Towards a General Architecture for the Integration of Reactive and Deliberative Behaviour:
Insights from Cognitive Neuroscience

Rick Cooper
Birkbeck, University of London

Birkbeck
UNIVERSITY OF LONDON

Overview

- Reactive behaviour versus deliberative behaviour
- A model of routine behaviour grounded in cognitive neuroscience
- Deliberative behaviour and cognitive neuroscience
  - Fractionating the “central executive”
- Towards an account of strategy generation
  - The “domino” architecture

Reactive Behaviour versus Deliberative Behaviour

Two Distinct Domains of Behaviour

- Reactive (routine) behaviour:
  - Fast
  - Habitual
  - Minimal cognitive effort, but …
  - Subject to slips and lapses
  - Requires deliberate / willed suppression
- Deliberative (nonroutine) behaviour:
  - Slow
  - Willed / Volitional
  - Requires attentional resources / cognitive effort
  - Comprised of sequences of reactive behaviours

Impairments of Reactive Behaviour

- Ideational apraxia (Liepmann, 1908):
  - Associated with left temporoparietal damage (De Renzi & Lucchelli, 1988)
  - The patient makes conceptual errors in simple tasks involving object use
  - Misuse of tools
  - Mislocation of actions
  - The patient can show no impairment of task knowledge, as assessed through picture sequencing tasks (Rumiati et al., 2001)

Impairments of Deliberative Behaviour

- Utilisation behaviour
  - The patient spontaneously uses objects in the environment in object-appropriate ways
- Anarchic hand syndrome
  - Similar to utilisation behaviour, but with one hand only
- Action disorganisation syndrome (?)
  - Disorganised goal-directed action
- Dysexecutive syndrome
  - Impairments on novel tasks requiring “flexible” thinking

The Contention Scheduling / Supervisory System Framework (Mark 1)

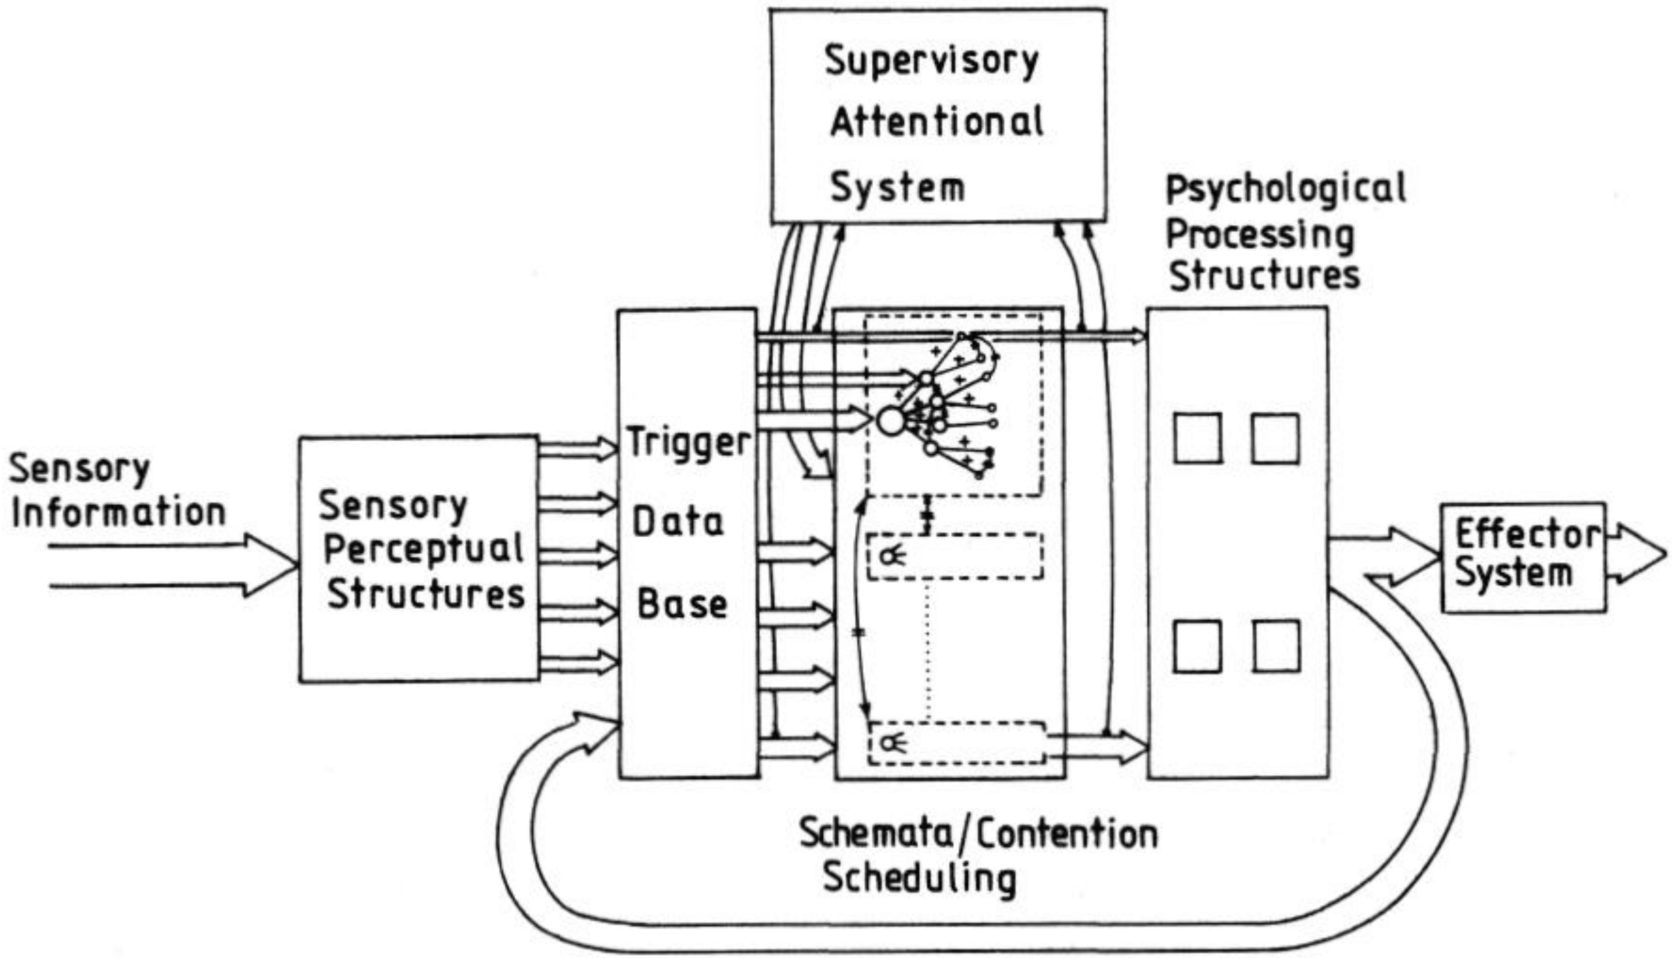

The Contention Scheduling / Supervisory System Framework (Mark 2)

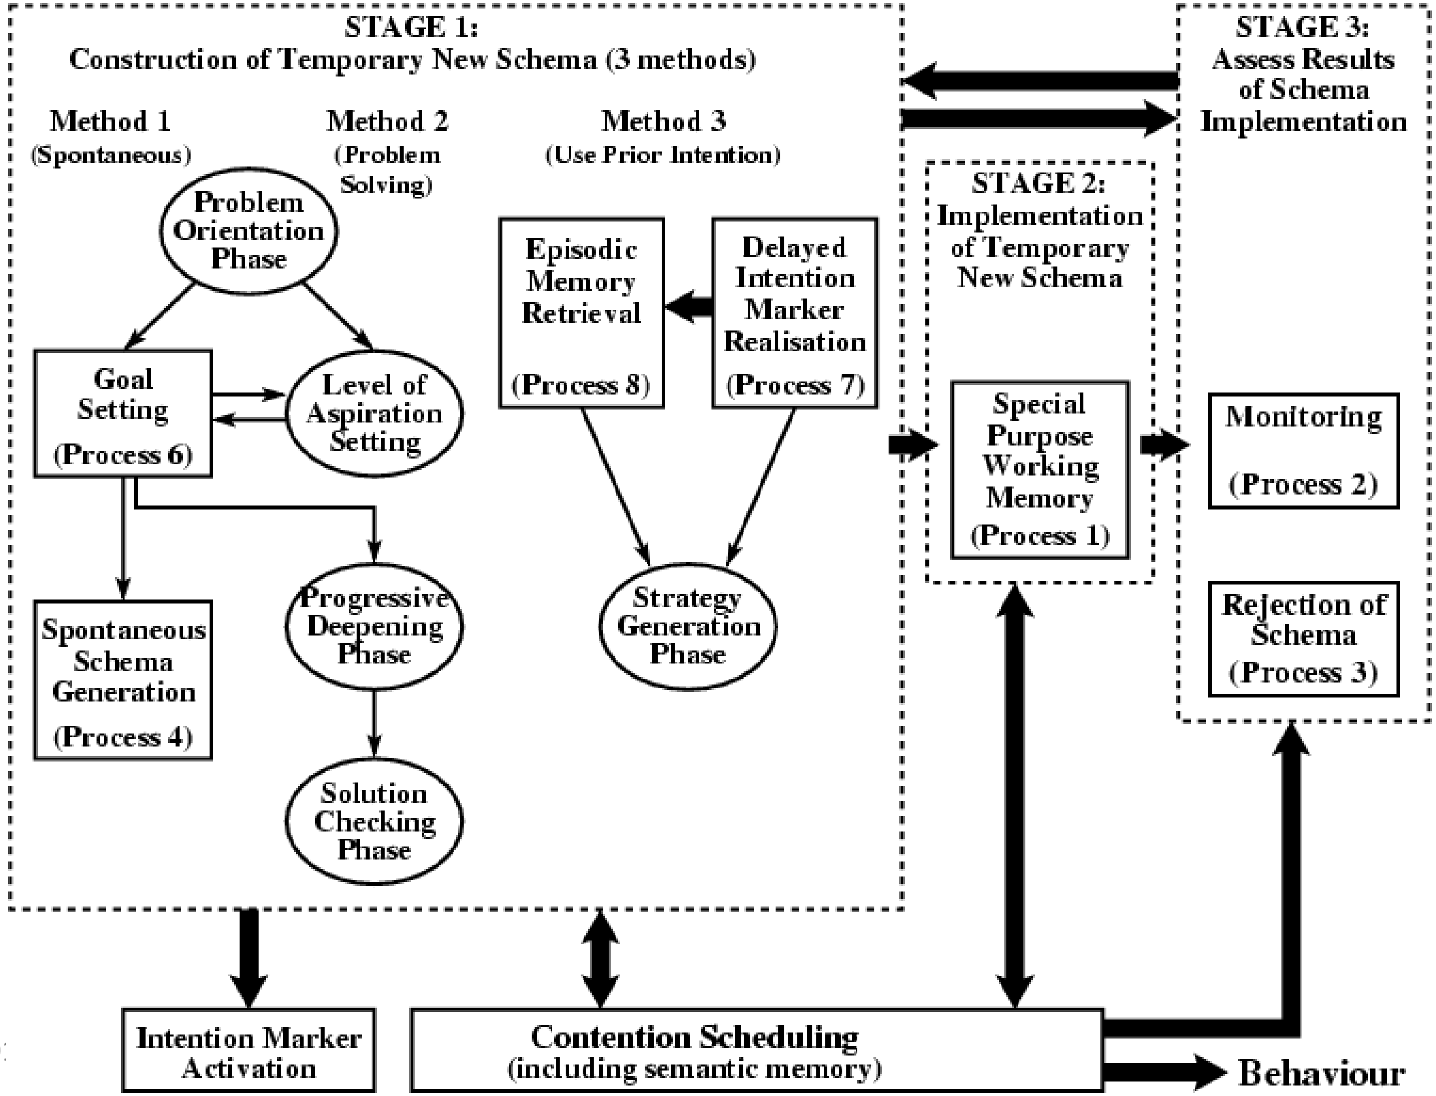

A Model of Routine Behaviour Grounded in Cognitive Neuroscience

Basic Assumptions

- Behaviour comprises performance of action schemas
- Action schemas are abstractions over commonly performed goal-directed sequences of action
- Action schemas are associated with nodes within a hierarchically structured interactive activation network
- Nodes have activation values that are:
  - Excited (top down) intentionally by SS or when their parent node is selected
  - Excited (bottom-up) when their triggering conditions match the representation of the environment
  - Inhibited by nodes corresponding to competing schemas

A Model of CS:
A Sample Schema Hierarchy

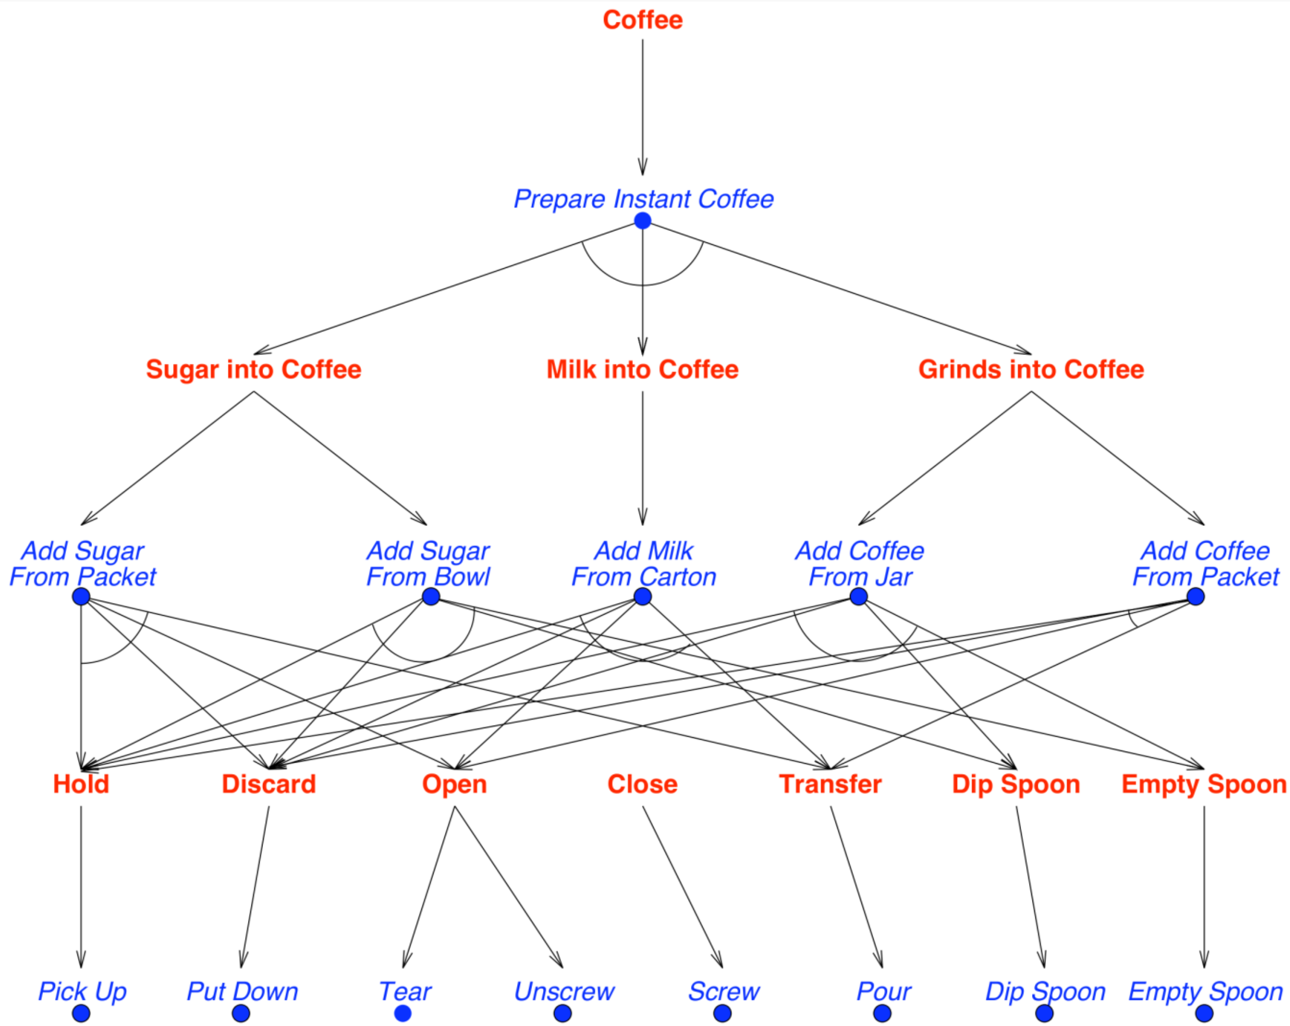

A Model of CS:
Processing and Selection of Actions

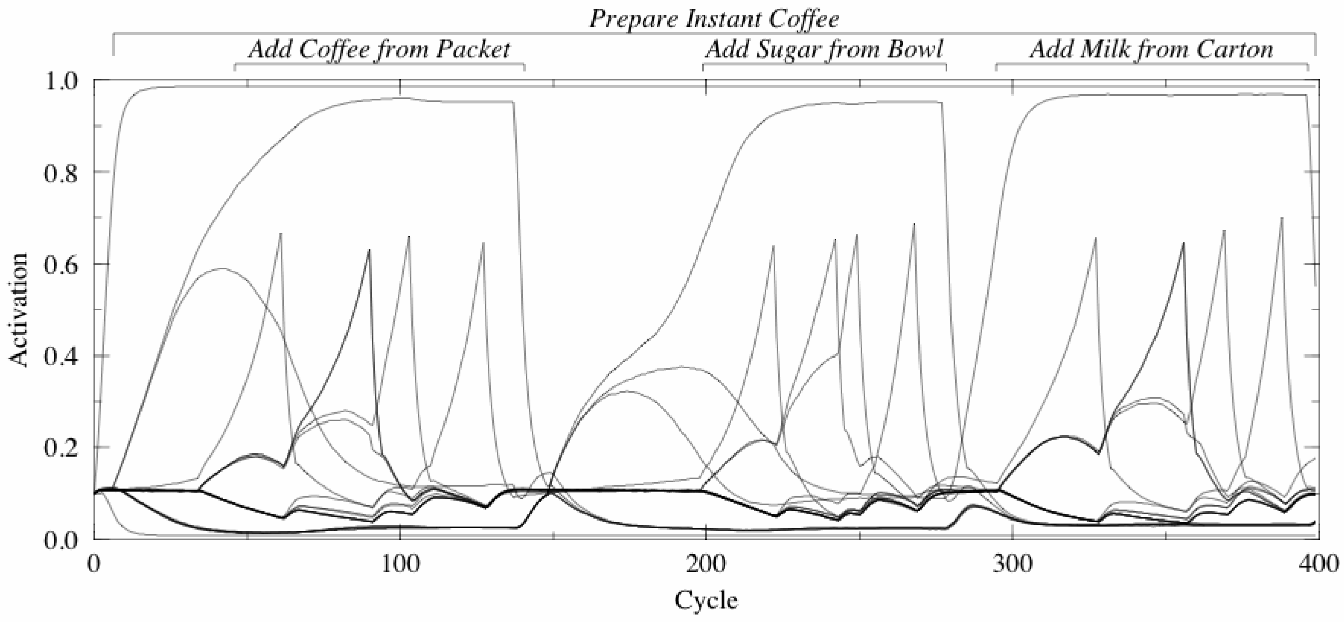

6 +*Prepare Instant Coffee*
34 +*Add Coffee from Packet*
61 +*Pick Up*
61 Picking up coffee packet with left hand
66 −*Pick Up*
90 +*Tear*
90 Tearing coffee packet (with right hand)
94 −*Tear*
103 +*Pour*
103 Pouring coffee packet into coffee mug
110 −*Pour*
127 +*Put Down*
127 Putting down coffee packet
137 −*Put Down*
148 −*Add Coffee from Packet*
198 +*Add Sugar from Bowl*
222 +*Pick Up*
222 Picking up spoon with left hand
226 −*Pick Up*
242 +*Dip Spoon*
242 Dipping spoon into sugar bowl
245 −*Dip Spoon*
249 +*Empty Spoon*
249 Emptying spoon into coffee mug
255 −*Empty Spoon*
268 +*Put Down*
268 Putting down spoon
277 −*Put Down*
285 −*Add Sugar from Bowl*
295 +*Add Milk from Carton*
327 +*Pick Up*
327 Picking up milk carton with right hand
332 −*Pick Up*
356 +*Tear*
356 Tearing milk carton (with left hand)
360 −*Tear*
369 +*Pour*
369 Pouring milk carton into coffee mug
376 −*Pour*
388 +*Put Down*
388 Putting down milk carton
396 −*Put Down*
397 Trial complete

A Model of CS: The Structure of Schemas

- The flexibility of routine behaviour suggests schemas are complex entities:
  - They are goal directed
  - They have an associated triggering condition, which governs activation from the internal model of the environment
  - They specify a set of sub-goals
  - Each sub-goal has pre-conditions and post-conditions, which determine if a sub-goal is optional and when a sub-goal is complete

A Model of CS:
Schema / Object-Representation Interactions

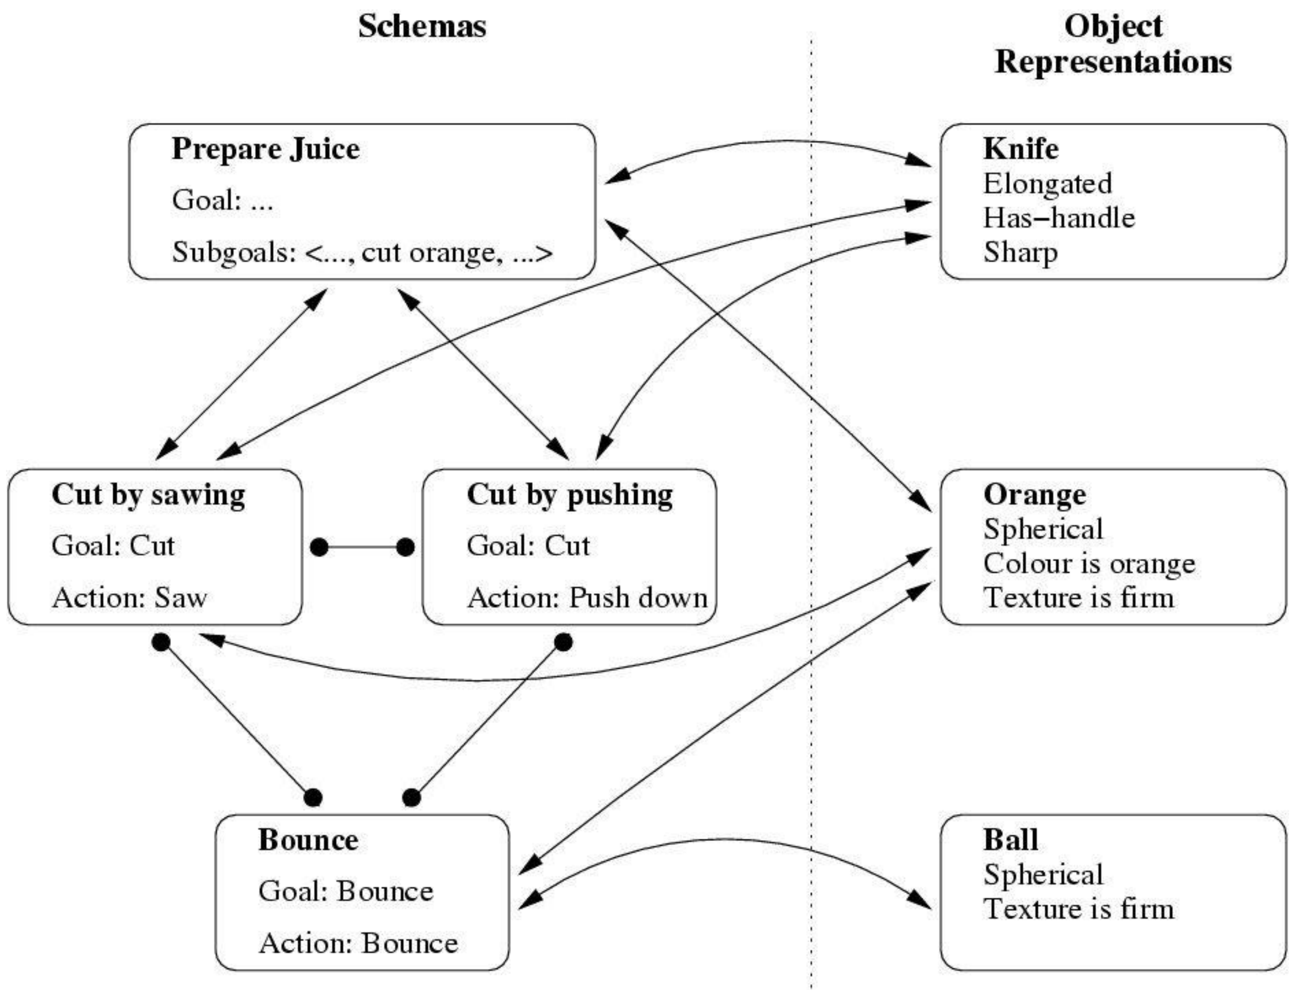

Support for the Model

- Cooper & Shallice (2000):
  - Qualitative simulations of normal slips and lapses, and of action disorganisation syndrome following closed head injury
- Cooper, Schwartz, Yule & Shallice (2005):
  - Quantitative simulations of ADS (error distribution; the distractor effect)
- Cooper & Shallice (2006):
  - An alternative Simple Recurrent Network (SRN) model based on chaining actions within context cannot account adequately for observed error types, or the flexibility of real-world action
- Cooper (2007):
  - Quantitative simulations of ideational apraxia (following left hemisphere lesions), and in particular two subtypes of IA reflecting disconnection of (i) object-representations from schemas and (ii) schemas from object-representations.

Implications for N/S Theory

- The model demonstrates that CS is, in principle, capable of the functions ascribed to it
- The model has clarified the structure of schemas (vis-à-vis goal direction, optionality of sub-goals, etc.)
- The model has clarified the interaction between CS and the representation of the environment
- The model demonstrates the proposed explanations of action disorganisation syndrome and ideational apraxia are viable
- Schemas provide an interface between SS and CS: SS may indeed modulate CS by exciting/inhibiting schema nodes

Applications in Cognitive Robotics

- Chernova and Arkin (2007):
  - Sony QRIO robot with six pre-specified behaviours
  - Transfer of behaviour from a deliberative to a routine system
  - The robot acquires complex behavioural routines by chaining subroutines
- Burattini, Finzi and Staffa (2010):
  - Controlling attention under conditions of distraction
  - Lateral inhibition between competing behaviours is used within a mobile robot to implement adaptive attention mechanisms

Deliberative Behaviour and Cognitive Neuroscience:

Fractionating the “Central Executive”

Some Context

- What is the nature of “the central executive”?
- Miyake et al. (2000):
  - Multiple components, including response inhibition, task shifting, memory updating, dual tasking, …
- Higher level functions:
  - Planning, dealing with novelty, prospective memory for events, adaptive thinking, …
- Many models of specific EFs in specific tasks
- Few (no?) accounts of how multiple EFs interact to control behaviour in complex tasks (Cooper, 2010)

The Approach

- Basic Issue:
  - What can be deduced (from Cognitive Neuroscience) about the organisation of the Supervisory System?
- Strategy:
  - What functions are required in addition to those of Contention Scheduling to support intelligent behaviour?
  - Does cognitive neuroscience support the existence of subsystems that carry out those functions?
  - We use COGENT, a kind of *Lingua Franca*, to support the theory development

The COGENT Notation (Cooper & Fox, 1998): Basic Contention Scheduling

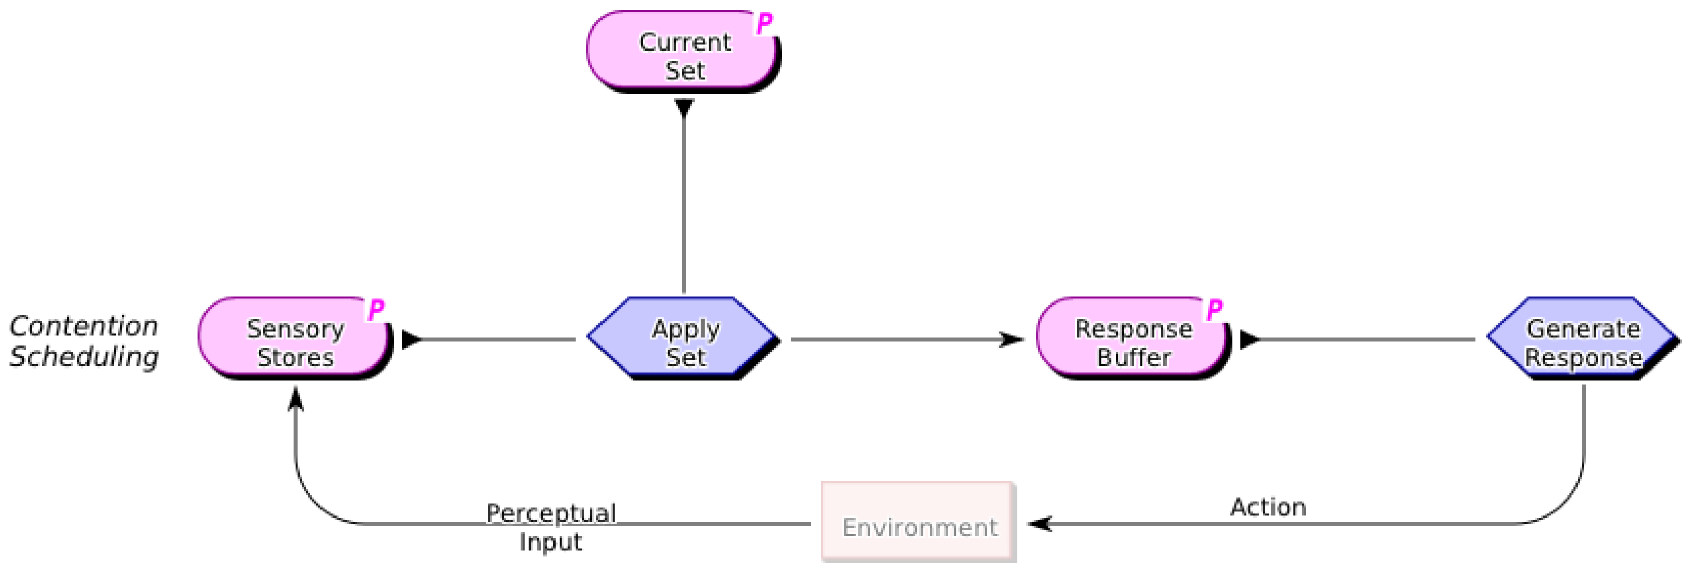

Augmenting Contention Scheduling 1: Prediction …

- Rapid motor control in which motor commands are generated before their effects are felt requires the generation and maintenance of internal models of the environment and the state of the motor system (Wolpert & Ghahramani, 2000)
- Two primary types of internal model:
  - Inverse model: Planning
  - Forward model: Predicted consequences
- Similar concepts apply at higher levels of the cognitive system (e.g., Alexander & Brown, 2010)
- Cognitive neuroscience:
  - Cerebellum is implicated in the learning of forward motor models
  - (Superior) parietal lobes are implicated in functions such as prediction

Augmenting Contention Scheduling 1: … and Monitoring

- If prediction is available, then:
  - Comparison of prediction and effect allows error detection
  - Comparison of prediction and intention allows one to pre-empt and prevent error (i.e., pro-active control)

- Cognitive neuroscience:
  - fMRI studies show ACC activation on erroneous trials of experimental tasks (e.g., Carter et al., 1998)
  - Behaviour of patients on a range of tasks with focal right LPFC damage can be understood in terms of a monitoring deficit (Shallice et al., 2008)

Augmenting Contention Scheduling 1: Prediction and Monitoring

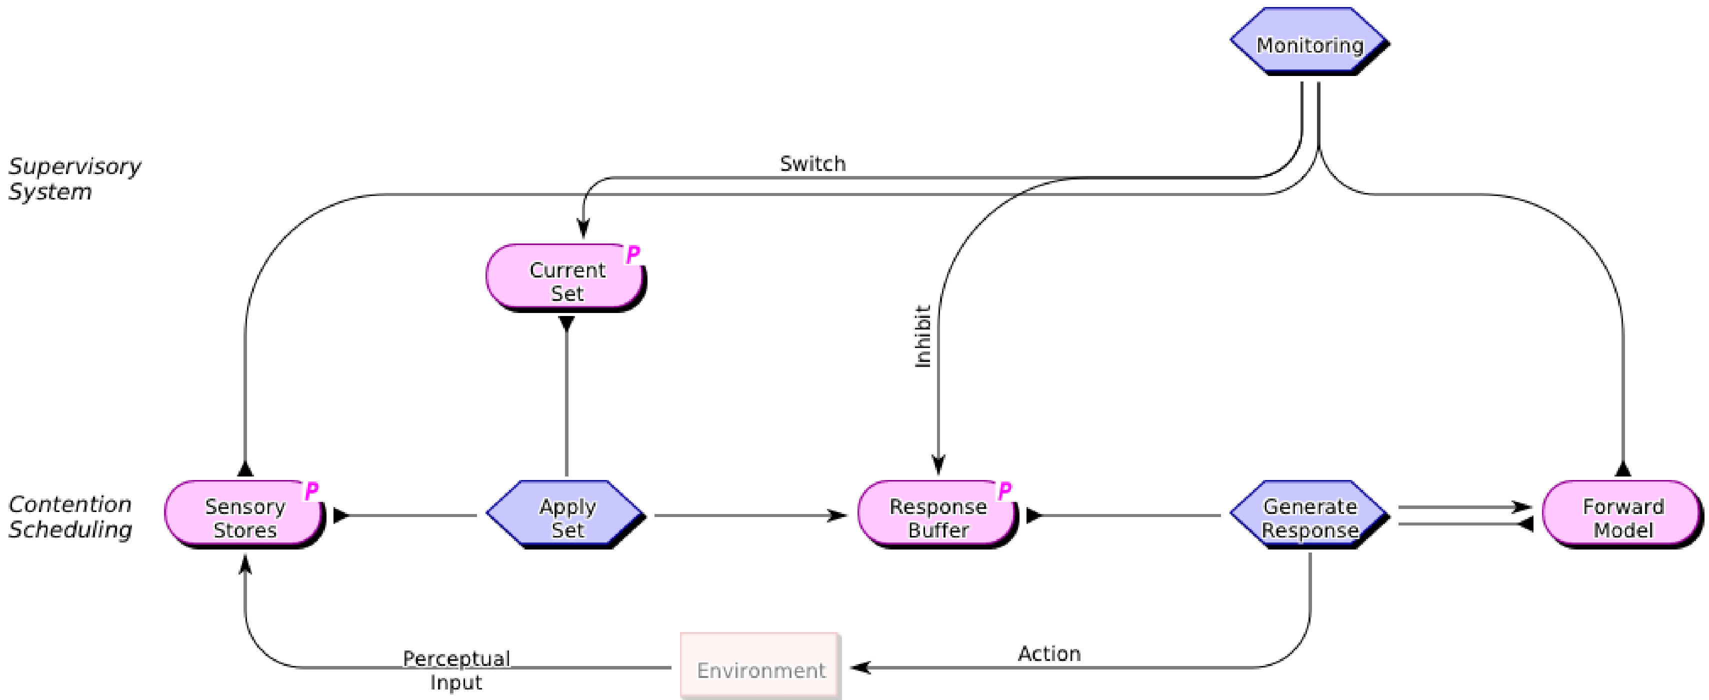

Augmenting Contention Scheduling 2: Working Memory

- Clearly there are neural mechanisms to support the short-term retention of information
- Such mechanisms are required if behaviour is to be sensitive to recent experience
- Cognitive neuroscience evidence:
  - Short term retention of verbal information via the phonological loop (e.g., Baddeley, 1986)
  - Frontal systems are implicated in the manipulation of information retained in short-term memory (e.g., frontal patients and backwards digit span)

Augmenting Contention Scheduling 2: Working Memory

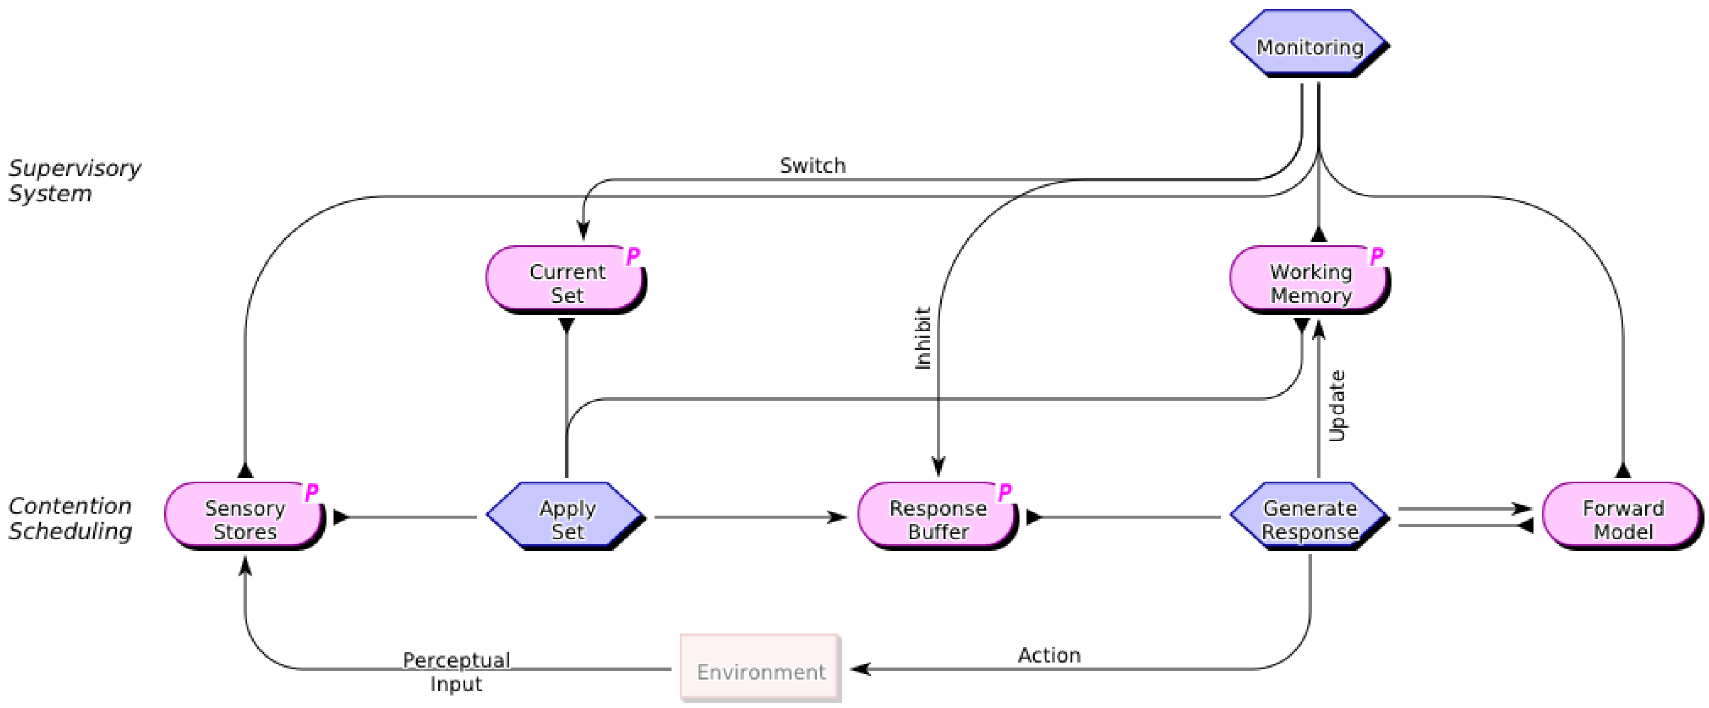

Augmenting Contention Scheduling 3: Attentiveness and Active Maintenance

- Working memory decays if not actively maintained:
  - O’Reilly and Frank (2006): Gating switches between WM maintenance and updating
- Working hypothesis: Task set behaves analogously
  - Altmann and Gray (2008): Decay is functional
  - *Attentiveness* serves the maintenance function for task set
- Cognitive neuroscience evidence:
  - Gating is implemented by the basal ganglia (O’Reilly & Frank, 2006)?
  - Attentiveness is a function of inferior medial PFC (Shallice et al., 2008)?

Augmenting Contention Scheduling 3: Attentiveness and Active Maintenance

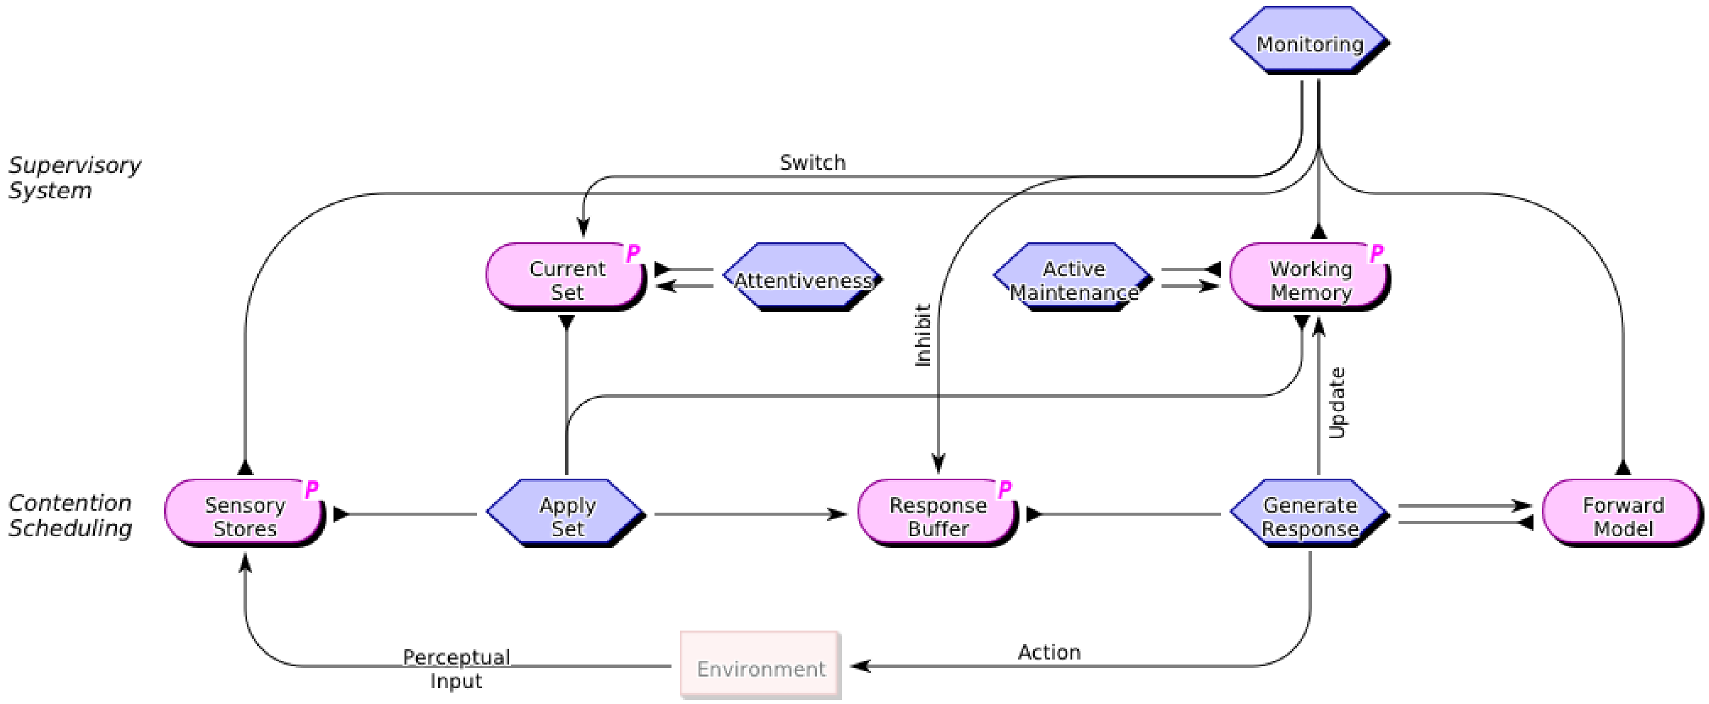

Augmenting Contention Scheduling 4: Strategy Generation

- Intelligent behaviour requires the generation of task-appropriate temporary schemas:
  - We view strategy generation as a separable process
- Strategy generation can involve:
  - Adaptation of an existing strategy
  - Reasoning over hypothetical states of the world
  - Planning
  - Induction and / or insight
- Is strategy generation a function of left LPFC (e.g., Shallice, 1982; Reverberi et al., 2005)?

Augmenting Contention Scheduling 4: Strategy Generation

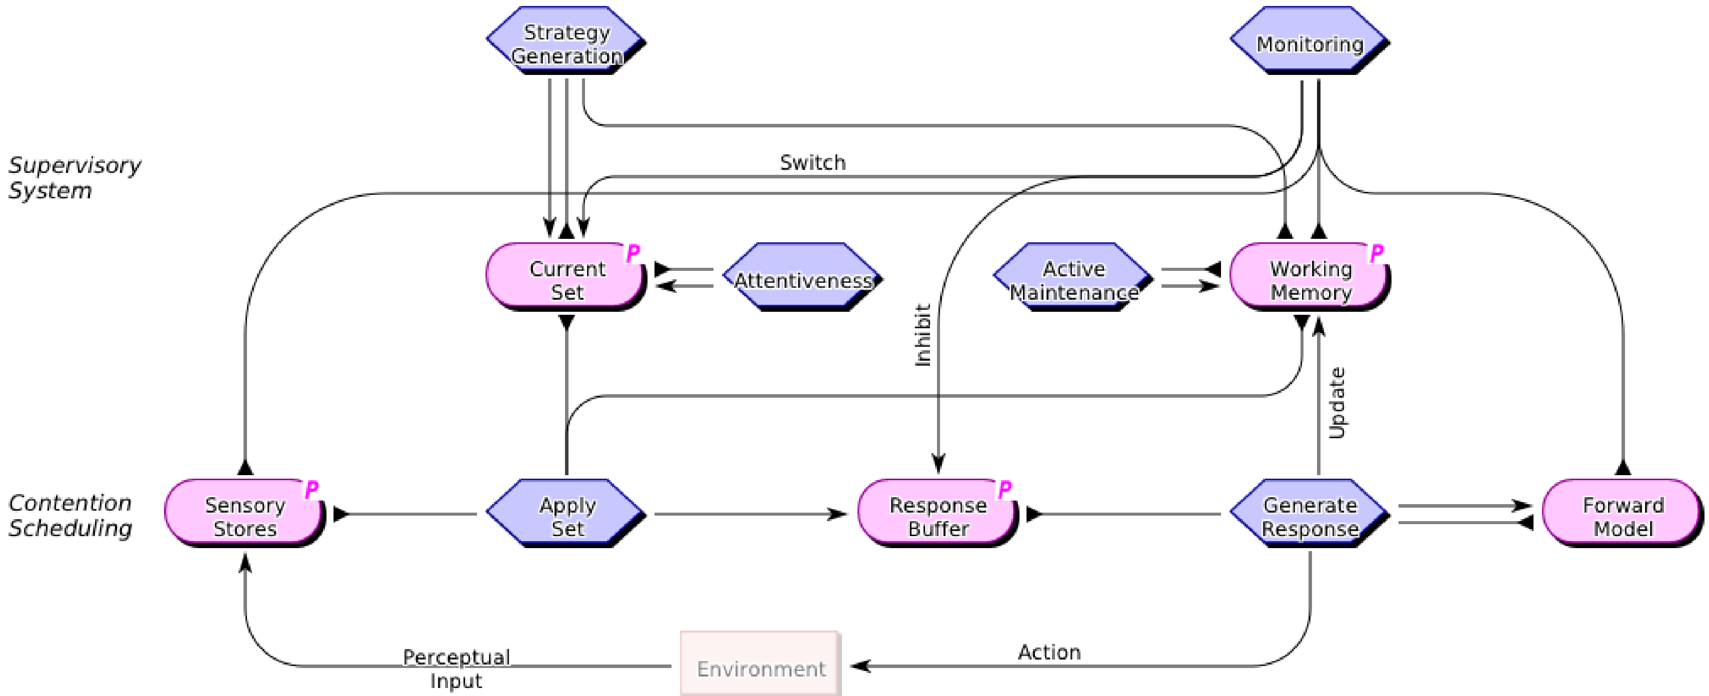

Relation to Standard Accounts of EF

- Separable functions of Miyake et al. (2000) are signals within the wider architecture:
  - Response inhibition, memory updating, task shifting
- Standard accounts fail to consider how EF might interact
- We view monitoring as a separate process with a more general remit (not just WM monitoring)
- More complex EF (e.g., planning) are subsumed by strategy generation:
  - But is strategy generation an homunculus?

Alternative Approaches to PFC Functioning

- Hierarchical approaches:
  - Fuster (1989): PFC sits at the apex of a pyramid rooted in perception and action (cf. Botvinick, 2007)
  - Koechlin et al. (2003): Neuroscience suggests increasing temporal abstraction as one moves forward in the PFC
  - Badre and D'Esposito (2009): Imaging and patient studies suggest a rostro-caudal axis of representational abstraction within PFC
- None of these approaches offers a concrete computational account of complex task performance
  - How can these ideas be cached out in any specific task?

Applications of the Architecture
1: Random Sequence Generation

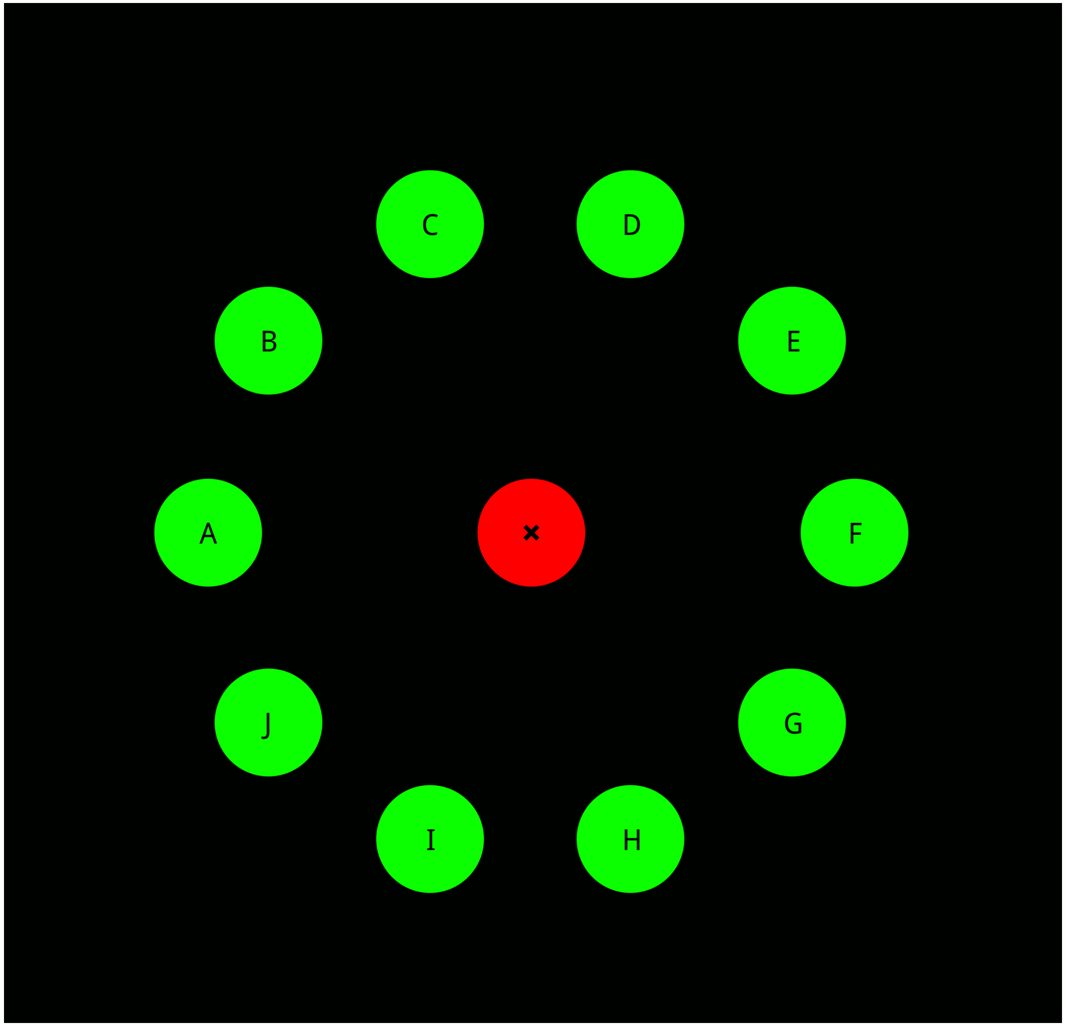

Applications of the Architecture 1: Random Sequence Generation

	R	RNG	RR	AA	OA
Ctrl	0.962	0.300	0.014	0.259	0.131
Varying half-life of *Working Memory*					
H/L=10	0.621	0.398	0.011	0.443	0.059
H/L=20	0.753	0.278	0.014	0.327	0.108
H/L=30	0.771	0.263	0.018	0.285	0.131
H/L=40	0.924	0.256	0.024	0.281	0.141

Applications of the Architecture 2: Wisconsin Card Sorting Test

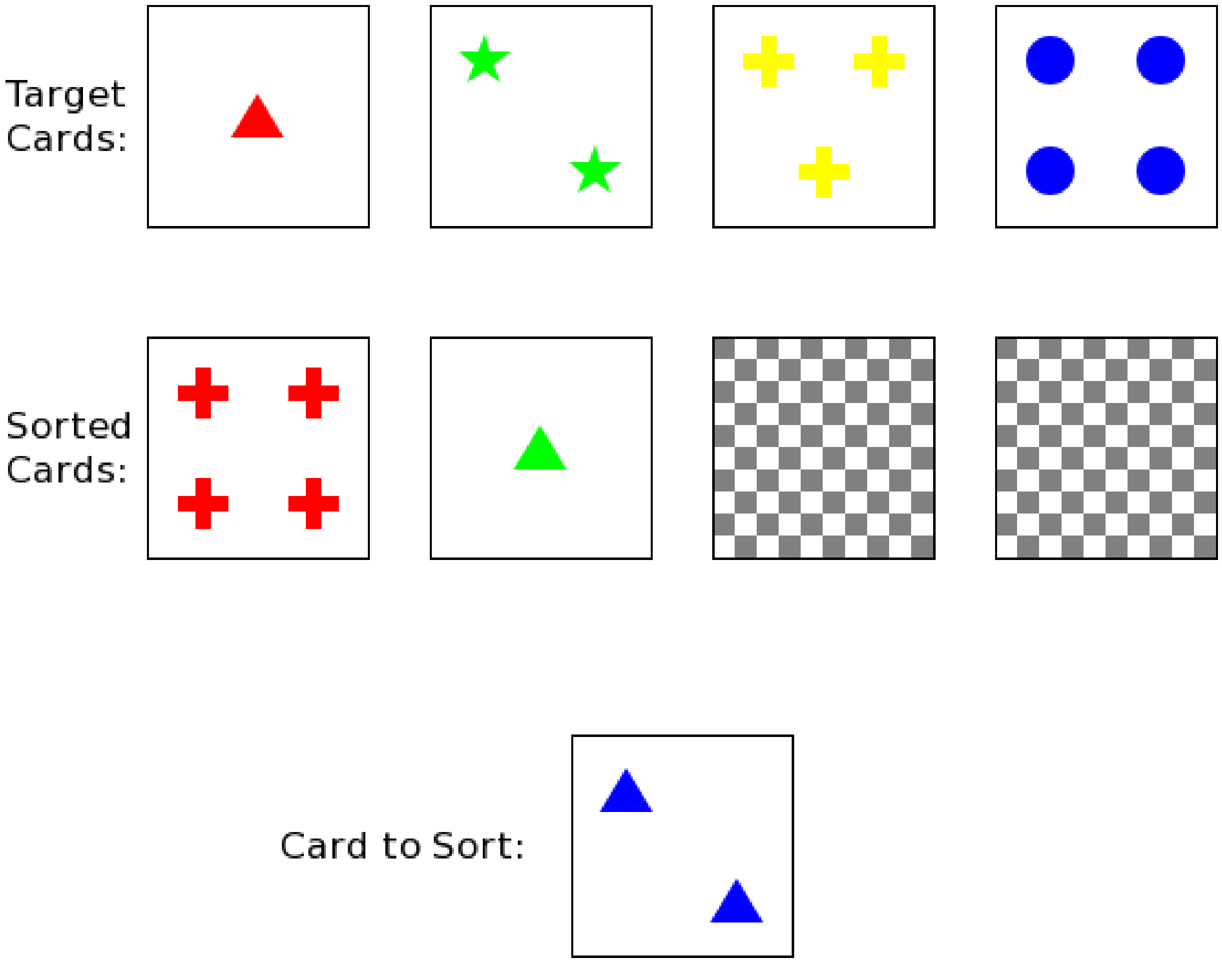

Applications of the Architecture
2: Wisconsin Card Sorting Test

	Cats.	PPC	PPR	Set Loss
Ctrl	3.9 (0.3)	7.4 (0.8)	1.3 (0.6)	1.0 (0.3)
IM	2.6 (0.6)	10.6 (1.7)	2.9 (0.9)	2.6 (0.7)
Decreasing *Attentiveness*				
Mon = 1.4 Att = 0.4	4.0 (0.7)	6.9 (1.4)	1.3 (0.9)	1.1 (1.1)
Mon = 1.4 Att = 0.3	2.3 (1.3)	5.6 (4.0)	1.7 (1.4)	3.3 (1.8)
Mon = 1.4 Att = 0.2	0.7 (0.8)	3.9 (3.9)	3.1 (1.6)	6.1 (1.7)

Applications of the Architecture 3: Tower of London

Current Configuration:

Desired Configuration:

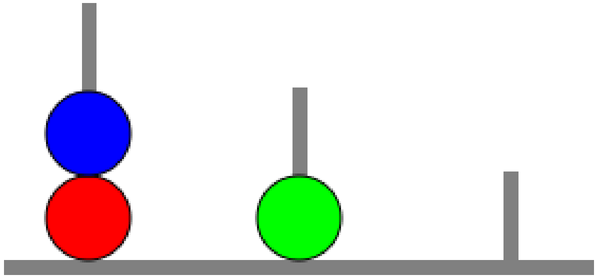

Towards an Account of Strategy Generation:

The “Domino” Architecture

The "Domino" Framework (Fox & Das, 2000)

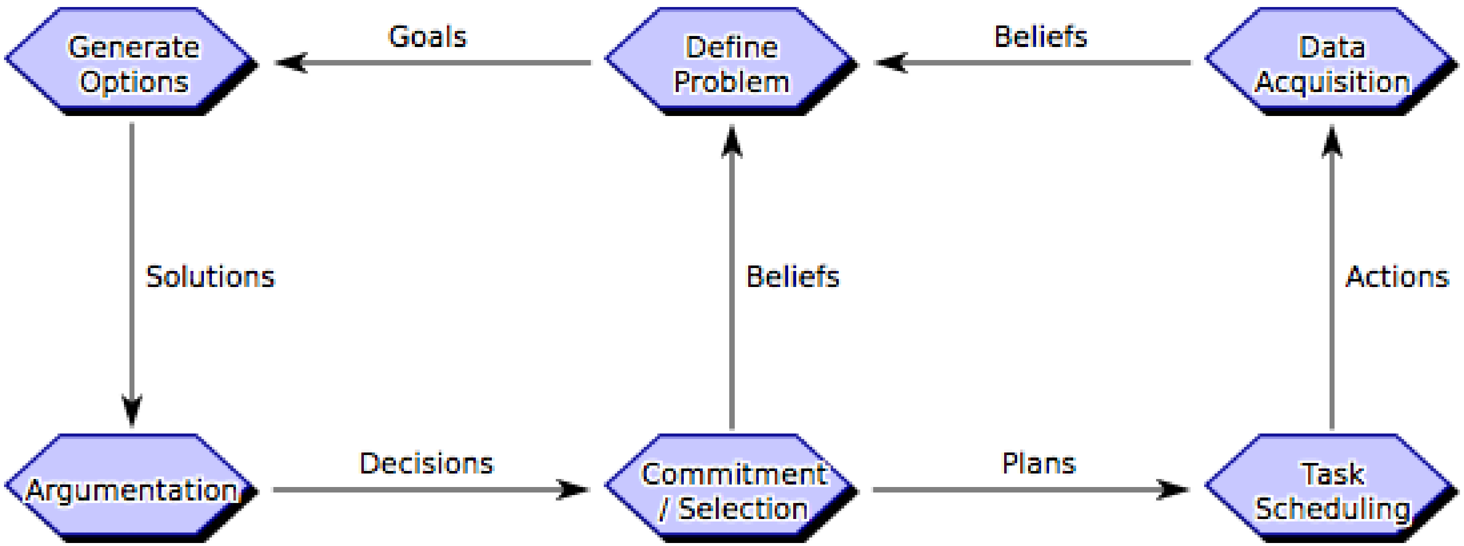

WCST Processing within the Domino: 1

- Initial processing:
  - Define problem $\Rightarrow$ problem is to sort the card
  - Generate options $\Rightarrow$ match to colour, number or form
  - No strong arguments either way $\Rightarrow$ choose at random

- Within the wider architecture:
  - Prediction of feedback may be associated with an action
  - Example: We think the rule is “match to colour”, so we place the current card (two blue triangles) under four blue circles and predict that feedback will be positive

WCST Processing within the Domino: 2

- Following positive feedback:
  - Feedback is as predicted - strategy generation is not invoked
- Following negative feedback:
  - Feedback is not as predicted - strategy generation is invoked and the current sorting method is rejected
  - Feedback provides an argument against the previous sorting method
  - Select a method with positive arguments (or no negative arguments)

Comparison with Other Approaches

- Comparison with ACT-R:
  - ACT-R is built on a production system model of memory and problem solving
  - It can simulate RNG and WCST performance, but …
  - No separable processes of monitoring, task setting, etc.,
  - More generally, no special-purpose mechanisms to support EF
- Comparison with EPIC (Executive Process - Interactive Control):
  - Built largely on requirements of PRP effects
  - Executive component consists of preprogrammed strategies to avoid simultaneous use of limited perceptual/motor resources
  - Again, no specific mechanisms to address response inhibition, etc.

Summary and Conclusions: 1

- The schema-based IA model of routine action selection provides a viable account of reactive behaviour

- The account is supported by error data from neurologically healthy individuals and a variety of groups of neuropsychological patients

- The IA model may be modulated by a structured higher-level system in order to generate non-routine, deliberative behaviour

Summary and Conclusions: 2

- Components of the higher-level system include:
  - Monitoring, working memory, strategy generation, and processes related to active maintenance of and attentiveness to buffer elements
- Strategy generation need not be considered a homunculus:
  - The domino framework provides a viable decomposition

Acknowledgements

- Thanks to:
  - Tim Shallice
  - John Fox
  - Pete Yule
  - Eddy Davelaar
  - Karolina Wutke
  - Hala Alireza