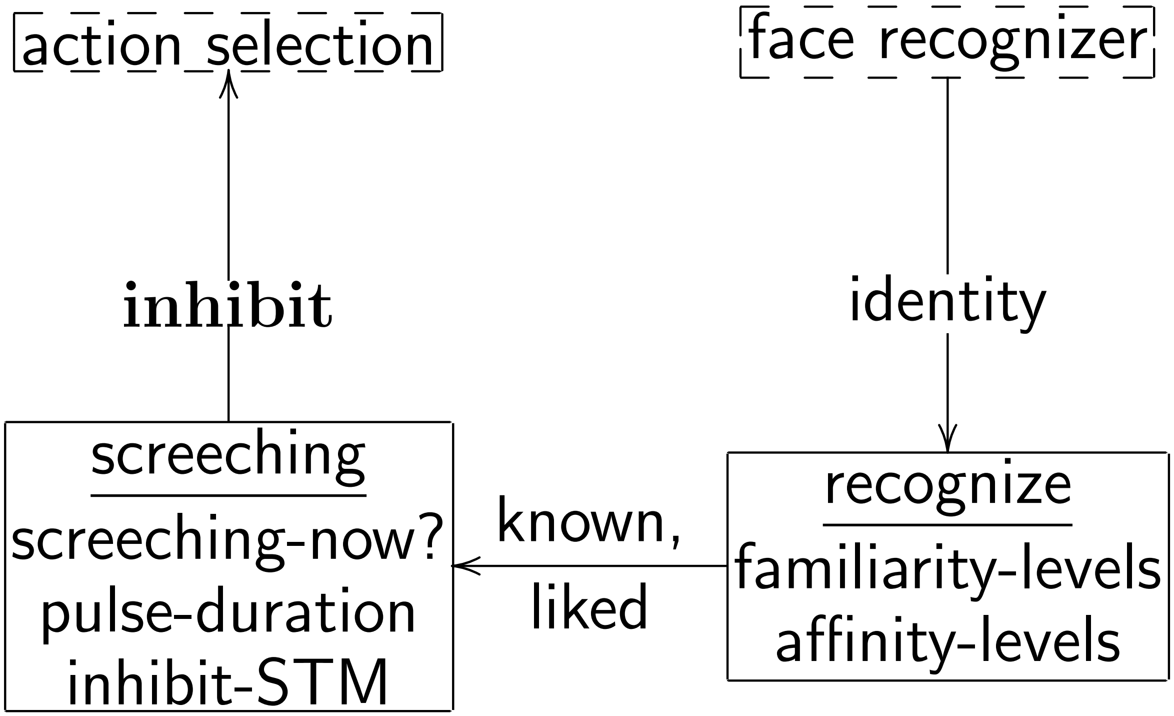

Some Things Follow: Action Patterns

$\langle$get a banana $\rightarrow$ peel a banana $\rightarrow$ eat a banana$\rangle$

Another Representation:
Production Rules

$$
\begin{aligned}
\text{(have hunger)} &\Rightarrow \text{get a banana} \\
\text{(have a banana)} &\Rightarrow \text{peel a banana} \\
\text{(have a peeled banana)} &\Rightarrow \text{eat a banana}
\end{aligned}
$$

Are Production Rules Better than Action Patterns?

$$\begin{aligned}\text{(have hunger)} &\Rightarrow \text{get a banana}\\ \text{(have a banana)} &\Rightarrow \text{peel a banana}\\ \text{(have a peeled banana)} &\Rightarrow \text{eat a banana}\end{aligned}$$

No — A Sequence is State

$\langle$get a banana from left $\rightarrow$ pass a banana to right$\rangle$

(left neighbor offers banana) $\Rightarrow$ get a banana from left

(have a banana) $\Rightarrow$ pass a banana to right

Basic Reactive Plans: State + Flexibility

$$
\text{(hungry)} \Rightarrow \uparrow \left\langle \begin{array}{rl} \text{(full)} \Rightarrow & \textit{goal} \\ \text{(have a peeled banana)} \Rightarrow & \text{eat a banana} \\ \text{(have a banana)} \Rightarrow & \text{peel a banana} \\ \Rightarrow & \text{get a banana} \end{array} \right\rangle
$$

Many different *expressed* plans (sequences of behavior) are determined by one *reactive* plan.

Drive Collections:
BRPs for Environment Monitoring

$$\textit{life} \Rightarrow \left\langle\left\langle \begin{array}{r} \text{(something looming)} \Rightarrow \text{avoid} \\ \text{(something loud)} \Rightarrow \text{attend to threat} \\ \text{(hungry)} \Rightarrow \text{forage} \\ \Rightarrow \text{lounge around} \end{array} \right\rangle\right\rangle$$

How Hard is Learning Everything?