

# Learning Flexible Goal-Directed Behavior

Christian Balkenius

Lund University Cognitive Science









**3rd ANNUAL INTERNATIONAL WORKSHOP**

**AUGUST 4 & 5 • 2003**

**EPIGENETIC ROBOTICS**

**MODELING COGNITIVE DEVELOPMENT IN ROBOTIC SYSTEMS**

**SPEAKERS**  
**GYÖRGY GERGELY**  
 INSTITUTE FOR PSYCHOLOGICAL RESEARCH, HUNGARIAN ACADEMY OF SCIENCES BUDAPEST, HUNGARY  
**ROD GRUPEN**  
 LABORATORY FOR PERCEPTUAL ROBOTICS UNIVERSITY OF MASSACHUSETTS AMHERST, MA, USA  
**DEB ROY**  
 MEDIA LAB MIT, USA

**www.epigenetic-robotics.org**  
**BOSTON UNIVERSITY BOSTON MA USA**

**BOSTON**

Sponsored by European Commission Research Laboratory, Japan Workshop Co-Sponsor, The Cognitive & Neural Systems Department, Boston University

Epilab 2003 will be held after the Cognitive Science Society Meeting

Deadline for submission of papers and poster abstracts March 14, 2003

**4th annual August Robotics**

Tatyana Serezhnikova  
 Isabella Fedige Biomedical  
 Glass van Nifflaen Psychol  
 Jürgen Krambeck Roman  
 Jacqueline Roubi CARS D

ally

Sponsored by European Commission Research Laboratory, Japan Workshop Co-Sponsor, The Cognitive & Neural Systems Department, Boston University

www.epigenetic-robotics.org

March 1st, 2004 Deadline for submission of papers and posters

**2002 August 10 • 11**

**University of Edinburgh Scotland**

**Luc Steels**  
WJL Brussels & Sony CS Lab, Paris

**Cohwyn Trevarthen**  
University of Edinburgh, UK

**John Weng**  
Michigan State University, USA

Workshop chair: Christophe Prater  
 University of Minnesota Duluth, USA  
 christophe@uic.edu

Program chair: Yuzuru Omura,  
 Imperial College, UK  
 y.omura@ic.ac.uk

Treasurer: Shōichi Kuroki,  
 Communications Research Laboratory, Japan  
 shokuroki@nri.go.jp

Proceedings: Christian Balkenius,  
 Lund University, Sweden  
 christian.balkenius@lth.se

Local arrangements: Yuzuru Omura,  
 University of Edinburgh, UK  
 y.omura@uic.edu

Organizer:  
 Communications Research Laboratory, Japan


**2nd International Workshop on EPIGENETIC ROBOTICS**

**MODELING COGNITIVE DEVELOPMENT IN ROBOTIC SYSTEMS**


**www.epigenetic-robotics.org**

**Rob**


- ✦ Stimulus-Approach
- ✦ Stimulus-Response
- ✦ Contextual Inhibition




E. L. Thorndike  
(1874-1949)



E. C. Tolman  
(1886-1959)




E. L. Thorndike  
(1874-1949)



E. C. Tolman  
(1886-1959)


# Reactive Behavior





E. L. Thorndike  
(1874-1949)

Reactive Behavior




E. C. Tolman  
(1886-1959)

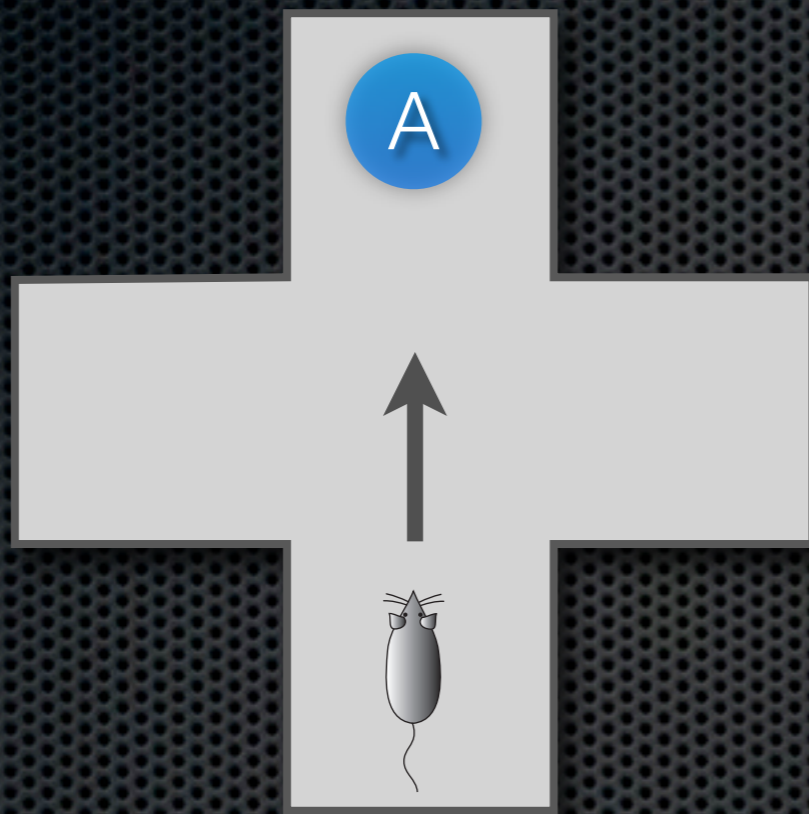
Purposive Behavior




Mackintosh, 1983



Mackintosh, 1983




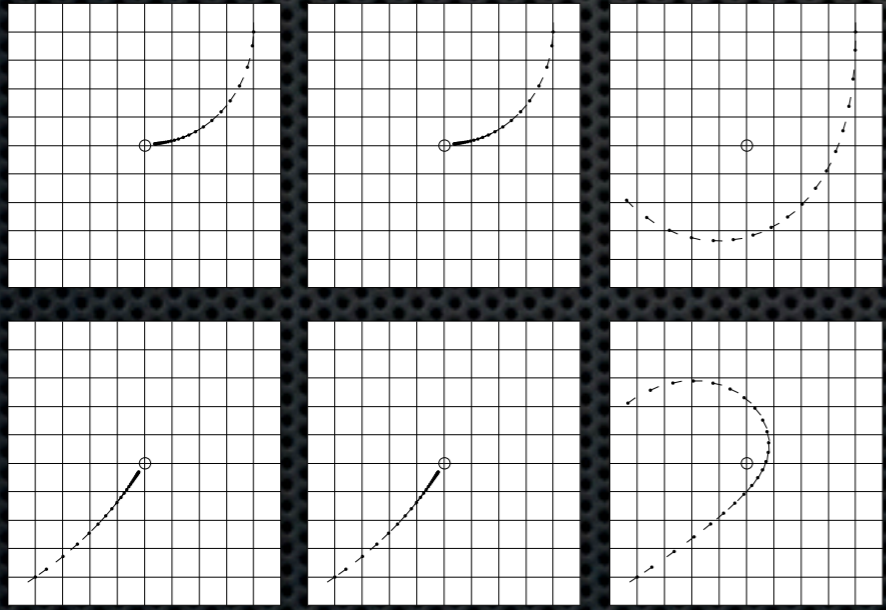

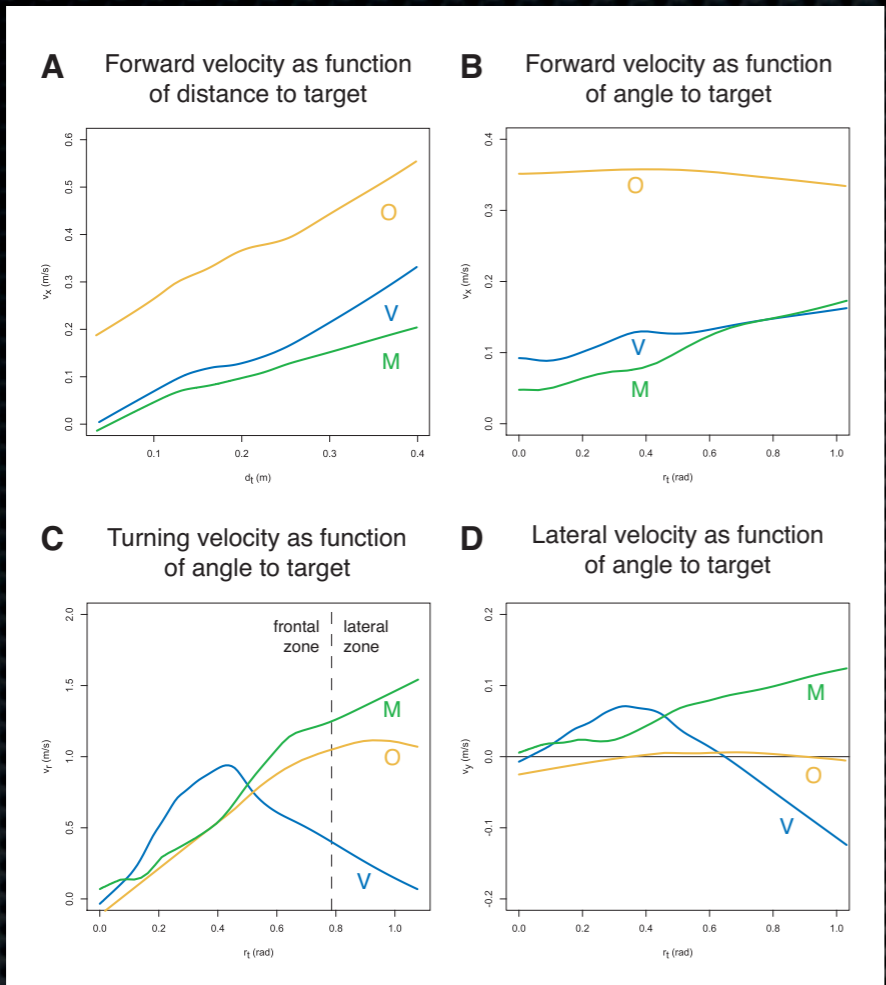
Stimulus-Response

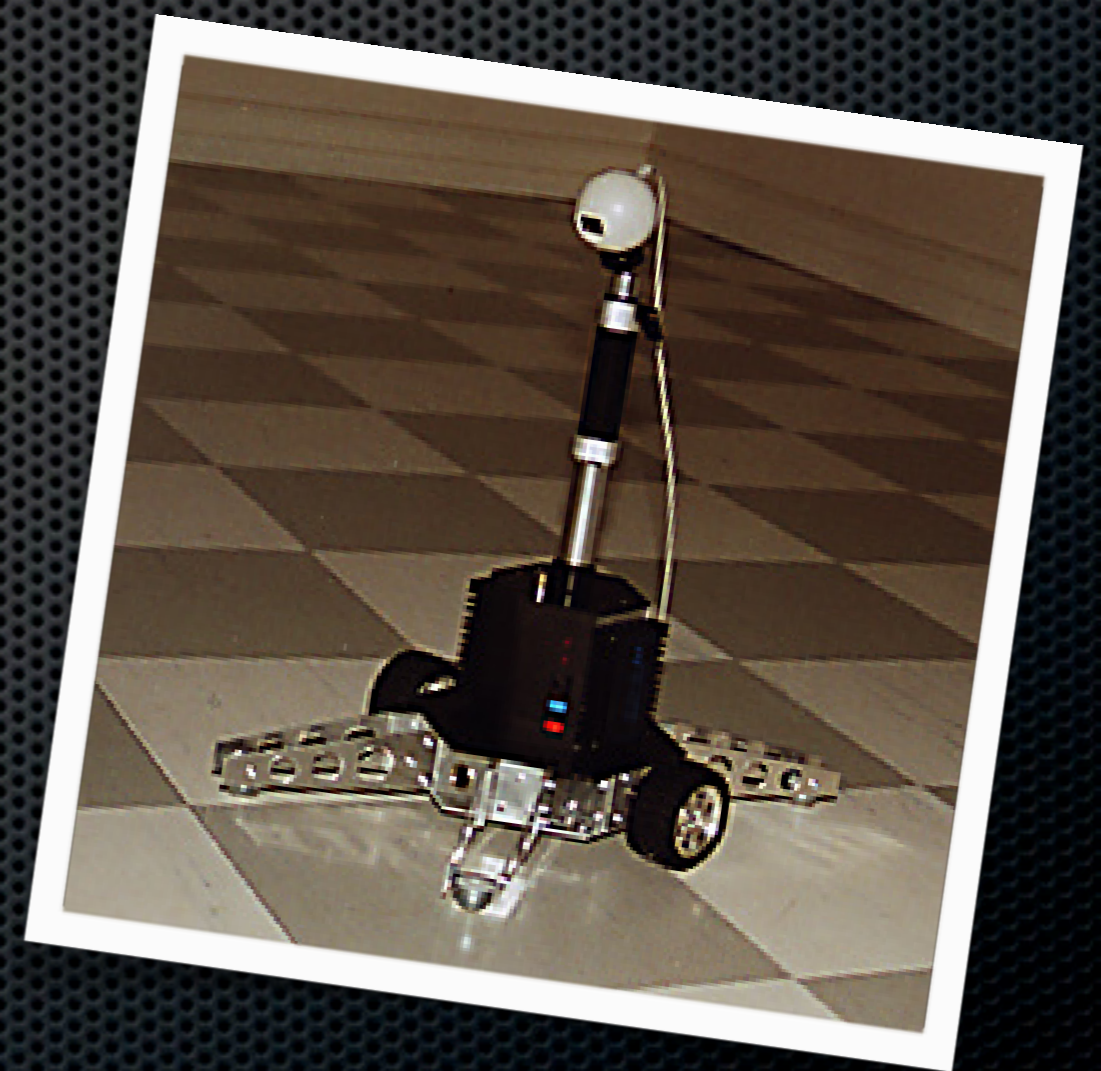



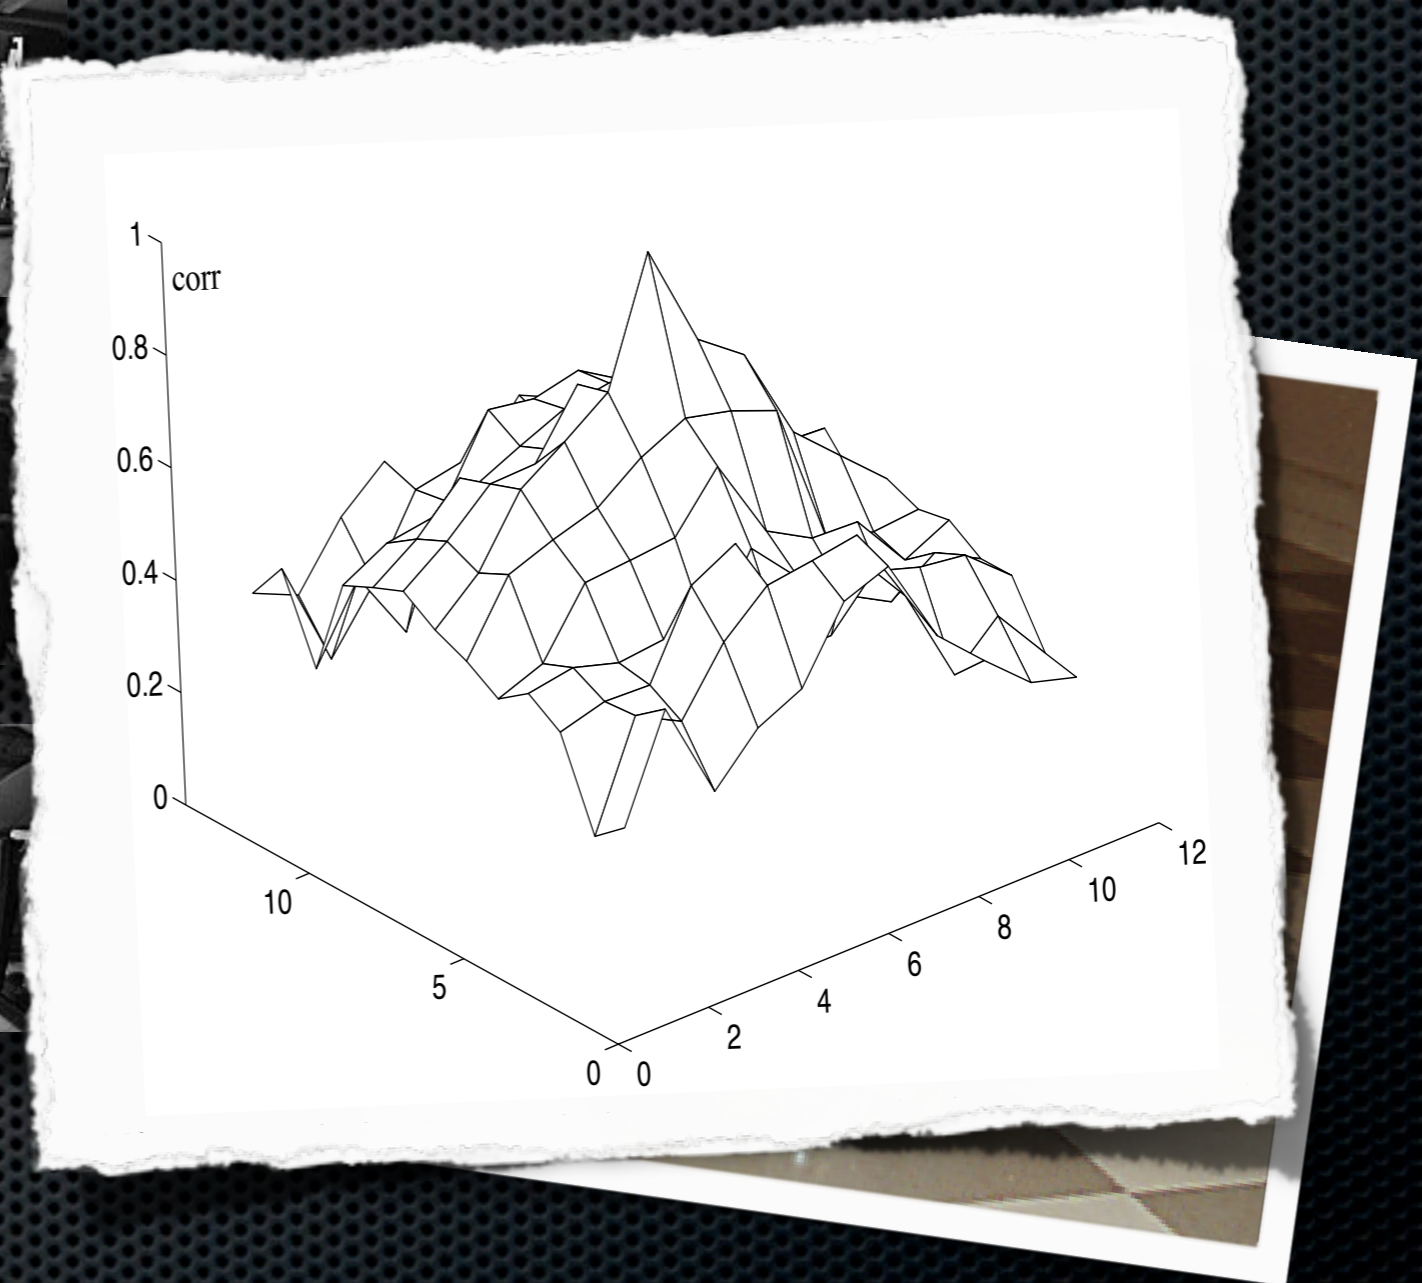

Stimulus-Approach

Mackintosh, 1983




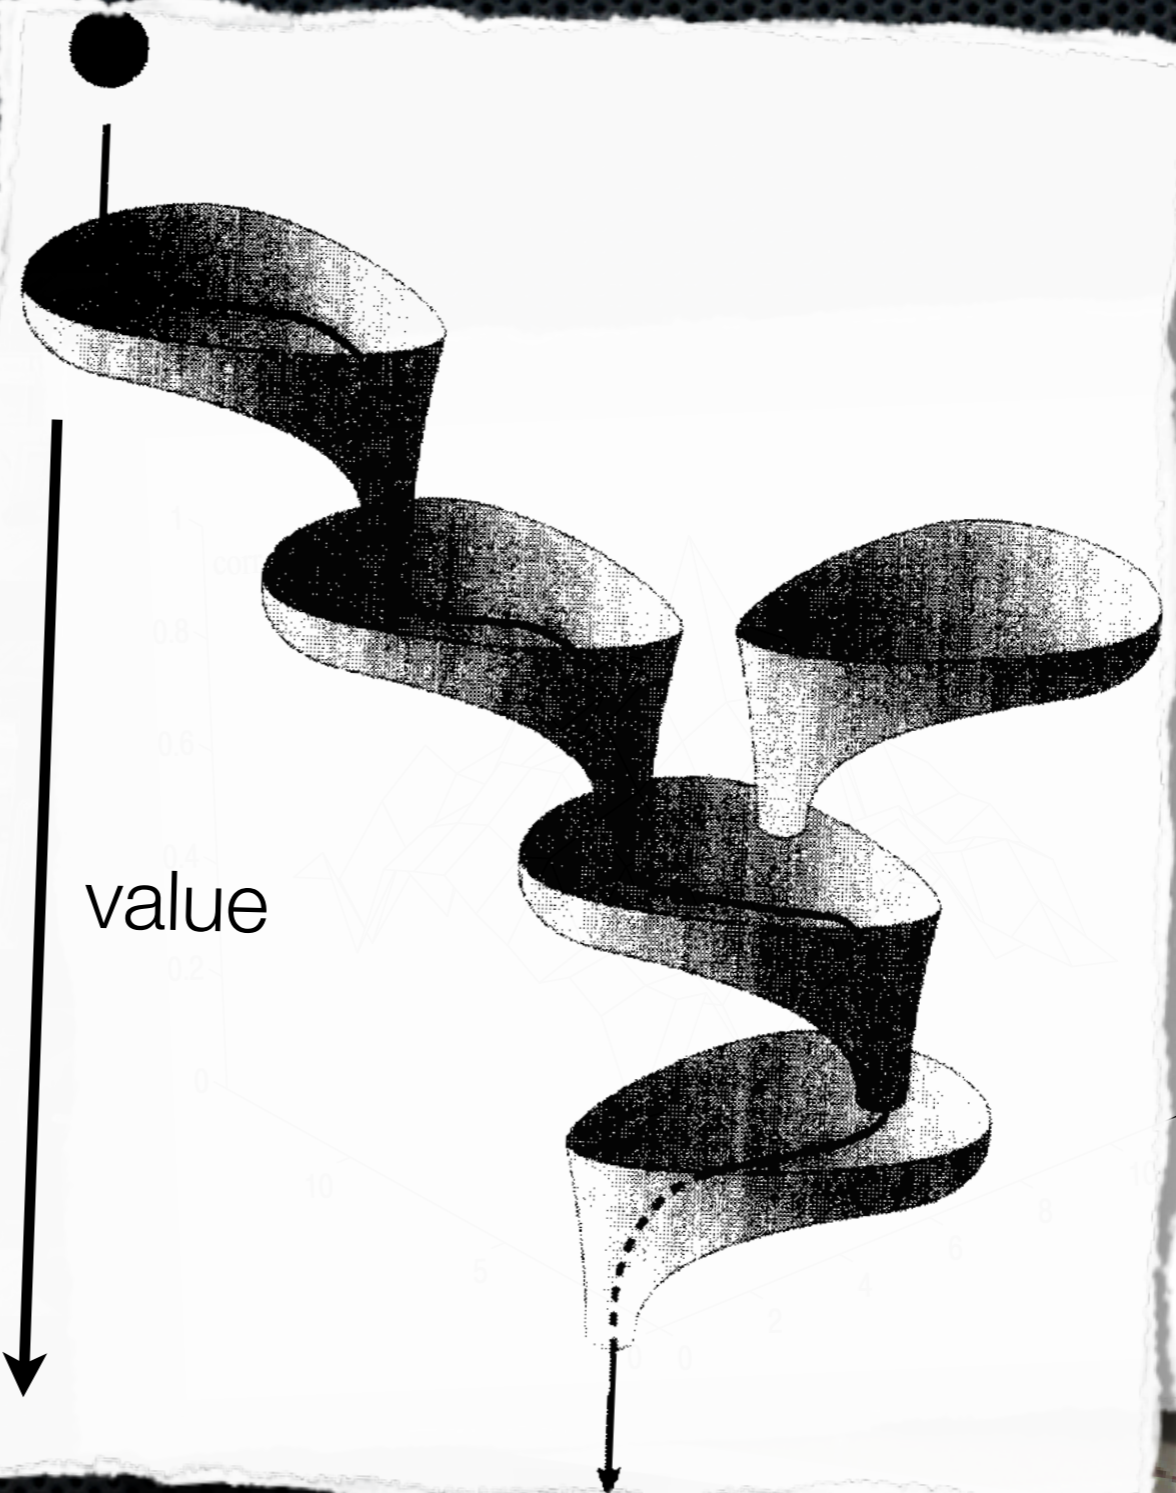








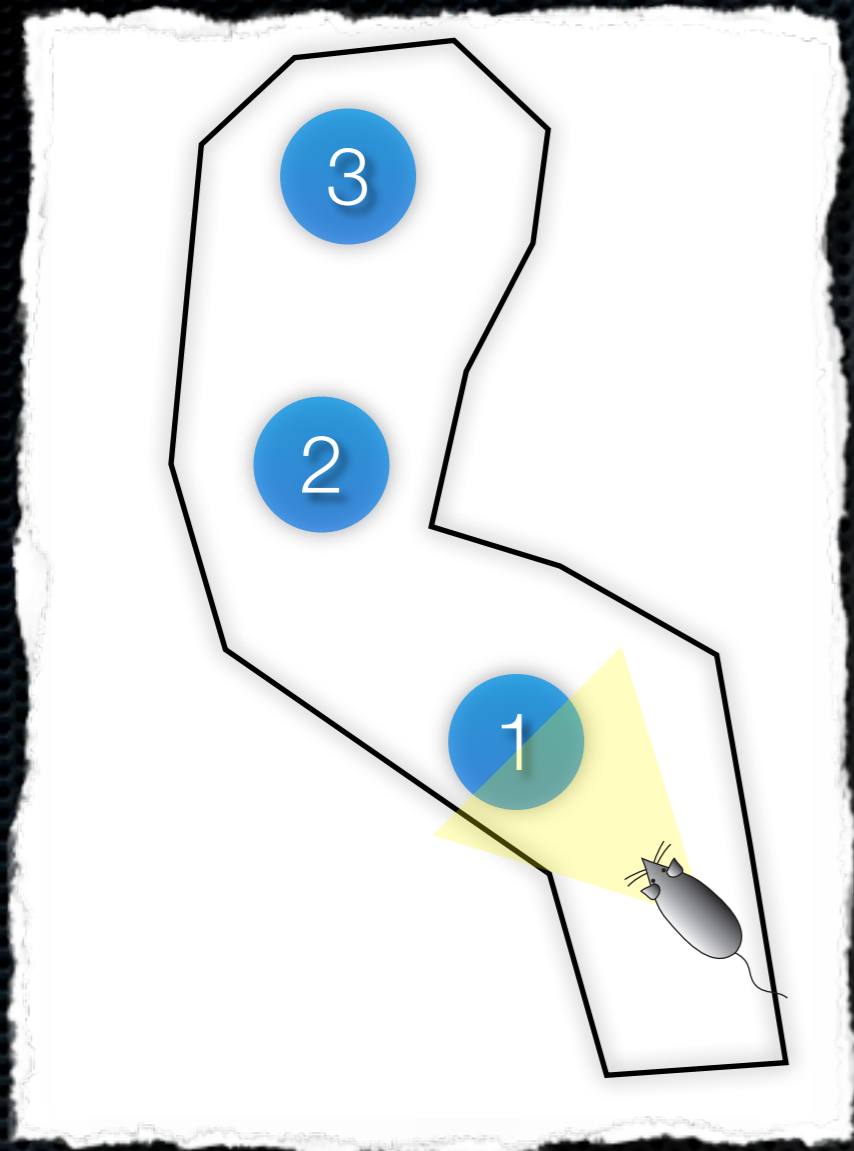




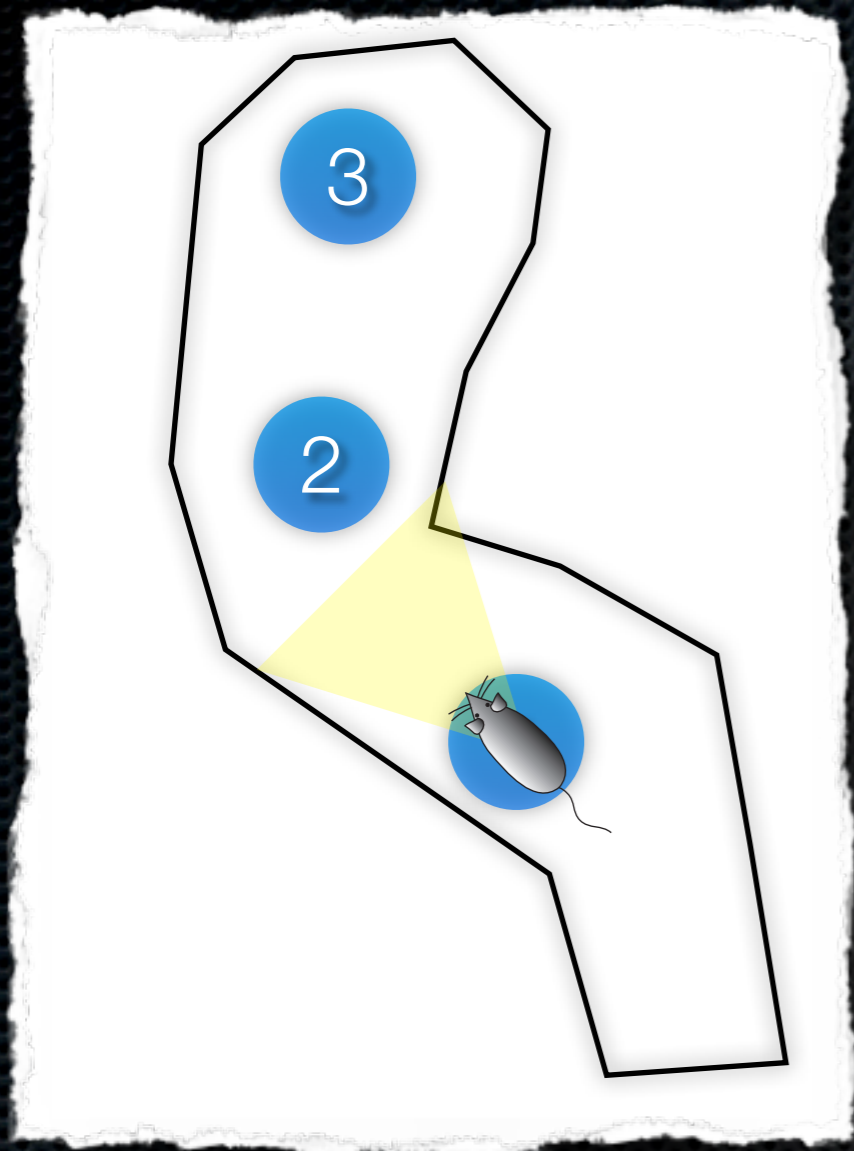





Stimulus-Approach



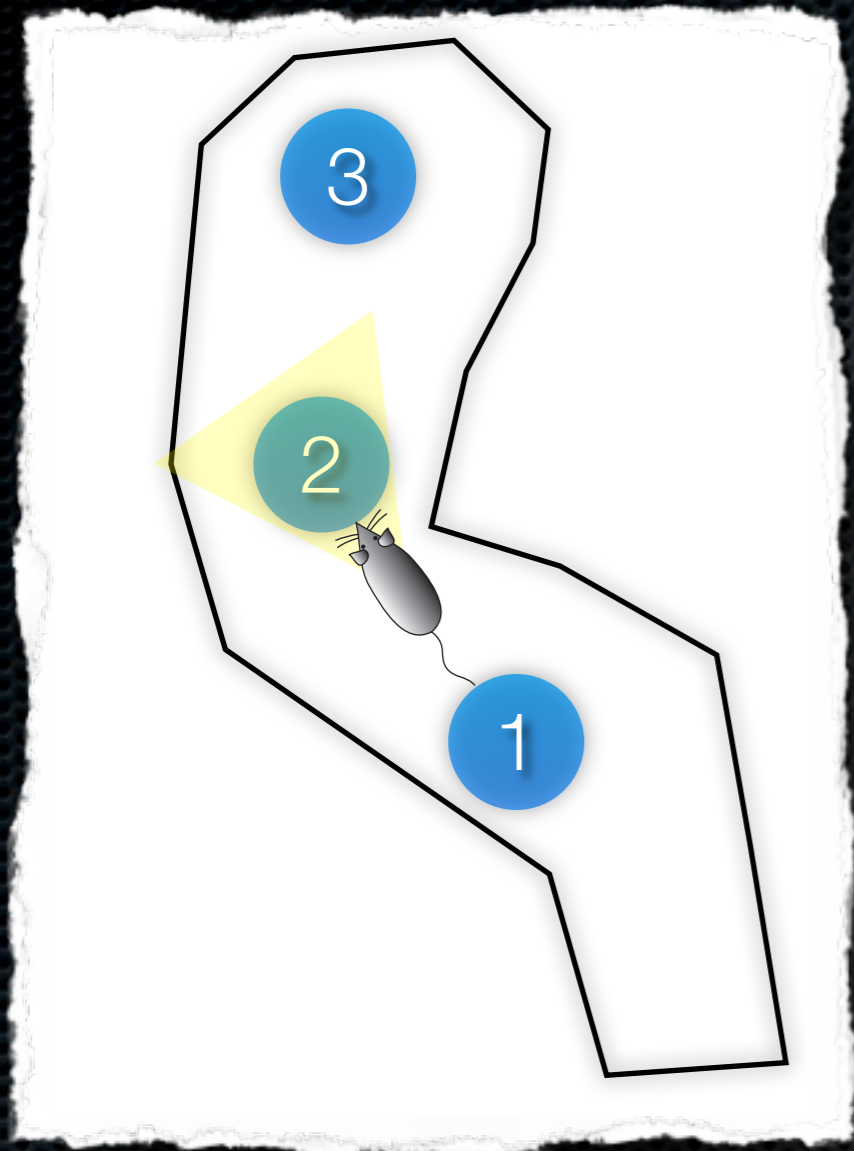
Stimulus-Approach




Stimulus-Approach



Stimulus-Approach




Stimulus-Approach



Stimulus-Approach

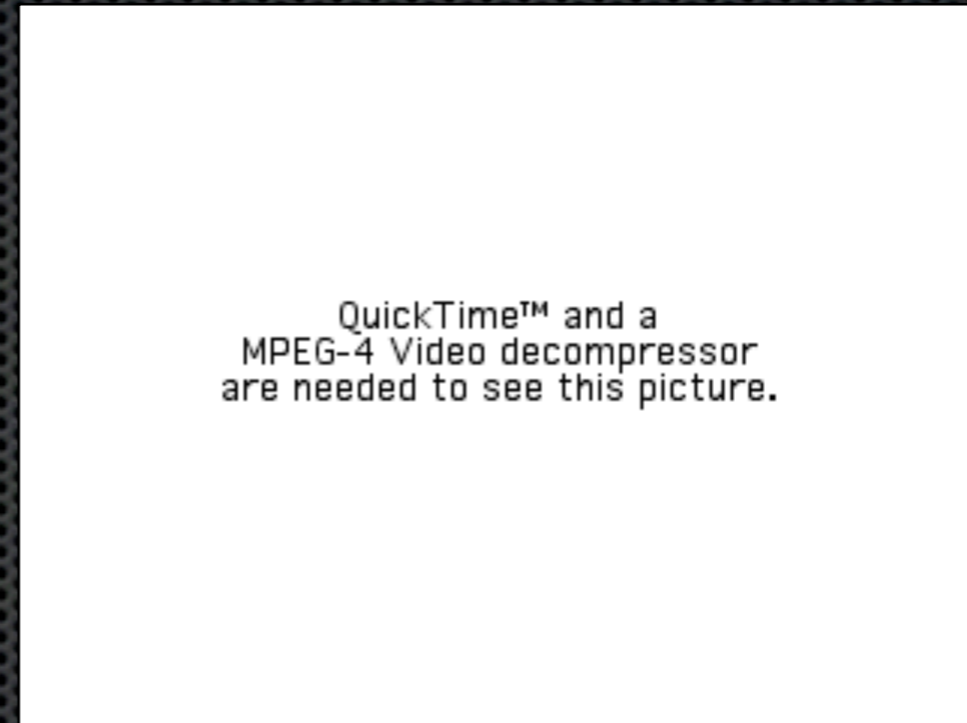





Stimulus-Approach

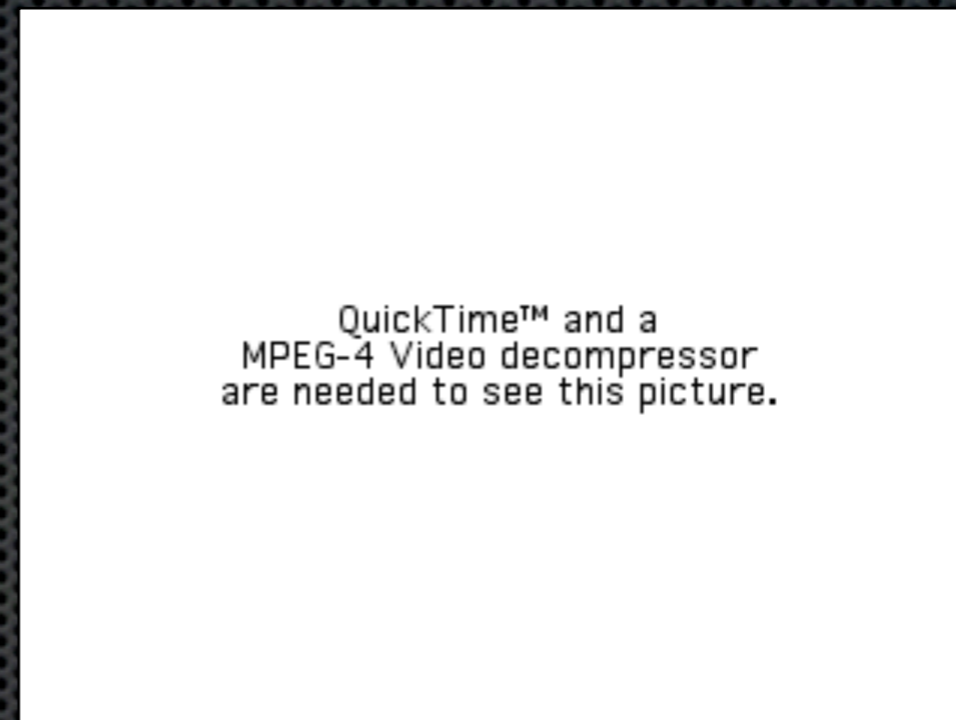
# Switching Control

eye regions



# Switching Control

eye regions




orienting

Balkenius & Kopp, 1997

# Switching Control

eye regions



saccades


orienting

QuickTime™ and a  
MPEG-4 Video decompressor  
are needed to see this picture.

Balkenius & Kopp, 1997

# Switching Control

eye regions



QuickTime™ and a  
MPEG-4 Video decompressor  
are needed to see this picture.

fixation / smooth pursuit

saccades

orienting

Balkenius & Kopp, 1997



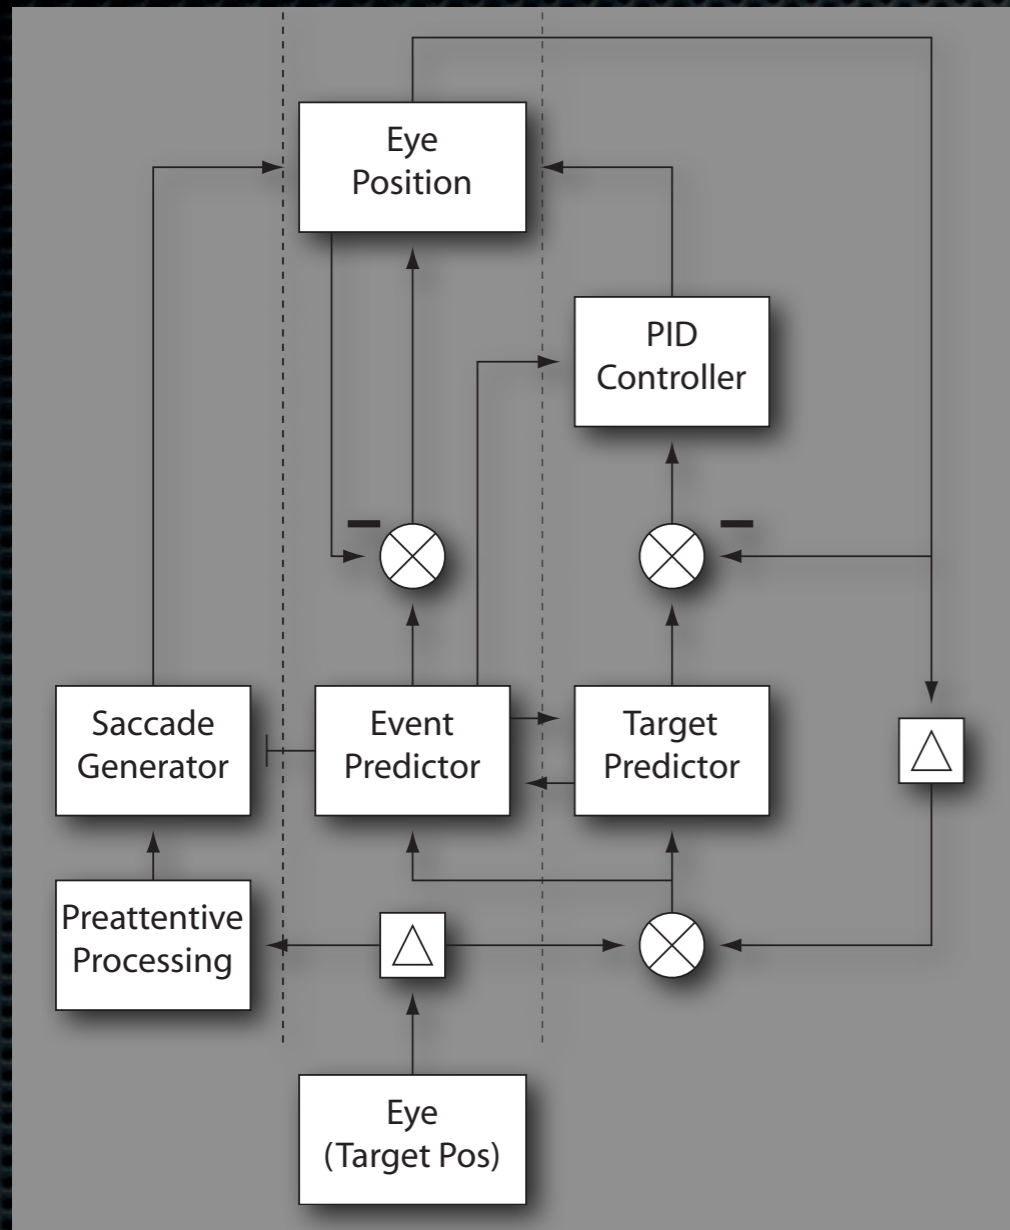


BabyLab

# Smooth Pursuit Eye-Movements




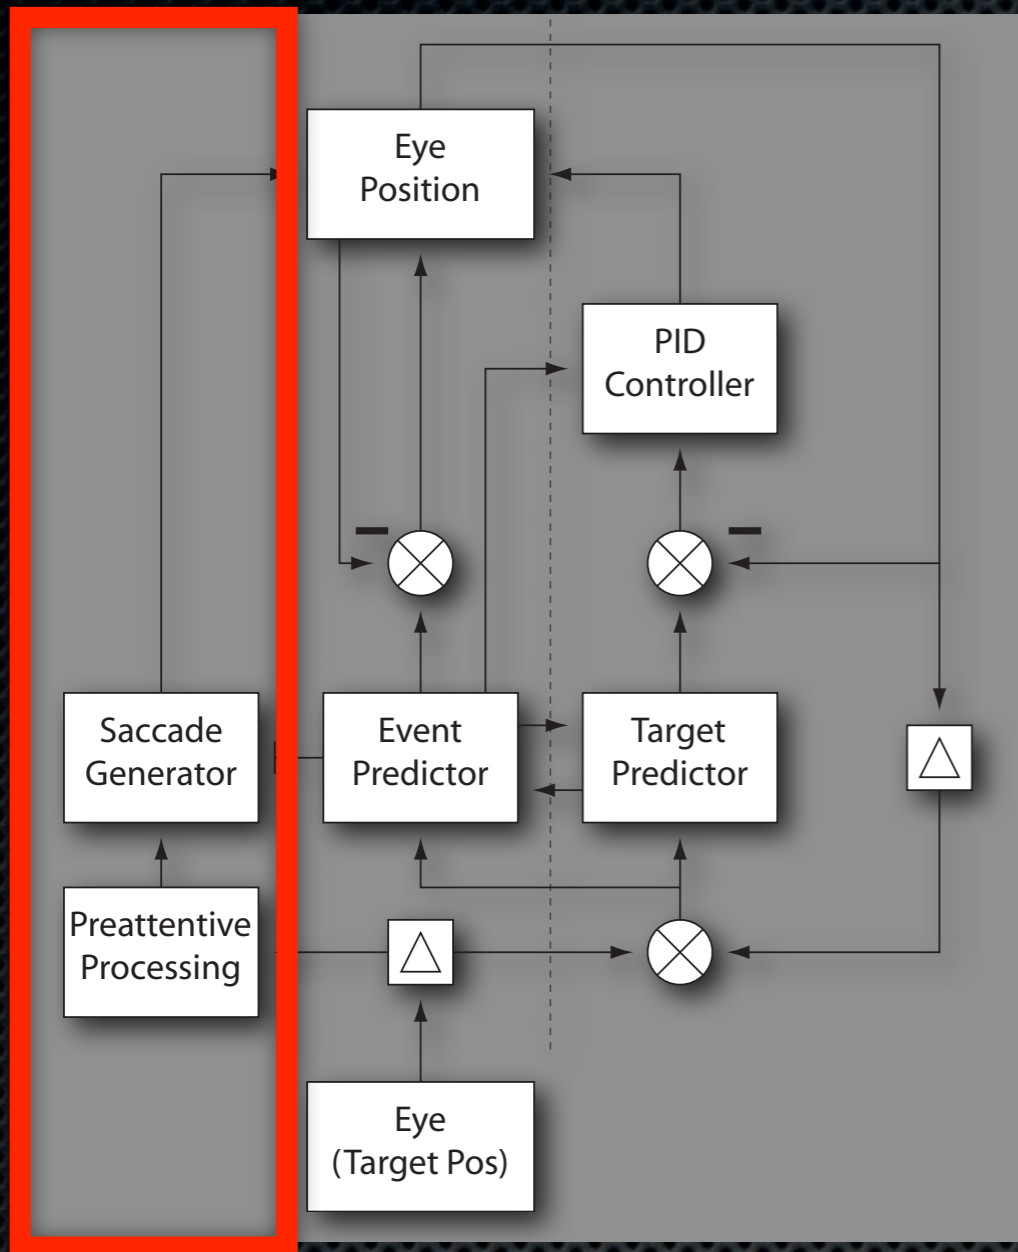




Balkenius & Johansson, 2005



# Orienting




$F_0$



$F_1$



$F_2$




$F_3$

$$S(x, y) = G(x, y) * \sum_m \theta_m F_m(x, y)$$

Balkenius, Åström, Eriksson, 2004

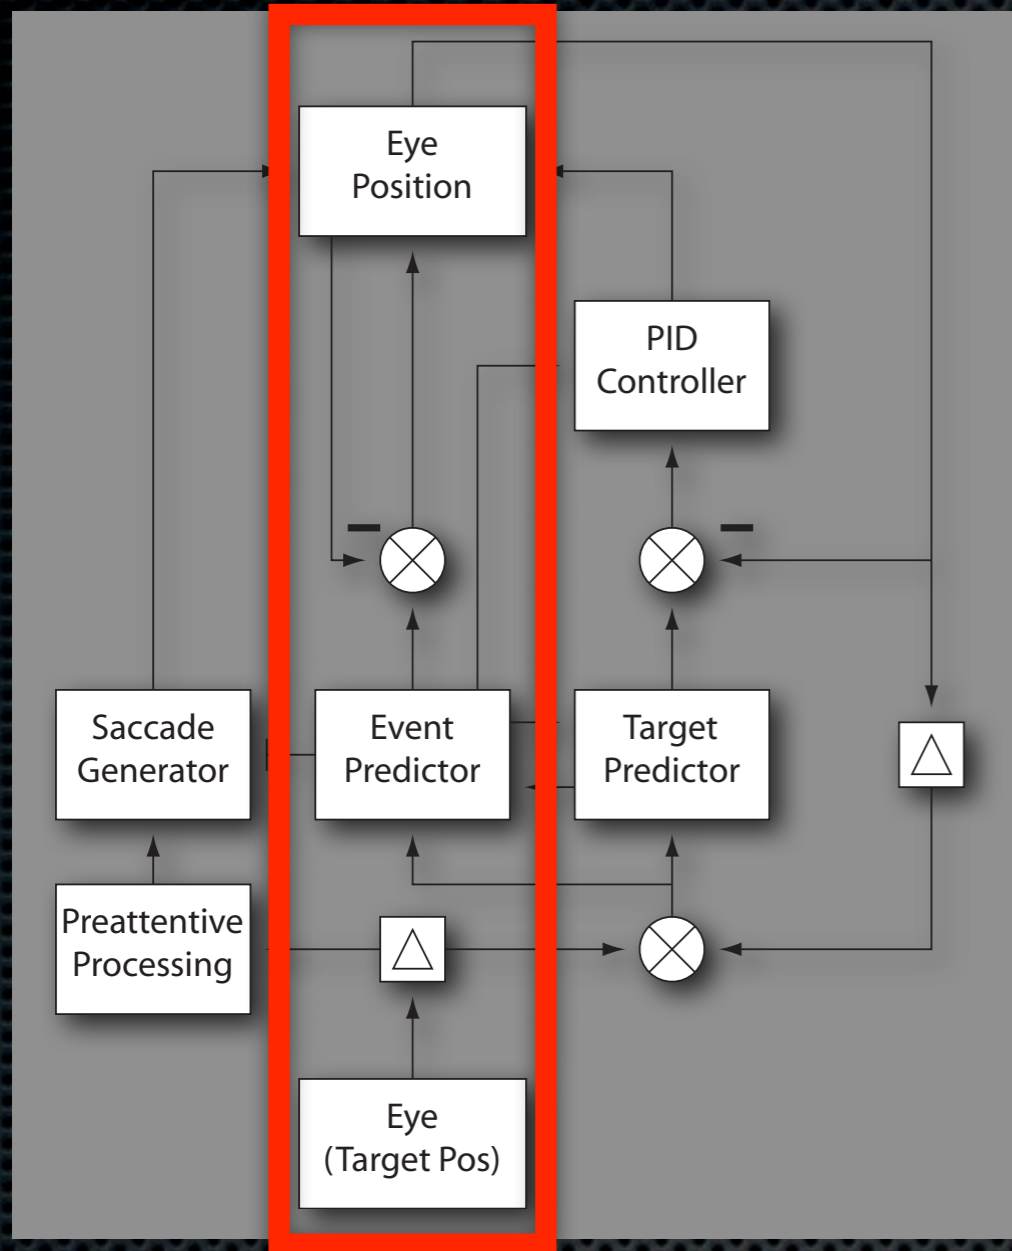
# Saccades



stimulus-response

Balkenius, 2000

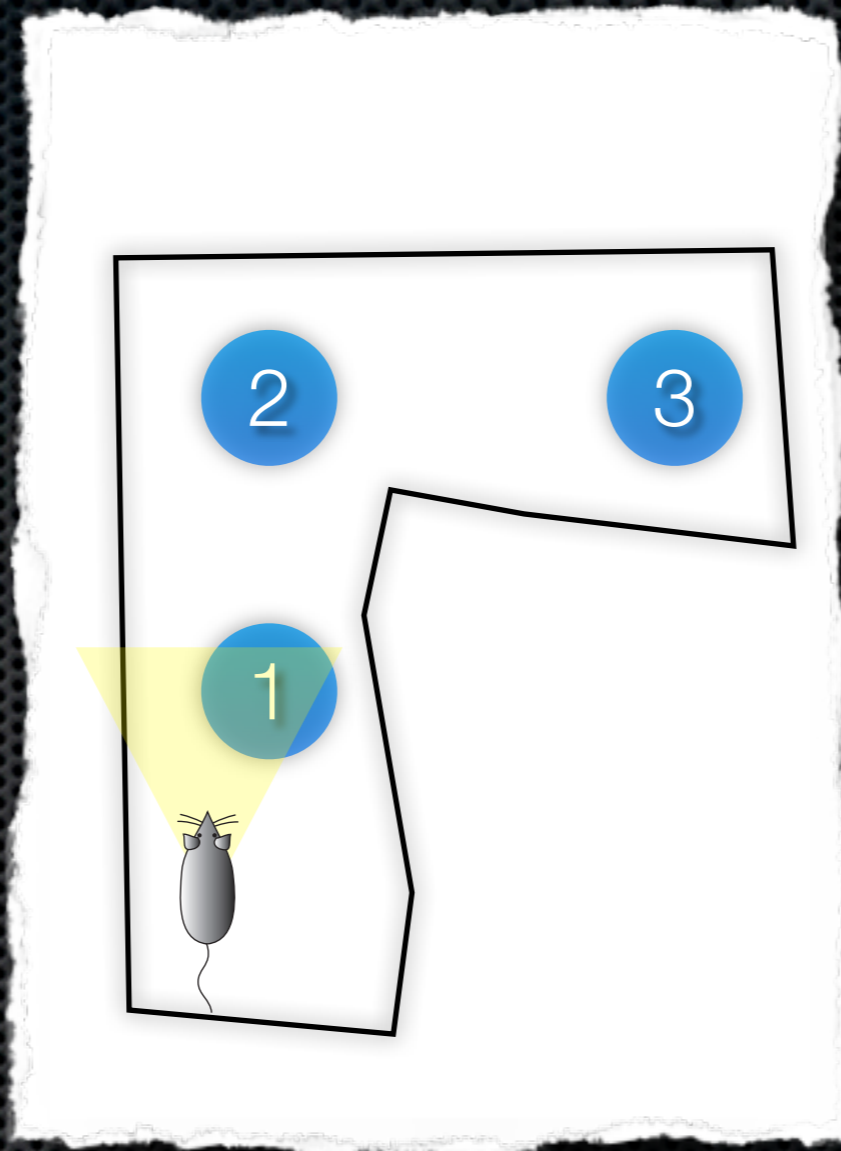

# Saccades



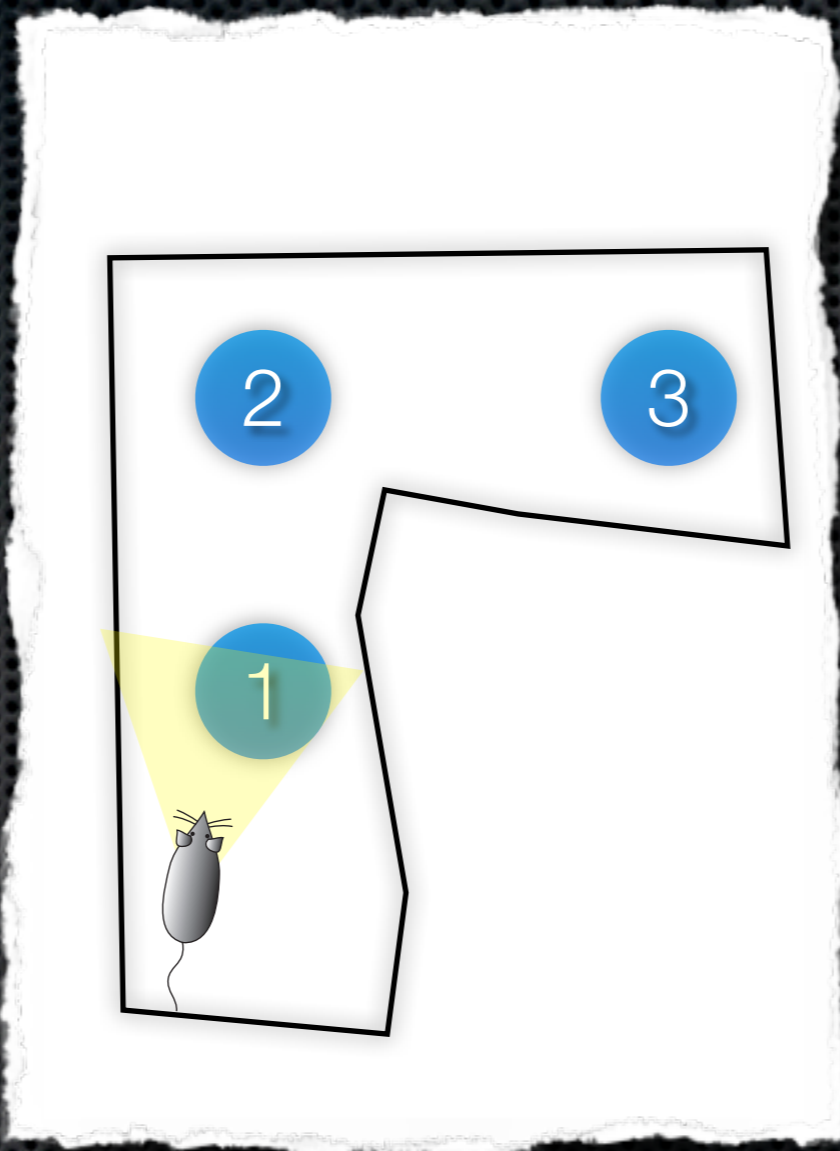

Stimulus + Gaze	Learned Saliency	Focus	Gain
+		⊕	
S*	R*	SG	
E(S*)	E(R*)	E	

stimulus-response



Balkenius, 2000



Stimulus-Approach

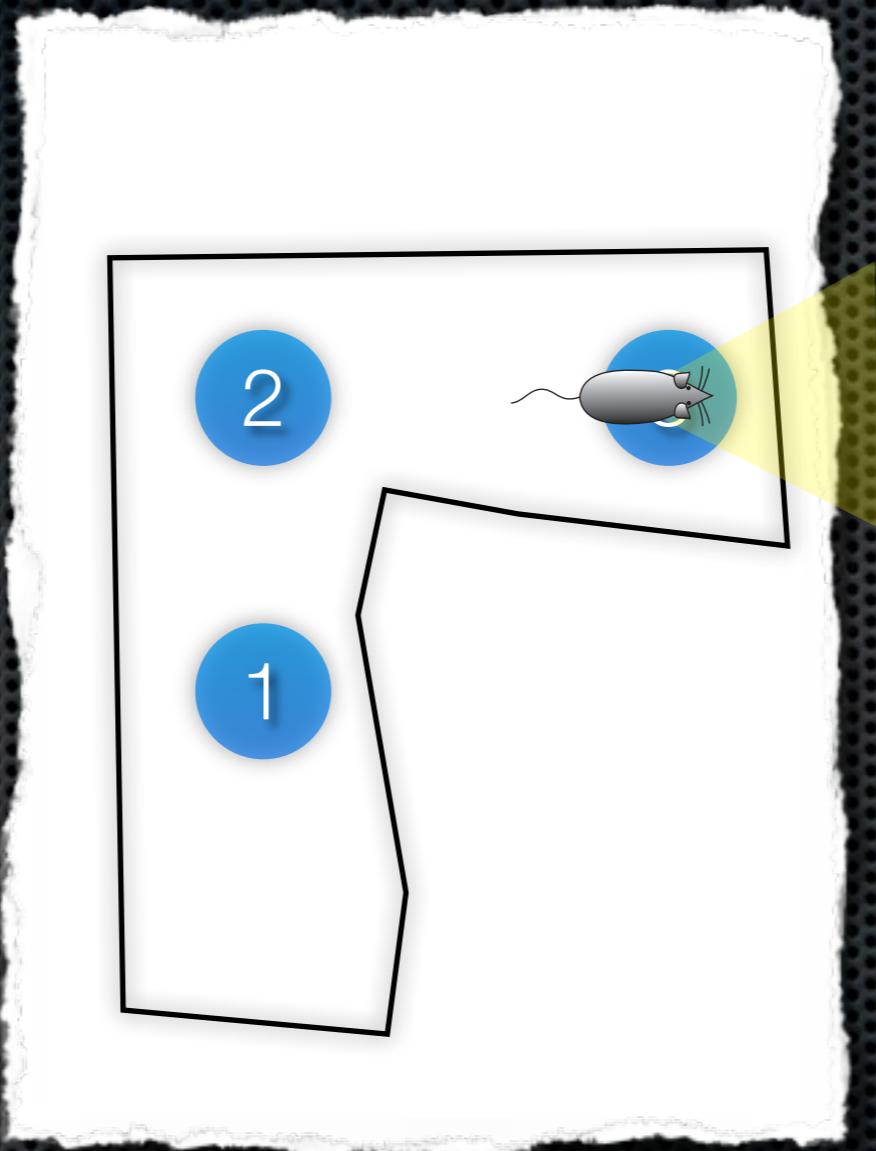



Stimulus-Approach




Stimulus-Approach

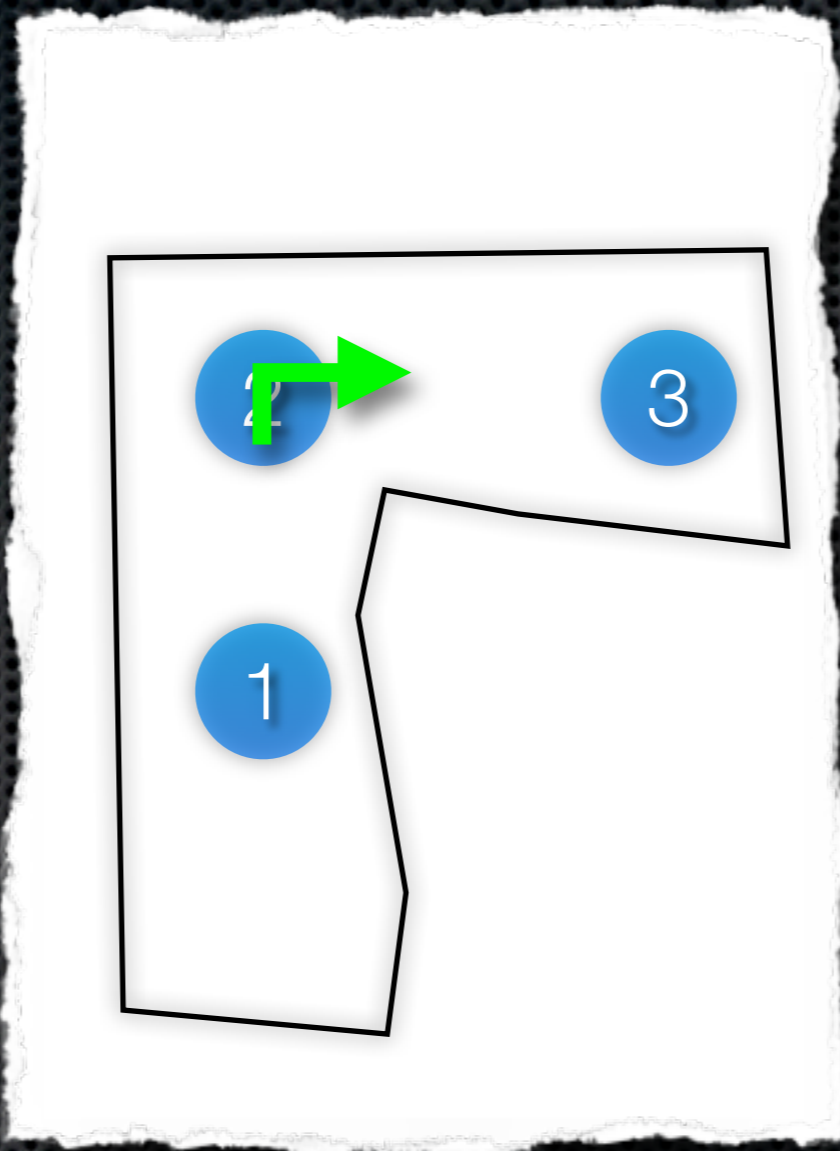




Stimulus-Approach



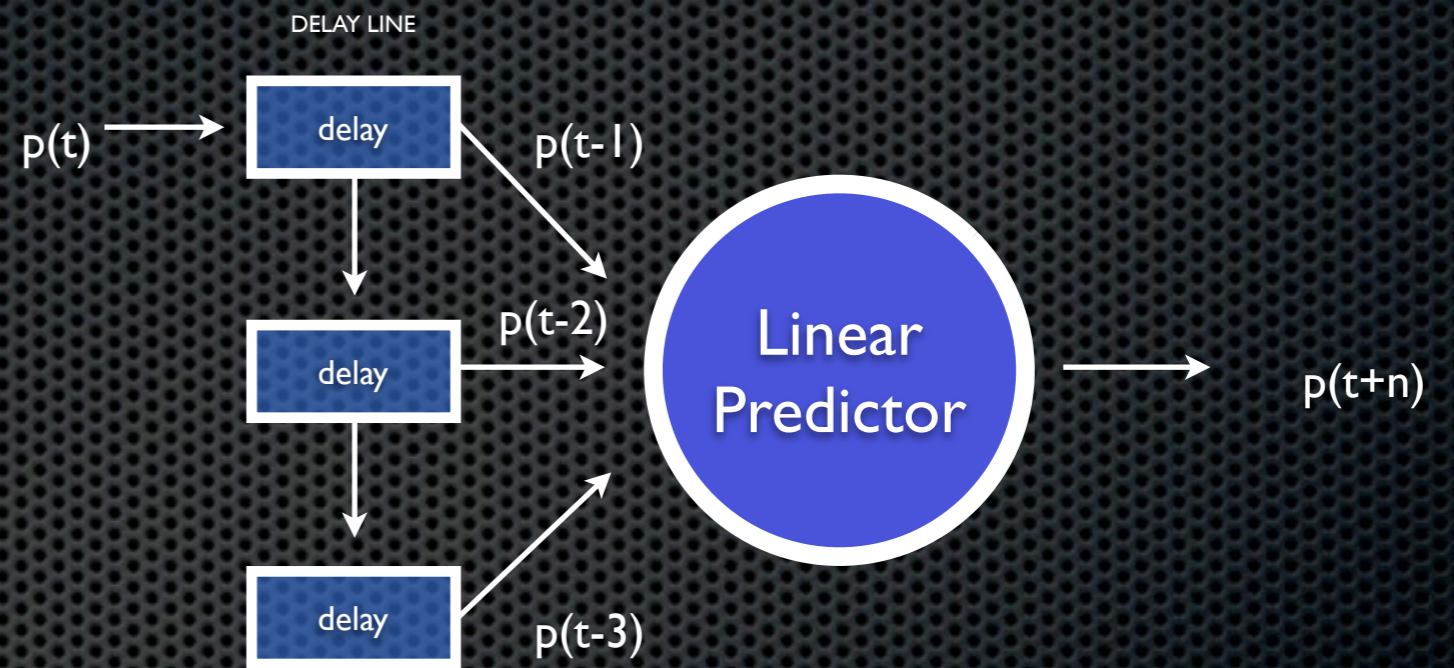
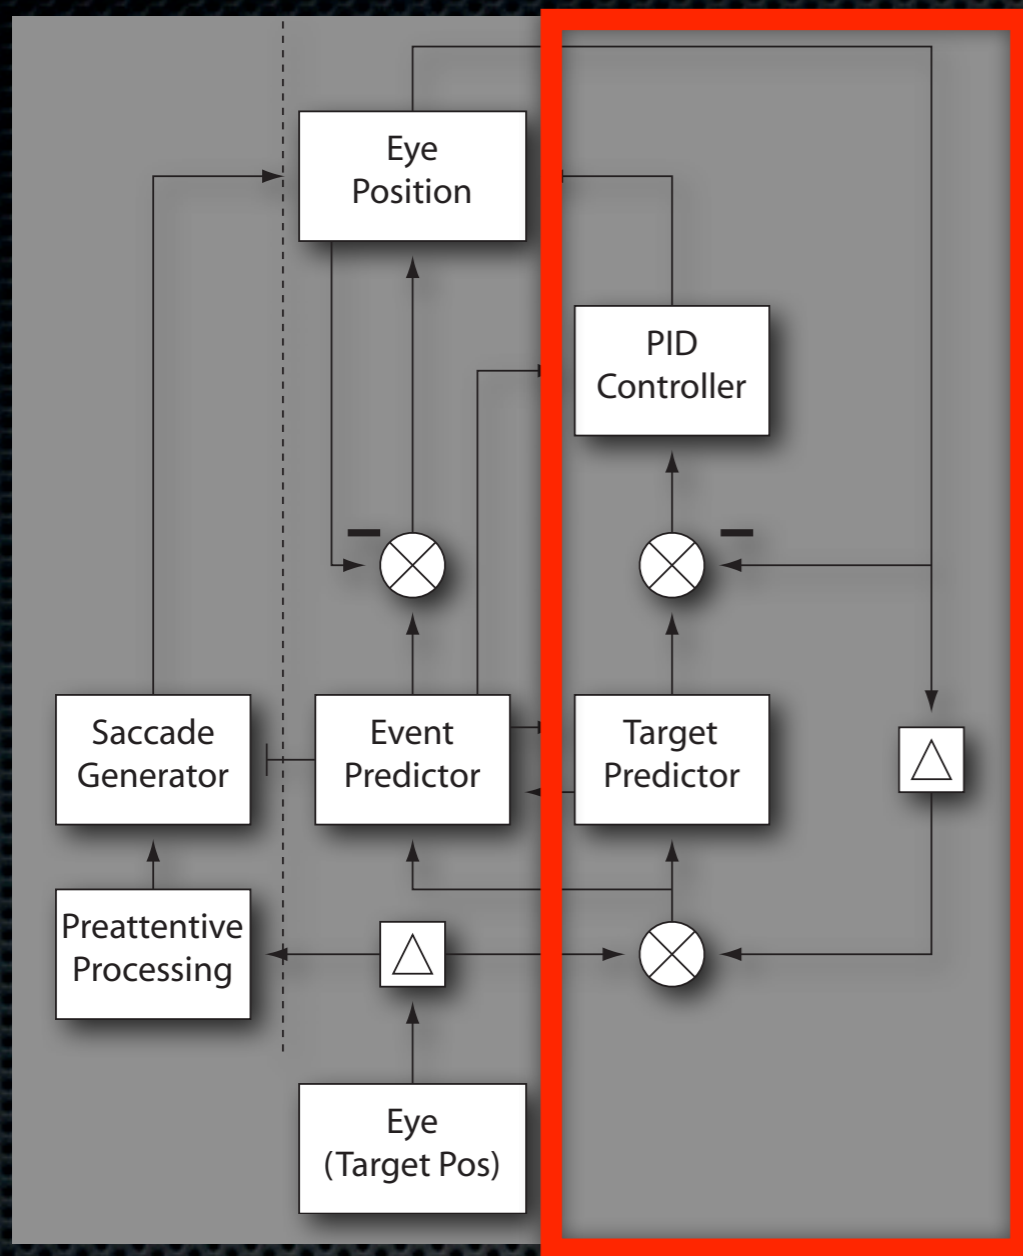
Stimulus-Approach



Stimulus-Response



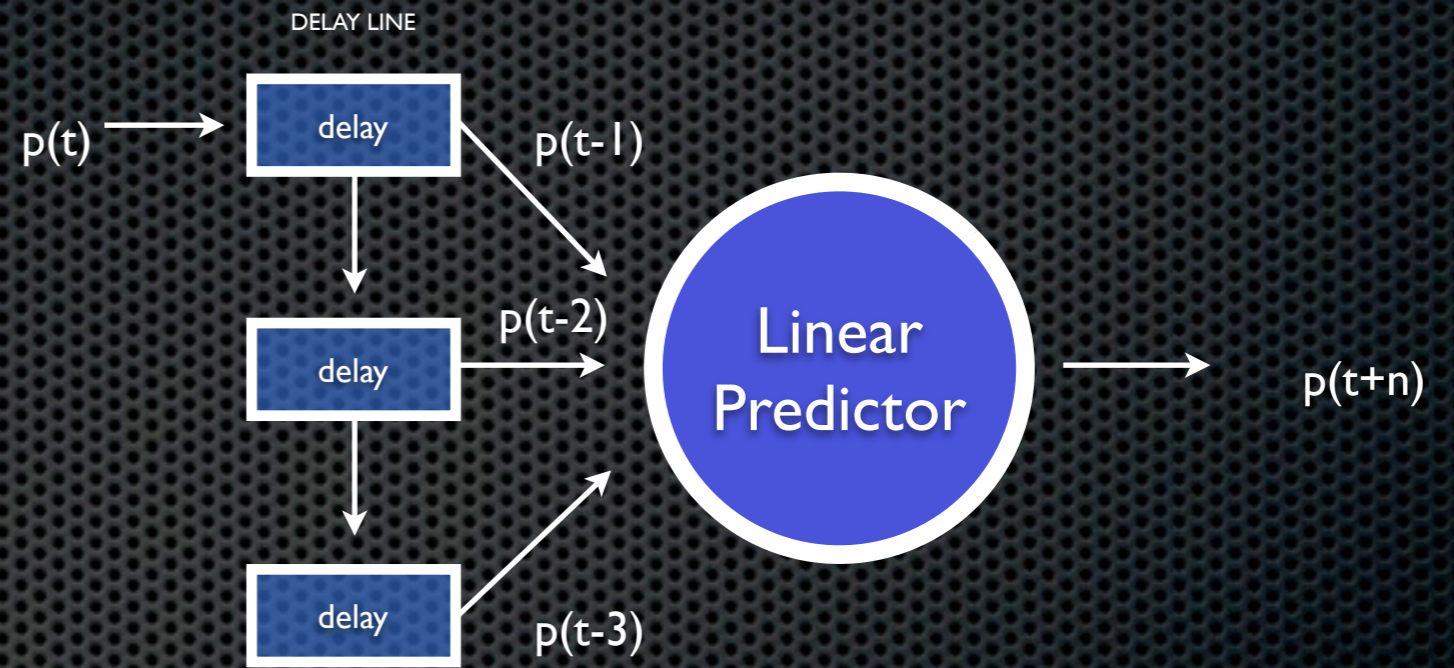
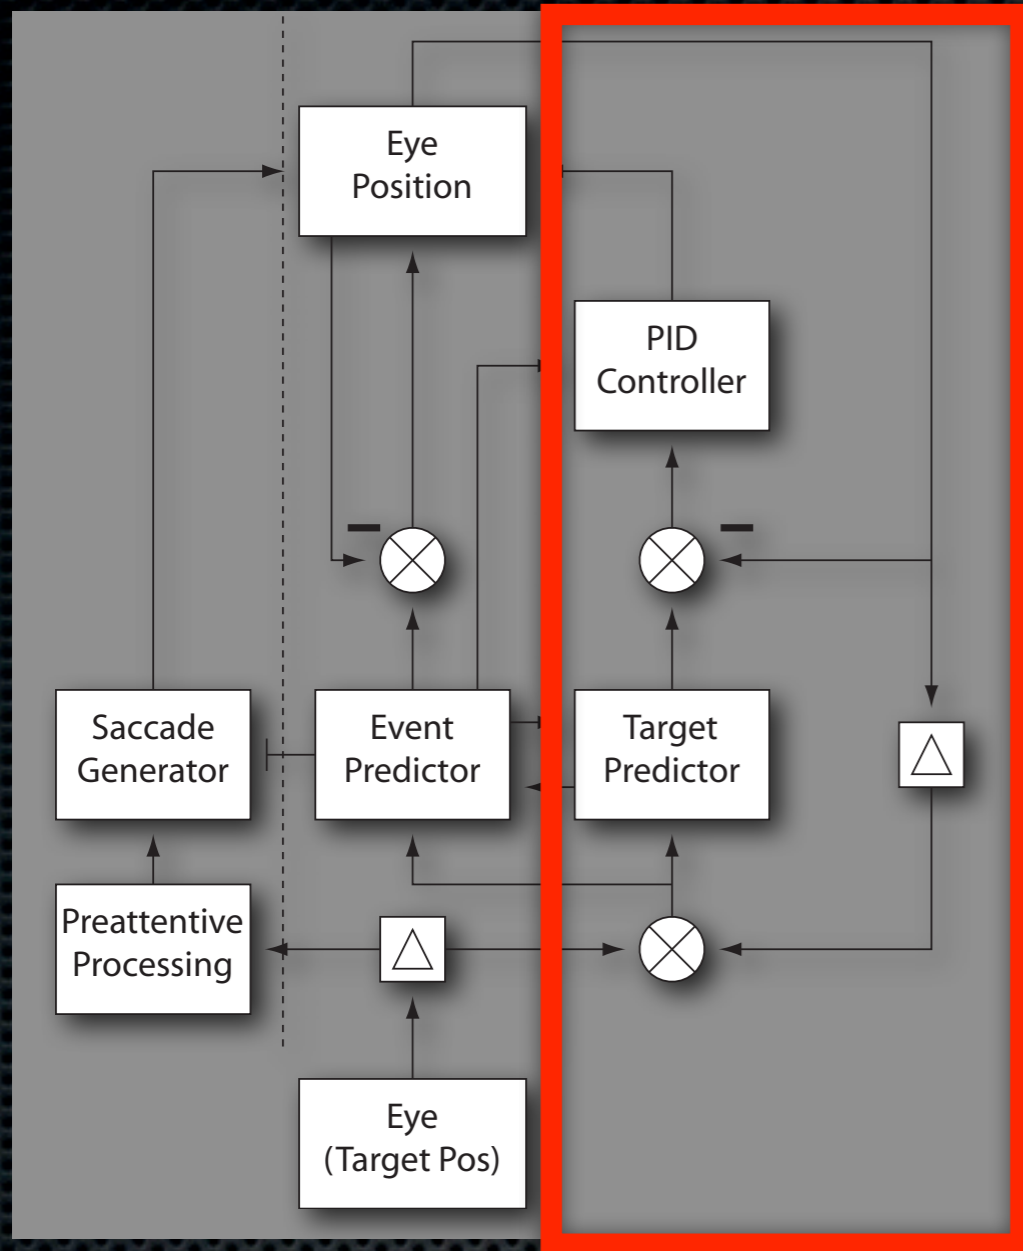
# Smooth Pursuit



stimulus-approach

Balkenius & Johansson, 2005

# Smooth Pursuit



Prediction confidence sets the gain of the controller


stimulus-approach

Balkenius & Johansson, 2005

# The Development of Smooth Pursuit


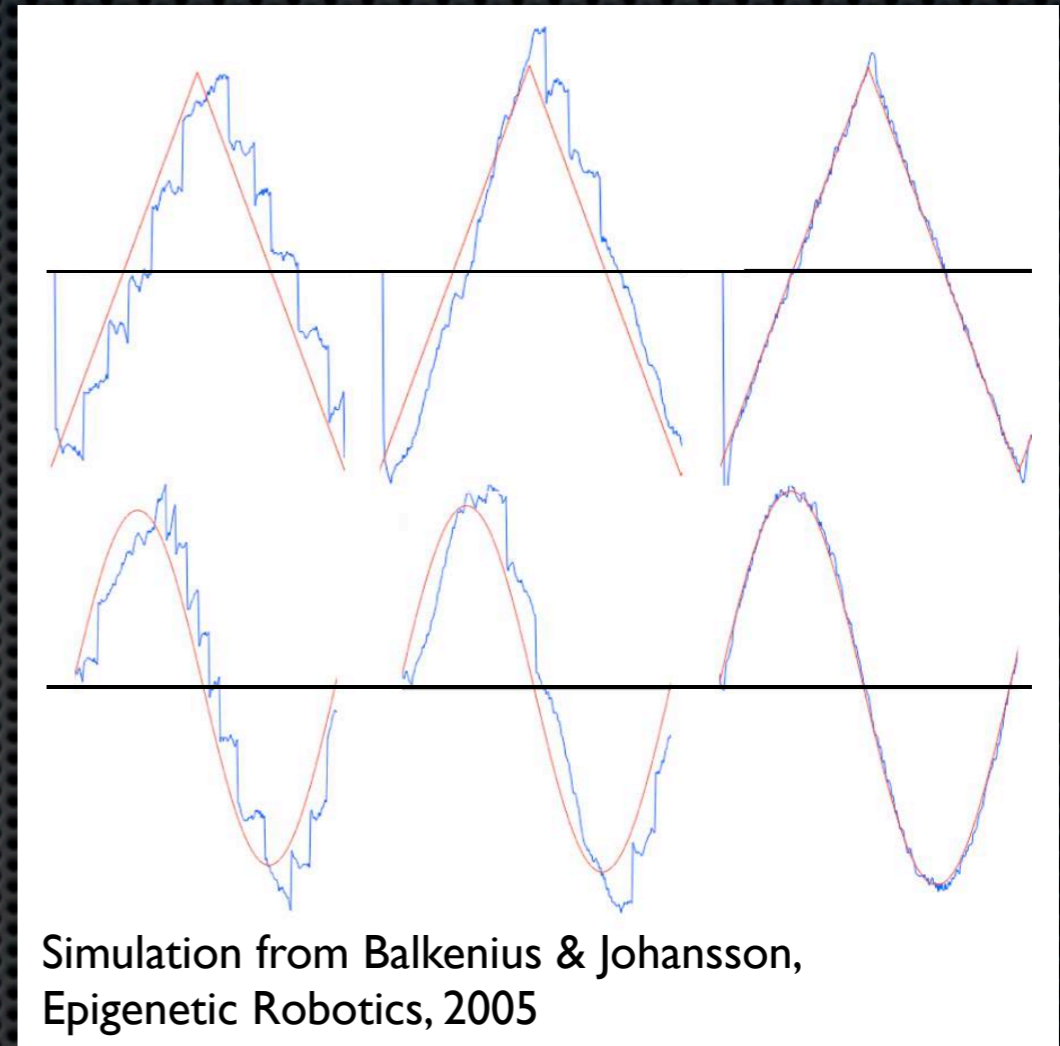
- Gradual development from catch up saccades to smooth pursuit from 0-4 month

Simulation from Balkenius & Johansson, Epigenetic Robotics, 2005


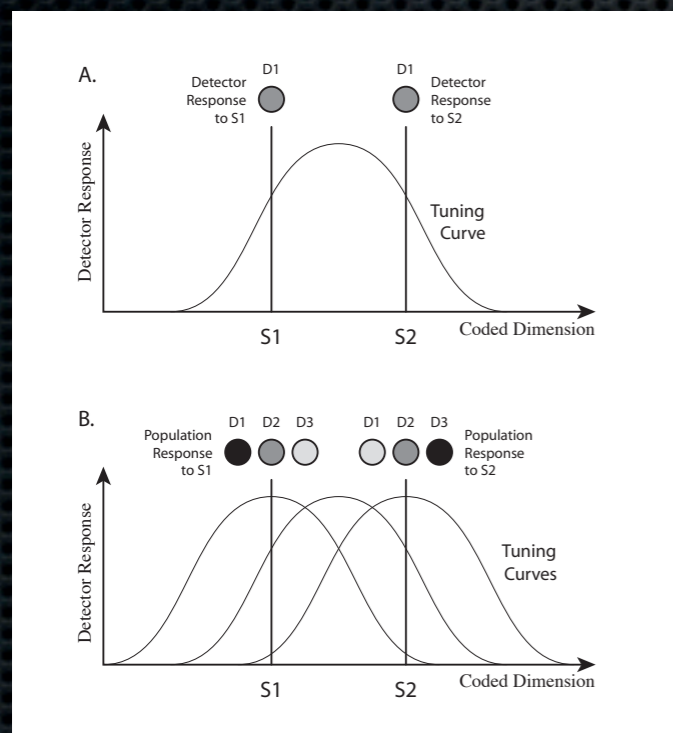


# The Development of Smooth Pursuit

- Gradual development from catch up saccades to smooth pursuit from 0-4 month


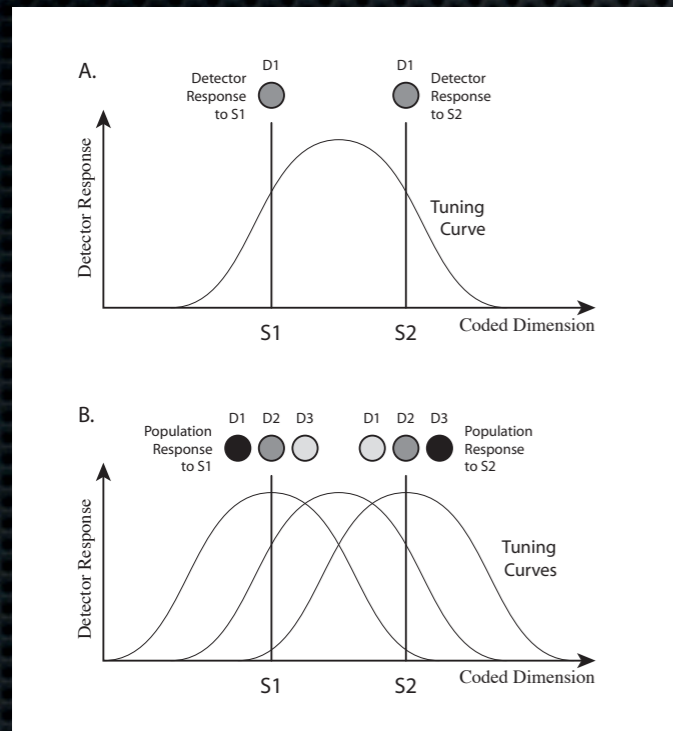


# Learning to Reach







# Learning to Reach



the system learns associations between retinal positions and joint angles



# Learning to Reach



the system learns associations between retinal positions and joint angles

not driven by error between hand position and target

# Learning to Reach




the system learns associations between retinal positions and joint angles


not driven by error between hand position and target

population coding supports interpolation and some extrapolation


# Learning to Reach



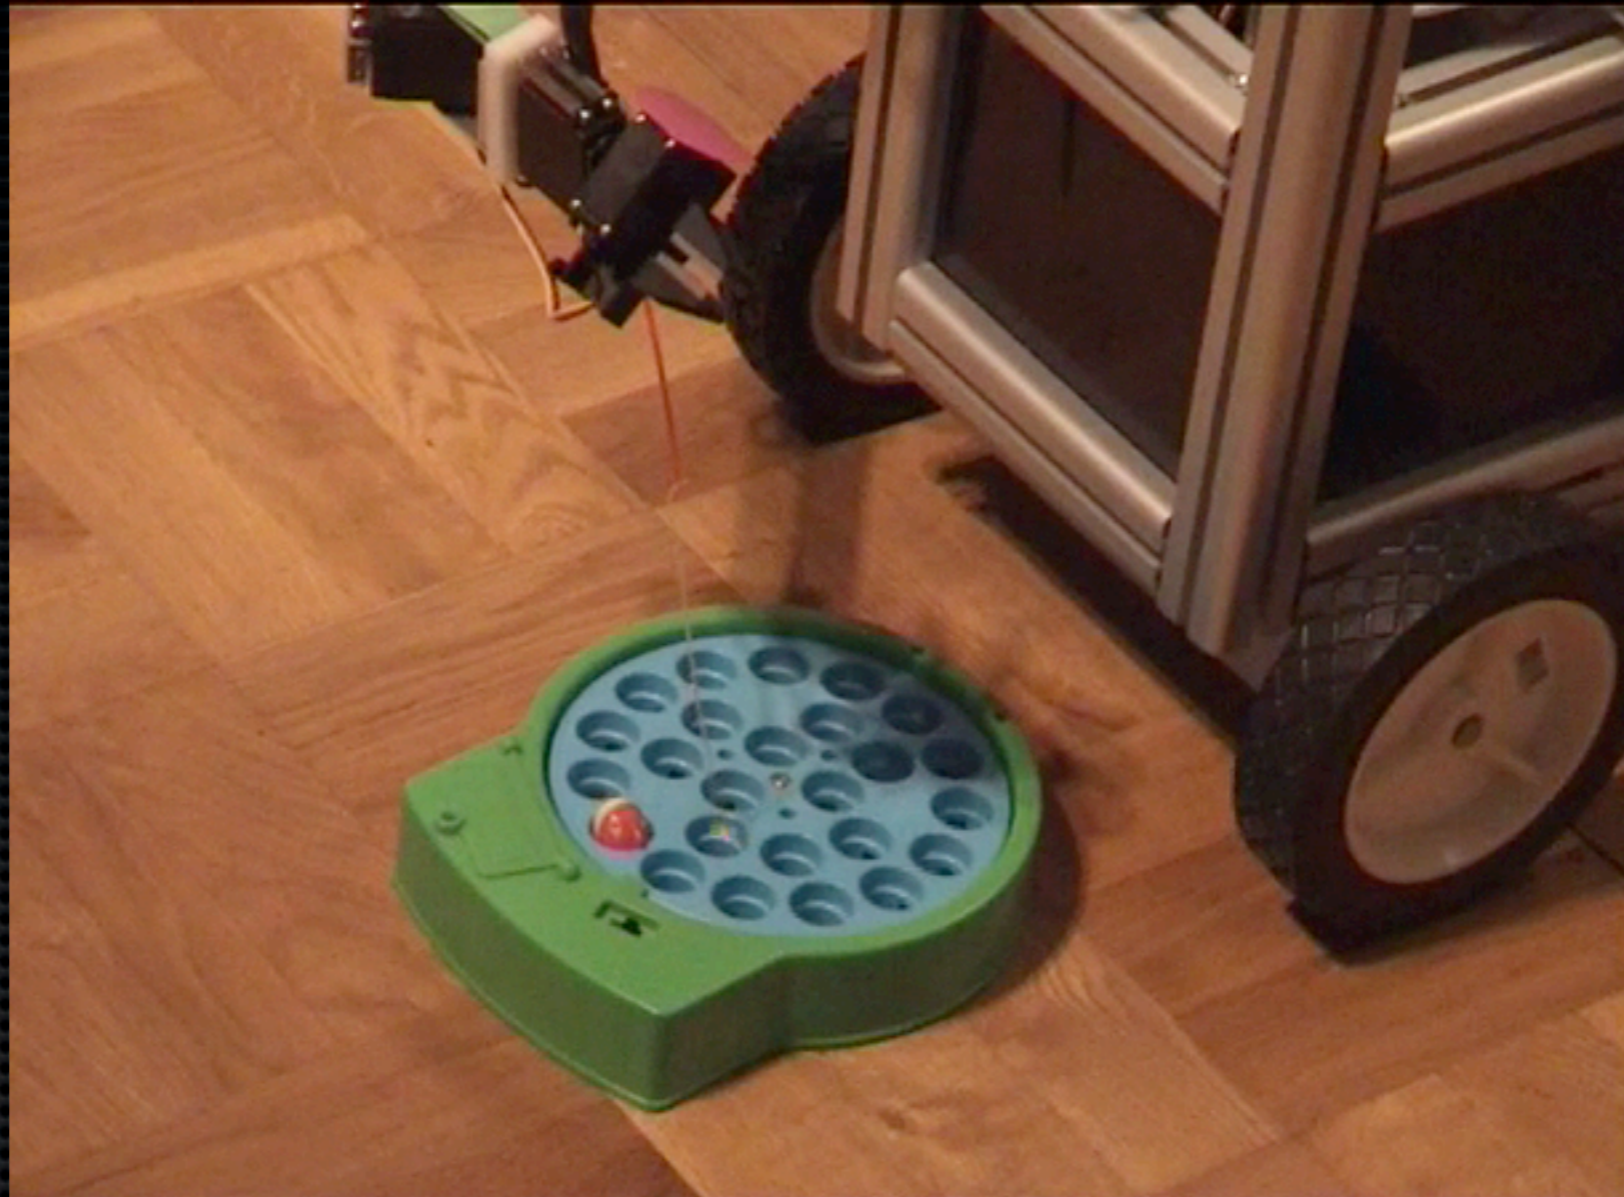
# Learning to Reach



# Learning to Reach




ongoing interaction



Fishing

0 min

500 ms sensory-motor delay



Fishing


2 min






Fishing

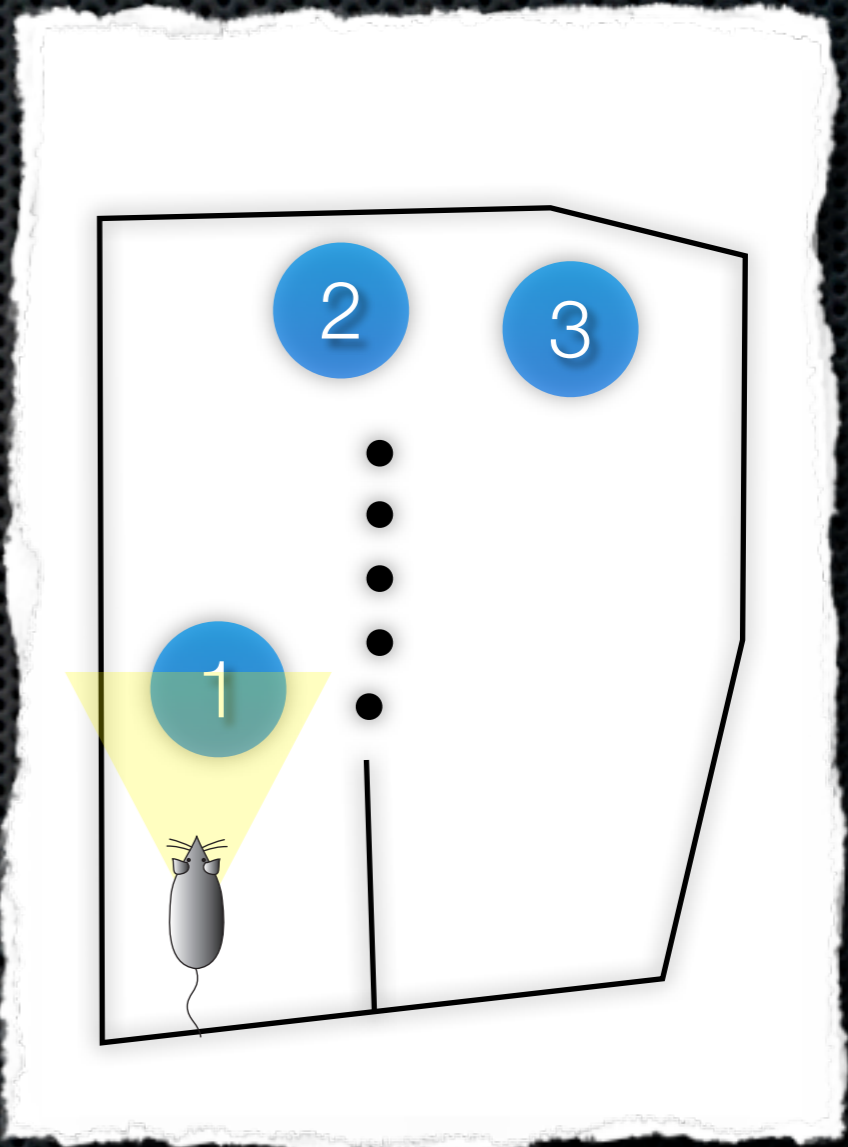
5 min




Stimulus-Approach




Stimulus-Response







Stimulus-Approach




Stimulus-Response



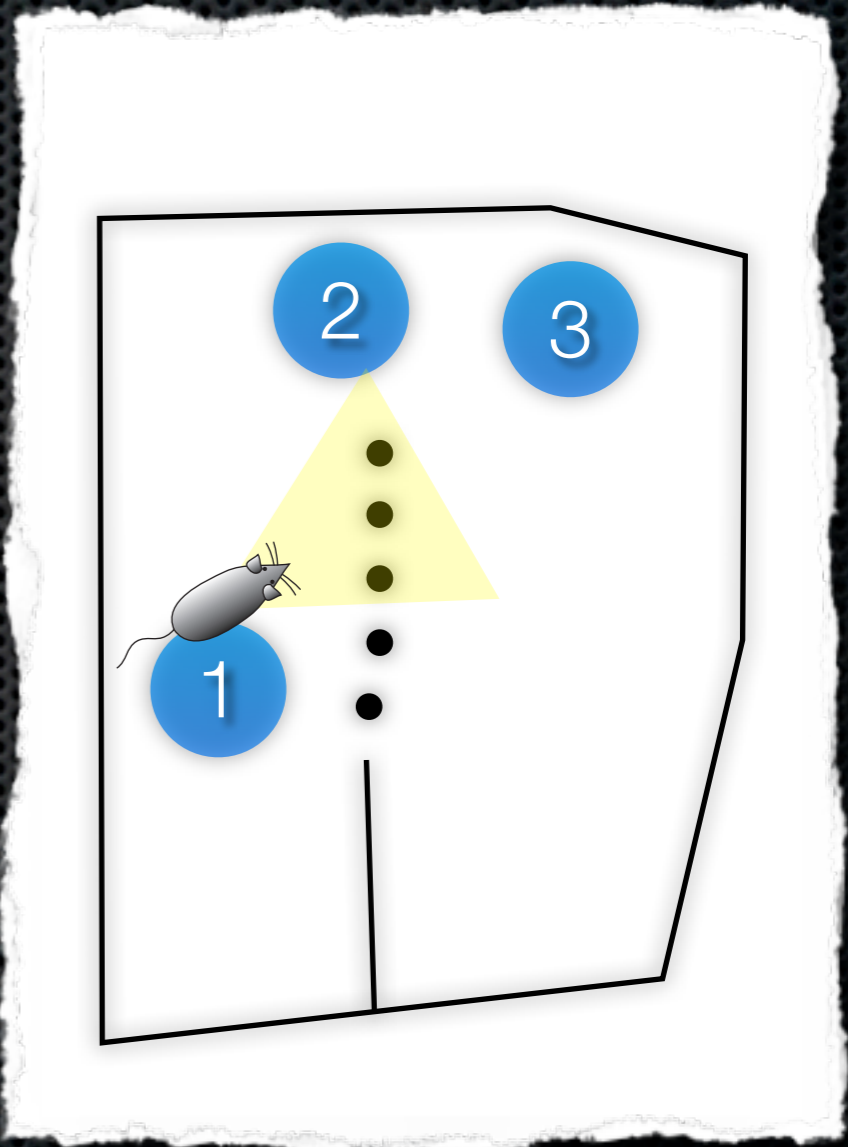





Stimulus-Approach




Stimulus-Response







Stimulus-Approach




Stimulus-Response







Stimulus-Approach




Stimulus-Response







Stimulus-Approach




Stimulus-Response



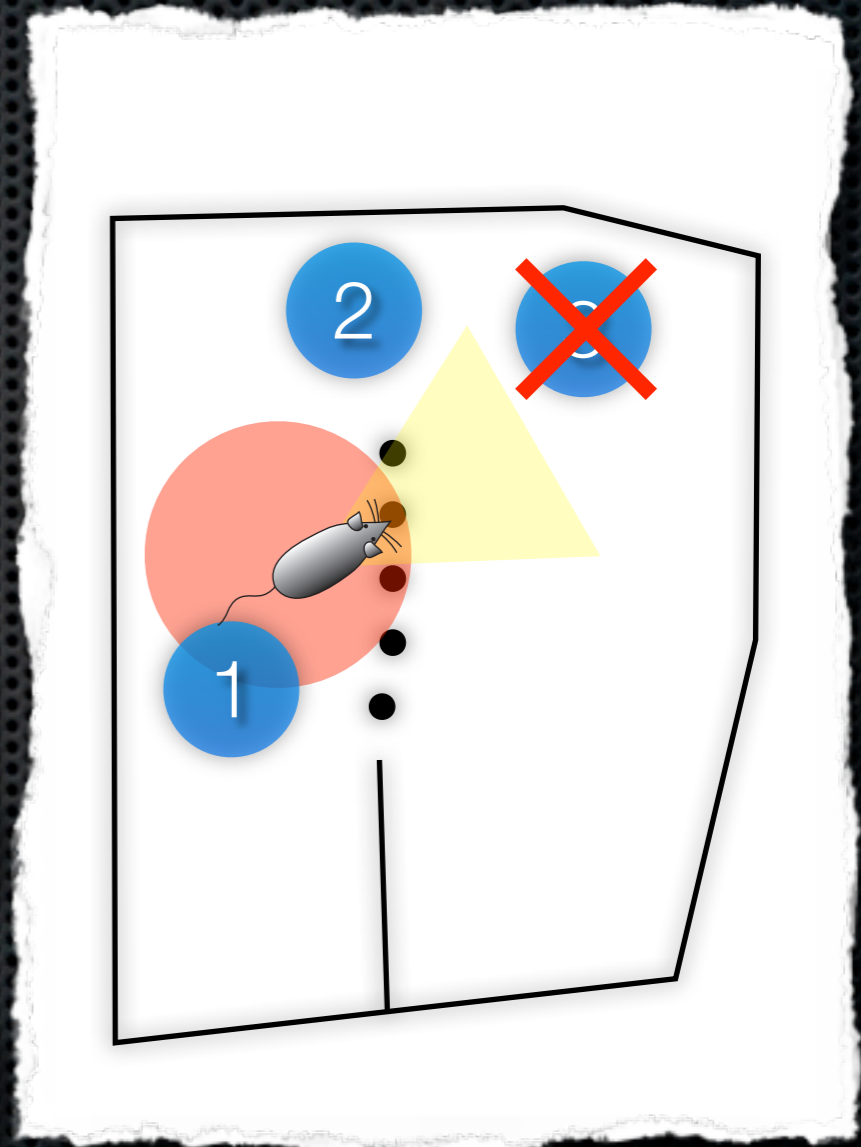





Stimulus-Approach




Stimulus-Response




Contextual Inhibition




Stimulus-Approach




Stimulus-Response



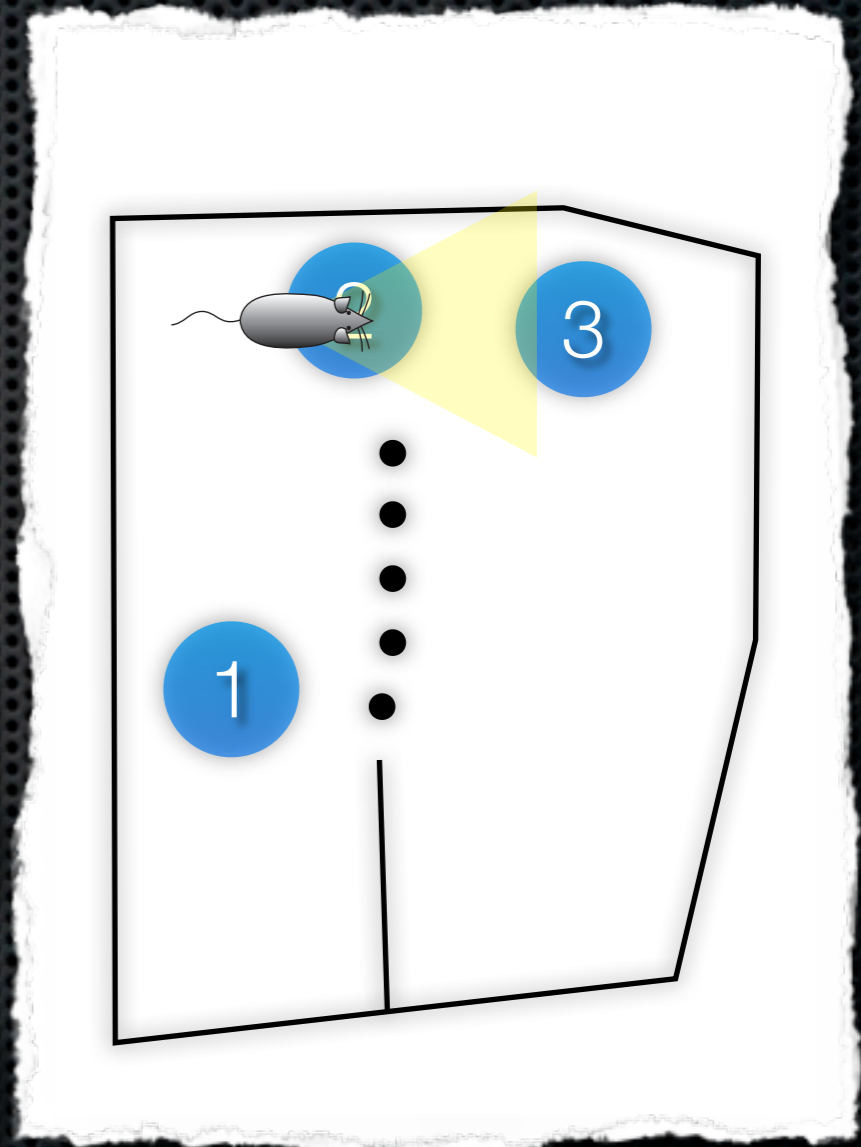
Contextual Inhibition




Stimulus-Approach




Stimulus-Response




Contextual Inhibition





Stimulus-Approach



Stimulus-Response



Contextual Inhibition



Neuronal population vectors are plotted every 10 ms vs time. C, onset of the delay; E, end of the waiting period. The *filled circle* on the abscissa indicates the time after the beginning of the delay (130 ms) at which the population vector reached statistical significance

# Context Effects

Phase 1	CXA : CS + US
Phase 2	CXA : CS
Test A	CXA : CS → no-CR
Test B	CXB : CS → CR

Extinction does not transfer to a new context (Bouton 1991, 1992)

# Contextual Inhibition

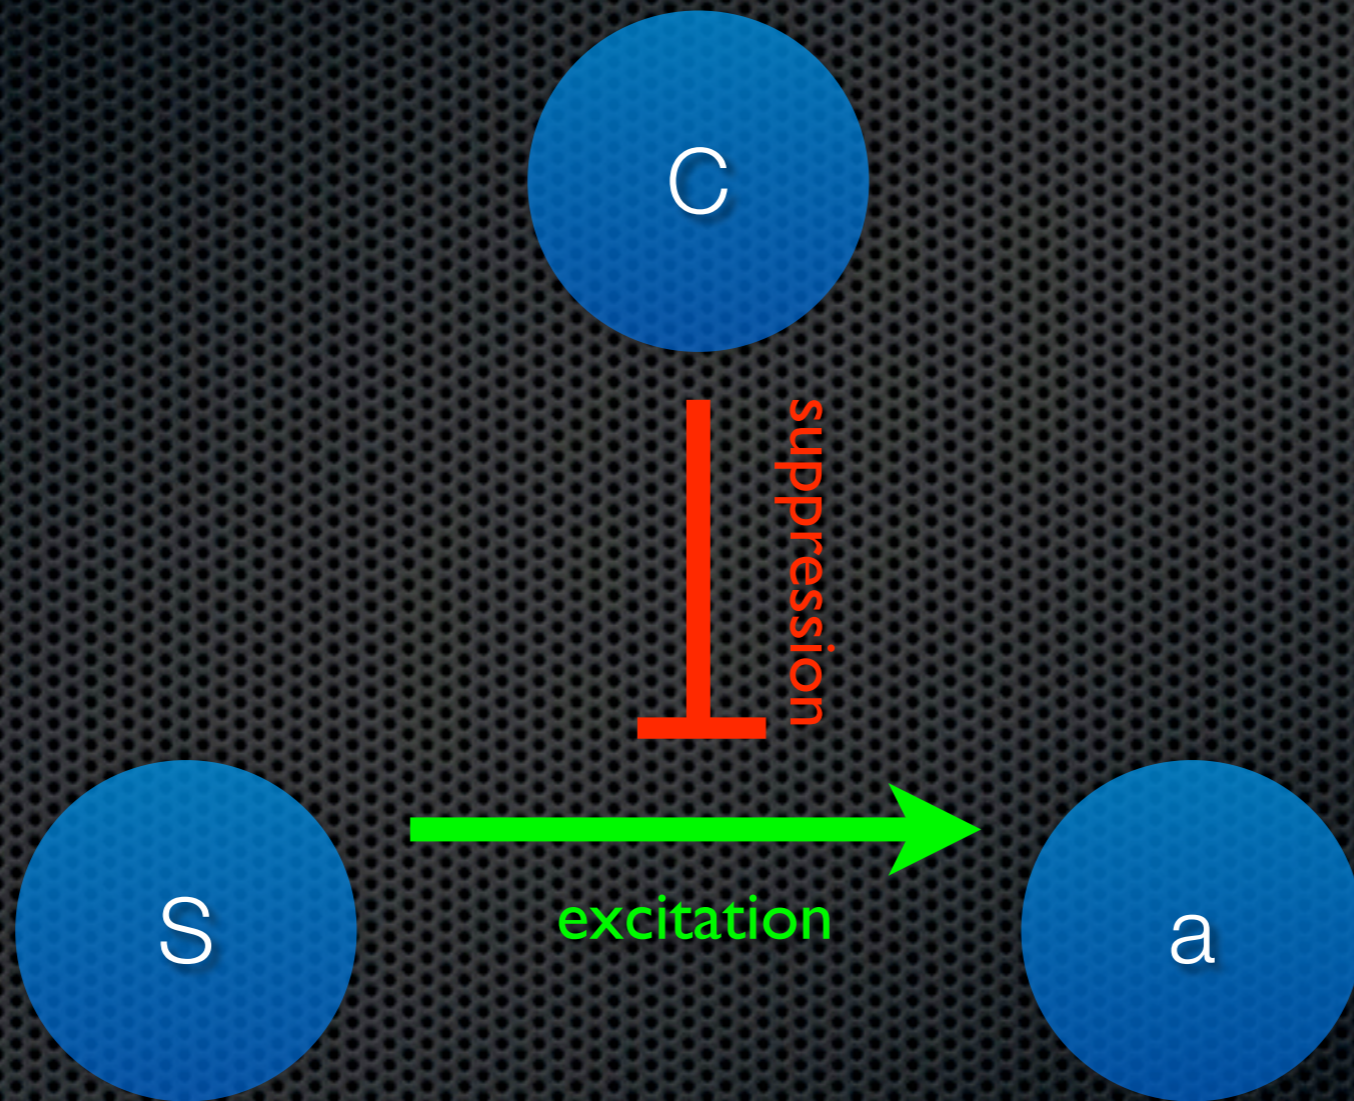


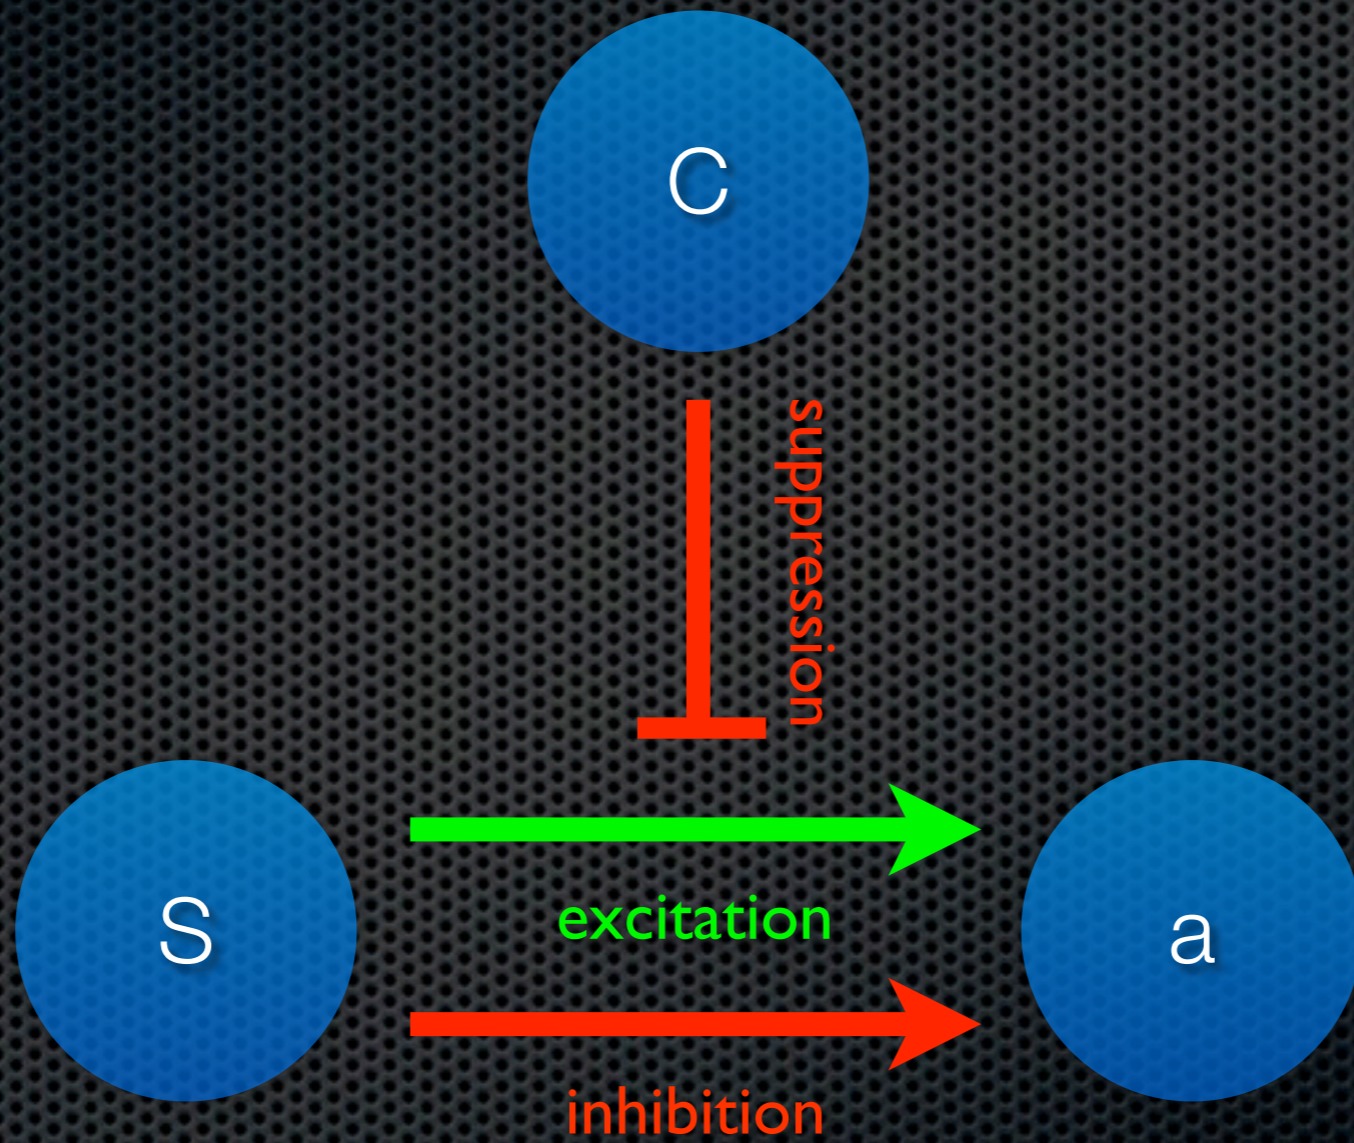
# Three Learning Conditions

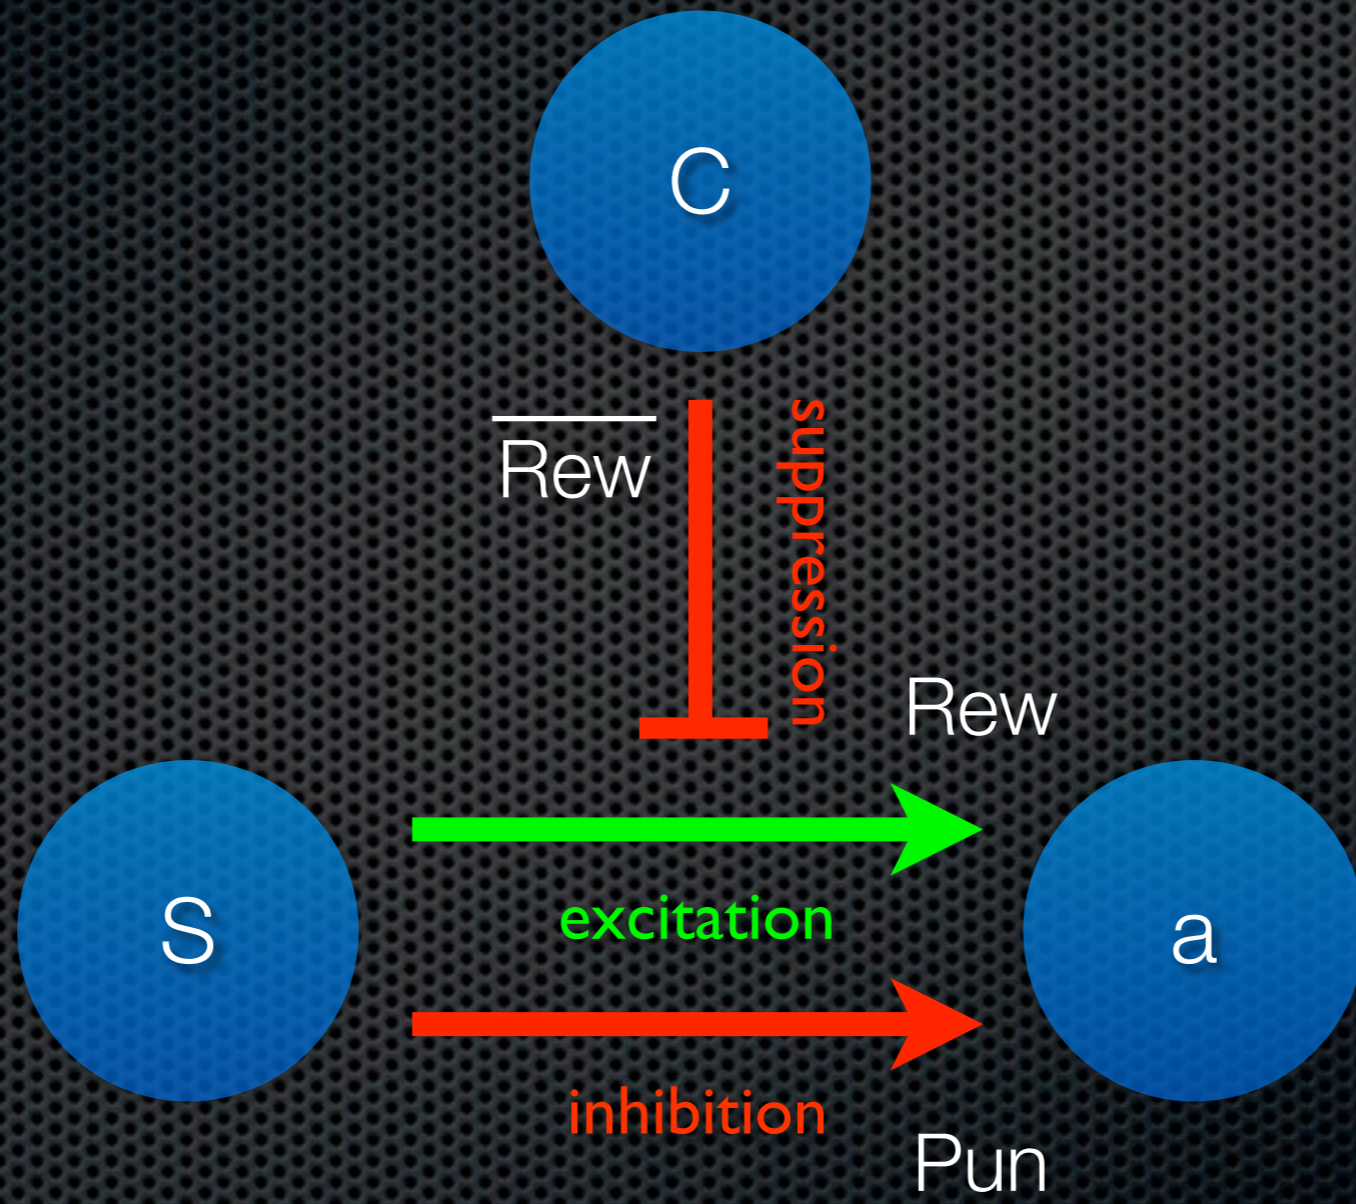
Rew	better than expected	maximal generalization
$\overline{\text{Rew}}$	worse than expected	contextual exception
Pun	bad	minimal generalization

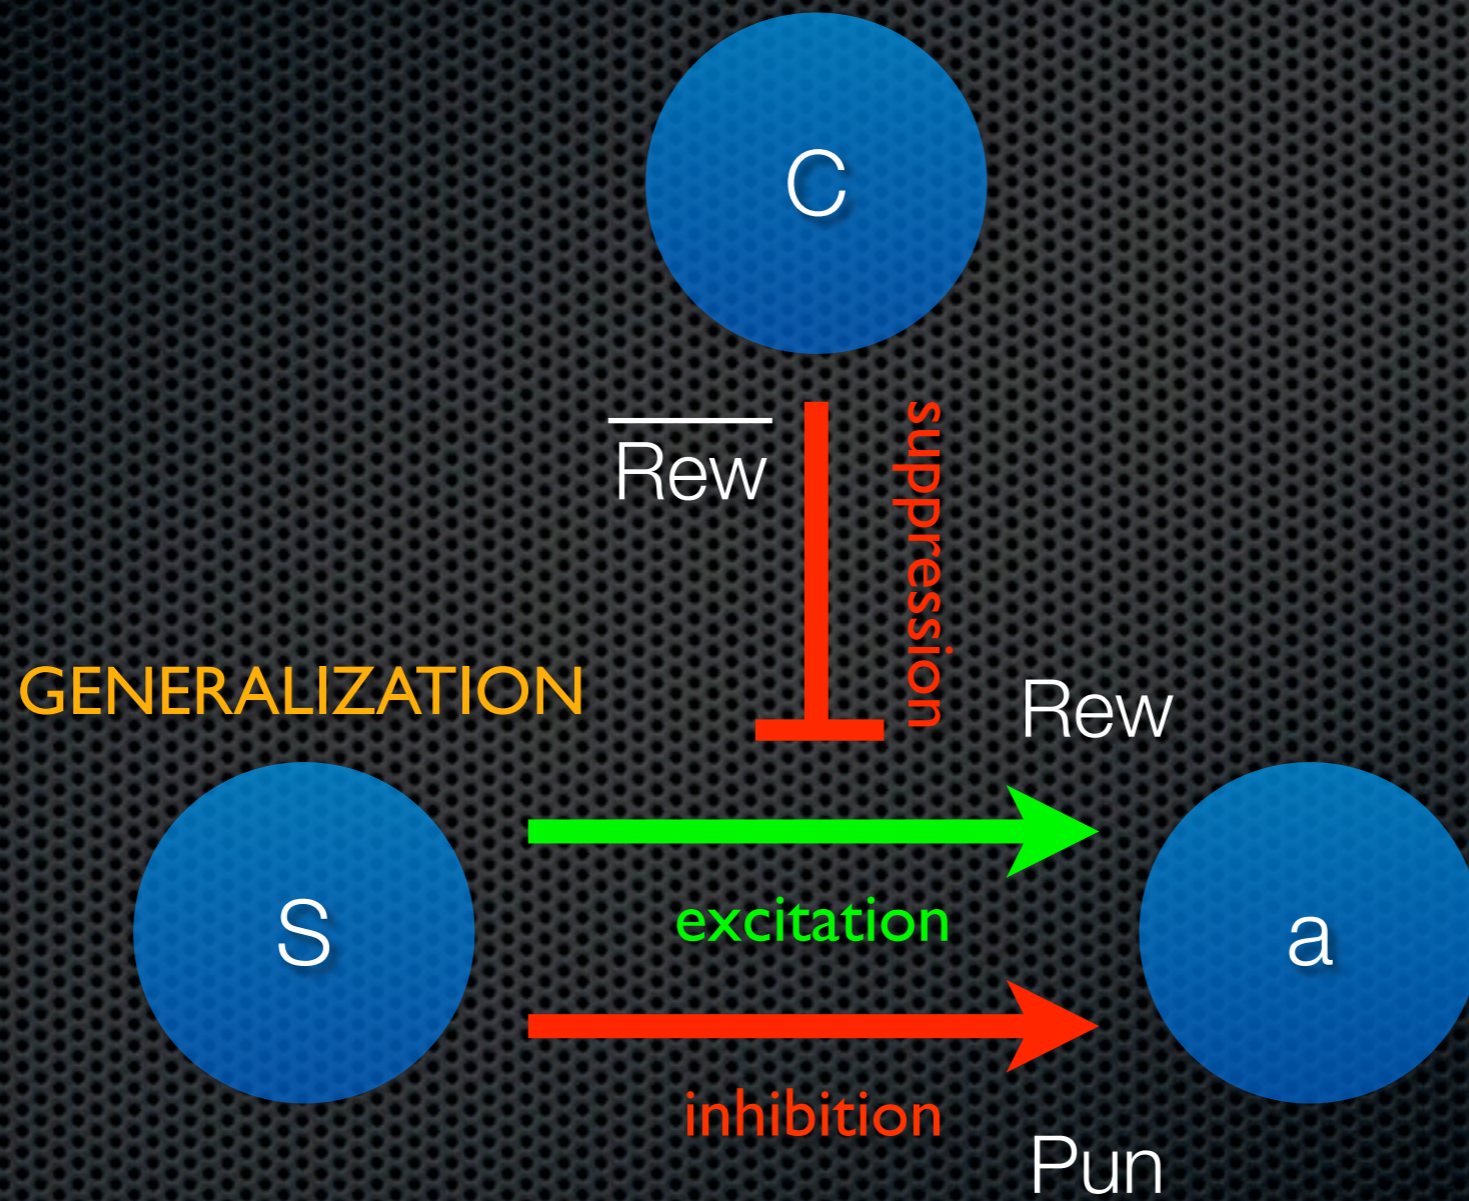












SPECIALIZATION



GENERALIZATION



excitation



inhibition



Rew

Pun

SPECIALIZATION



suppression

GENERALIZATION



SPECIFIC



excitation



inhibition

Rew



Pun

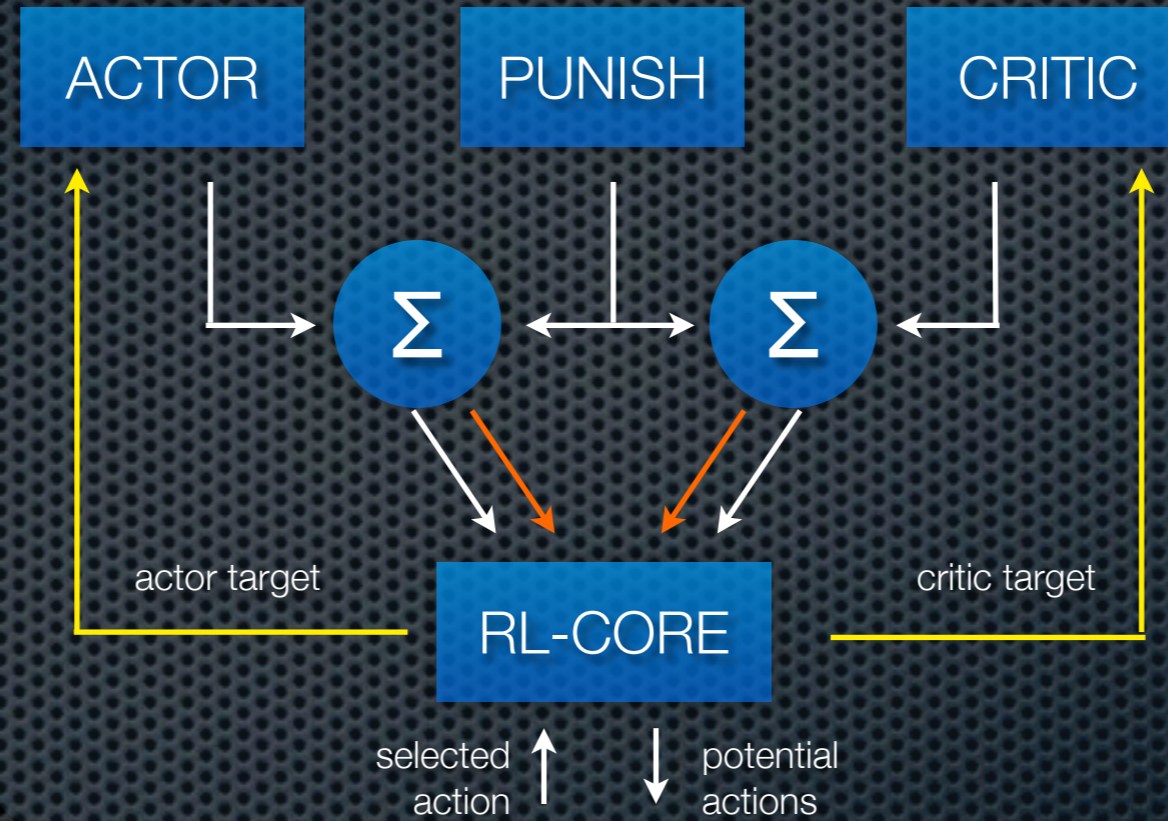
Sensory Coding

State & Action  
Evaluation

Action Selection

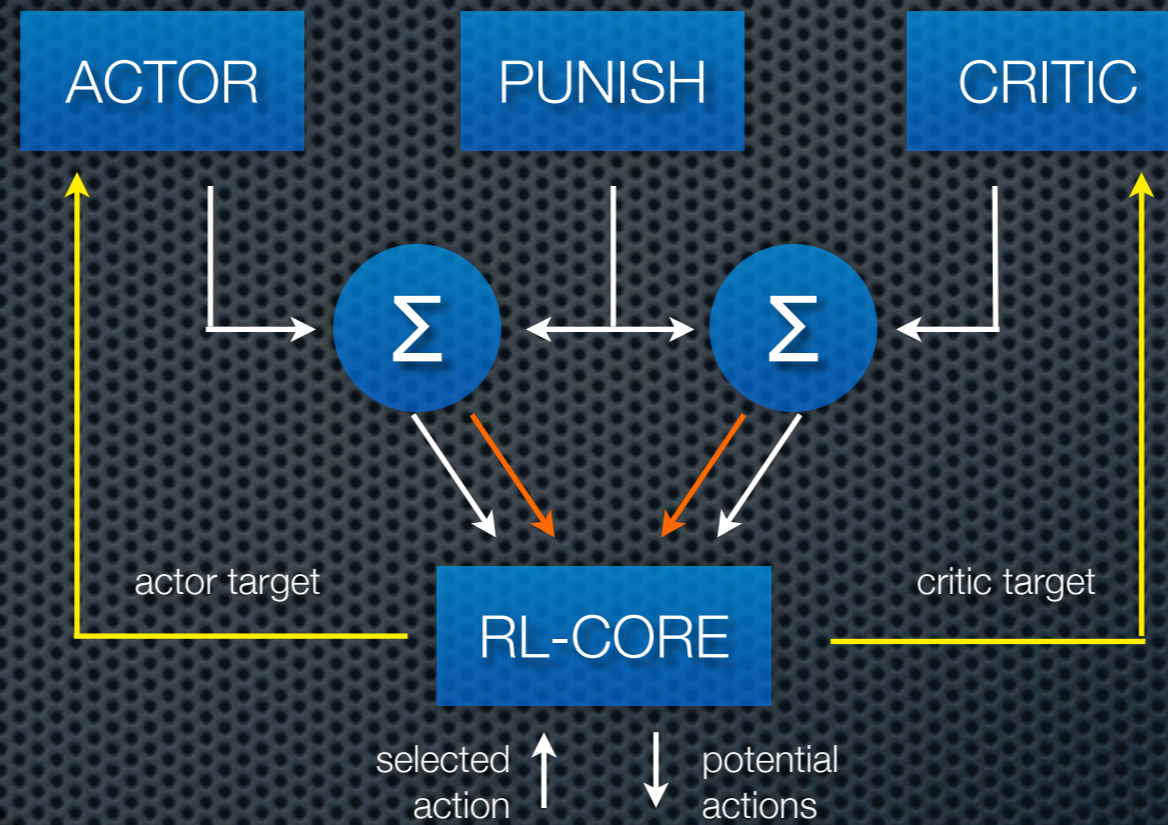


# Sensory Coding




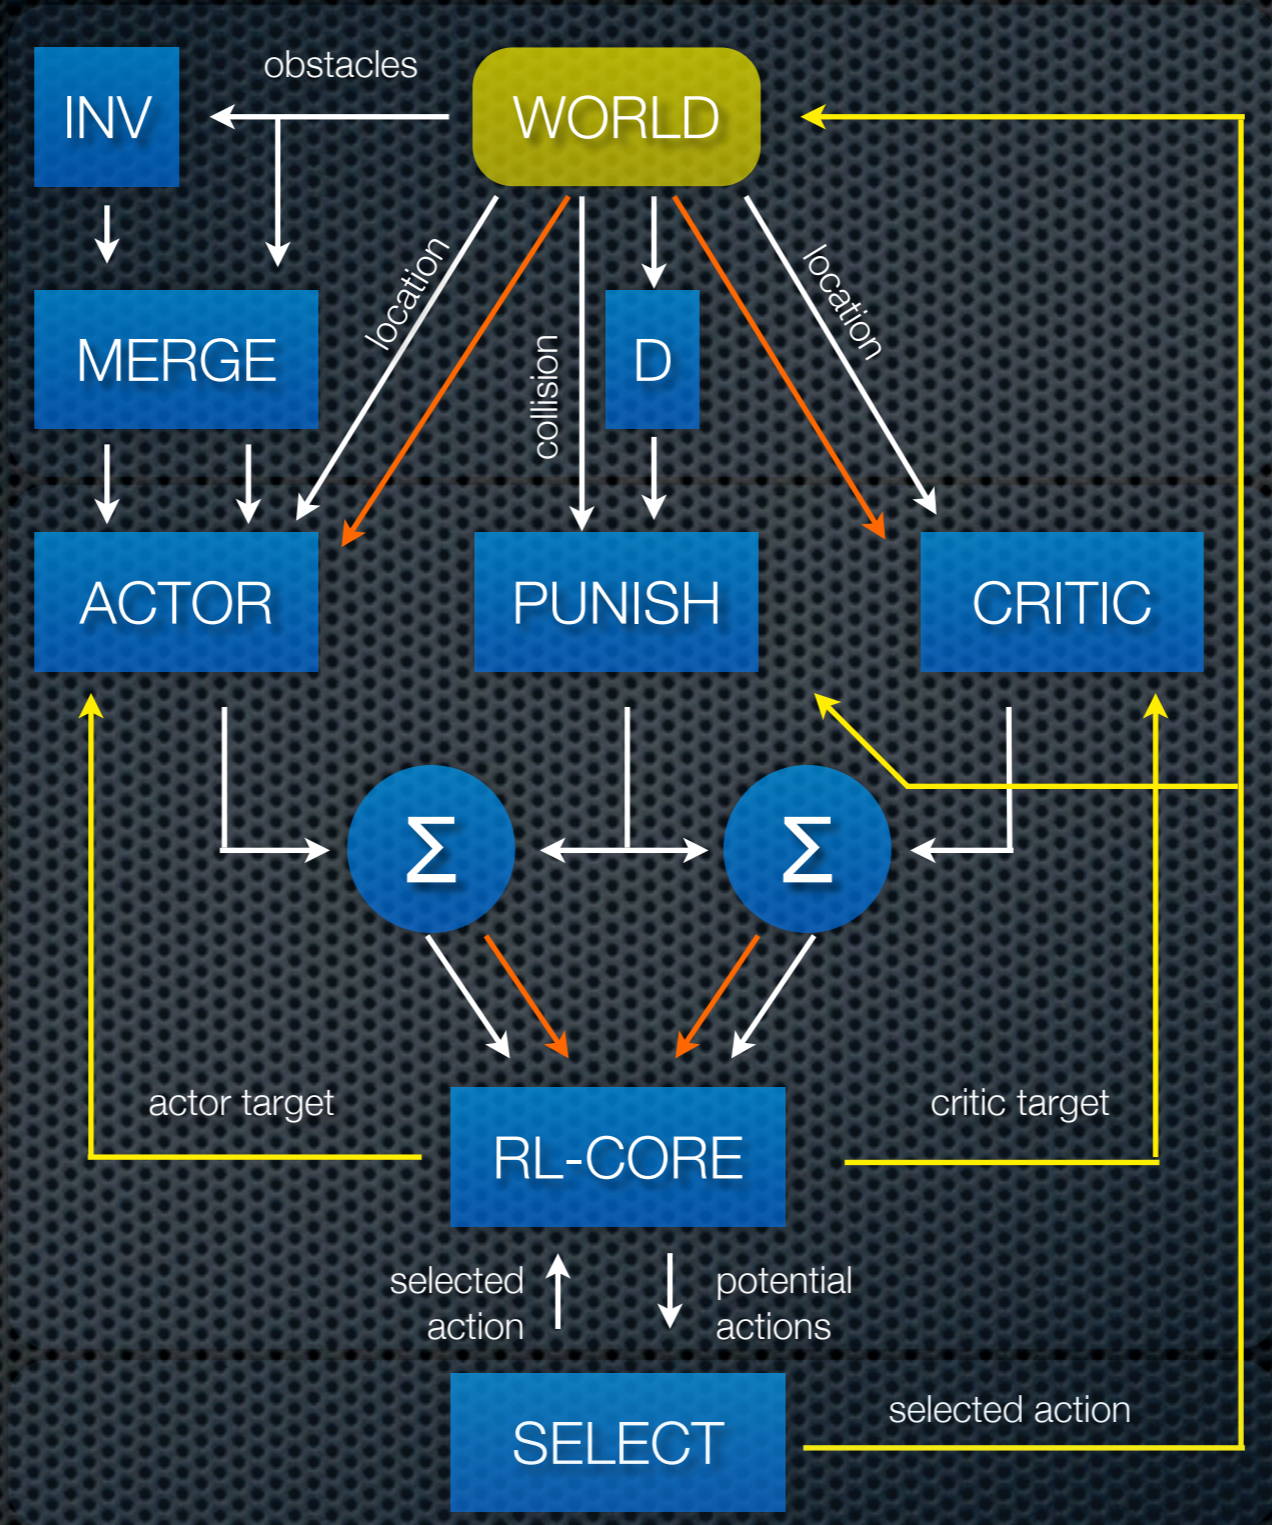
# Action Selection

# Sensory Coding



# Action Selection







# Learning Algorithm

$$Q(c, s, a_j) = \sum_{i=0}^n s_i w_{ij} I_{ij},$$











$$I_{ij} = \prod_{k=0}^p (1 - c_k u_{ijk})$$

$$w_{ij}^{(t+1)} = w_{ij}^{(t)} + \alpha \frac{s_i a_j}{|s|} \Delta Q_t \quad \Delta Q_t > 0$$









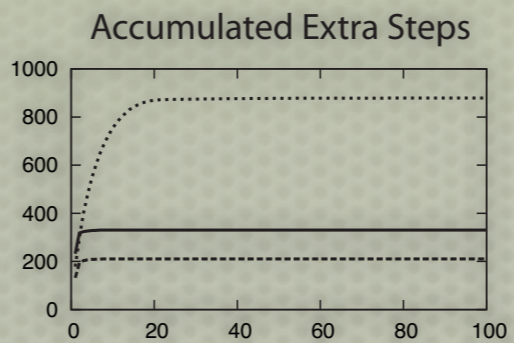
$$u_{ijk}^{(t+1)} = u_{ijk}^{(t)} - \beta (1 - u_{ijk}^{(t)}) \frac{s_i a_j c_k}{|s| w_{ij}} \Delta Q_t \quad \Delta Q_t < 0$$

# Maze



# Accumulated Extra Steps



### Maze




# A More Complex Example







# Context Prevents Catastrophic Forgetting



Q-Learning with a regular linear network



Context-Q



# Four Algorithms

Q	ContextQ	ContextAC	ContextACP
$Q(s, a)$	$Q(c, s, a)$	$Q(c, s, a)$ $V(s)$	$Q(c, s, a)$ $V(s)$ $P(s, a)$
'stimulus generalization'	contextual specialization	progress separate from state that controls action	learns to avoid doing bad things

# Four Algorithms

Q	ContextQ	ContextAC	ContextACP
$Q(s, a)$	$Q(c, s, a)$	$Q(c, s, a)$ $V(s)$	$Q(c, s, a)$ $V(s)$ $P(s, a)$
'stimulus generalization'	contextual specialization	progress separate from state that controls action	learns to avoid doing bad things

general  less general

- ✦ Stimulus-Approach
- ✦ Stimulus-Response
- ✦ Contextual Inhibition

