Open-Ended Robot Learning about Objects and Activities

(IROS'2019 Workshop on Open-Ended Learning for Object Perception and Grasping: Current Successes and Future Challenges, Macau, Nov. 8th, 2019)

Luís Seabra Lopes(*)

Universidade de Aveiro, Portugal

(*) Participation funded by FCT project UID/CEC/00127/2019 and Fundação Oriente

Abstract

- If robots are to adapt to new users, news tasks and new environments, they will need to conduct a long-term learning process to gradually acquire the knowledge needed for that adaptation.
- One of the key features of this learning process is that it is open-ended.
- The Intelligent Robotics and Systems group of the University of Aveiro has been carrying out research on open-ended learning in robotics for more than a decade.
  - Different learning techniques were developed for object recognition, grasping and task planning.
  - These techniques build upon well established machine learning techniques, ranging from instance-based learning and bayesian learning to abstraction and deductive generalization.
  - Our approach includes the human user as mediator in the learning process. Key features of open-ended learning will be discussed.
  - New experimental protocols and metrics were designed for open-ended learning and will also be presented.

2000

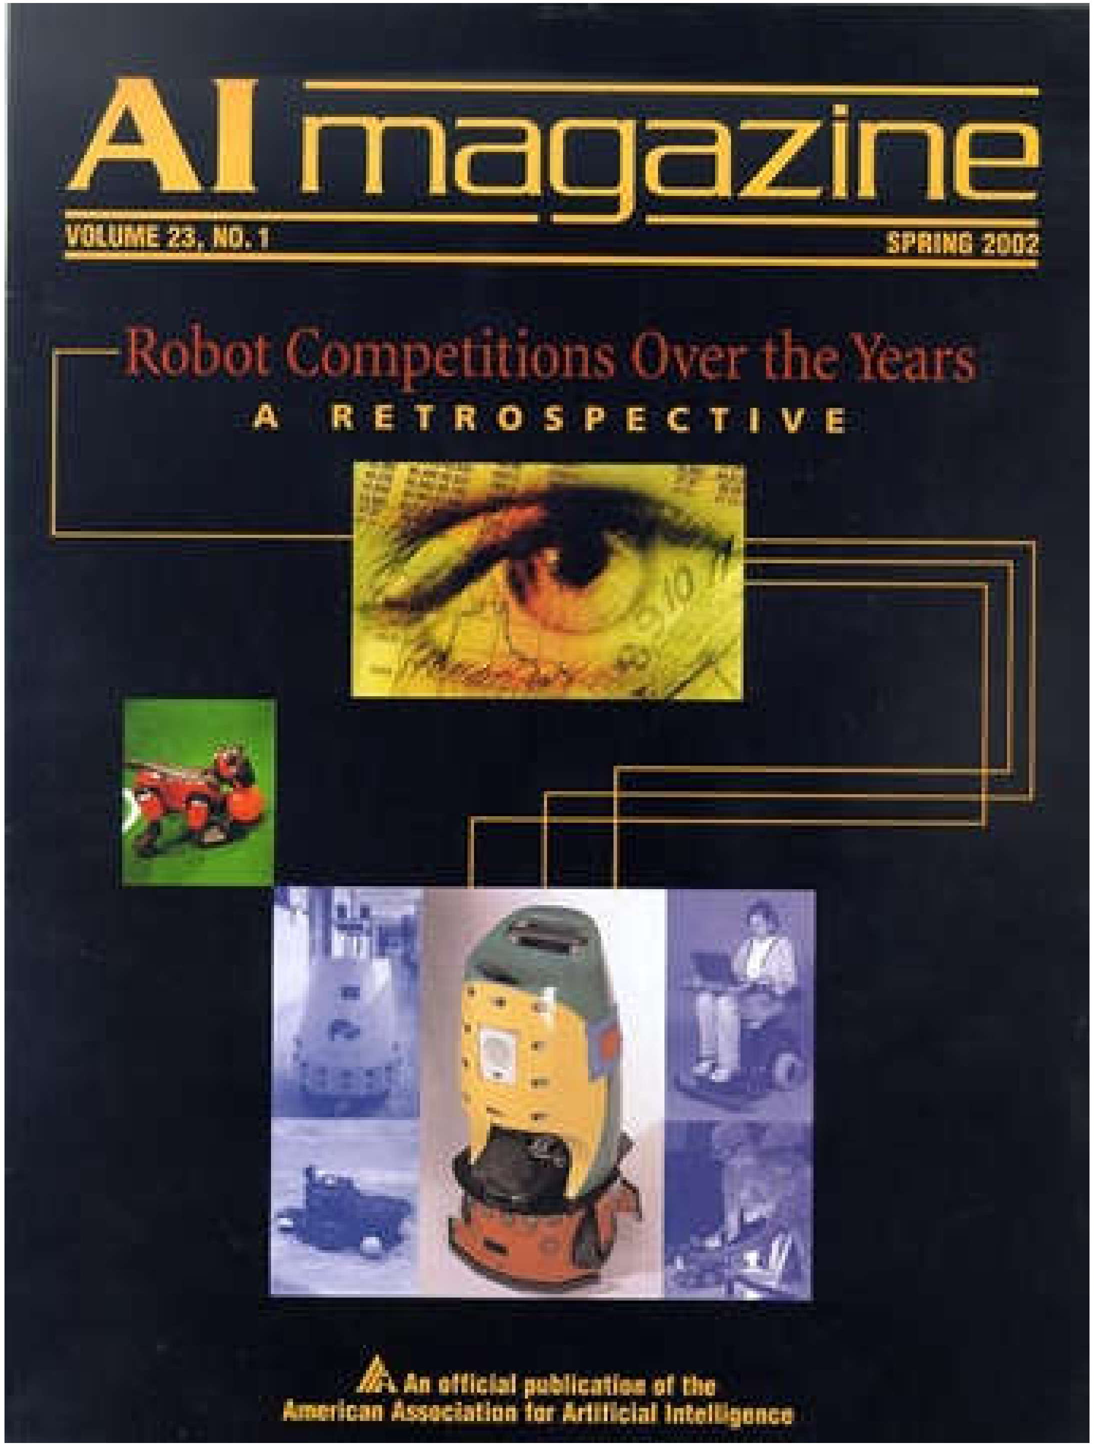
AI magazine
VOLUME 23, NO. 1
SPRING 2002
Robot Competitions Over the Years
A RETROSPECTIVE
An official publication of the
American Association for Artificial Intelligence

2002

Intelligent Robots

- **Animate** - Respond in real-time to changing conditions in the environment
- **Adaptable** - adapt to different tasks, users and environments
- **Accessible**- explain beliefs, motivations and intentions; be easy to instruct

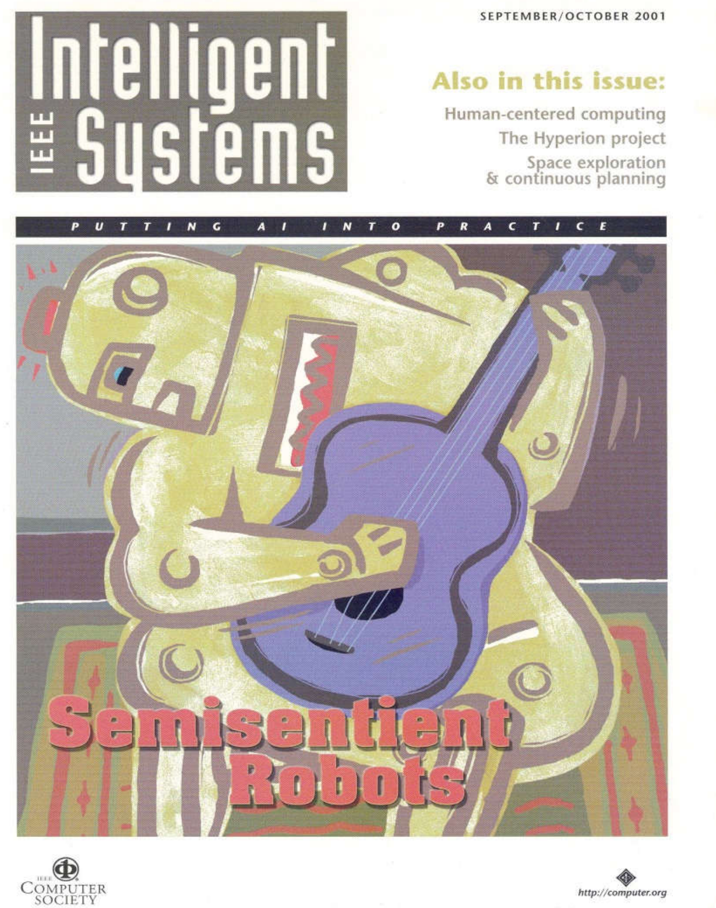

Seabra Lopes, L. Connell & JH (2001) Semisentient Robots: Routes Integrated Intelligence, *IEEE Intelligent Systems*, 16 (5), 10-15.

Capabilities of Intelligent Robots

- A tight coupling of perception and action, to be animate
- Reasoning, to be adaptive
  - High-level interpretation capabilities, for updating a model of the state of the world
  - Planning, to determine sequences of action to achieve the given goals
- Learning, to be even more adaptive
- Memory, to store the world model and learned knowledge
- Speaking the language of the user, to be accessible

Learning for intelligent robots

- **Supervised** — to include the human instructor in the learning process.
- **On-line** — so that learning takes place while the agent is running.
- **Opportunistic** — the system must be prepared to accept a new example when it is observed or becomes available, rather than according to a pre-defined training schedule.
- **Incremental** — it is able to adjust the learned descriptions when a new example is observed.
- **Open-ended** — it is able to acquire and incorporate new concepts.
- **Concurrent** — it is able to handle multiple learning problems at the same time.
- **Meta-learning** — it is able to determine which learning parameters and configurations are more promising for different problems, ensuring each problem is handled in the best way.

Seabra Lopes, L. and A. Chauhan (2007) **How many Words can my Robot learn? An Approach and Experiments with One-Class Learning**, *Interaction Studies*, 8(1), p. 53-81.

Intrinsic Motivation and Meta-Learning

- Intrisic motivation is a set of hard-wired internal rewards and mechanisms that drive the behavior of an agent
  - Especially when exploring and learning about the environment
  - Intrinsic meta-learning mechanisms can be essencial to maximize learning success

Seabra Lopes, L. and Q.H. Wang (2002) **Towards Grounded Human-Robot Communication**, *Proc. 11th IEEE Int'l Workshop on Robot and Human Interactive Communication (ROMAN'2002)*, Berlin, Germany, p. 312-318.

Oudeyer P-Y, Kaplan , F. and Hafner, V. (2007) **Intrinsic Motivation Systems for Autonomous Mental Development**, IEEE Transactions on Evolutionary Computation, 11(2), p. 265-286

Architecture (2007)

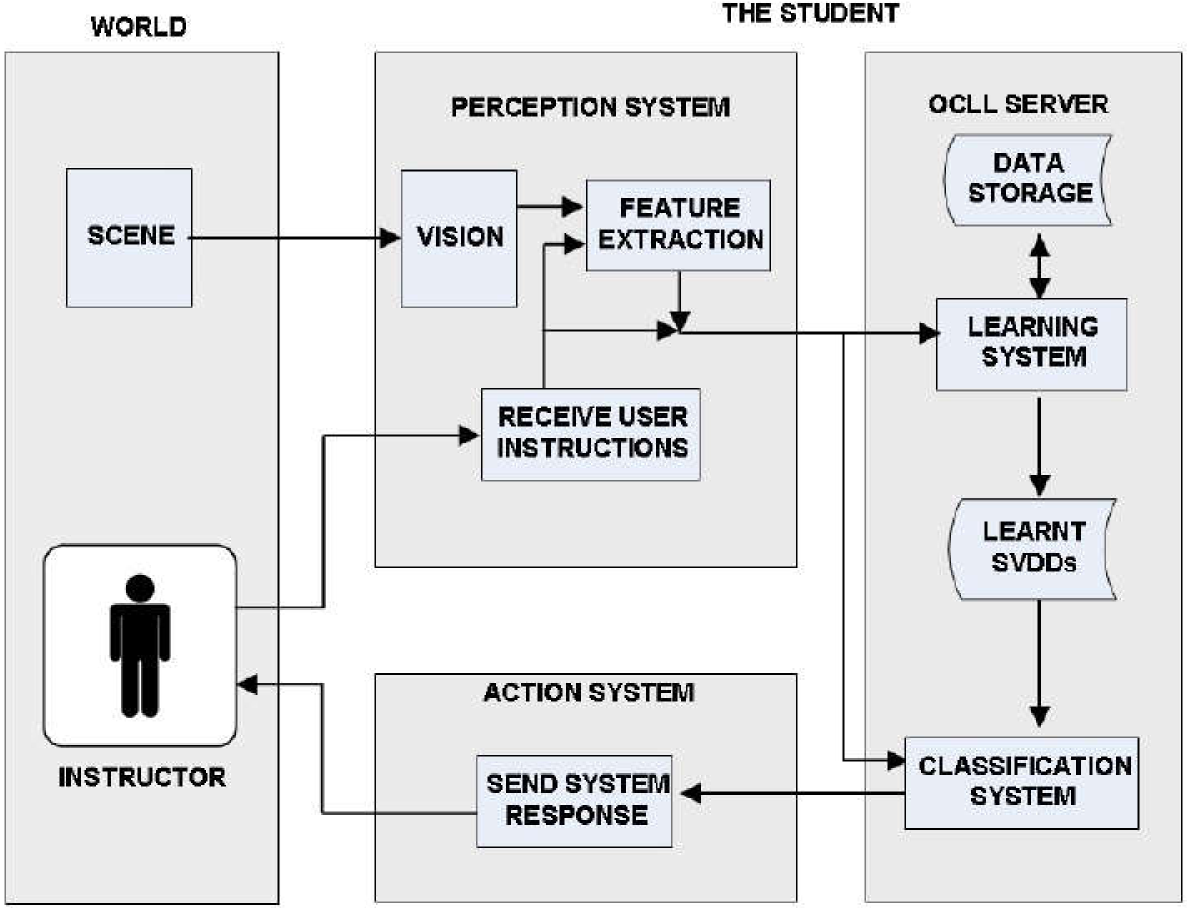

Instructor actions

- Point
- Teach
- Ask
- Correct

Experimental setup (2007)

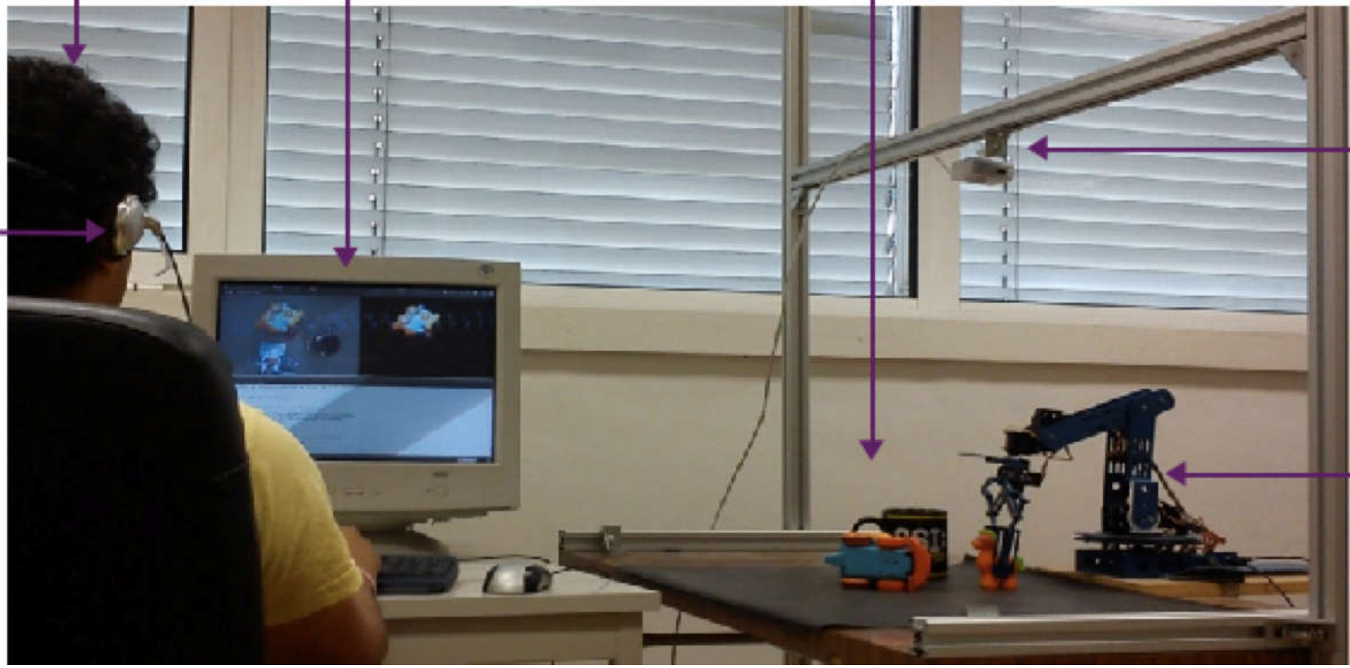

Evolution – Recovery – Breakpoint

- **Evolution** - Depends on the ability of the learner to adjust category representations when a new word is introduced.
- **Recovery** - The discrimination performance will generally deteriorate with the introduction of a new word. The time spent in system evolution until correcting and adjusting all current categories defines recovery. Recovery is based on classification errors and corresponding corrective feedback.
- **Breakpoint** - Inability of the learner to recover and evolve when a new category is introduced.

Experiments: teaching protocol

```
Teach(Category_1)
n ← 1
repeat
        n ← n + 1
        Teach(Category_n)
        k ← 0
        c ← 1
        repeat
                x ← previously unseen instance of Category_c
                C_x ←  Ask(x)
                if C_x != Category_c:
                        Correct(x,Category_c)
                c ← c + 1 if c<n else 1
                k ← k+1 if k ≤ 3n else k
                A ← average accuracy in last k question/correction iterations
        until ((A > A_min and k ≥ n) or
                (teacher sees no improvement in protocol accuracy))
until (teacher sees no improvement in protocol accuracy)
```

Metrics and plots

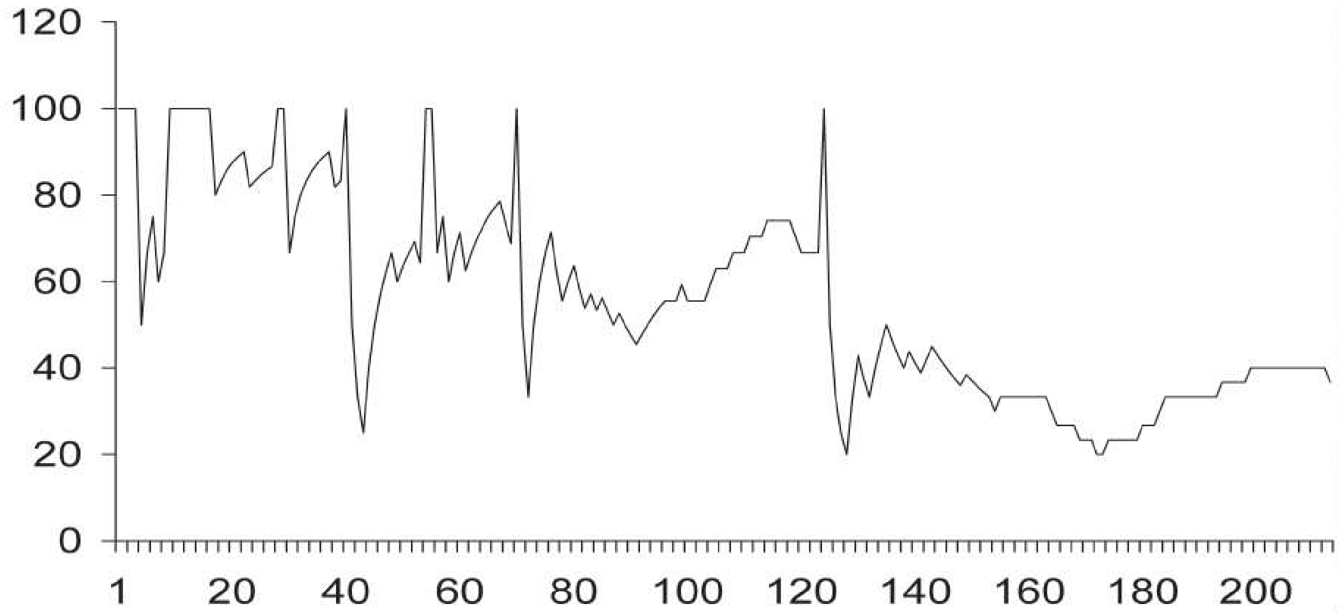

Figure 7. Evolution of classification precision versus number of question/correction iterations

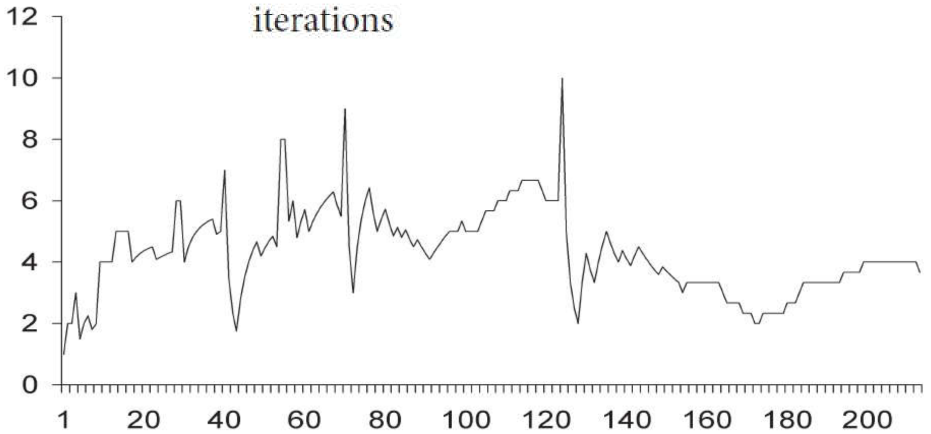

Figure 8. Evolution of learning efficiency versus number of question/correction iterations

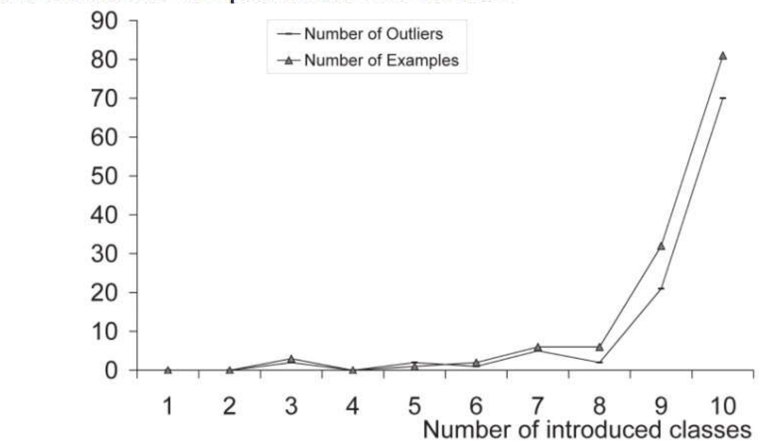

Figure 9. Number of outliers and positive examples added after each new class was introduced in the experiment of Figures 7–8

Architecture (2008)

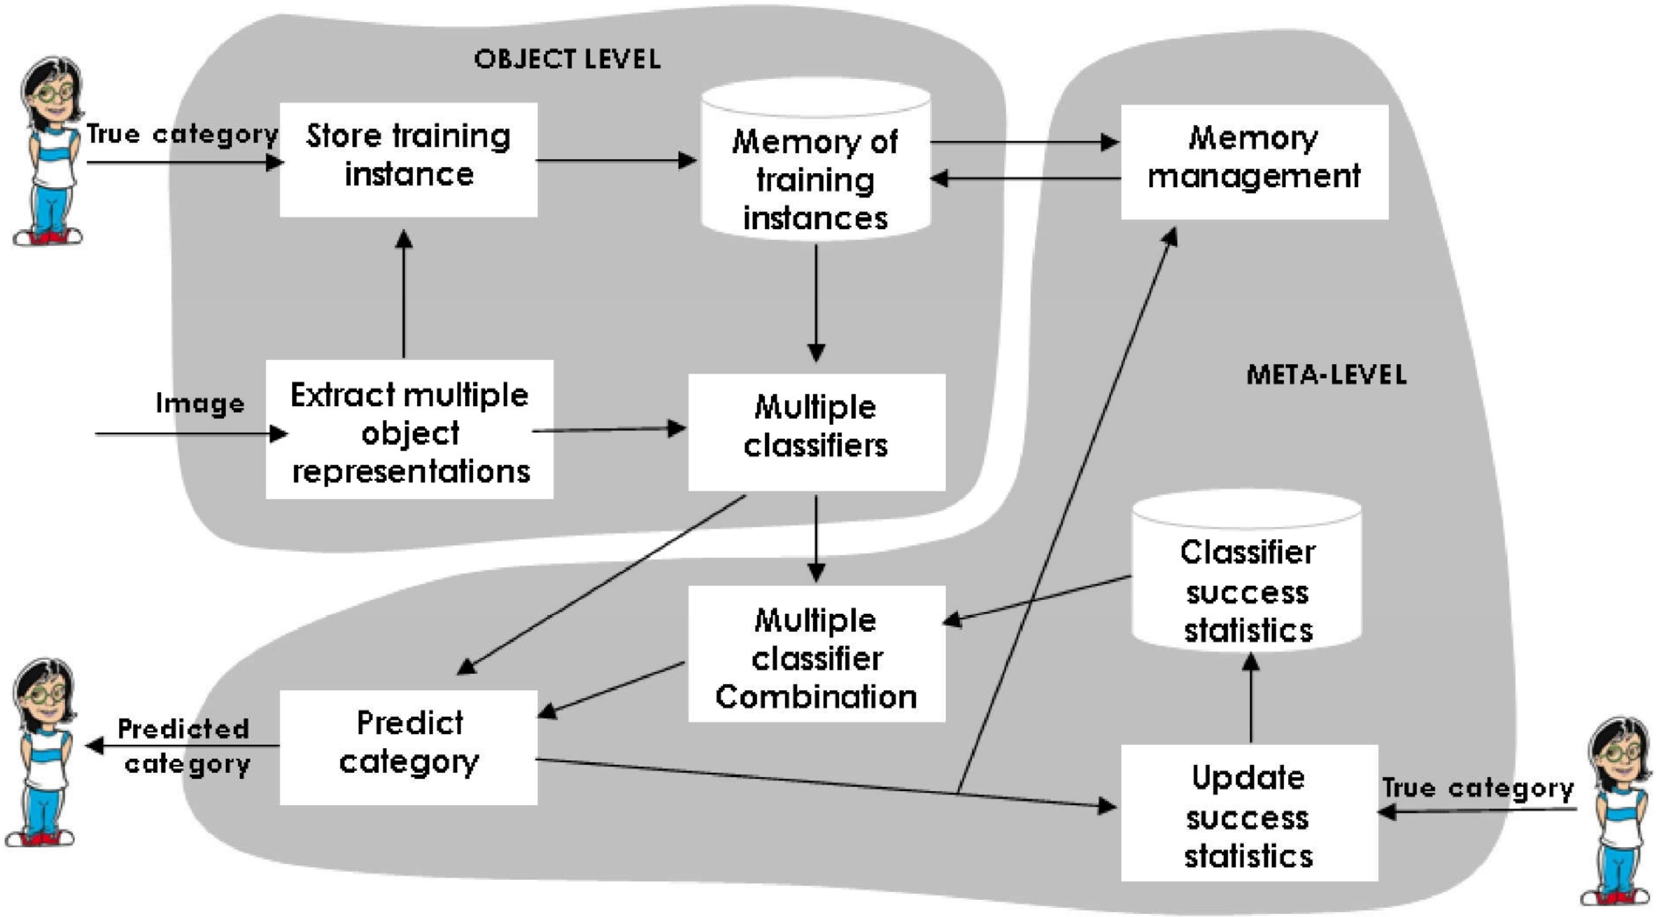

Seabra Lopes, L. and A. Chauhan (2008) **Open-Ended Category Learning for Language Acquisition**, _Connection Science_, 20 (4), 277-297.

A longer experiment – 68 objects

Evolution of classification precision versus number of question/correction iterations

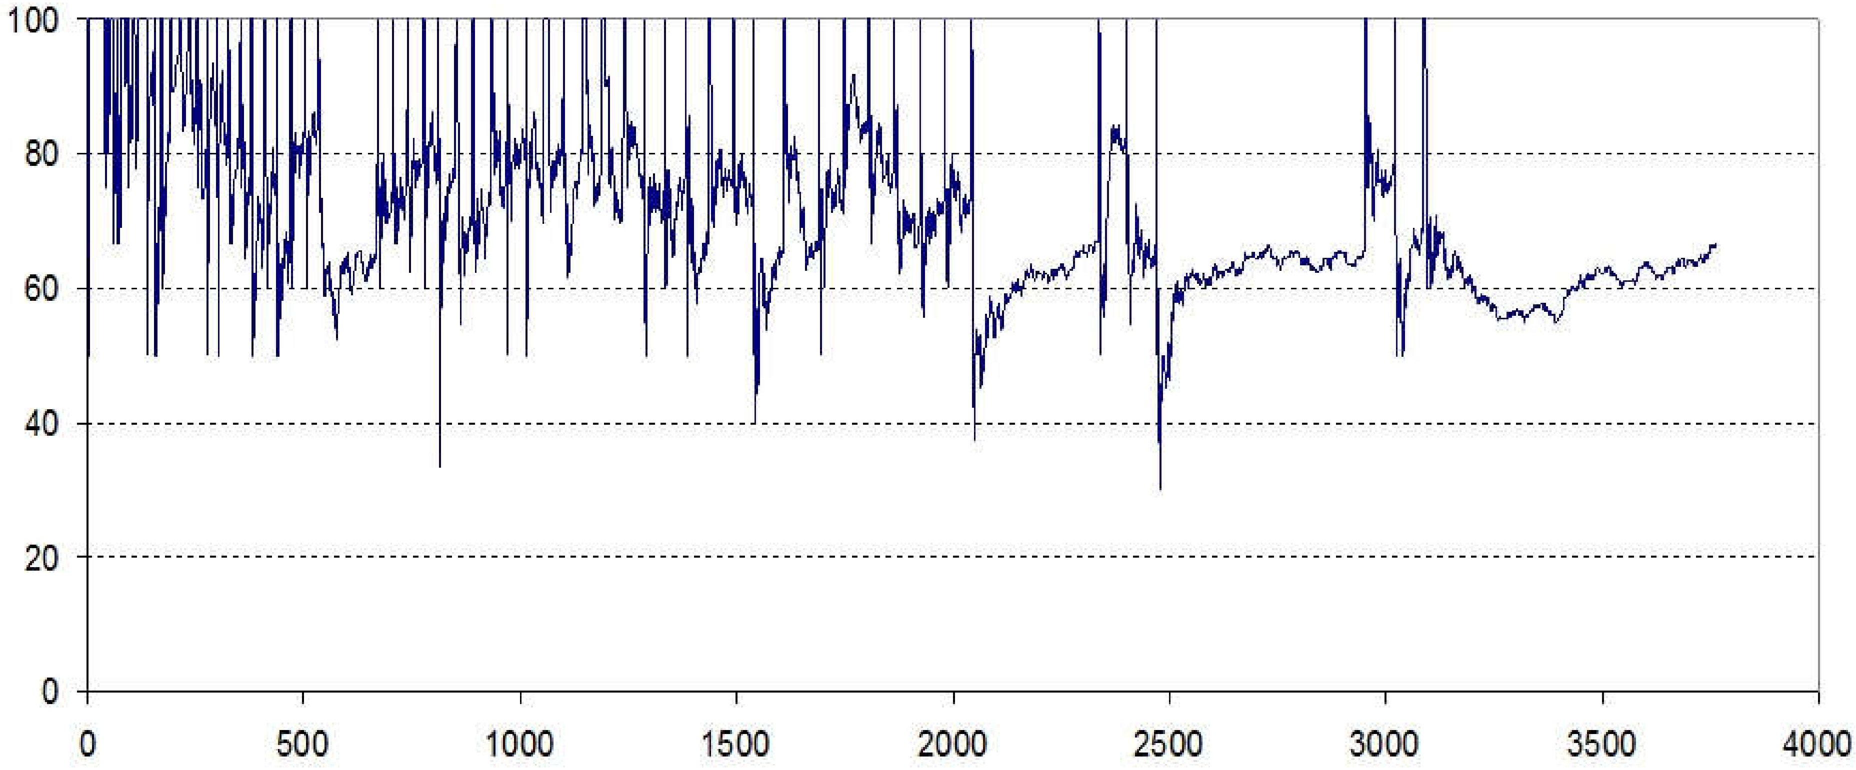

Simulated Teacher

- Take over the task of the human teacher
- Follow the **teaching protocol** and interact with the learning agent using the **teach**, **ask** and **correct** actions
- Given an object-images **dataset**, the simulated teacher picks images randomly for interaction with the learning system and uses each stored image at most once
- When the learning agent is ready to learn a new category, the simulated teacher randomly selects and teaches the next category

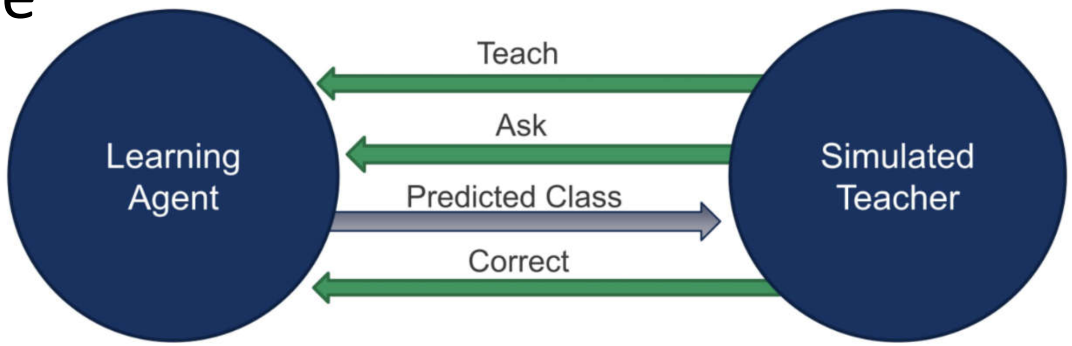

Grounding spoken words

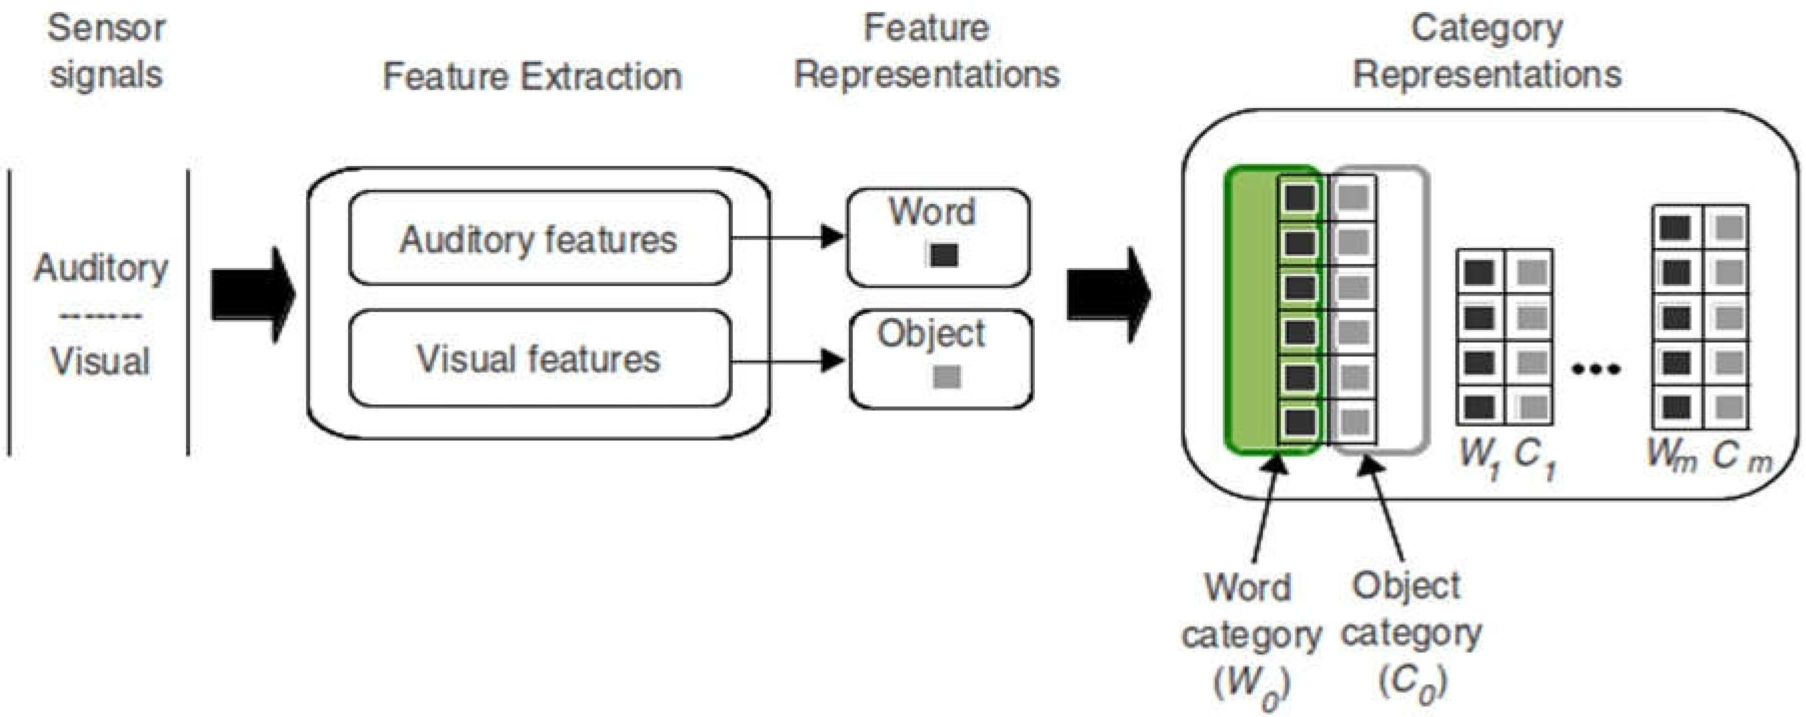

Chauhan, A., and L. Seabra Lopes (2011) **Using Spoken Words to guide Open-ended Category Formation**, <u>*Cognitive Processing*</u>, vol. 12(4), p. 341-354.

A long experiment (293 cats.)

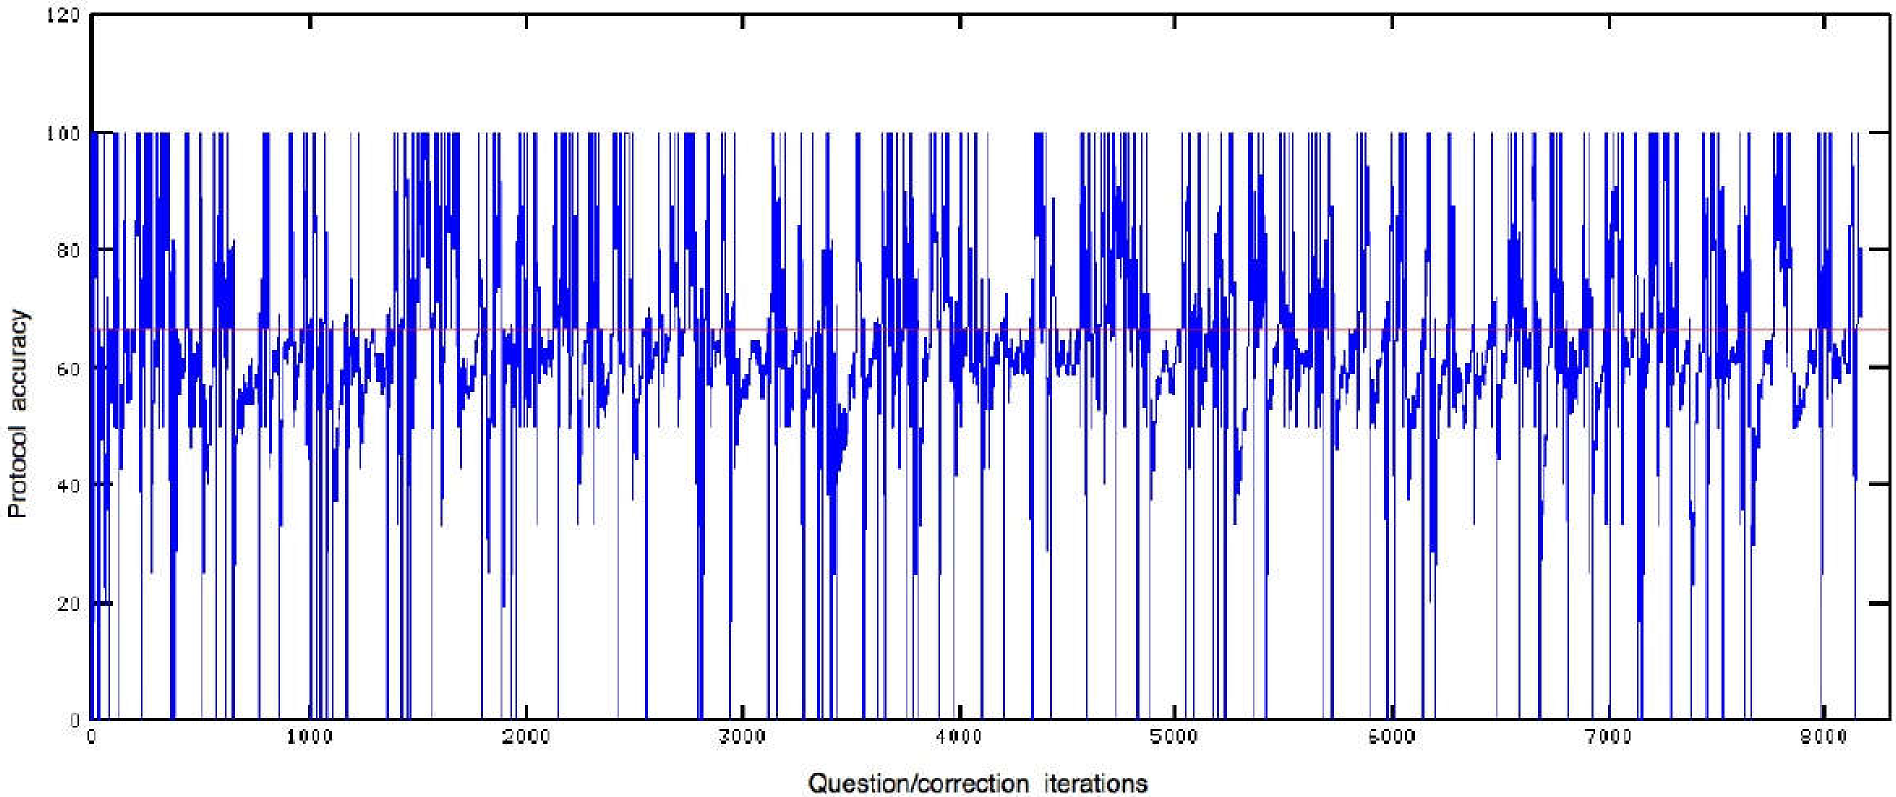

Chauhan, A., L. Seabra Lopes (2015) **An Experimental Protocol for Evaluation of Open-Ended Category Learning Algorithms**, *IEEE Conference on Evolving and Adaptive Intelligent Systems (EAIS'2015)*, Douai, France.

Results

Dataset	Learning approach	Stopping condition	# exps.	#cats learned	#iters	#inst/cat	Avg. protocol acc.(%)	Global acc.(%)
LANGG68	SVDD	Lack of data	5	25.0	1070.4	19.5	53.2	56.7
	COMP	Lack of data	2	14.5	761.5	26.1	50.4	52.2
		Breakpoint	3	8.3	140.3	8.22	52.5	60.4
	MCML	All cats. learned	5	68.0	855.6	3.8	77.0	78.1
COIL-100	SVDD	Lack of data	3	9.0	322.0	19.9	46.0	49.6
		Breakpoint	2	7.5	141.5	10.0	47.9	53.9
	COMP	Lack of data	4	17.5	530.0	13.3	54.7	58.6
		Breakpoint	1	10	158.0	7.5	48.5	57.6
	MCML	All cats. learned	1	100.0	2319.0	8.3	63.2	68.4
		Lack of data	4	64.3	1385.0	7.8	63.6	68.1
ETH80	SVDD	Breakpoint	5	3.6	35.0	5.4	49.9	59.0
	COMP	Breakpoint	5	4.2	66.4	8.1	46.1	50.4
	MCML	All cats. learned	5	8.0	141.6	8.3	53.7	59.9
ALOI-1000	SVDD	Breakpoint	5	3.4	54.2	8.5	41.3	47.3
	MCML	Lack of data	5	259.4	7789.0	10.6	64.4	68.0

Axes of evaluation

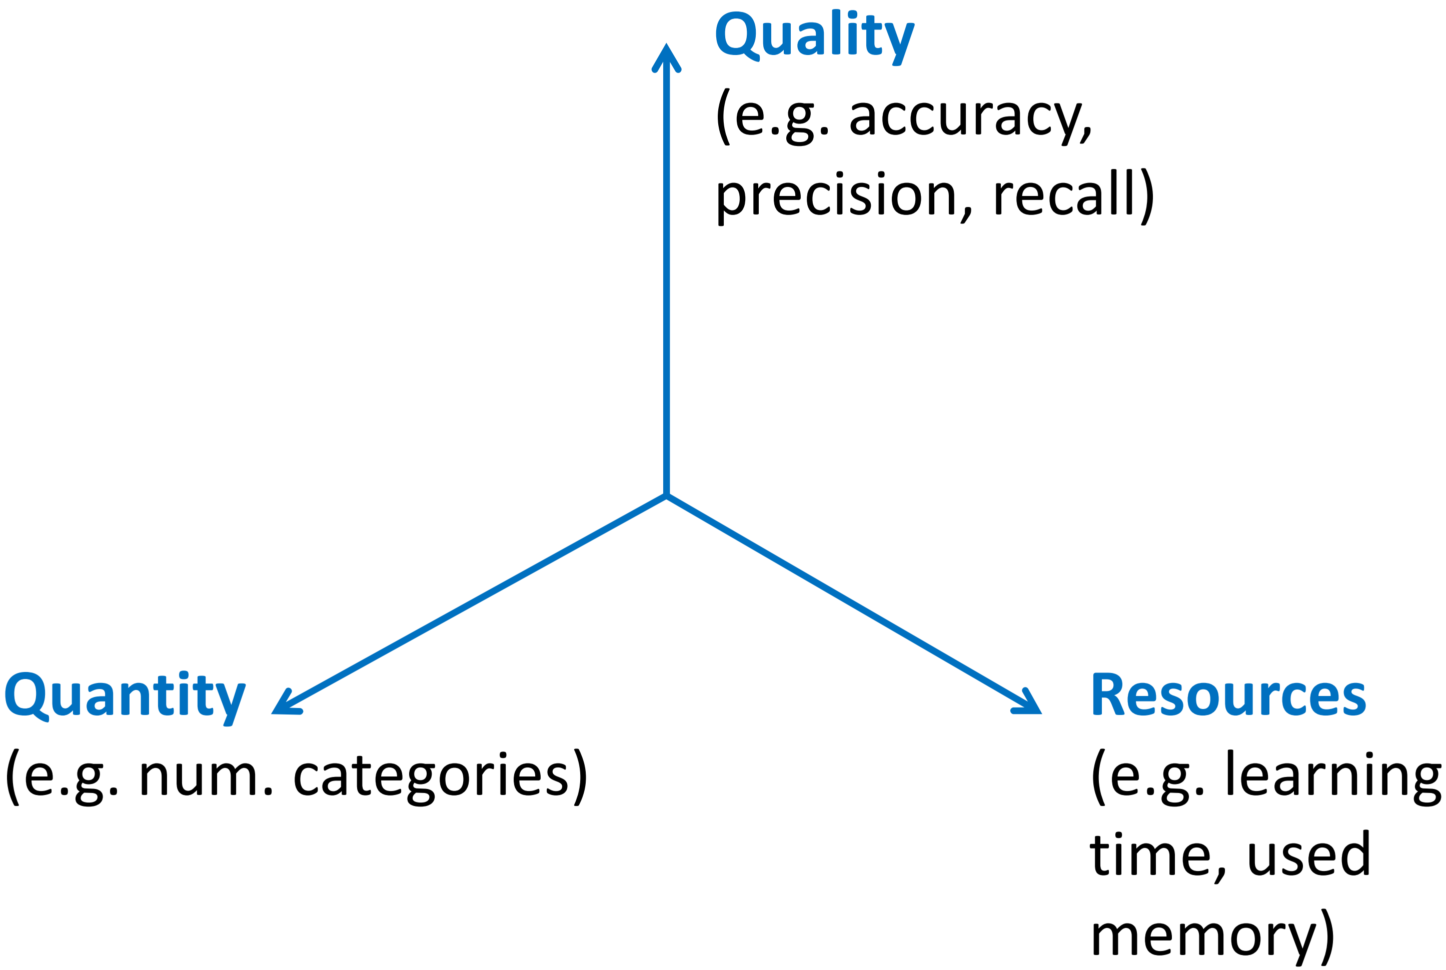

RACE – 2011-2014

- RACE = Robust Autonomous Competence Enhancement
- FP7
- Hamburg, Leeds, Orebro, Osnabrück, Aveiro
- Aveiro:
  - Interactive open-ended learning about objects and activities

Experience Extraction

- Experiences typically abstract from low-level or irrelevant data, given the intended use of experiences
- The information in a raw experience can be subject to
  - Representation change
  - Downsampling
  - Heuristics
    - Temporal segmentation
    - Filtering

Teaching the category of an object

- The instructor points to an object (a Mug) with the hand
- In the graphical interface, a menu can be used to label de object as Mug
- After that, the object is recognized as Mug

Teaching how to serve a coffee

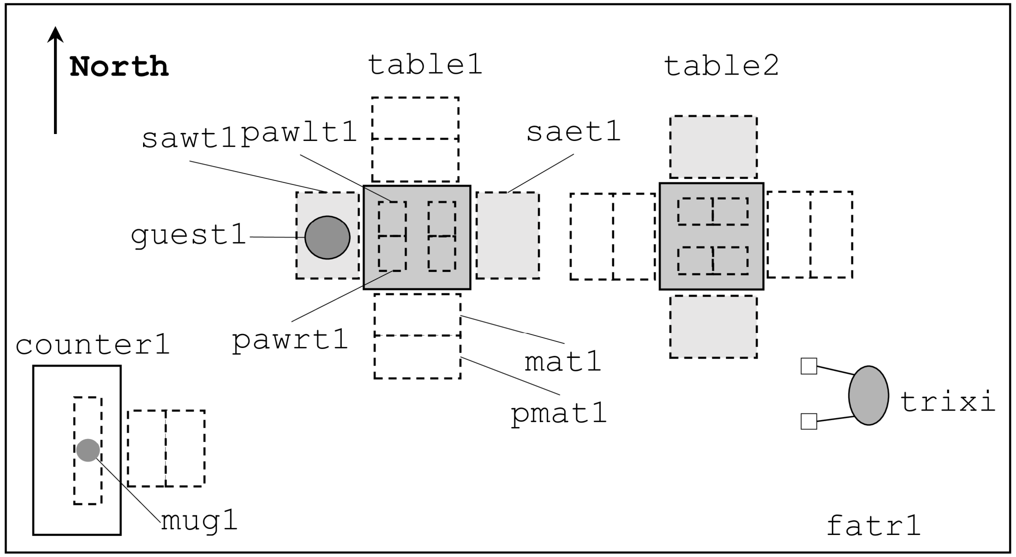

- **Achieve** drive_robot preManipulationAreaEastCounter1
- **Achieve** grasp_object_w_arm mug1 rightArm1
- **Achieve** drive_robot preManipulationSouthTable1
- **Achieve** put_object mug1 placingAreaWestRightTable1
- **Teach_task** serve_coffee guest1

Very different experiences

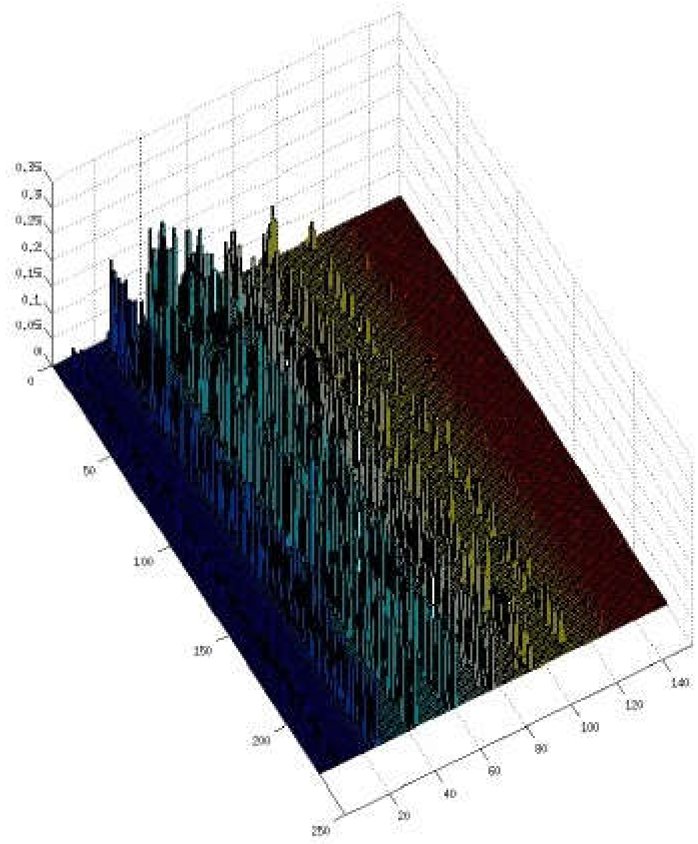

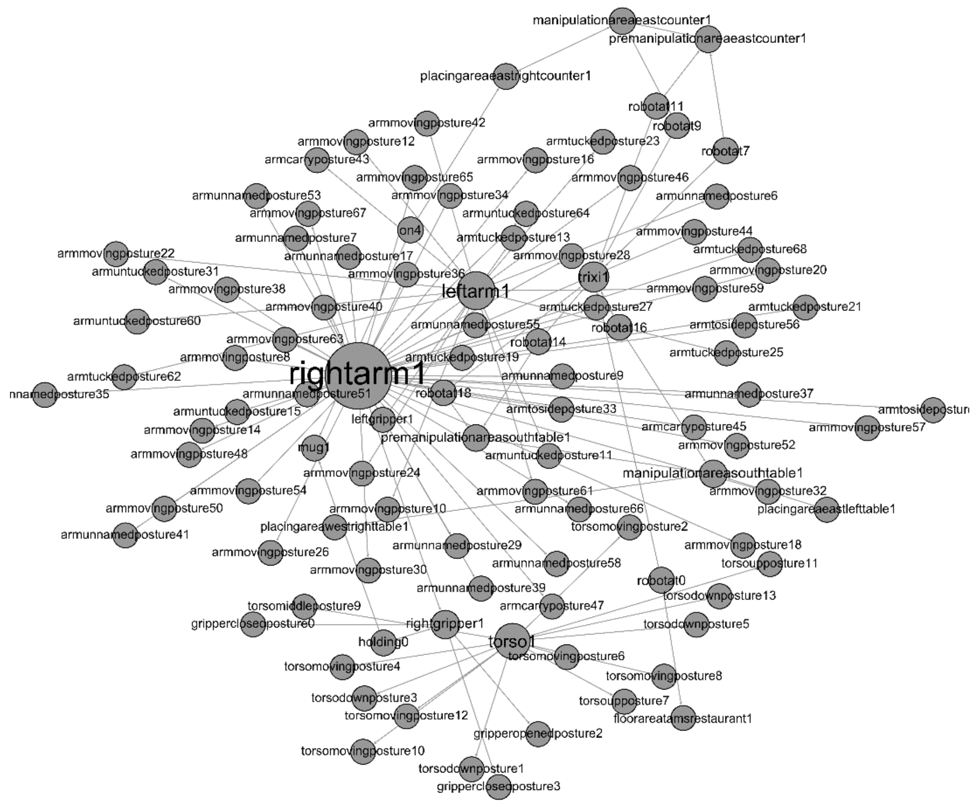

Concepts

- Conceptualization is the process of forming a new concept by assigning a description and a name to a real-world pattern
  - Concepts can be generated from one or multiple experiences
- Learned concepts enable to:
  - Recognize instances, e.g. recognize an object as being a Mug
  - Generate solutions to concrete problems, e.g. Generate a plan to correctly serve a coffee to a guest

Semantic versus perceptual information

- Semantic information
  - Symbolic
  - Explicit / declarative / relational
  - Slow processing
- Perceptual information
  - Numeric
  - Implicit / pattern based
  - A lot of data
  - Requires fast processing
- Connecting language / semantic information to perception
  - Grounding concept names
  - Anchoring object names

RACE – architecture discussion

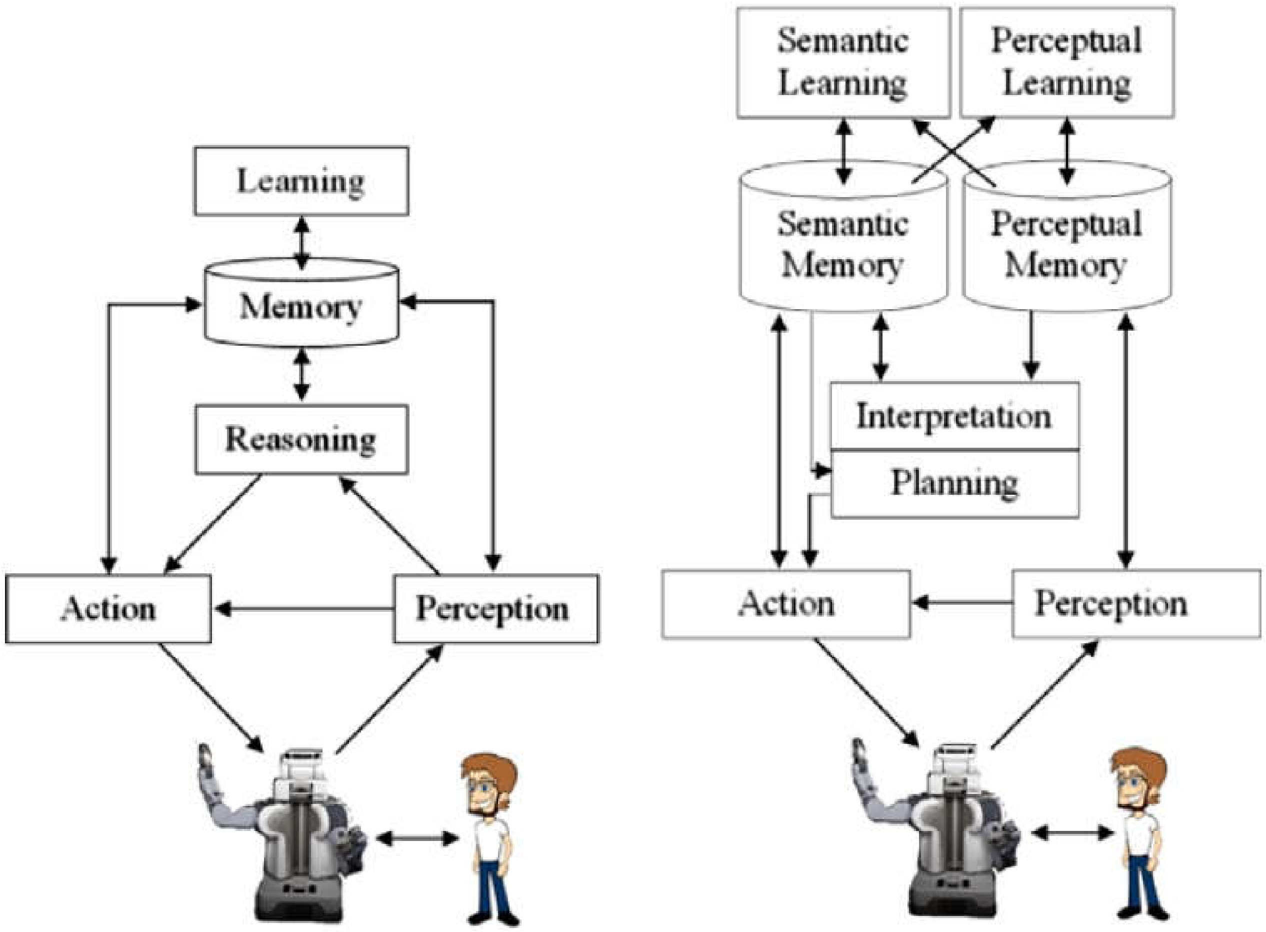

Oliveira, M., L. Seabra Lopes, H. Kasaei, G.H. Lim, A.M. Tomé, A. Chauhan (2016) **3D Object Perception and Perceptual Learning in the RACE Project**, <u>*Robotics and Autonomous Systems*</u>, 75, p. 614-626.

Perception and perceptual learning

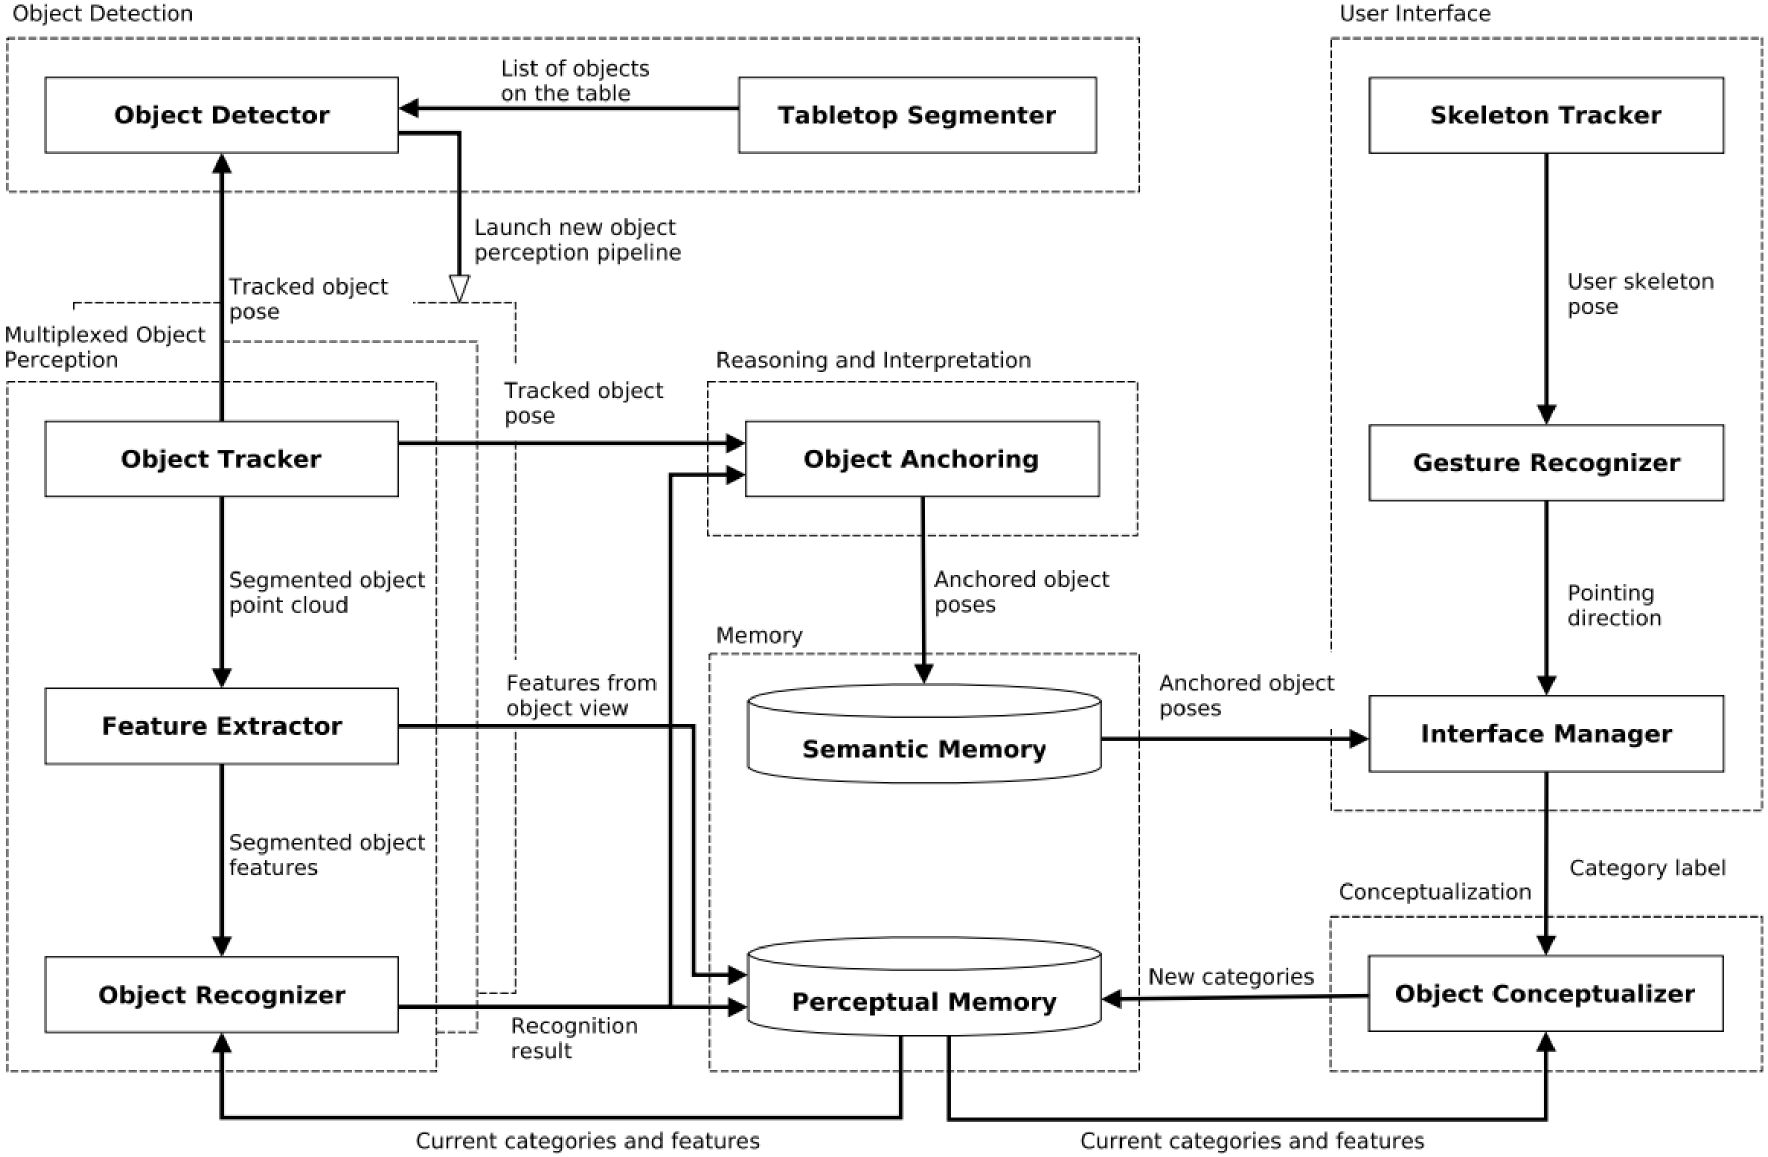

Objects at multiple levels of abstraction

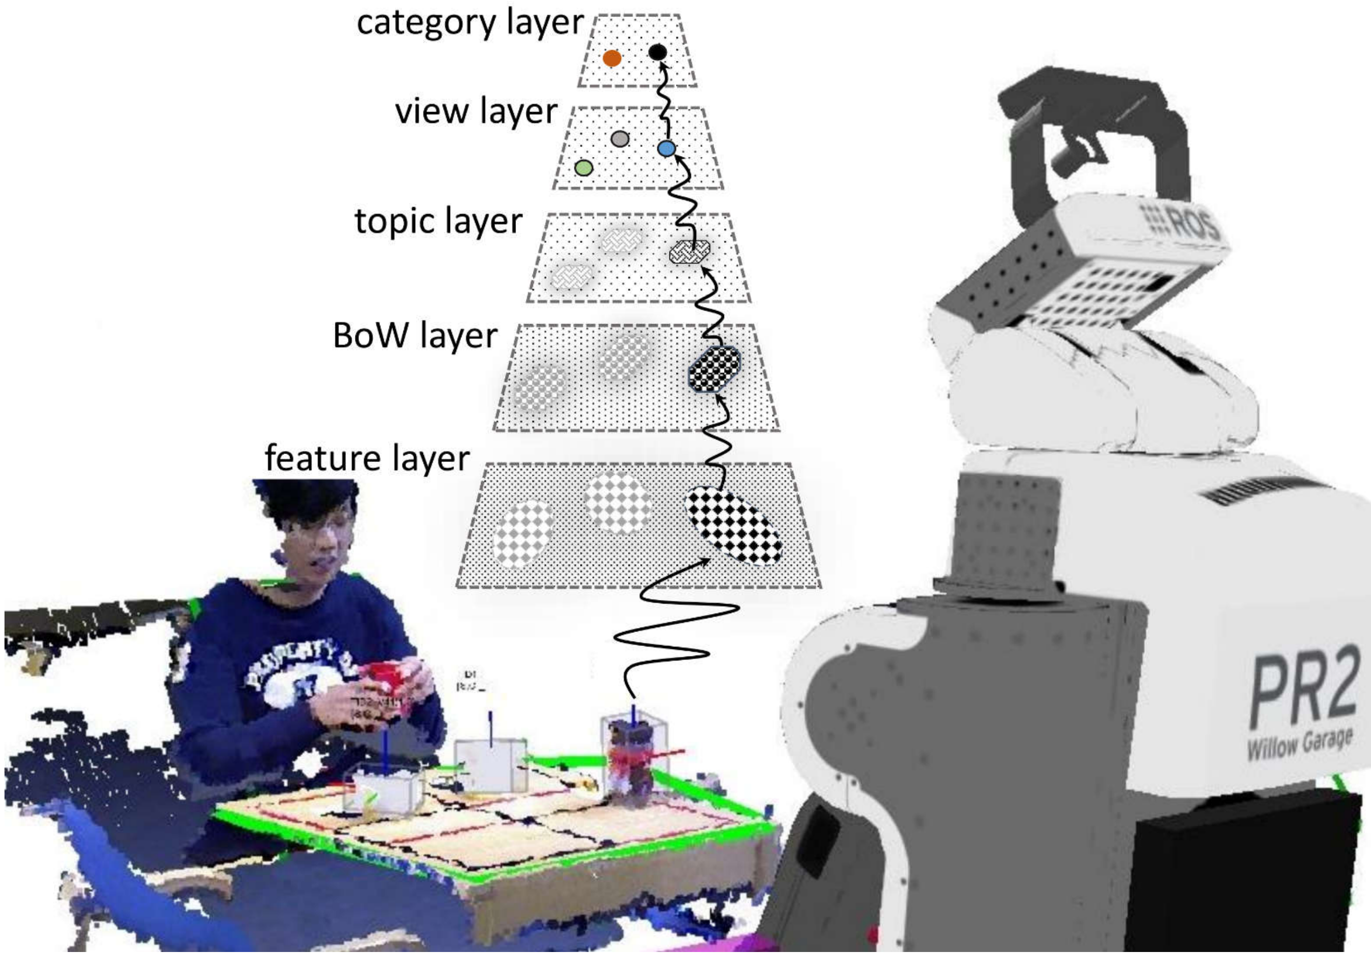

Objects at multiple levels of abstraction

- Feature layer: spin-images (local features)
- Bag of visual words dictionary
- Latent Direchlet Allocation: local topics, per object category (Local-LDA)
- View layer: representations of object views
- Category layer: each category represented by a set of object views

GOOD:
Global Orthographic Object Descriptor

- Provides a good tradeoff between
  - Descriptiveness
  - Robustness to noise
  - Computation time
  - Memory usage
- Integrated into PCL 1.9.x

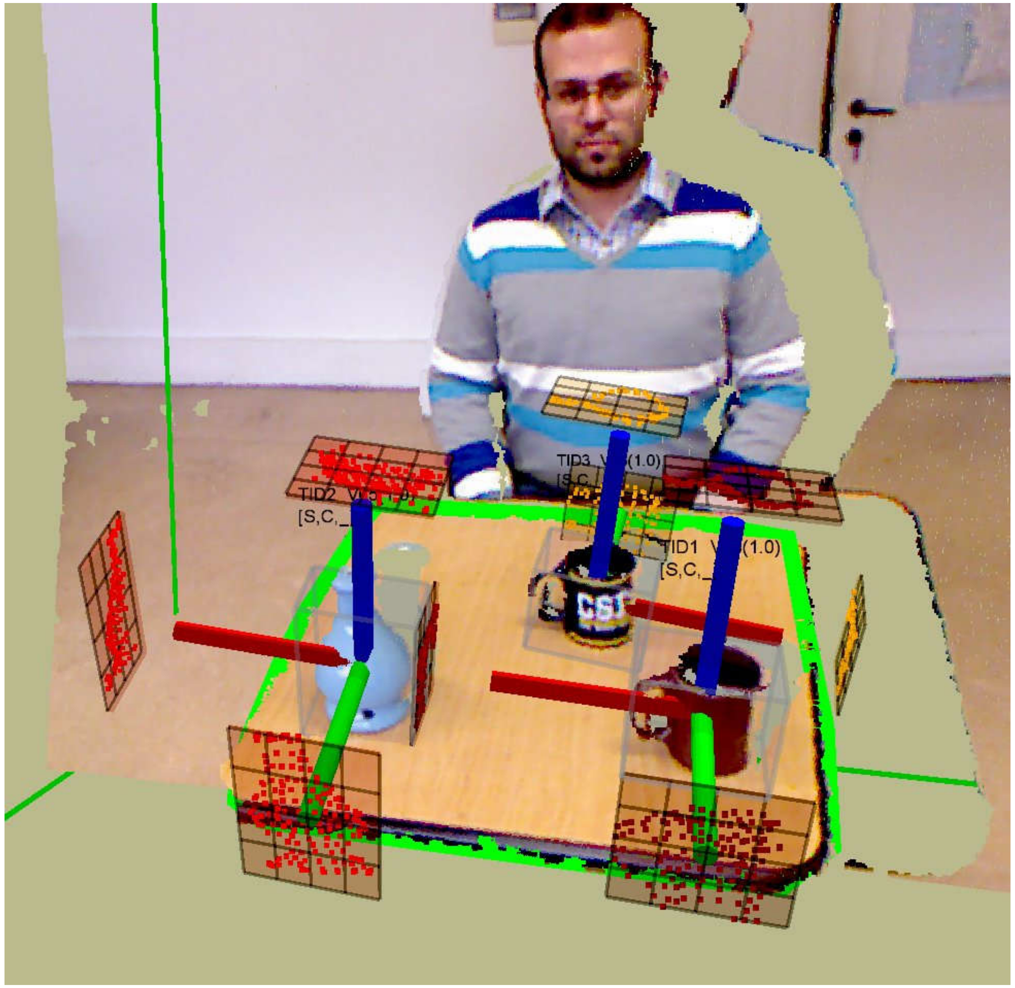

GOOD

- Video: https://www.youtube.com/watch?v=iEq9TAaY9u8

Open-ended object category learning

- “Open-ended” means
  - The set of target categories is not known in advance
  - The training instances are online experiences of a robot, and are not available in advance
- Learning approaches
  - Instance-based
  - Model-based (naïve Bayes)
- Evaluation using a simulated teacher following the open-ended experimental protocol

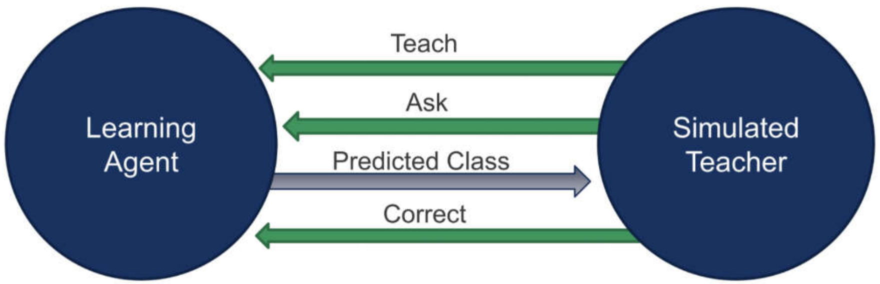

Additional selected publications

- Kasaei, S.H.M., L. Seabra Lopes, A.M. Tomé (2019) **Local-LDA: Open-Ended Learning of Latent Topics for 3D Object Recognition**, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, to appear
- Kasaei, H., L. Seabra Lopes, A.M. Tomé, J. Sock, T.-K. Kim (2018) **Perceiving, Learning, and Recognizing 3D Objects: An Approach to Cognitive Service Robots**, 32nd Conf. Artificial Intelligence (AAAI-18), USA, p. 596-603.
- Kasaei, S.H.M., M. Oliveira, G.H. Lim, L. Seabra Lopes, A.M. Tomé (2018) **Towards Lifelong Assistive Robotics: a Tight Coupling between Object Perception and Manipulation**, *Neurocomputing*, vol. 291, p. 151-166.
- Kasaei, S.H.M., A.M. Tomé, L. Seabra Lopes, M.R. Oliveira (2016) **GOOD: A Global Orthographic Object Descriptor for 3D Object Recognition and Manipulation**, *Pattern Recognition Letters*, 83, 312-320.
- Kasaei, H., M. Oliveira, G.H. Lim, L. Seabra Lopes, A.M. Tomé (2015) **Interactive Open-Ended Learning for 3D Object Recognition: An Approach and Experiments**, *Journal of Intelligent & Robotic Systems*, Springer, vol. 80, p. 537-553.

Adaptability to context change

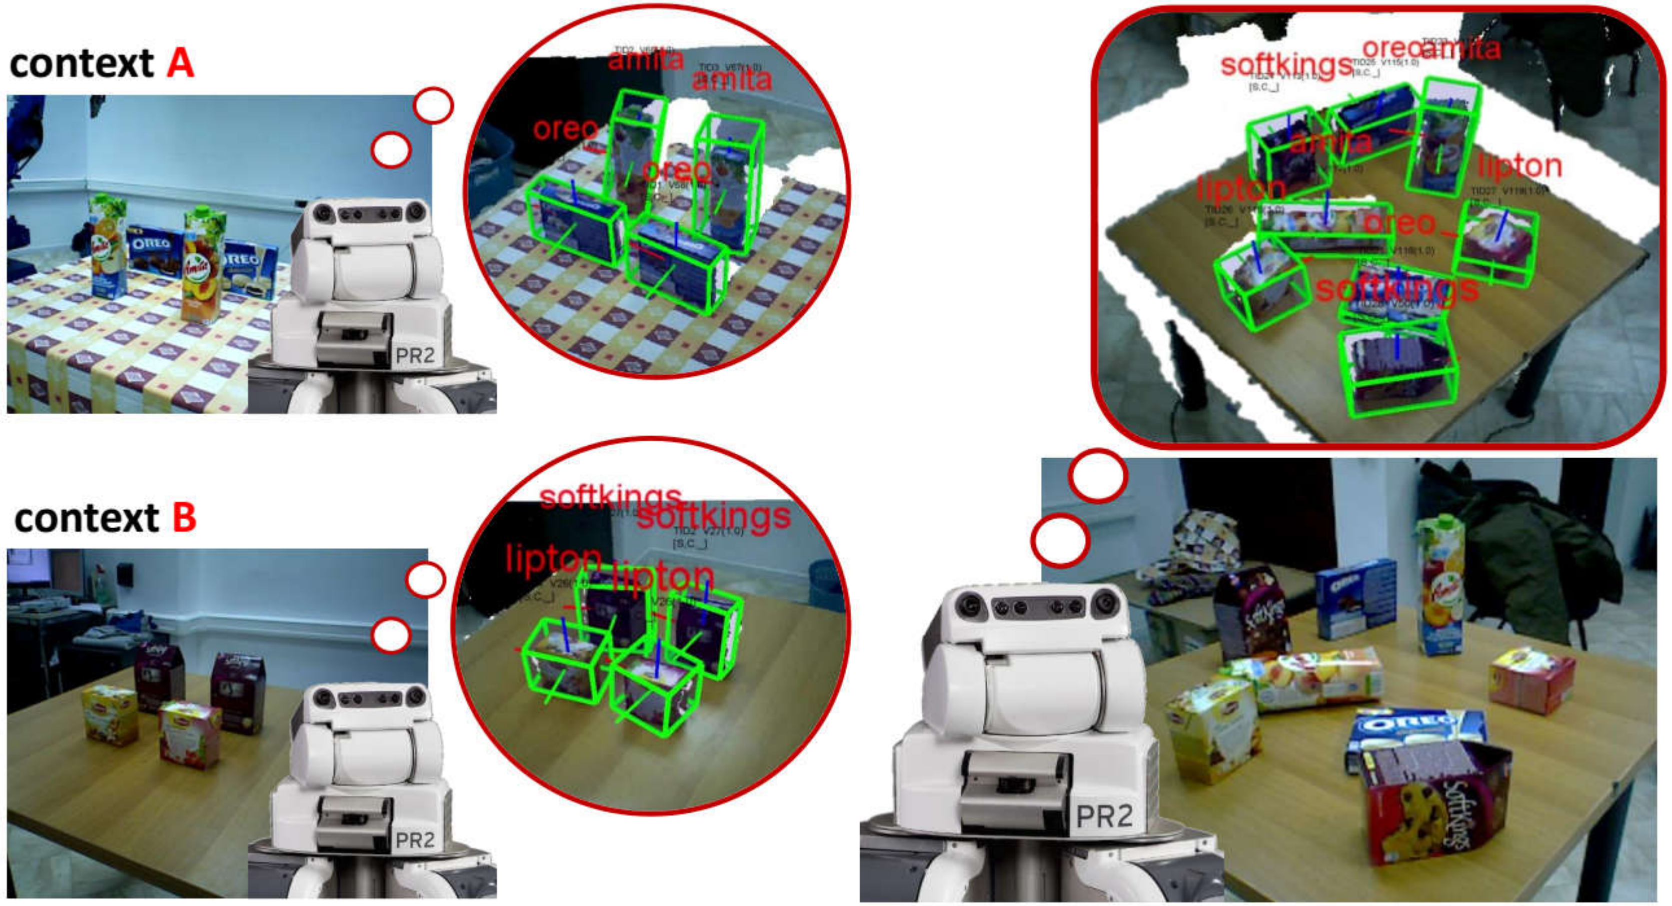

Adaptability to context change (I)

- Evaluate the intrisic ability of different learning/recognition approaches to comply with context change
  - Without explicit context information
- Two rounds
  - Evaluation using the standard teaching protocol
  - Modified teaching protocol, simulating a change of context

Adaptability to context change (II)

- 1st round
  - Evaluation using the standard teaching protocol
  - Compute the Average Learned Categories (ALC)
- 2nd round
  - Modified teaching protocol, simulating a change of context
  - Context transition point generated randomly in the interval [0.65,0.85]
- Compute adaptability

$$A = \frac{ALC_2}{ALC_1}$$

  - Where: ALC1 and ALC2 are the average learned categories in the first and second contexts

Adaptability to context change (III)

Approach	ALC_1	ALC_2	Adaptability
RACE (sets of local features)	14.9	9.4	0.63
Bag of Words	16.3	16.5	**1.01**
Standard LDA	10.9	6.3	0.58
Local LDA	**21.7**	18.6	**0.86**
GOOD	**27.3**	13.6	0.50

Kasaei, H., L. Seabra Lopes, A.M. Tomé (2018) **Coping with Context Change in Open-Ended Object Recognition without Explicit Context Information**, *Proc. 2018 IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS'2018)*, p. 6806-6812.

Open-ended learning for grasping

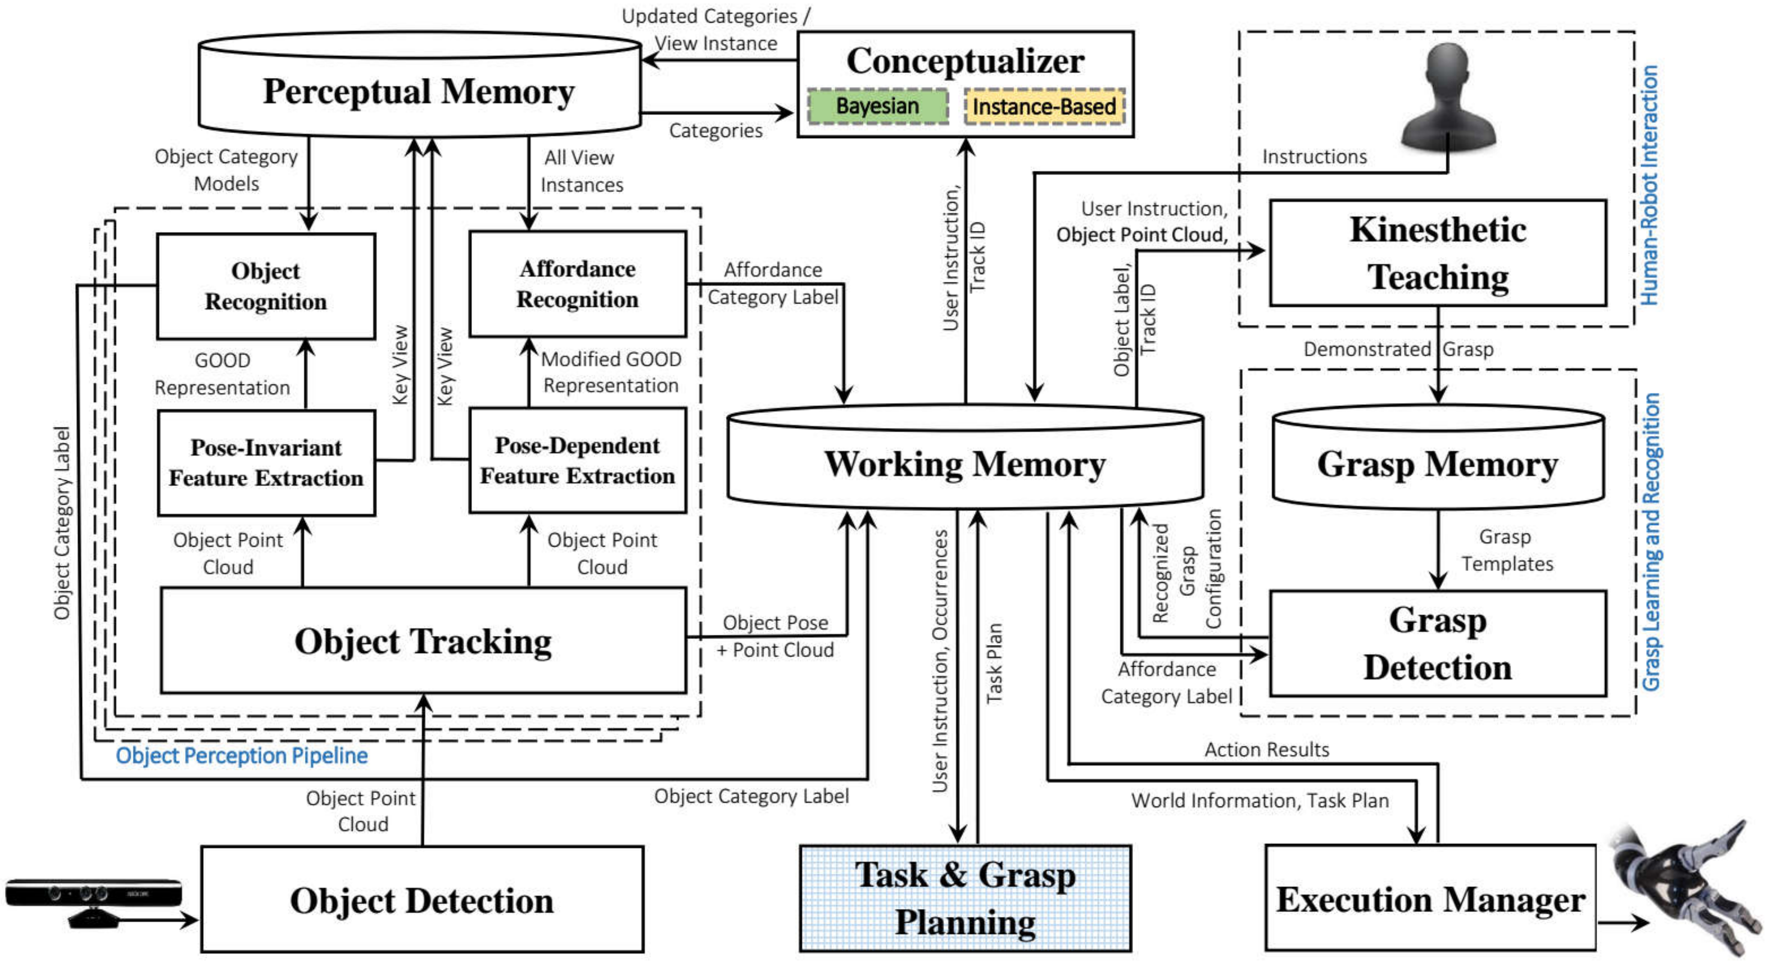

Compare to the dorsal (“what”) and ventral (“where” and “how”) pathways in the brain

Grasp teaching actions

- **Point**: point to the target object
- **Teach-category**: teach the object category or the affordance category of the selected object
  - Each stable pose of an object on the table may map to a different affordance category.
- **Ask-category**: inquire the object category or the affordance category of the target object, which the learning agent will predict based on previously learned knowledge
- **Correct-category**: if the agent could not recognize a given object or its affordance correctly, the user can teach the correct one
- **Teach-grasp**: using kinesthetic teaching, teach a grasp configuration of the robotic arm to grasp the target object
- **Grasp**: command the robot to grasp the target object

Kasaei, H., N. Shafii, L. Seabra Lopes, A.M. Tomé (2019) **Interactive Open-Ended Object, Affordance and Grasp Learning**, *Proc. IEEE Int. Conf. Robotics and Automation* (*ICRA'2019*).

Shafii, N., S.H.K. Kasaei, L. Seabra Lopes (2016) **Learning to Grasp Familiar Objects Using Object View Recognition and Template Matching**, *Proc. 2016 IEEE/RSJ Int. Conf. Intelligent Robots and Systems* (*IROS'2016*), p. 2895-2900.

Learning and grasping

- Video: https://www.youtube.com/watch?v=HoEjJJOynmY

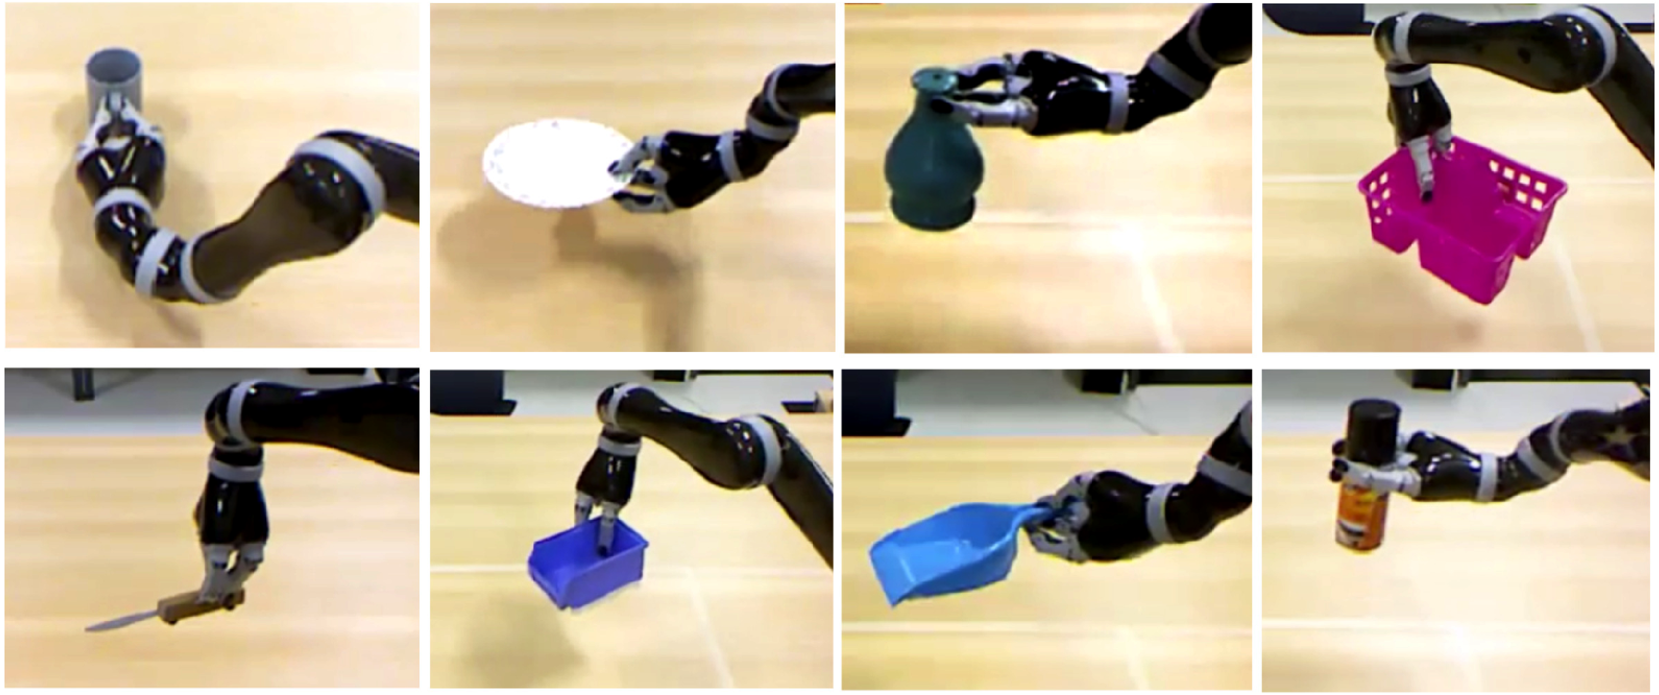

- 40 objects
- First round
  - Teach how to grasp the first 6 objects
  - Try to grasp all 40 objects
- Second round
  - Teach how to grasp 2 additional objects
  - Try again to grasp all 40 objects

	1st round	2nd round
W/o affordance recognition	58%	65%
Shafii et al, 2016	55%	70%
Kasaei et al., 2019	65%	95%

Teaching how to serve a coffee

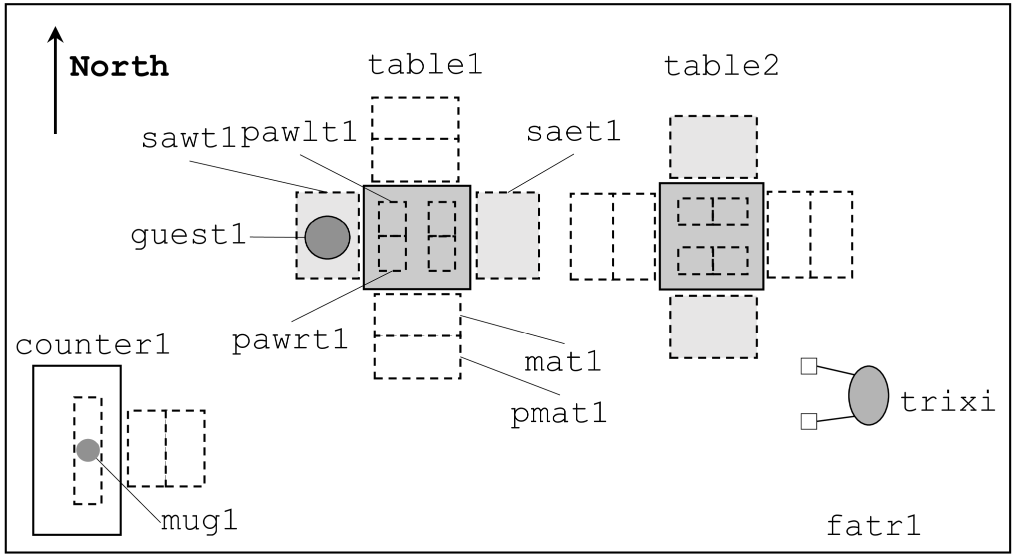

- **Achieve** drive_robot preManipulationAreaEastCounter1
- **Achieve** grasp_object_w_arm mug1 rightArm1
- **Achieve** drive_robot preManipulationSouthTable1
- **Achieve** put_object mug1 placingAreaWestRightTable1
- **Teach_task** serve_coffee guest1

Learning to achieve plan-based tasks

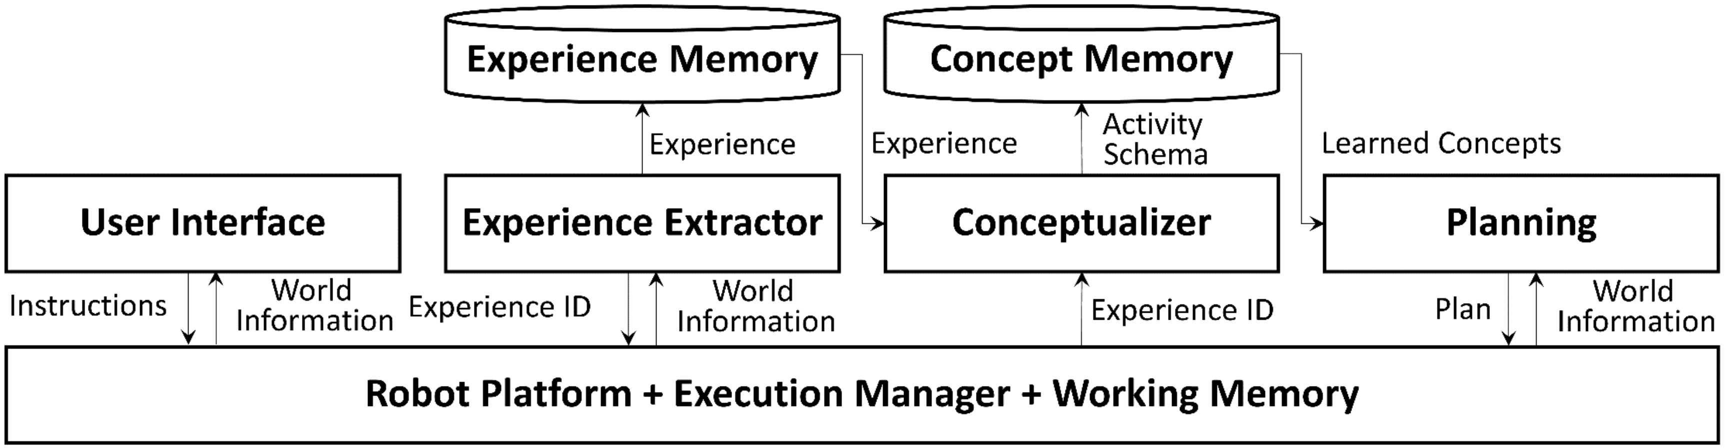

- Teaching robots how to achieve tasks
- Extracting robot activity experiences
- Learning activity models from experiences
- Problem solving using the learned models

Experience-Based Planning Domains (EBPDs)

- An **experience-based task planning domain** is a tuple [Mokhtari et al., 2016]

$$\Delta = (D_a, D_c, A, E, M),$$

  - $D_a = (\Sigma_a, S_a, O_a)$ is an **abstract planning domain**
  - $D_c = (\Sigma_c, S_c, O_c)$ is a **concrete planning domain**
    - Σ is the static world information
    - S is the set of possible states
    - O is the set of planning operators
  - A is a set of **abstraction hierarchies** relating D_c to D_a
    - $o_c \in O_c \rightarrow o_a \in O_a \quad , \quad p_c \in (\Sigma_c \cup S_c) \rightarrow p_a \in (\Sigma_a \cup S_a)$
  - E is a set of plan-based robot activity **experiences**
  - M is a set of learned methods (i.e., **activity schemata**)

- A **task planning problem** is a tuple $P = (t, \sigma, s_0, [g])$, where t is the task to be acheived, $\sigma \in \Sigma$, $s_0 \in S$ and $g \in S_g$ $(S_g \subseteq S)$

Conceptualizing robot activity experiences

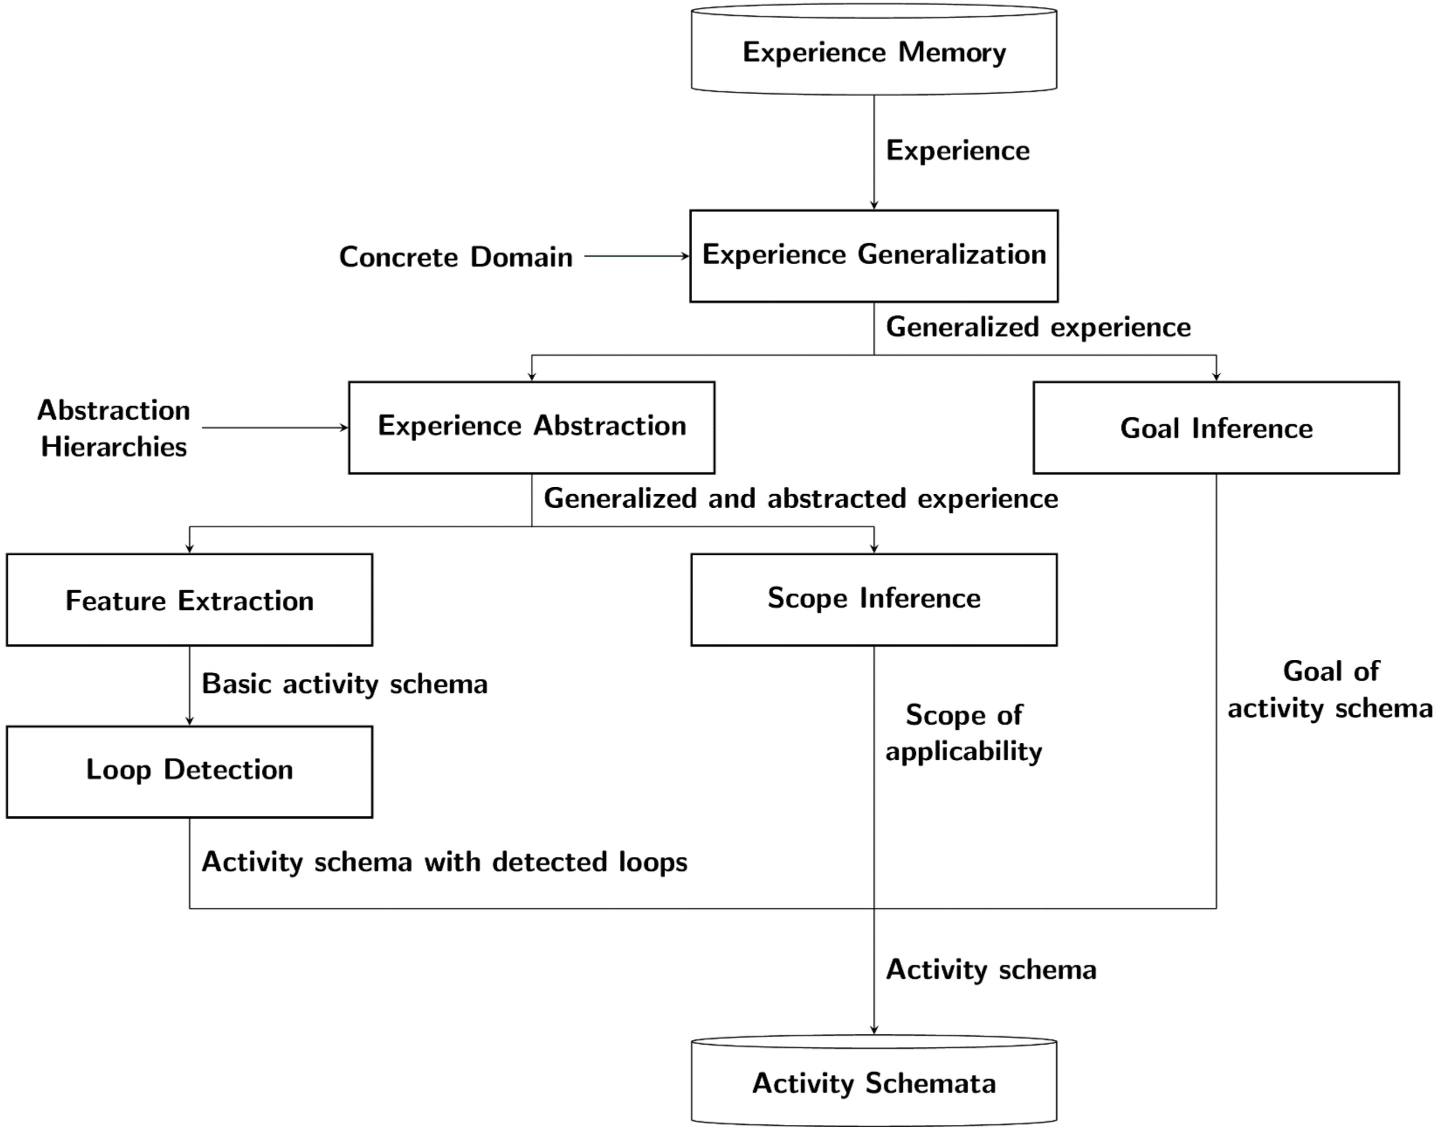

Planning using the learned knowledge

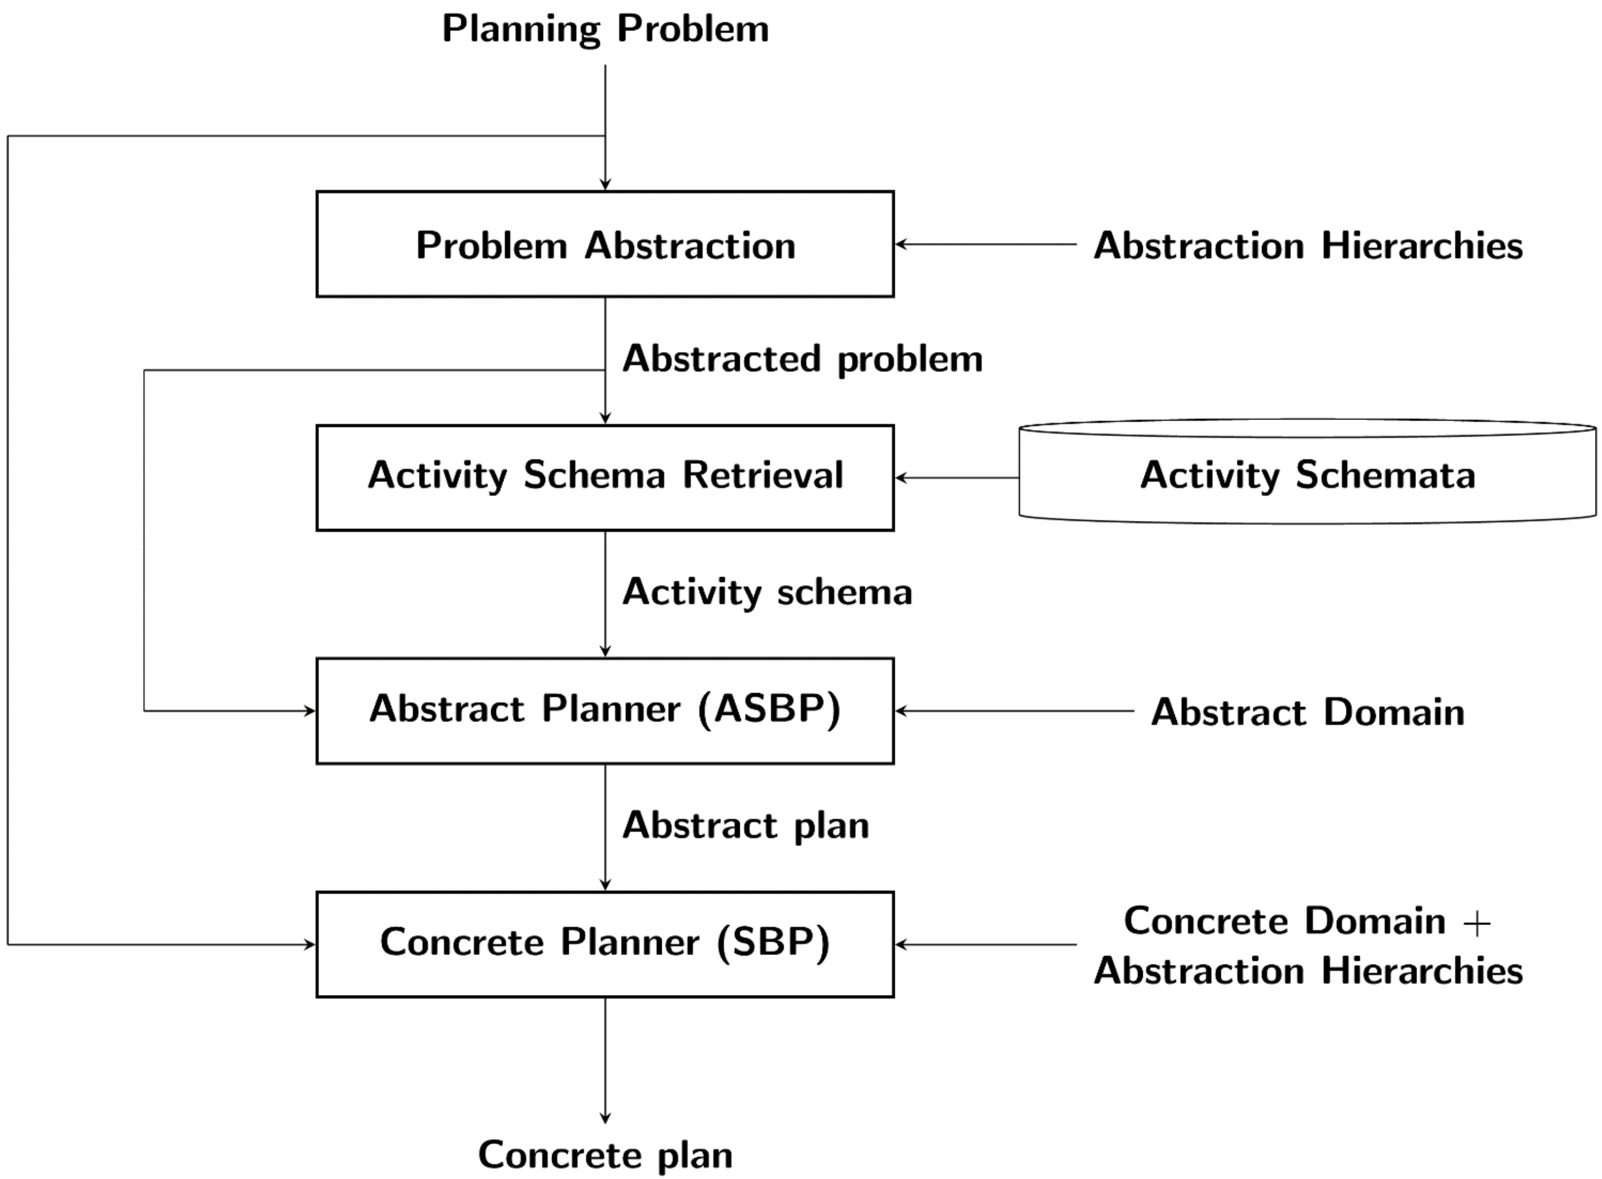

Learning and planning: selected recent publications

Mokhtari, V., R. Manevich, L. Seabra Lopes, and Armando J. Pinho (2019) **Learning the Scope of Applicability for Task Planning Knowledge in Experience-Based Planning Domains**, *Proc. 2019 IEEE/RSJ Int. Conf. Int. Robots and Systems* (*IROS'2019*).

Mokhtari, V., L. Seabra Lopes, A. Pinho (2017) **Learning Robot Tasks with Loops from Experiences to Enhance Robot Adaptability**, *Pattern Recognition Letters*, 99, p. 57-66.

Mokhtari, V., L. Seabra Lopes, A. Pinho (2017) **An Approach to Robot Task Learning and Planning with Loops**, *Proc. 2017 IEEE/RSJ Int. Conf. Int. Robots and Systems* (*IROS'2017*), p. 6033-6038.

Mokhtari, V., L. Seabra Lopes, A. Pinho (2016) **Experience-Based Planning Domains: an Integrated Learning and Deliberation Approach for Intelligent Robots**, *Journal of Intelligent & Robotic Systems*, 83 (3), p. 463-483.

Mokhtari, V., L. Seabra Lopes, A. Pinho (2016) **Experience-Based Robot Task Learning and Planning with Goal Inference**, *Proc. 26th Int. Conf. Automated Planning and Scheduling* (*ICAPS'2016*), p. 509-517.

Thanks to team members

Aneesh Chauhan

Hamidreza Kasaei

Vahid Mokhtari

Ana Tomé

Armando Pinho

Miguel Oliveira

Gi Hyun Lim

Nima Shafii