INSTITUTO SUPERIOR TÉCNICO

Institute for Systems and Robotics

On Learning and Using Affordances with Humanoid Robots

José Santos-Victor, Alexandre Bernardino,
Luis Montesano and Manuel Lopes

Computer and Robot Vision Laboratory
http://vislab.isr.ist.utl.pt

Object-Action Complexes Workshop
Humanoids, Paris, Dec. 2009

Outline

- Motivation: affordances
  - Modeling
  - Learning
- Using the model
  - Imitation games
  - Task learning
  - Mirror an canonical neurons
- Extensions: multimodal perception (words)
- Discussion

Outline

- **Motivation: affordances**
  - Modeling
  - Learning
- Using the model
  - Imitation games
  - Task learning
  - Mirror an canonical neurons
- Extensions: multimodal perception (words)
- Discussion

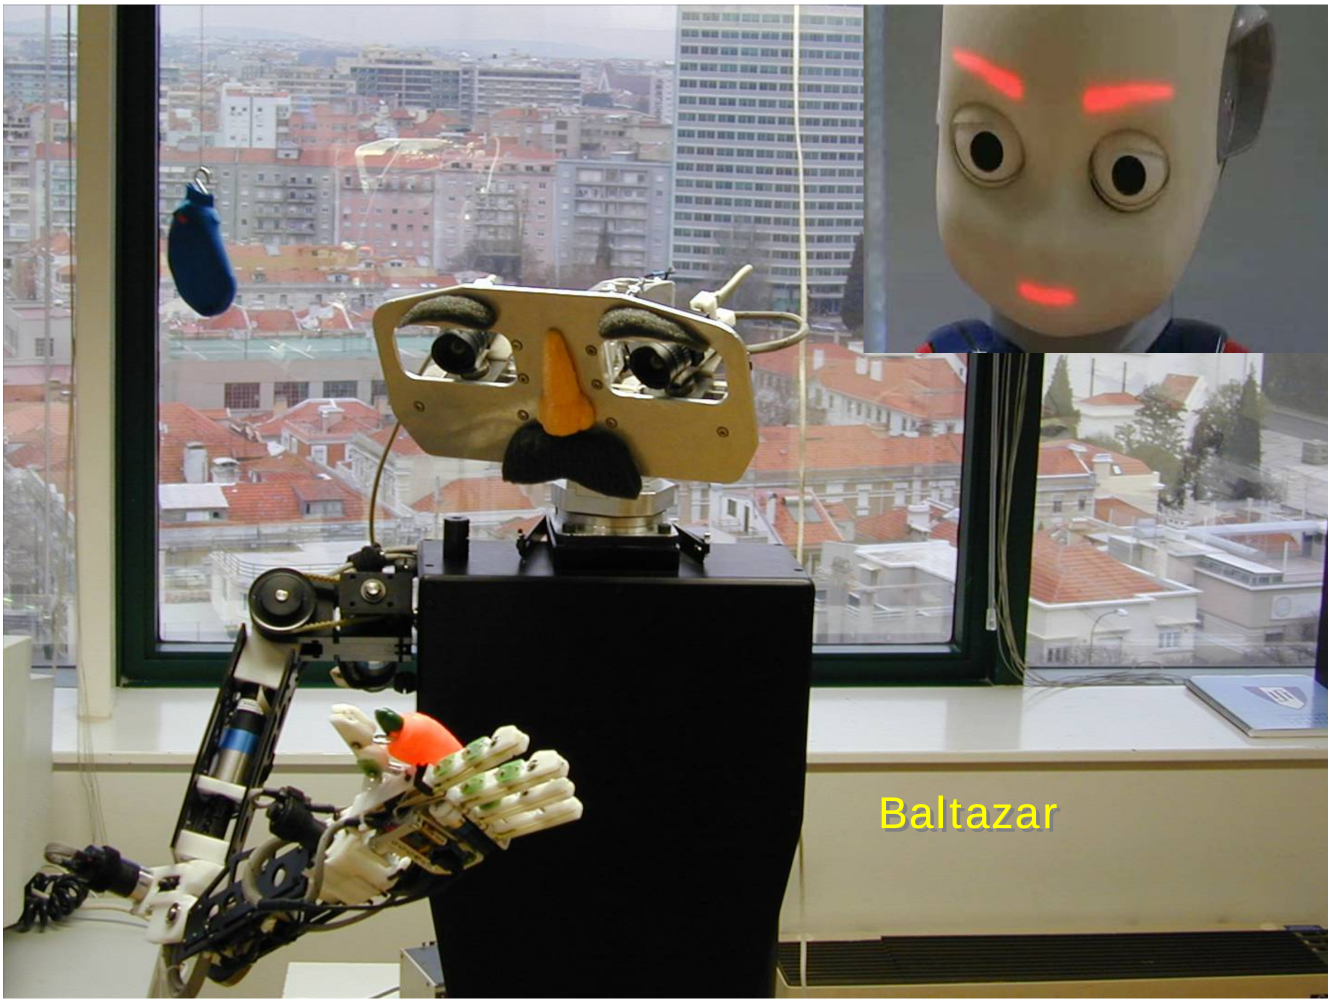

What are affordances?

Affordances Definition

"The Affordances of the environment are what it offers the animal, what it provides or furnishes, either for good or ill" (Gibson, 1986)

What are affordances?

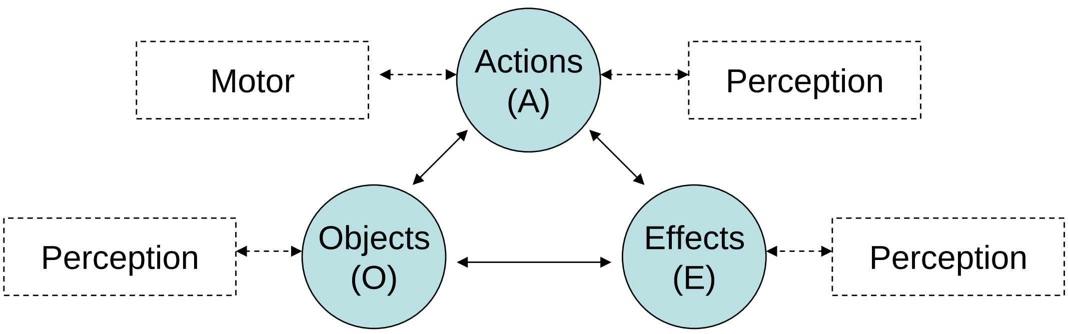

What are affordances?

What are affordances?

Robotic Affordances

- Relations between actions, objects and effects
- Involves action and perception
- **Bayesian networks to model relations**
- **Learn through experimentation or observation**

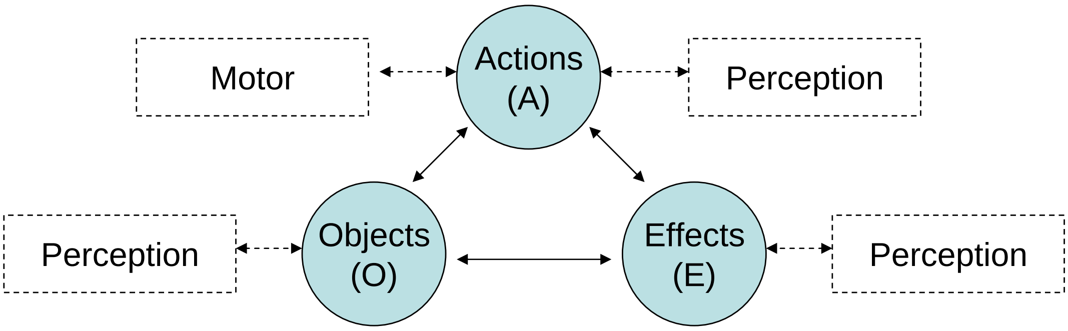

Instituto Superior Técnico

Mathematical model: BN

- We use Bayesian Networks to represent affordances
- Nodes are:
  - Actions and action parameters
  - Object properties
  - Resulting effects
- BNs provide a unified and sound probabilistic framework for learning and using affordances

Institute for Systems and Robotics
Computer and Robot Vision Lab

Learning Bayesian Networks

- Use a set of acquired data D to learn the network

Prior knowledge

Structure Learning

Parameter Learning

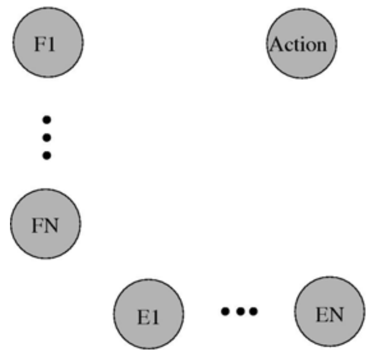

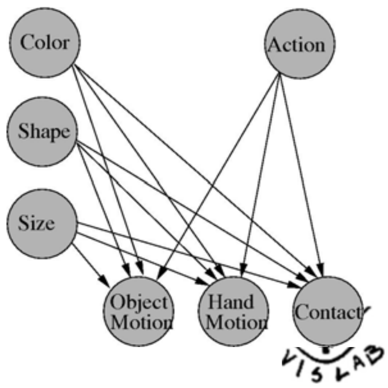

Instituto Superior Técnico

Learning Bayesian Networks

- Previous development phases provide the Actions, the object descriptors and the effect categories
- Bayesian structure learning

$$G^* = \arg max_{\mathcal{G}} p(G \mid D) = \arg max_{\mathcal{G}} \eta p(D \mid G) p(G)$$

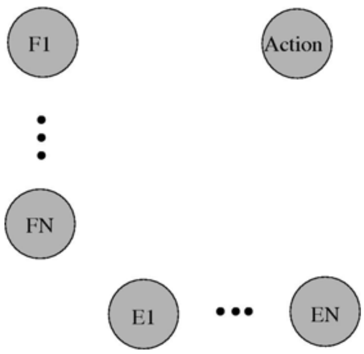

- Likelihood of the data given the model

$$p(D \mid G) = p(X^{1:N} \mid G) = \prod_i p(x_i^{1:N} \mid x_{Pa(x_i)}^{1:N})$$

- Prior over models

$$p(G)$$

- Interventional data

Institute for Systems and Robotics
Computer and Robot Vision Lab

Object and effects description

- Percepts of object features and effects are clusterized in unsupervised manner
- The categories form the space of

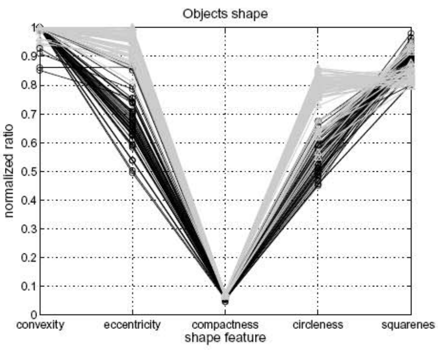

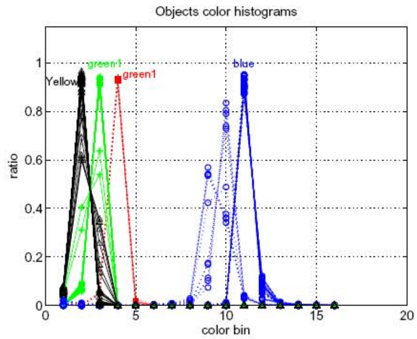

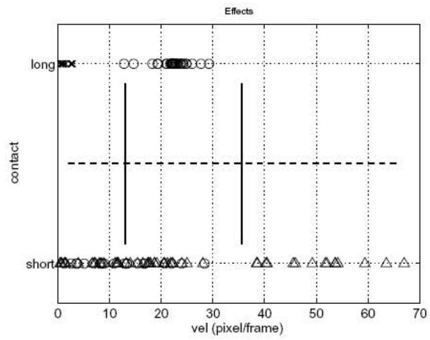

Object and effects description

- Object features and effects are described using unsupervised learned clusters from

Symbol	Description	Values
A	Action	*grasp, tap, touch*
H	Height	discretized in 10 values
C	Color	*green$_1$,green$_2$, yellow, blue*
Sh	Shape	*ball, box*
S	Size	*small, medium, big*
V	Object velocity	*small, medium, big*
HV	Hand velocity	*small, medium, big*
Di	Object-hand velocity	*small, medium, big*
Ct	Contact duration	*none, short, long*

Learning affordances: Baltazar

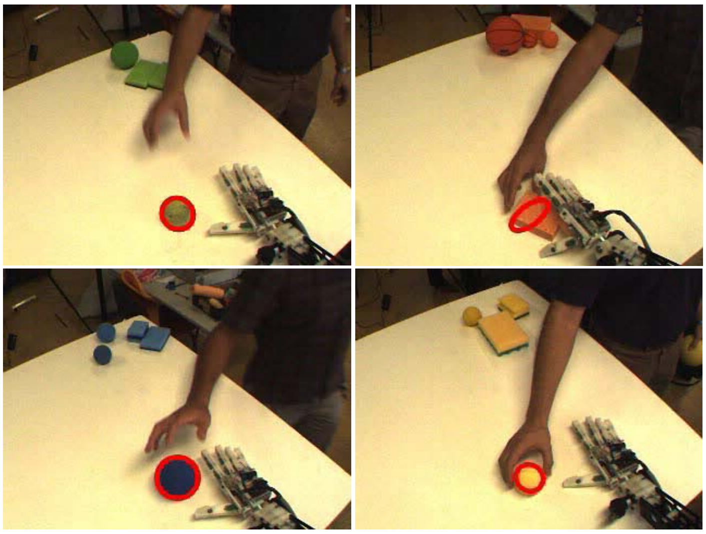

Learned object affordances

- 300 trials using random exploration, i.e. object action pair
- The structure network learnt by the MC3 algorithm (*Markov Chain Monte Carlo Model Composition algorithm*)

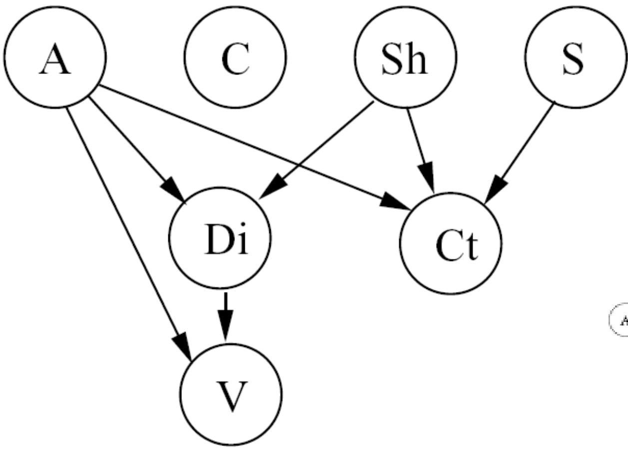

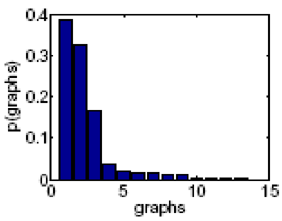

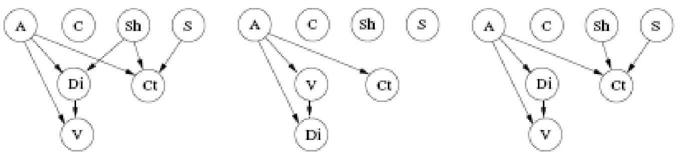

Not all network structures explain the data equally well

Robotic Affordances

Affordances as the link between sensory-motor coordination and social interaction

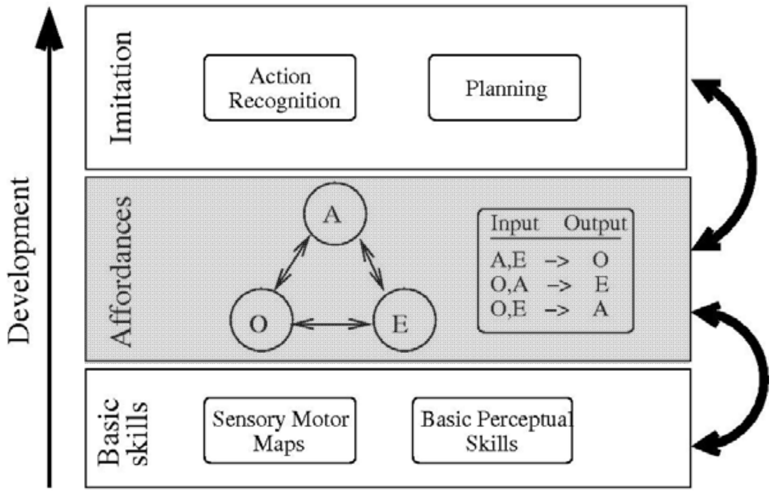

Outline

- Motivation: affordances
  - Modeling
  - Learning
- Using the model
  - **Imitation games**
  - Task learning
  - Mirror an canonical neurons
- Extensions: multimodal perception (words)
- Discussion

From SMM to social interaction

- Affordances can answer the following questions:

inputs	outputs	function
(O, A)	E	Predict Effect
(O, E)	A	Recognize action & Planning
(A, E)	O	Object recognition & selection

- Core capabilities for social interaction
- Emulation: achieve the same effect

Simple emulation

- One step Bayesian decision

$$< a^*, o^* >= \underbrace{\arg max}_{a \in A, o \in \mathcal{O}} \mathtt{E}\left[r(a^d, f^d, e^d, a, f^o, e^o)\right]$$

- It is possible to define different simple reward functions to obtain different behaviors:
  - Matching an effect:

$$r(e^d) = \begin{cases} 1, & if\ E^i = \hat{e}^d \\ 0, & otherwise \end{cases}$$

  - Matching effect and Object:

$$r(e^d, f^d, f^i) = \begin{cases} 1, & if\ E^i = \hat{e}^d \wedge F^i = \hat{f}^d \\ 0, & otherwise \end{cases}$$

Imitation games

Objective: select **action and object** to obtain the same effect on **a similar object**

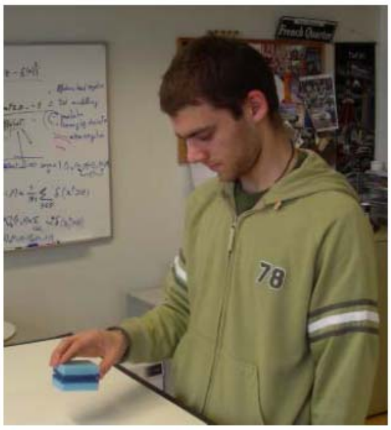

Demonstration

(grasp on small box)

Which action gives the same effect?

The reward now also includes information about the object features

Instituto Superior Técnico

Imitation games

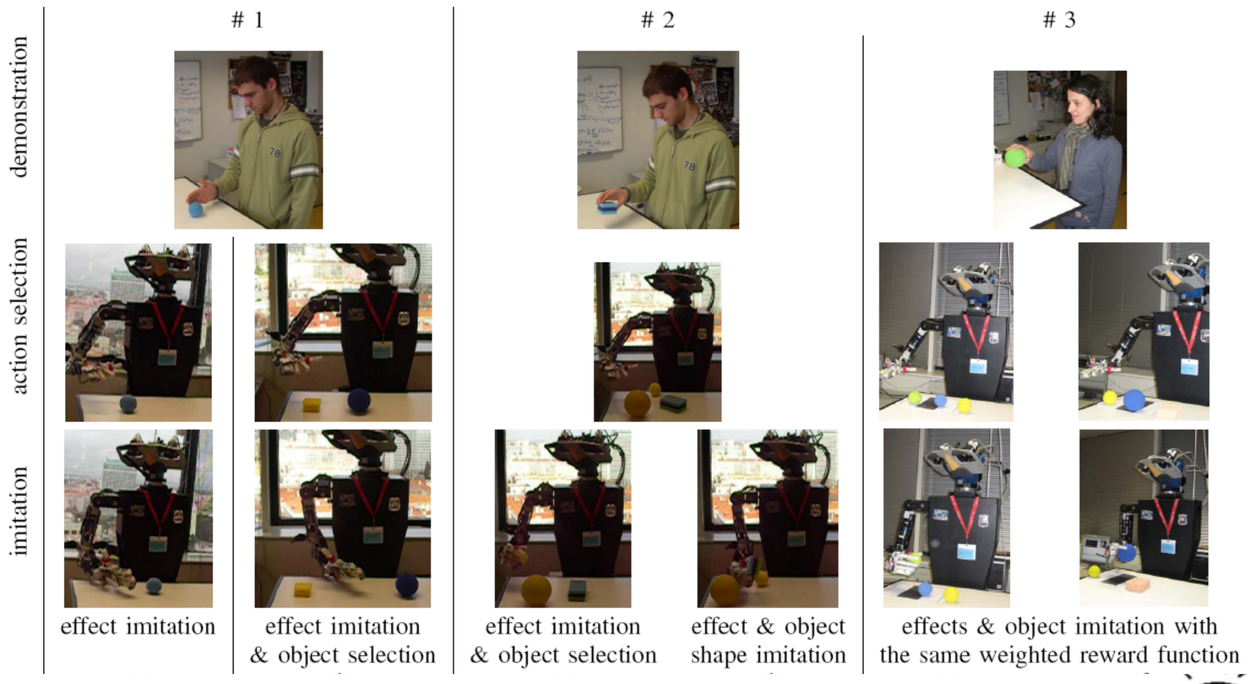

Institute for Systems and Robotics
Computer and Robot Vision Lab

Outline

- Motivation: affordances
  - Modeling
  - Learning
- Using the model
  - Imitation games
  - **Task learning**
  - Mirror an canonical neurons
- Extensions: multimodal perception (words)
- Discussion

Imitation (task)

Affordance based imitation

-World model (cause-effect)
-Action recognition

Task interpretation (BIRL)
-Bayesian Inverse Reinforcement Learning
-(sequences of actions, estimating the
task reinforcement function.)

Imitation (emulation)

Demonstration
{(s1,e1), (s2,e2), … }
↓
Demo Interpretation
(affordances+BIRL)
↓
Imitation
(execution)

Experiments

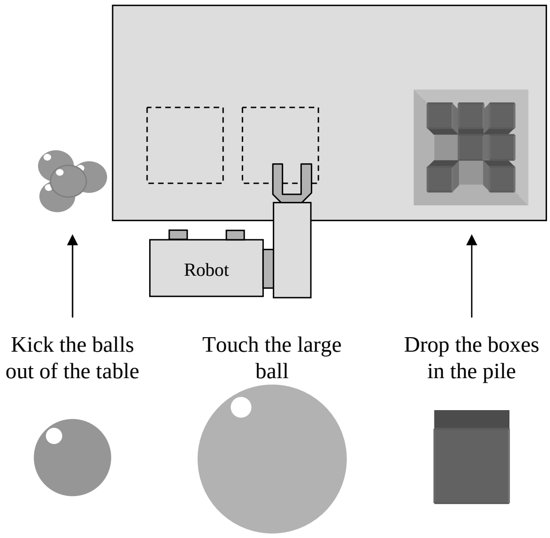

Inaccurate and incomplete demonstration

EXPERIMENT 2: INACCURATE, INCOMPLETE DEMONSTRATION (DEMONSTRATED AND LEARNED POLICIES).

State	Demo	Learned
(∅, BBall)	$\boxed{-}$	TouchR
(∅, Box)	GraspR	GraspR
(∅, SBall)	TapR	TapR
(BBall, ∅)	TouchL	TouchL
(BBall, BBall)	$\boxed{\text{GraspR}}$	TouchL
(BBall, Box)	TouchL	TouchL
(BBall, SBall)	TouchL	TouchL
(Box, ∅)	GraspL	GraspL
(Box, BBall)	GraspL	GraspL
(Box, Box)	GraspL	GraspL
(Box, SBall)	GraspL	GraspL
(SBall, ∅)	TapL	TapL
(SBall, BBall)	TapL	TapL
(SBall, Box)	TapL	TapL
(SBall, SBall)	TapL	TapL

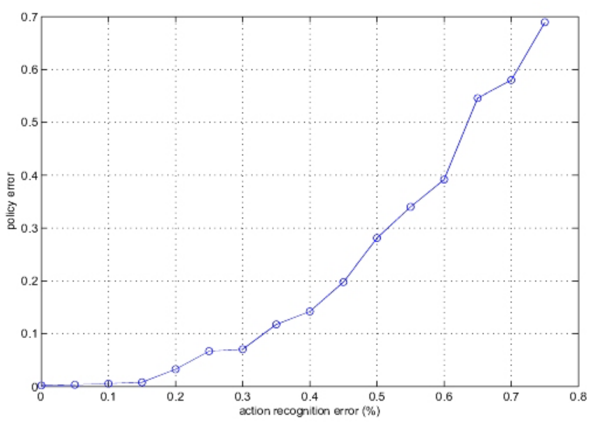

Imitation

Instituto Superior Técnico

Outline

- Motivation: affordances
  - Modeling
  - Learning
- Using the model
  - Imitation games
  - Task learning
  - **Mirror an canonical neurons**
- Extensions: multimodal perception (words)
- Discussion

Institute for Systems and Robotics

Computer and Robot Vision Lab

Mirror Neurons

[Gallese, Fadiga, Fogassi and Rizzolati, Brain, 1996]

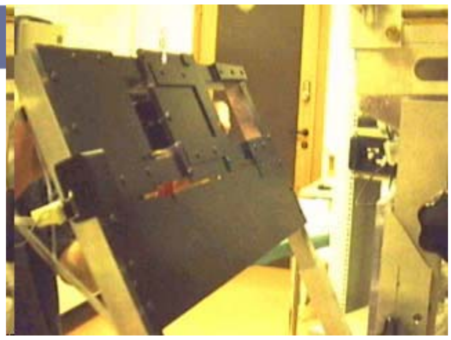

Active during observation of another monkey's or experimenter's hands interacting with objects.

Observed & executed actions are the same:

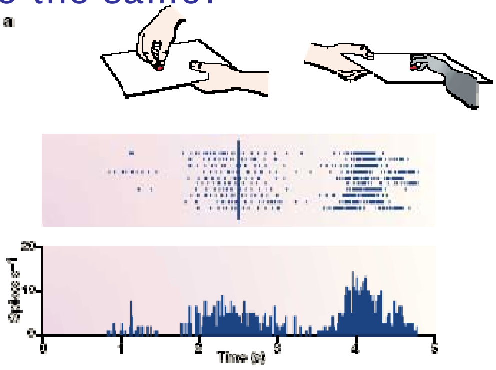

Observed & executed action are NOT the same (tool):

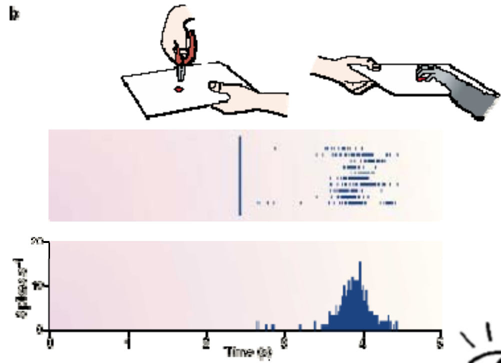

Canonical Neurons (affordances)

Motor neurons

Respond also to the presentation of

- food or
- graspable 3D objects,
- even in the absence of subsequent movement.

Object specific

(size and shape must be congruent with the type of grip coded by the neuron).

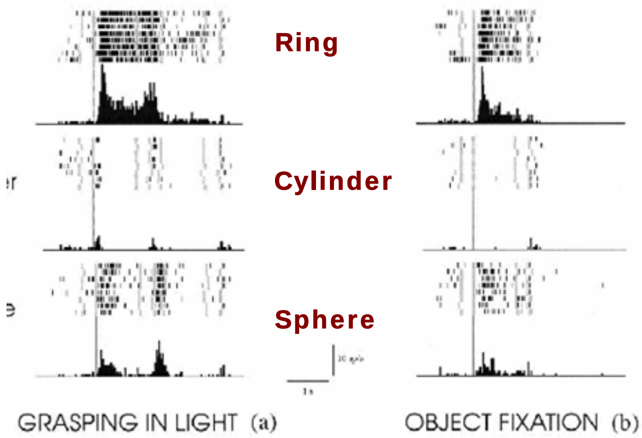

Responses for a selective neuron for ring shapes.

[L. Fadiga et al. Intl. Journal of Psychophysiology, 2000]

Action observation / execution ressonance

Individual A **Individual B**

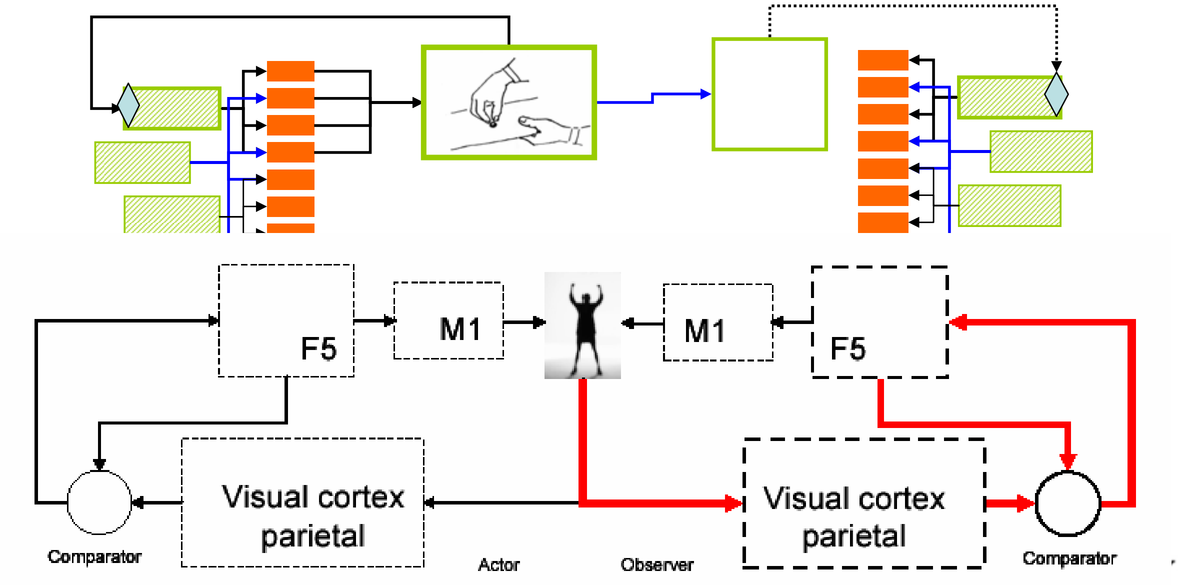

Instituto Superior Técnico

Affordances vs mirror / canonical

inputs	outputs	function
(O, A)	E	Predict Effect
(O, E)	A	Recognize action & Planning
(A, E)	O	Object recognition & selection

Institute for
Systems and Robotics
Computer and Robot Vision Lab

iCub example

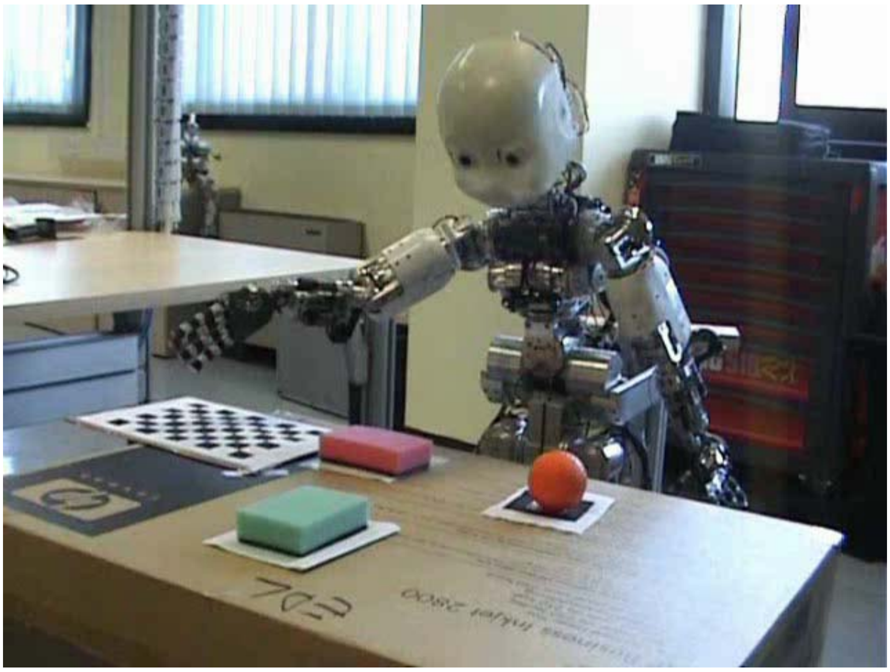

Instituto Superior Técnico

Outline

- Motivation: affordances
  - Modeling
  - Learning
- Using the model
  - Imitation games
  - Task learning
  - Mirror an canonical neurons
- **Extensions: multimodal perception (words)**
- Discussion

Institute for
Systems and Robotics
Computer and Robot Vision Lab

The setup

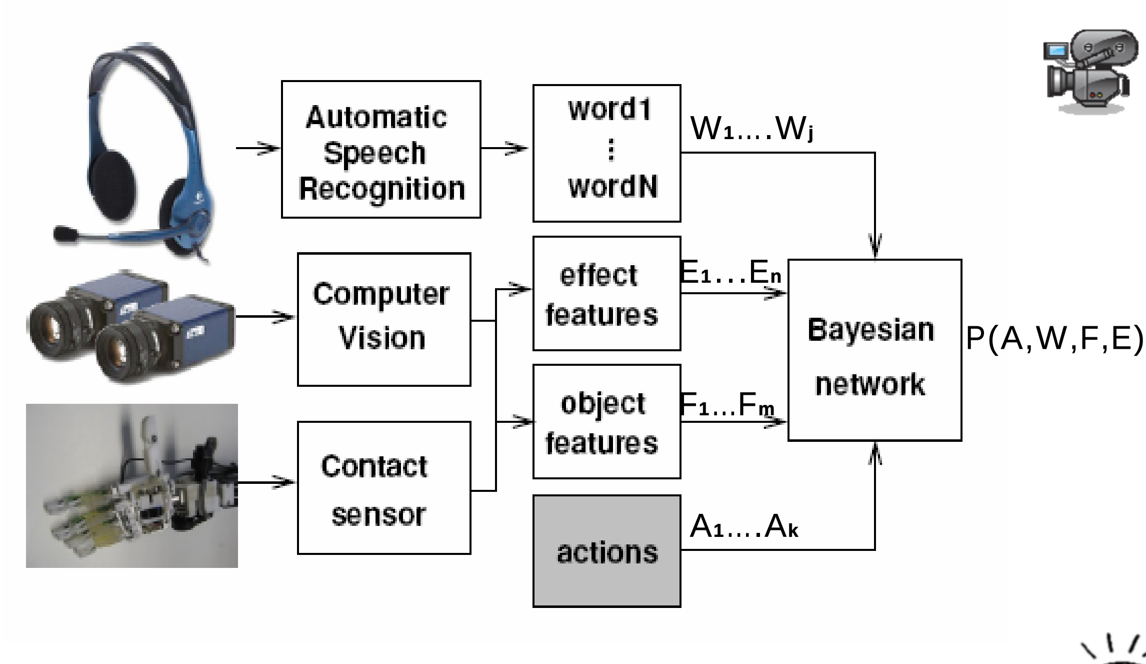

Adding words to the model

- p(X) represents the world behaviour, i.e. what the robot has learned through experience
- W a set of words, description of one experiment,
- Goal: to find mapping between W and X, achieved by estimating p(X,W)

$$p(X,W) = \prod_{\omega_i \in W} p(\omega_i \mid X_{\omega_i})\, p(X) \quad (1)$$

- X_{ω_i} is a subset of nodes X which are parents of word ω_i
- strong assumption is the independence among words, a “bag of words”
- to choose from all models described by (1) we use variation of the simple greedy approach – K2 algorithm.

Results

• Comparing results obtained with ideal speech recognizer (100%) and with the real one (81%), it is visible that mistakes from the real recognizer have only a small influence on the model performance

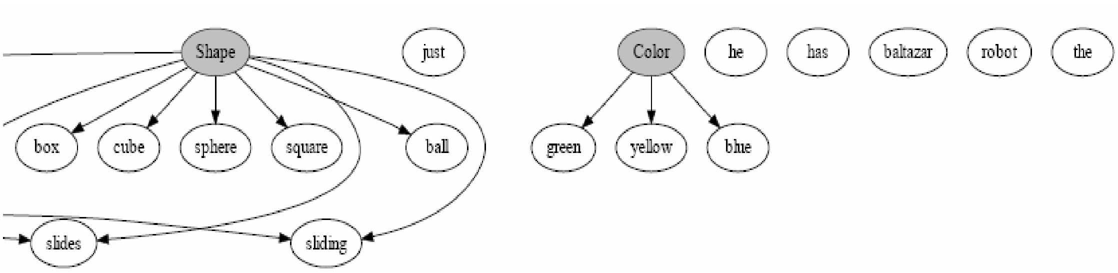

• model distinguishes non-referential words

• model sensitive to unbalanced training data

Results

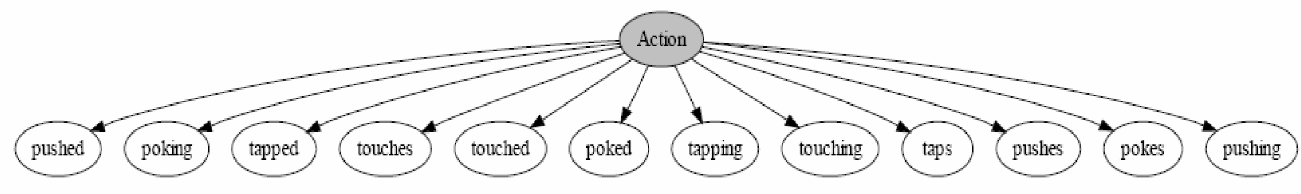

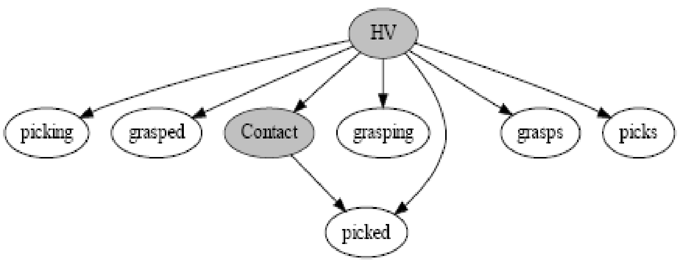

PLAY VIDEO

• grasp action and its association

Results – Instructing the Robot

Examples of using the Bayesian network to select actions and objects

objects on the table	Verbal input "small grasped"	"moving green"	"ball sliding"	"big rolling"	"has rising"	"sliding small"	"rises yellow"
light green circle big	-	grasp, p=0.034	-	**tap, p=0.227**	grasp, p=0.019	-	-
yellow circle medium	-	-	-	-	**grasp, p=0.073**	-	**grasp, p=0.3**
dark green box small	grasp, p=0.122	grasp, p=0.041	-	-	grasp, p=0.037	**tap, p=0.25**	-
blue box medium	-	-	-	-	grasp, p=0.037	-	-
blue box big	-	-	-	tap, p=0.022	grasp, p=0.017	-	-
dark green circle small	**grasp, p=0.127**	**tap, p=0.127**	-	-	grasp, p=0.064	-	-

- incomplete instructions
- both factors taken into consideration: available objects and verbal task assignment

Grasping (where & how)

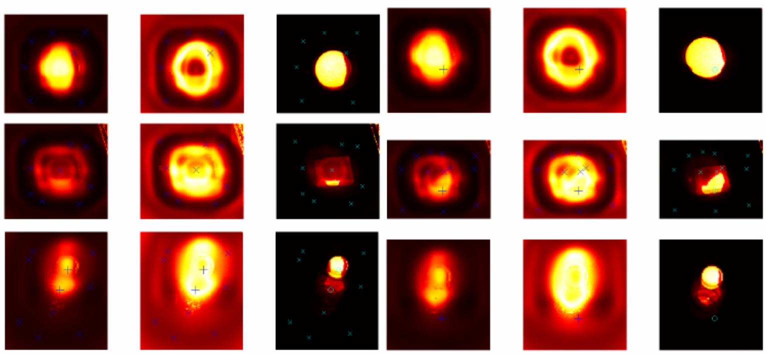

Figure: (left) predicted grasping probability , (center) variance of parameter p, (right) pixels $p > 0.5$ crosses (x) failed grasps, plus (+) successful ones.

Outline

- Motivation: affordances
  - Modeling
  - Learning
- Using the model
  - Imitation games
  - Task learning
  - Mirror an canonical neurons
- Extensions: multimodal perception (words)
- **Discussion**

Discussion

- Bridging sensorimotor layers and “higher cognitive behavior”
- Relationship to the “mirror” and “canonical” neurons
- Extension to other perceptual modalities (word/meaning)
- Future developments
  - Scale (structural learning with large dimensional graphs)
  - Batch versus incremental
  - Continuous versus discrete variables (parameterized is ok)
  - Application to grasping and handling

Credits (http://vislab.isr.ist.utl.pt)

Head, face, cover design: Ricardo Beira, M. Praça, L. Vargas
Vision & attention: Alexandre Bernardino, J.Ruesch
Learning, affordances; Manuel Lopes, Luis Montesano
Language acquisition: Jonas Hornstein, Cláudia Soraes
G. Salvi, Verica Krunic – Vision-speech association
Matteo Tajana – model based tracking
Ricardo Nunes – robot life support!

http://vislab.isr.ist.utl.pt
José Santos-Victor – jasv@isr.ist.utl.pt

Credits (http://vislab.isr.ist.utl.pt)