Goal-Leaders
Goal-directed, Adaptive Builder Robots

GIOVANNI PEZZULO
NATIONAL RESEARCH COUNCIL, ROME, ITALY

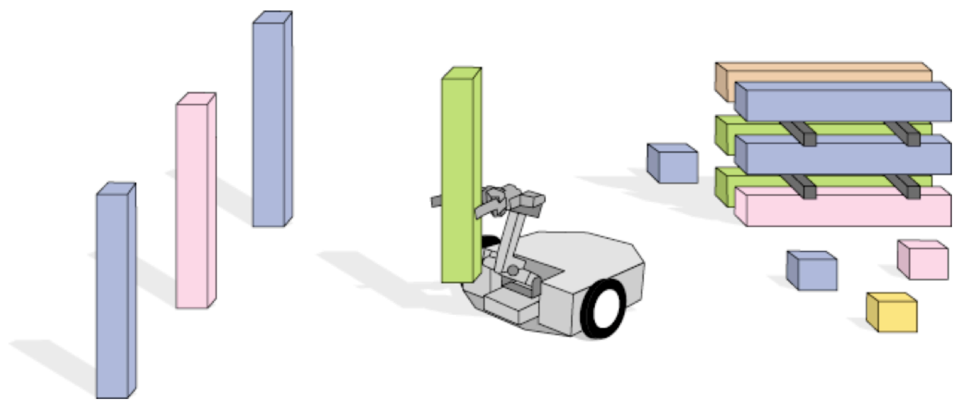

Goal-Leaders Consortium

National Research Council, Italy (CNR)	Universitat Pompeu Fabra, Spain (UPF)	Universiteit van Amsterdam, the Netherlands (UVA)	Lund University, Sweden (ULUND)
			
Giovanni Pezzulo	Paul Verschure	Cyriel Pennartz	Christian Balkenius
Computational modeling of anticipation and goal-directedness	Computational neuroscience, bio-constrained robot models, DAC	Neuroscience of spatial navigation, prediction and goal-directedness	Epigenetic robotics, attention control, navigation and motivation

Goal-Leaders: enhancing robots' goal-directedness and proactivity

Brains, goal-directedness and autonomy

- Brain “information processing” is dominated by endogenously determined motivations, predictions and goals, and processes that prepare to action
  - Cascading effects on perception, memory, attention, monitoring, behavior, etc.
  - In turn, this sensorimotor loop affects internal processing (drive and goal setting, action selection, prediction, etc.)

Brains, goal-directedness and autonomy

Affordance competition hypothesis

Free energy principle and predictive coding

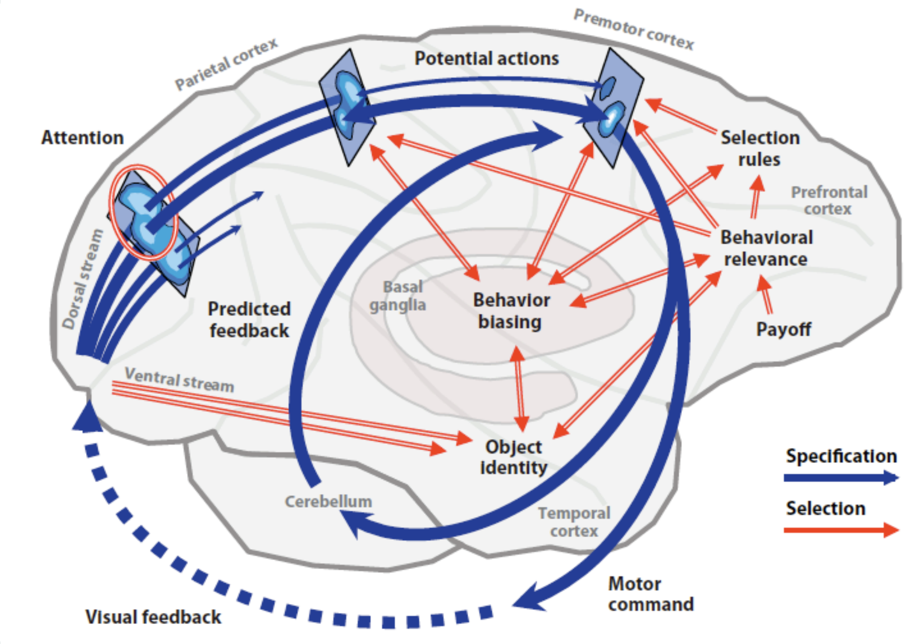

Cisek and Kalaska, 2010

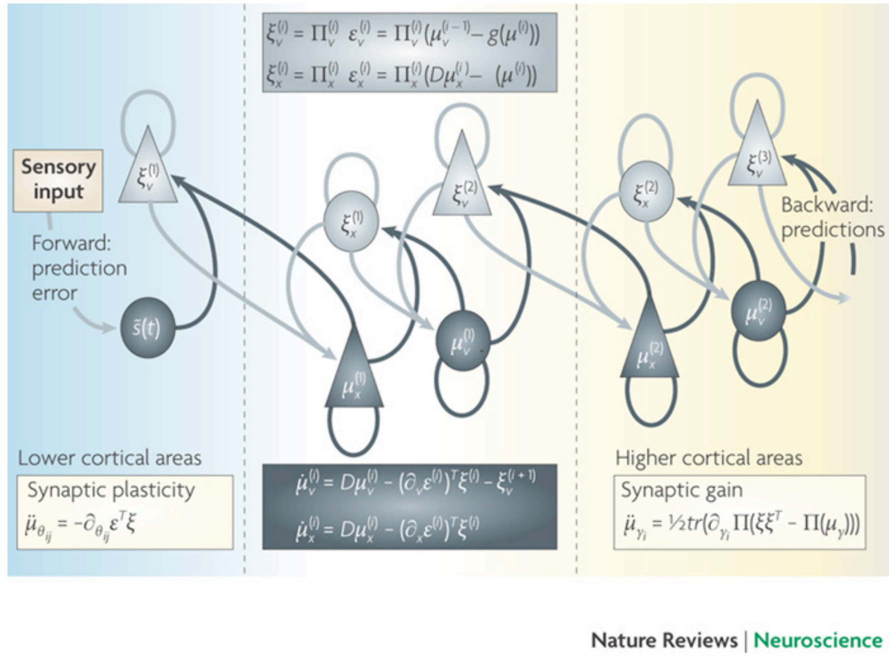

Friston, 2010

Many others: predictive brain, Bayesian brain, ideomotor principle, sensorimotor theories, model-based reinforcement learning in neuroeconomics, etc.

Goal-directed vs. habitual behavior

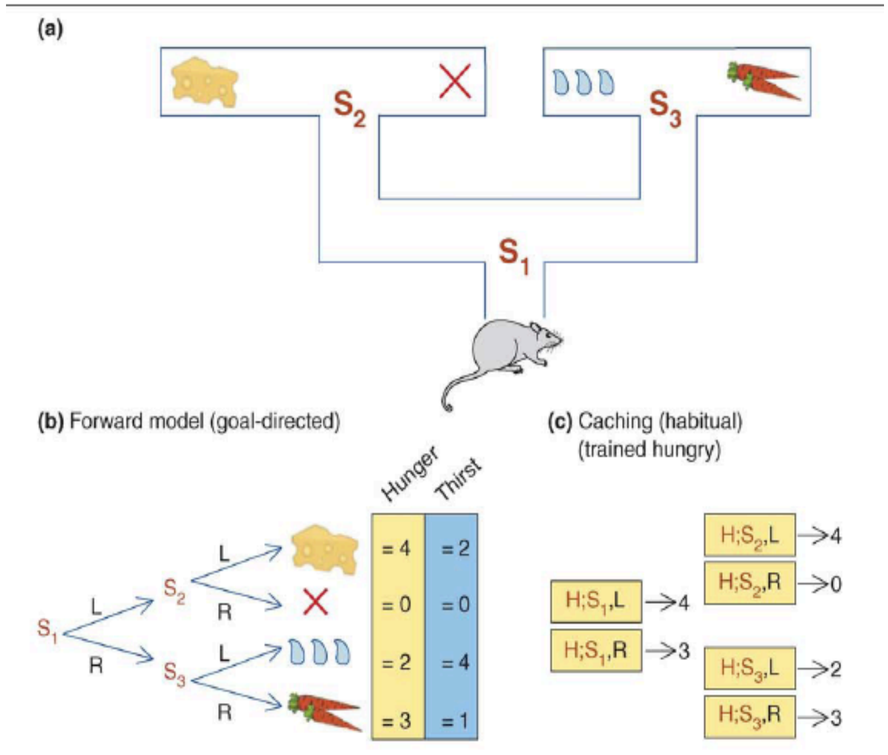

Model-based RL **Model-free RL**

Dickinson, Balleine (1998); Daw et al., (2005)

Project Objective:
Adaptive Builder Robots

Adaptive Builder Robots

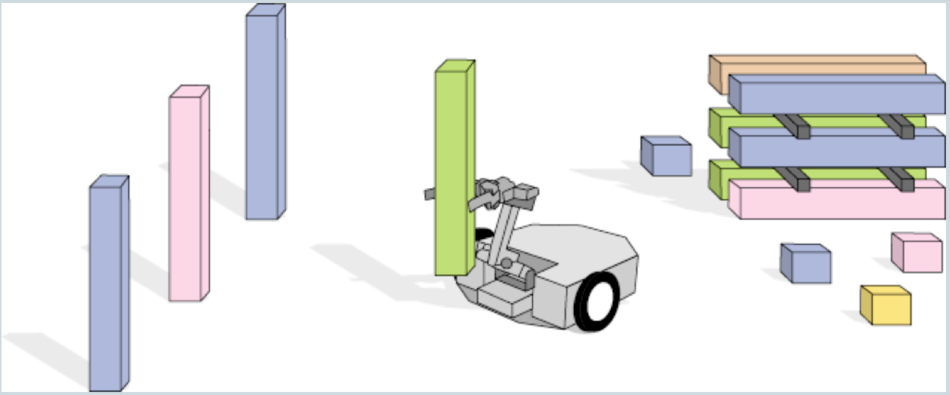

The robot is required to assemble a construction by selecting, reaching and assembling materials having different size and colour.

- Realize a set of externally assigned tasks (e.g., fetching objects, clearing an area, composing building parts)
- Maintain homeostatic drives in safe range (e.g., never remain without energy, not get hurt)

- Combinatorial tasks, subgoaling, cognitive control (e.g., finding and stacking objects to compose a given construction)
- Proactivity (e.g., store useful building parts, predict loss of energy and recharge before starting a long task)

We need breakthroughs in robots' goal-directedness and proactivity

1. Enhancing robots' goal-directedness

Learning increasingly sophisticated goals and goal-achieving strategies

...

More complex goals

More flexible actions

Goal states, objects, events

Drives

Stereotyped behaviors

Our starting point: Distributed adaptive control (DAC)

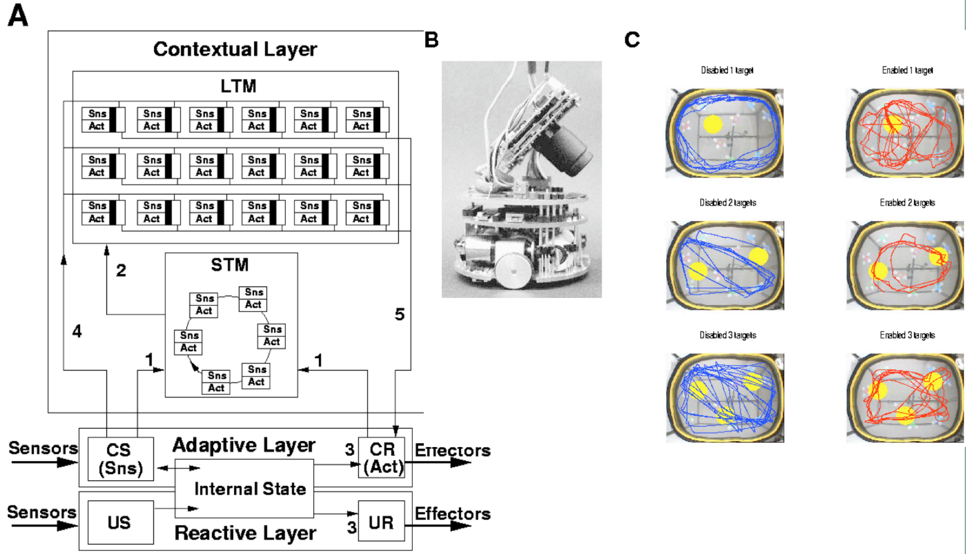

Verschure, Kröser et al.(1993); Verschure, Voegtlin et al. (2003

Reactive	Adaptive	Contextual
Drive: having stable structure, tall structure Reactive behavior: stacking	Represents objects in terms of how much they support other objects (cubes vs. sphere; big objects down, small objects on top). Objects / place value	Goal states: "Towers", "bridges" Strategies for efficient assembling of constructions, storing of useful objects, subgoaling

2. Enhancing robots' proactivity

From Marc Jeannerod's Lab

**Prediction,
mental simulation**

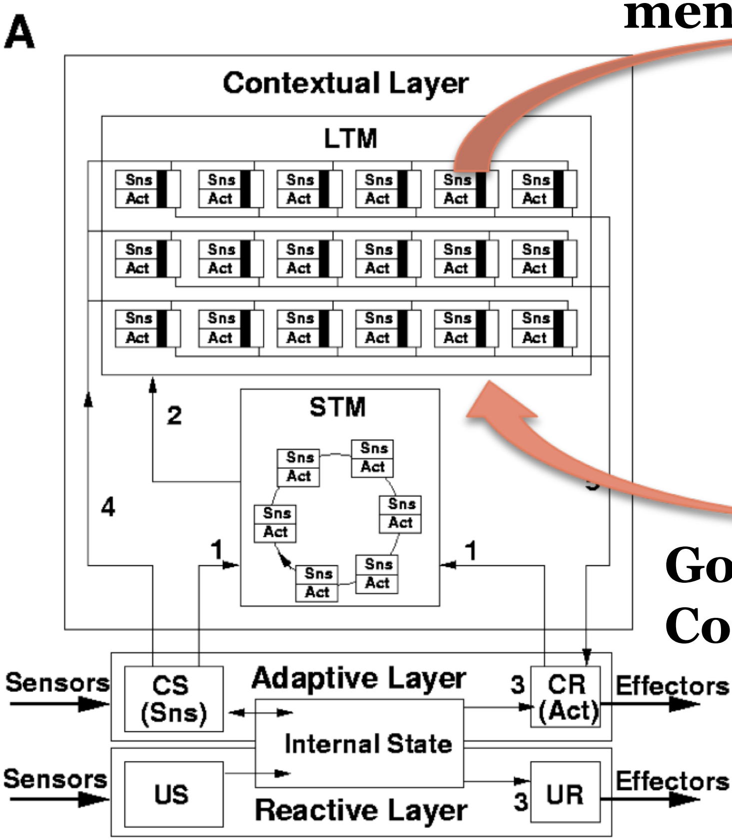

**Evaluate outcomes
Predict dangers
Set distal goals**

**Goal-directedness,
Cognitive control**

A few achievements so far

Hippocampal-striatal loops for mental simulation

Prospective coding in the rat hippocampus (CA3): forward sweeps at decision points

Johnson and Redish (2007)

The rats position is indicated by the white circle. The reddish areas indicate the firing rate of hippocampal neurons with place fields at that particular point of the maze.

Reward-predictive cues modulate firing patterns of hippocampal and striatal neurons

MISSING

Lansink et al. (in prep)

Goal-directed decision-making: hippocampal replay for accessing reward info in the striatum

Neural circuit for the affective labeling of spatial representations

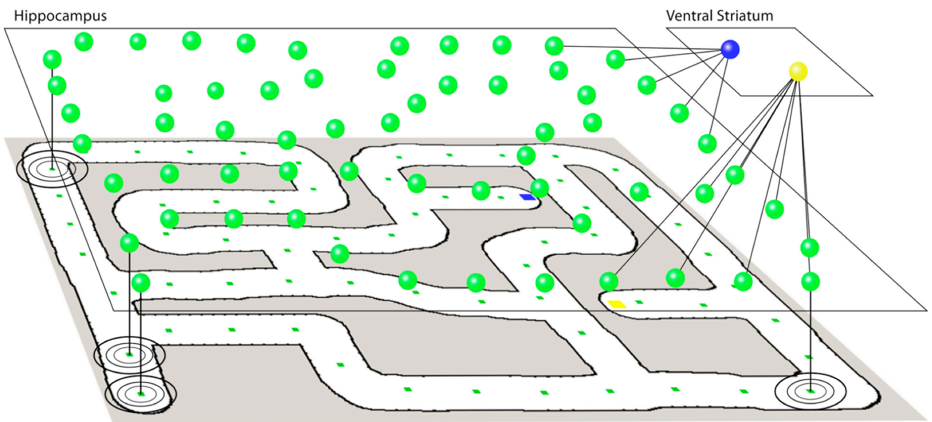

Chersi, Pezzulo (in prep.)

Goal-directed decision-making: hippocampal replay for accessing reward info in the striatum

Neural circuit for the affective labeling of spatial representations

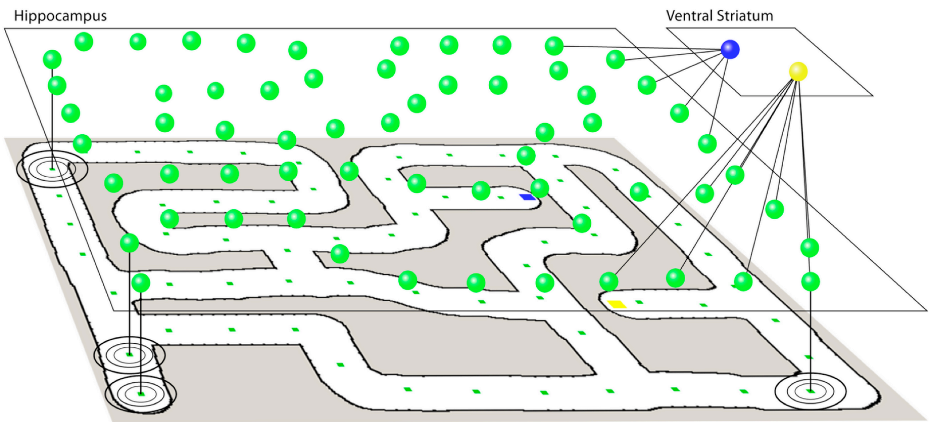

Chersi, Pezzulo (in prep.)

The neural circuit at work

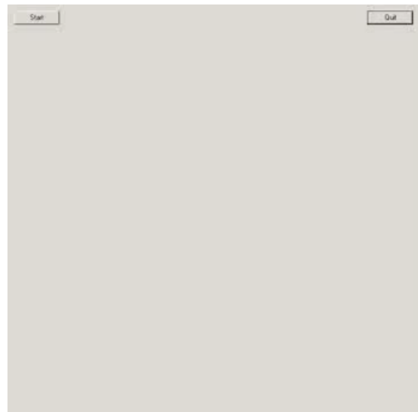

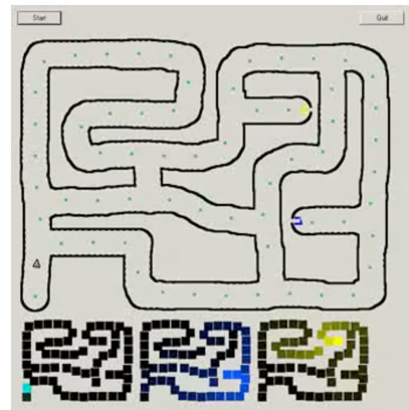

Balance between goal-directed vs. habitual choice

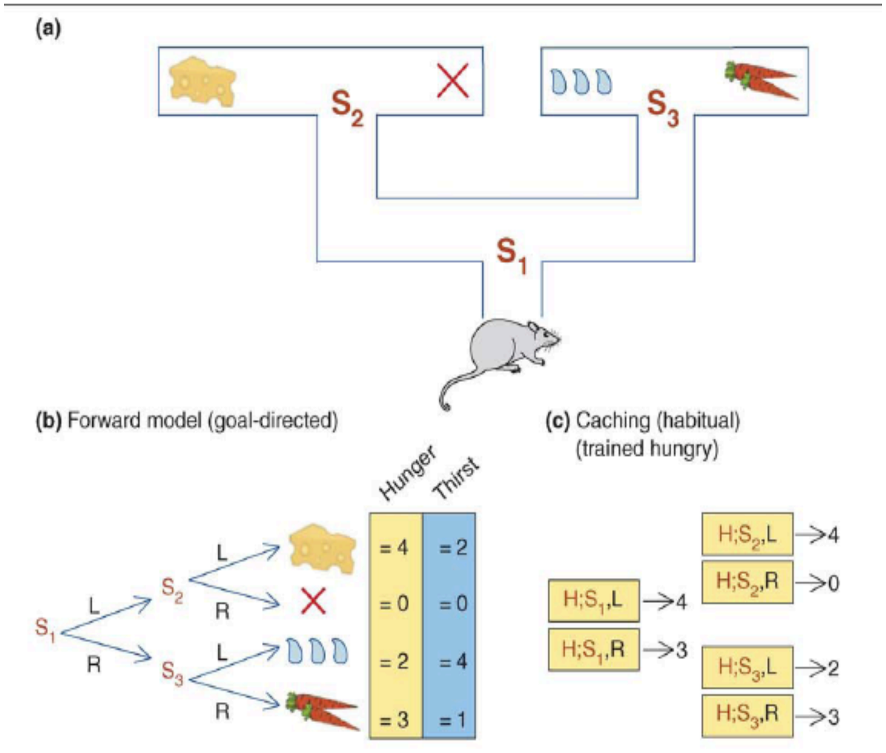

Model-based RL **Model-free RL**

Dickinson, Balleine (1998); Daw et al., (2005)

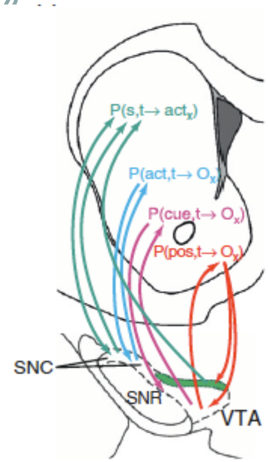

Different types of predictive learning associated with striatal sectors in rat brain.

Pennartz et al., 2011

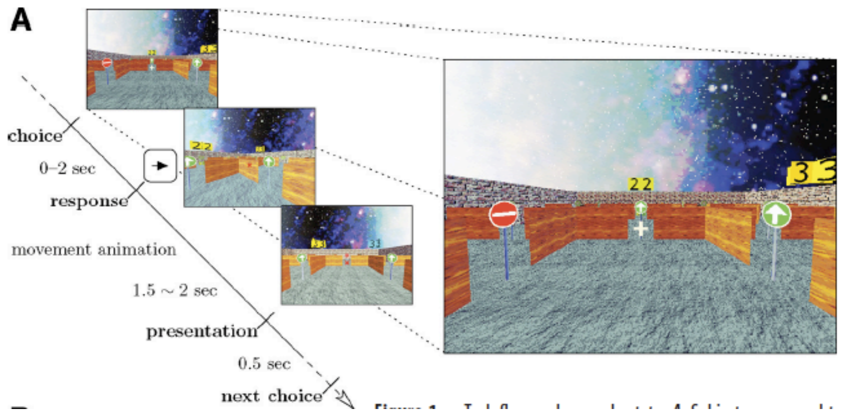

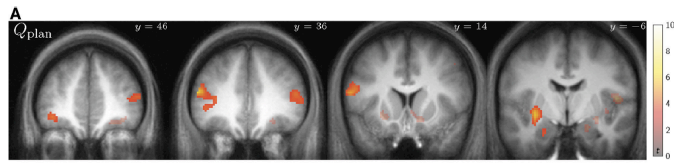

Simon and Daw., 2011

Balance between goal-directed vs. habitual choice

A mixed instrumental controller that solves an **exploration-exploitation dilemma**, trading off the costs of mental simulations with the value of information it produces

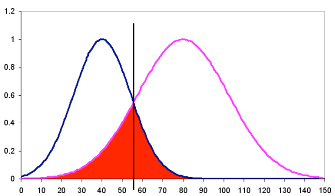

Uncertainty and variance of action value estimates

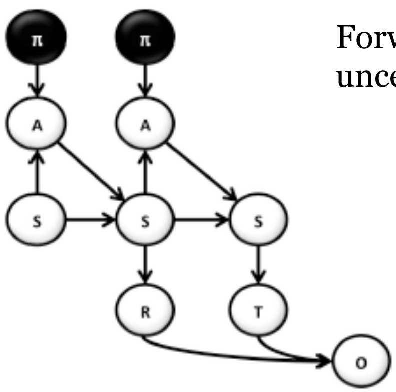

Forward sweeps to reduce uncertainty when necessary

Pezzulo et al. (sub)

Stable environment: habituation

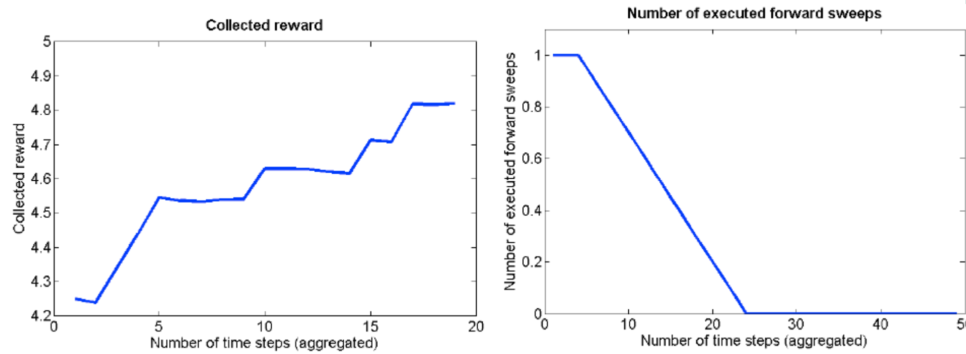

Changing reward contingencies

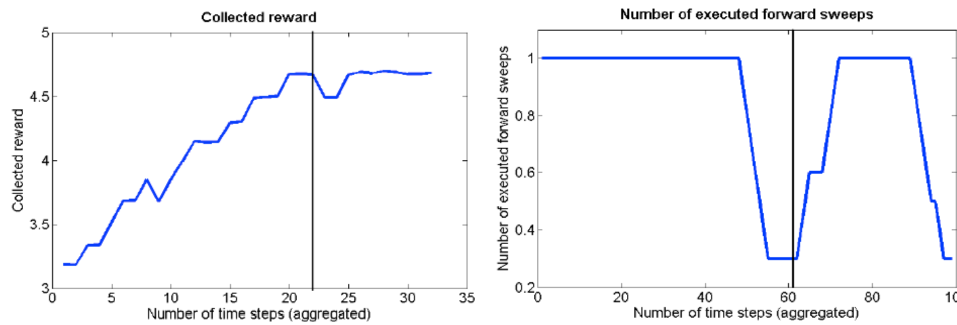

Multifunctionality:
Internal drive regulation of sensorimotor reflexes

Demo @ FET '11, Warsaw:
a catering assistant autonomous robot

Renno-Costa, Verschure

Integration of two loops: goal-directed navigation and face detection

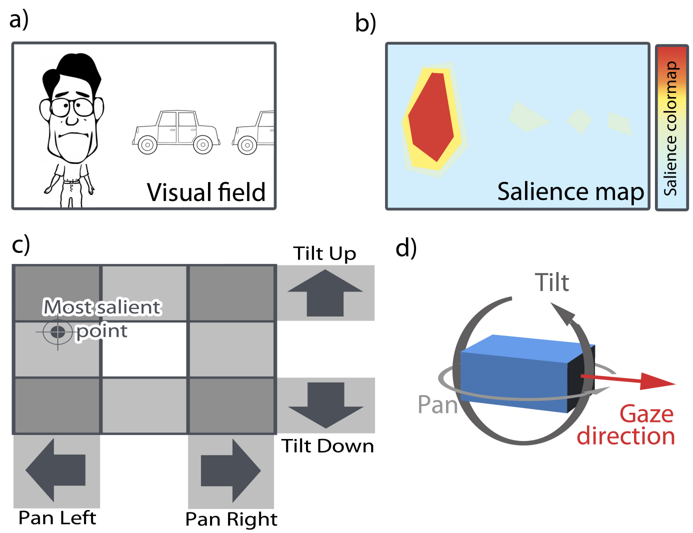

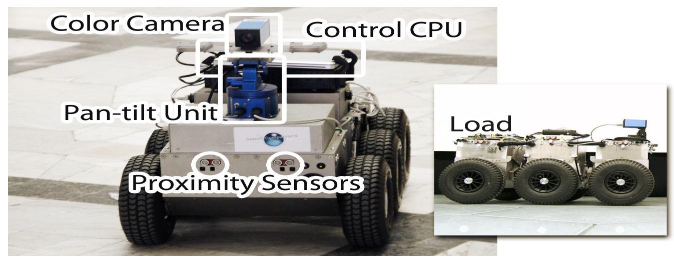

Hippocampal place fields modeling

POSTER No 97

Emerging single place field without plasticity using a spiking model

César Rennó-Costa[1], John E. Lisman[2], Paul F.M.J Verschure[1,3]

SPECS, Univesitat Pompeu Fabra, Barcelona, Spain; 2. Brandeis University, Boston, USA; 3. ICREA, Barcelona, Spain

Introduction

- Mean-field model emerges place fields instantaneously - and without plasticity - from the input of grid cells, found upstream in the medial entorhinal cortex.
- Model based on two mechanisms: the integration of massive convergent input and feedback inhibition, which is translated to the E%-MAX-WTA competition rule.
- Hierarchical topology can explain the formation of multiple place fields in the dentate gyrus and single place fields in the downstream region CA3.
- Can explain how place fields are affected by environmental changes in the rate remapping phenomena when non-spatial i[illegible]t is considered.
- To extend the [illegible]plica[illegible]bility of this model to more complex dynamics, incl[illegible]ding [illegible]me-dependence, we've developed an imp[illegible]eme[illegible]tion of its principles with spiking neurons using [illegible]egrat[illegible]nd-fir[illegible] units.

The Model

- Three layered network based on the hippocampal an[illegible]tomy: the entorhinal cortex, the dentate gyrus and [illegible]
- Each layer compromises a set of standard inte[illegible]rate-and-fire neurons.
- The forward connections are[illegible] conv[illegible]gent and random, with the weights set according to [illegible] distribution, and are elicited by spikes.
- At the dentate [illegible] a[illegible]d the [illegible] levels there are set associated i[illegible]ern[illegible]on ne[illegible]works capable of global inhibition[illegible]ese inte[illegible]rons are triggered after a specific interv[illegible] fro[illegible] the first spike detected in the population du[illegible] a cyc[illegible] This kind of circuitry generates overall population oscil[illegible]on as observed in these brain regions.

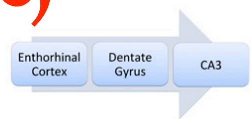

Figure 1: model hierarchy: from grid cells to place cells

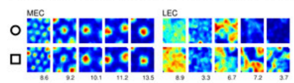

Figure 2: rate maps from grid cells and non-spatial cells

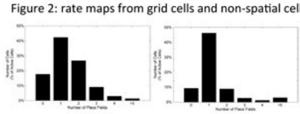

Figure 3: multiple place fields in the DG and single in the CA3

Results

- [illegible]he activity of the entorhinal cortex neurons at a certain time is set accordingly to the position of the virtual rat in a pre-defined trajectory. The activity of each cell mimics a grid cell with specific scale and spatial and angular offset. Though the model, the activity of the DG and CA3 neurons is computed.
- Spike activity is confronted with the position to obtain standard place cell analysis.
- Rate map analysis shows that the entorhinal cells exhibit grid cell formation. Place cells can be observed in both dentate gyrus and CA3 populations.
- A higher number of place fields per cell have been observed on the dentate gyrus.

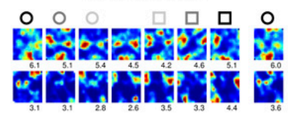

Figure 4: rate remapping in two DG cells

Conclusion

- Results show that the spike model produces similar results to the mean-field model.
- This establishes a checkpoint for further studies that might depend on spiking data such as well-timed features of the spatial code like phase precession and spike coincidence and to include other physiological facts such as the auto-associative connections in the CA3.

Acknowledgement

Supported by GoalLeaders project (FP7-ICT-97732)

Reference

L. de Almeida, M. Idiart, and J. E. Lisman, "A second function of gamma frequency oscillations: an E%-max winner-take-all mechanism selects which cells fire.," *The Journal of Neuroscience*, 29(23), pp. 7497-503, 2009.

L. de Almeida, M. Idiart, and J. E. Lisman, "The input-output transformation of the hippocampal granule cells: from grid cells to place fields.," *The Journal of Neuroscience*, 29 (23), pp. 7504-12, 2009.

L. de Almeida, M. Idart, and J. E. Lisman, "The single place fields of CA3 cells: A two-stage transformation from grid cells.", Hippocampus, 22(2), pp 200-8, 2010

C. Rennó-Costa, J. E. Lisman, and P. F. M. J. Verschure, "The mechanism of rate remapping in the dentate gyrus.," *Neuron*, 68(6), pp. 1051-8, 2010.

Renno Costa, Lisman, Verschure

Predicting objects dynamics

POSTER No 123

Learning to Simulate the Behavior of a Dynamical Object

Stefan Winberg and Christian Balkenius
Lund University Cognitive Science
Lund, Sweden

Lund University Cognitive Science

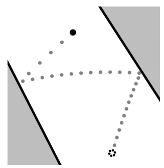

An example of prediction in a two dimensional environment. The large black circle is the ball being tracked, gray dots are the predictions and the black dots are a graphical representation of how future obstacles are detected.

The model predicts how a ball will fall down while bouncing on some obstacles.

Hesslow (2002) suggested that the brain simulates the future by first learning the expected outcome of an action or event and then uses this as input for a new prediction. As the process is repeated multiple times, the brain is able to simulate future consequences of an action. Here we apply this idea to the understanding of dynamical scenes where a system for prediction is converted into one for simulation.

To investigate how a system for prediction can be used for simulation we implemented the framework illustrated in the figure below. This framework consists of four modules that interact with each other and with the world in order to anticipate the temporal unfolding of an observed dynamic system.

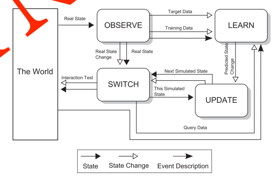

The model was implemented using the Ikaros system (Balkenius et al. 2010) where each of the components of the model were implemented as a separate module. A module reads data from its input connections in discrete time and generates new output at discrete intervals, usually referred to as ticks. These ticks can be locked to real-time when the system is used to control a robot and thus supports a seamless transition from simulation to a real robot.

We tested the accuracy of the predictions in a computer simulation using the open source physics engine Box2d (www.box2d.org). The simulated world was a two dimensional environment observed form the horizontal perspective, meaning that gravity pulled any objects down. A ball was launched from one of the edges of the world and as soon as the ball moved outside the boundaries of the world it was placed at a new starting position.

C. Balkenius, J. Morén, B. Johansson, and M. Johnsson, Ikaros: Building cognitive models for robots, Advanced Engineering Informatics, 24, 1, 40-48, 2010.
G. Hesslow, Conscious thought as simulation of behaviour and perception, Trends in Cognitive Sciences, 6, 6, 242-247, 2002.

GOAL LEADERS   SEVENTH FRAMEWORK PROGRAMME   LUND UNIVERSITY

Winberg, Balkenius

Vision hierarchies modulated by value

Hierarchical Models of Vision and Attention

Christian Balkenius and Birger Johansson
Lund University Cognitive Science
Lund, Sweden

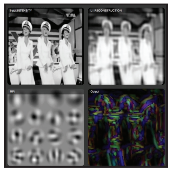

Top Left. Input image. Top Right. Reconstructed input from top-down attention signals. Bottom Left. Learned receptive fields. Bottom Right. Activity pattern in the first layer of the hierarchy. The color indicates the most active feature and the intensity represents the level of activation.

Image reconstructed from top-down expectations from the third hierarchial layer through layer 2 and 1.

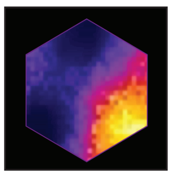

Activity pattern in one of the cortical modules.

We present a system-level computational model of visual learning and attention. The hexagons represent different cortical regions that are controlled by a central system consisting of subsystems for stimulus evaluation (S*), response evaluation (R*) as well as contextual modulation of learning and procesing (CX). The central system R is responsible for the production of 'innate' responses to visual stimuli.

The cortical part of the model learns on line to code for visual stimuli using hiearchical principal component analysys (HPCA). By combining features of self-organizing maps, convolutional networks and principal components analysis, a hierarchical network with an arbitrary number of layers [illegible]ganize from input data consisting of natural images. The hierarchical system is capable of both bottom-up analysis and top-down reconstruction of the visual input. The processing in the heirarchical system can also be modulated by value as coded by the S* system.

The S* system uses a form of classical conditioning while the R* system uses reinforcement learning. Finally, the CX system learns both to integrate stimuli over time to form a context and to modulete processing in the S*, R and R* systems.

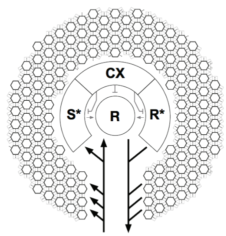

The model has been implemented within the Ikaros framework and exploits the duality between multiple self-organizing maps with weight-sharing and multiple parallel convolution operations to obtain real-time processing through the efficient utilization of multiple processor cores and hardware accelerated convolution operations.

POSTER No 7

Balkenius, Johansson

The Robot Builder

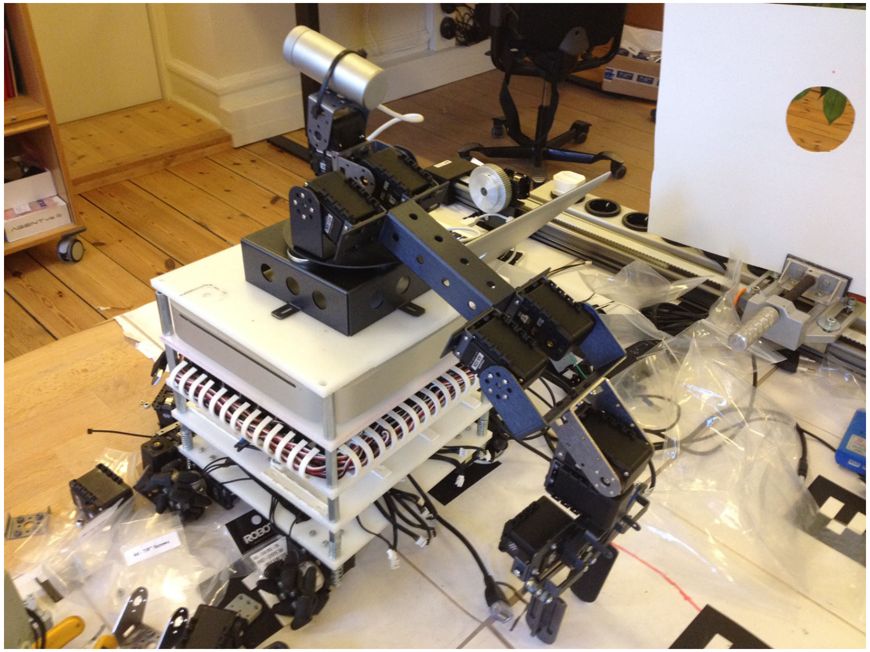

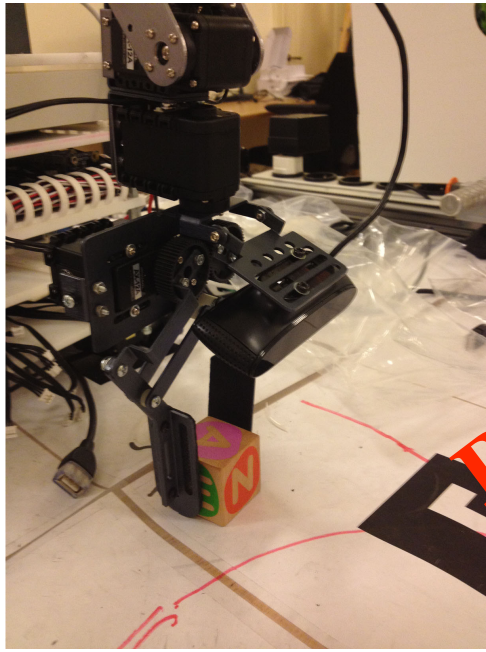

The Ikaros framework

A few planned achievements:

Action sequences, distal goals, subgoals, cognitive control

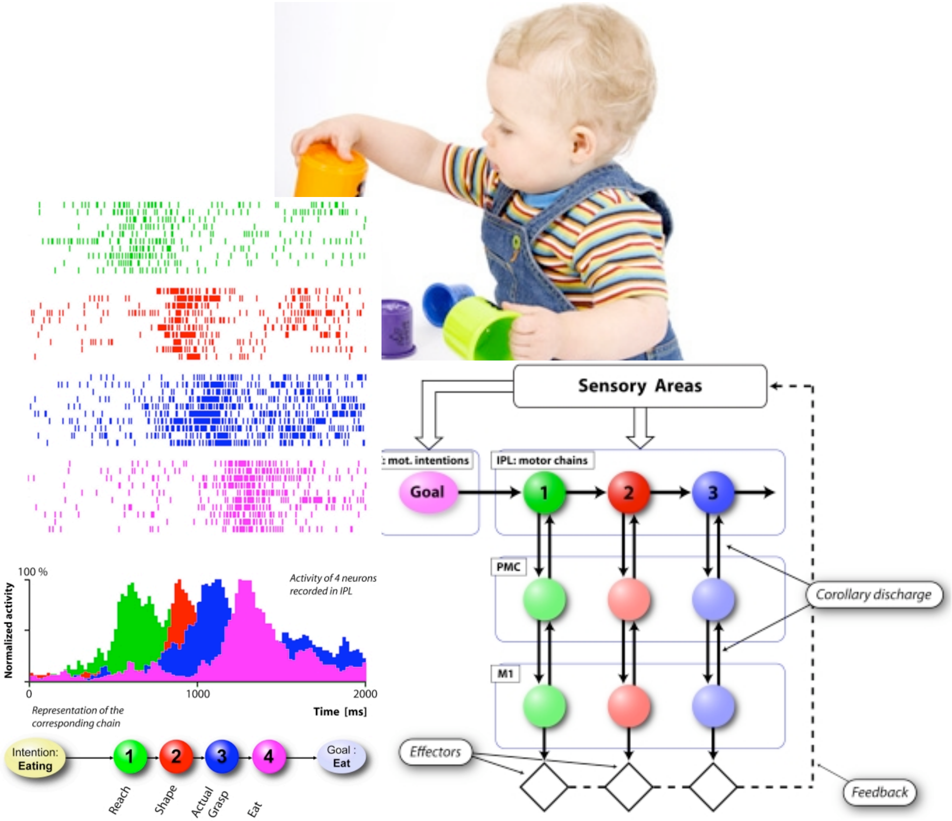

Fogassi et al., 2005 *Chersi et al, 2011*

Mental simulation Embodied problem solving

Thanks!

Our Advisory Board: Matthew Botvinick, Neil Burgess, Martin Butz, Michael Hasselmo, Bjorn Merker, Tony Prescott, Peter Redgrave

For these and other results, check:
www.goal-leaders.eu

Newsflash

Goal-Leaders Meeting @ Cogsys 2012
http://cogsys2012.acin.tuwien.ac.at/

Home

search...

Main Menu

- Home
- Aims of Goal-Leaders Project
- Project Partners
- Research & Products
- Publications
- News
- External links
- Contact us

Login Form

Username

Password

Remember Me

Login

- Forgot your password?
- Forgot your username?

Latest News

- Goal-Leaders Integration Meeting in Lund, March 7th 2012
- Goal-Leaders Meeting @ Cogsys 2012
- Goal-Leaders Y1 review meeting
- Goal-Leaders first semestral meeting
- Project Presentations

Popular

- IQR-DAC meeting
- Goal-Leaders first semestral meeting
- Kick-off Meeting
- Project Presentations
- Deliverables

Welcome to Goal-Leaders

Aims of Goal-Leaders project

The Goal-Leaders project aims at developing biologically-constrained architectures for the next generation of **adaptive service robots**, with unprecedented levels of **goal-directedness and proactivity**. Goal-Leaders will realize builder robots able to realize externally assigned tasks (e.g., fetching objects, composing building parts) and, at the same time, keeping their homeostatic drives within a safe range (e.g., never end up without energy or get hurt), by operating autonomously and for prolonged periods of time in open-ended environments.

To this aim, Goal-Leaders will pursue a combined neuroscientific, computational and robotic study of three key sets of competences:

-the biological and computational mechanisms behind **an agent's goal system**, which integrates somatic, affective, and cognitive elements, and realizes the setting and selection of the agent's goals;

-the biological and computational mechanisms that support the **assignment of situation-dependent value to an agent's state and action representations**, and therefore realize the link between the agent's goal system and perceptual-motor processes;

-the biological and computational mechanisms behind **an agent's anticipatory and mental simulation abilities.**

Who's Online

We have 1 guest online

Robot Builder Scenario

The Goal-Leaders achievements beyond the state of the art will be assessed against behavioral and neuroscientific data, and by realizing three demonstrators in which robots will perform **autonomous navigation and construction tasks**, and will readapt without reprogramming to novel task allocations or changes in their environment.

The Robot Builder scenario

Goal-Leaders Project

Goal-Leaders is a 3-year project funded by the European Union within the "Cognitive Systems and Robotics" initiative.

Short Facts:

Consortium: National research Council (Rome, Italy), University Pompeu Fabra (Barcelona, Spain), Universiteit van Amsterdam (Amsterdam, the Netherlands), Lund University (Lund, Sweden)

Start: March 2011; End: February 2013

Project acronym: Goal-Leaders

Project full title: Goal-directed,

Advertisement: Open PhD Position

PhD subject: Visual perception and motor anticipation in biological and artificial systems. Keywords: eye tracking, attention, prediction, neural networks. Driving fields: robotics, psychology, computational neuroscience

Starting date: position already available,

Affiliation: LAPSCO, Pascal Institute (Clermont-Ferrand, Fr)
Contact me or j-charles.quinton@univ.bpclermont.fr