From learning to search to finding what others miss: Developing a theory of mind

Dimitri Ognibene

Università degli Studi di Milano Bicocca

A Timeless Classic: Passive Vision

Why does this first-person clip seem unnatural?

Computer Vision Success

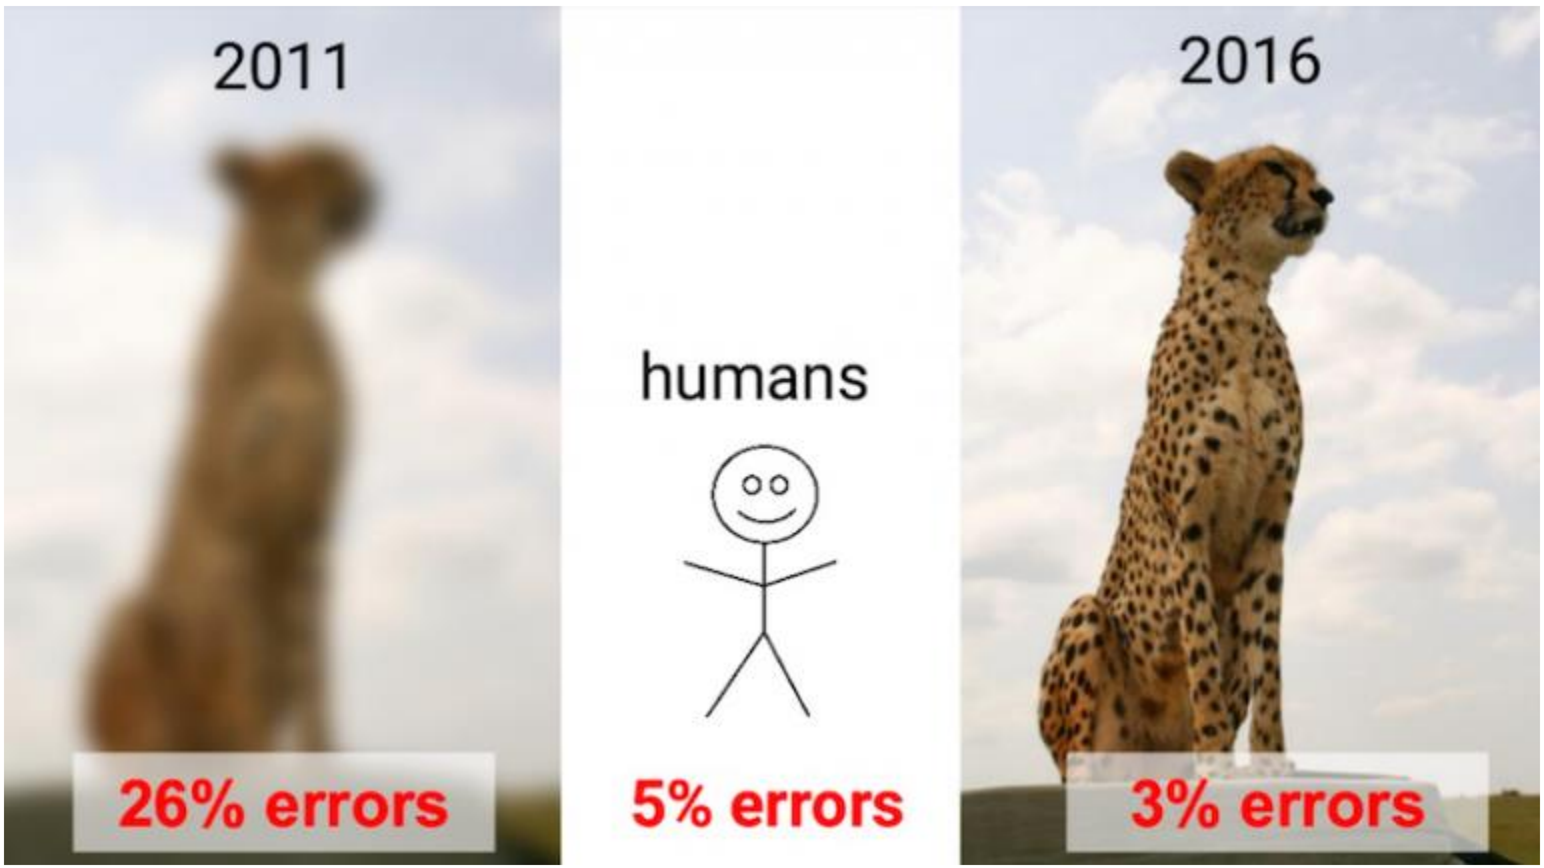

Robotic Success

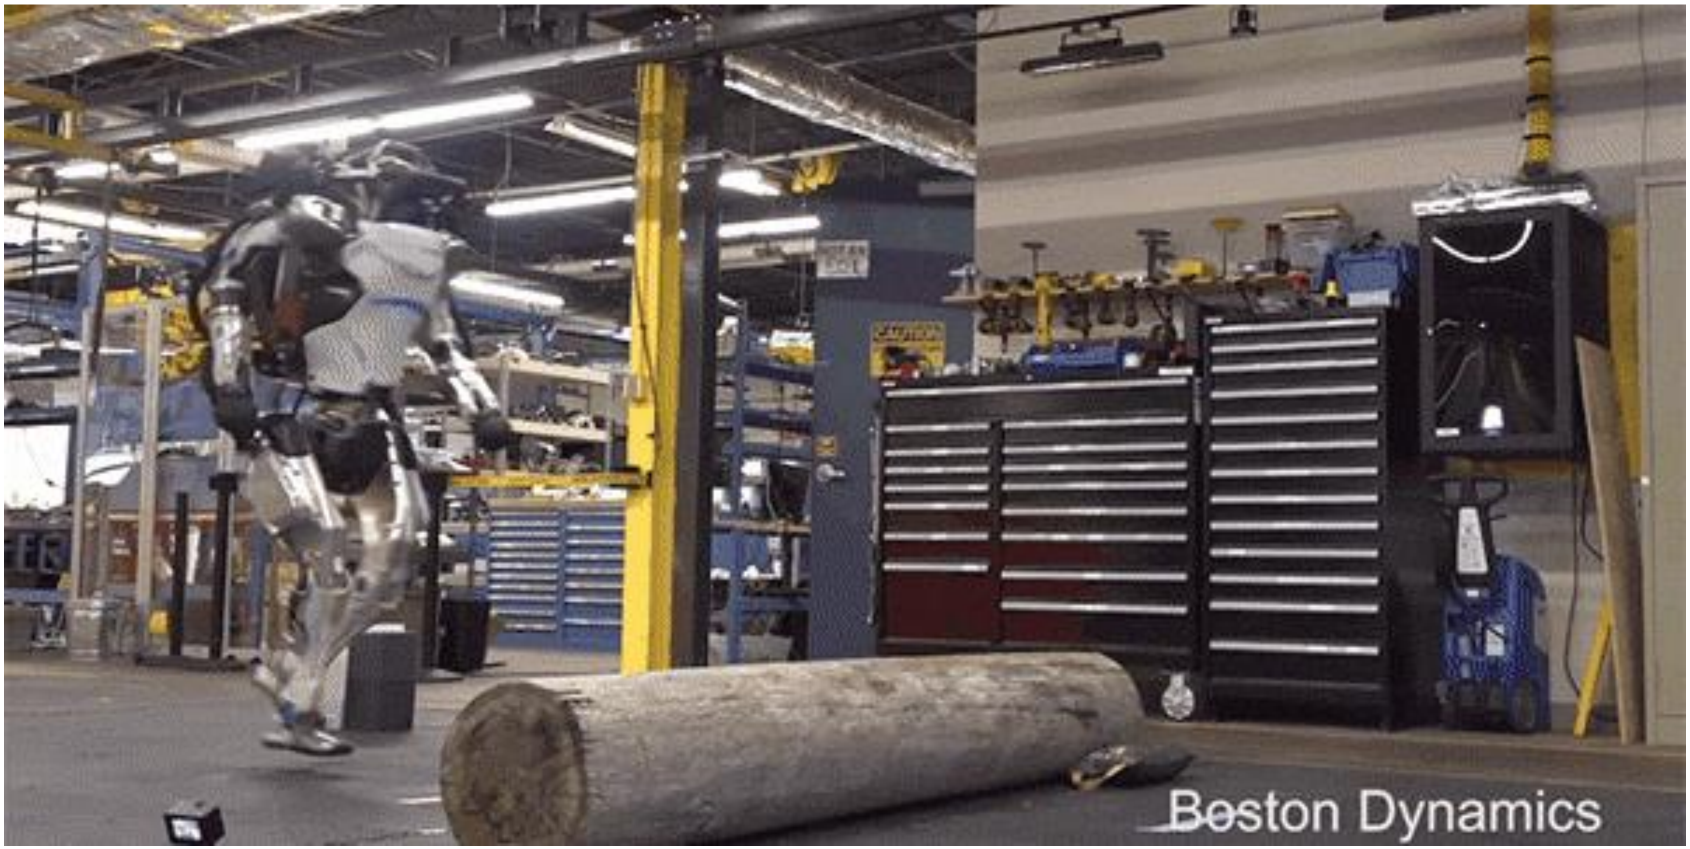

https://www.youtube.com/watch?v=Jky9I1ihAkg

Robotic Success

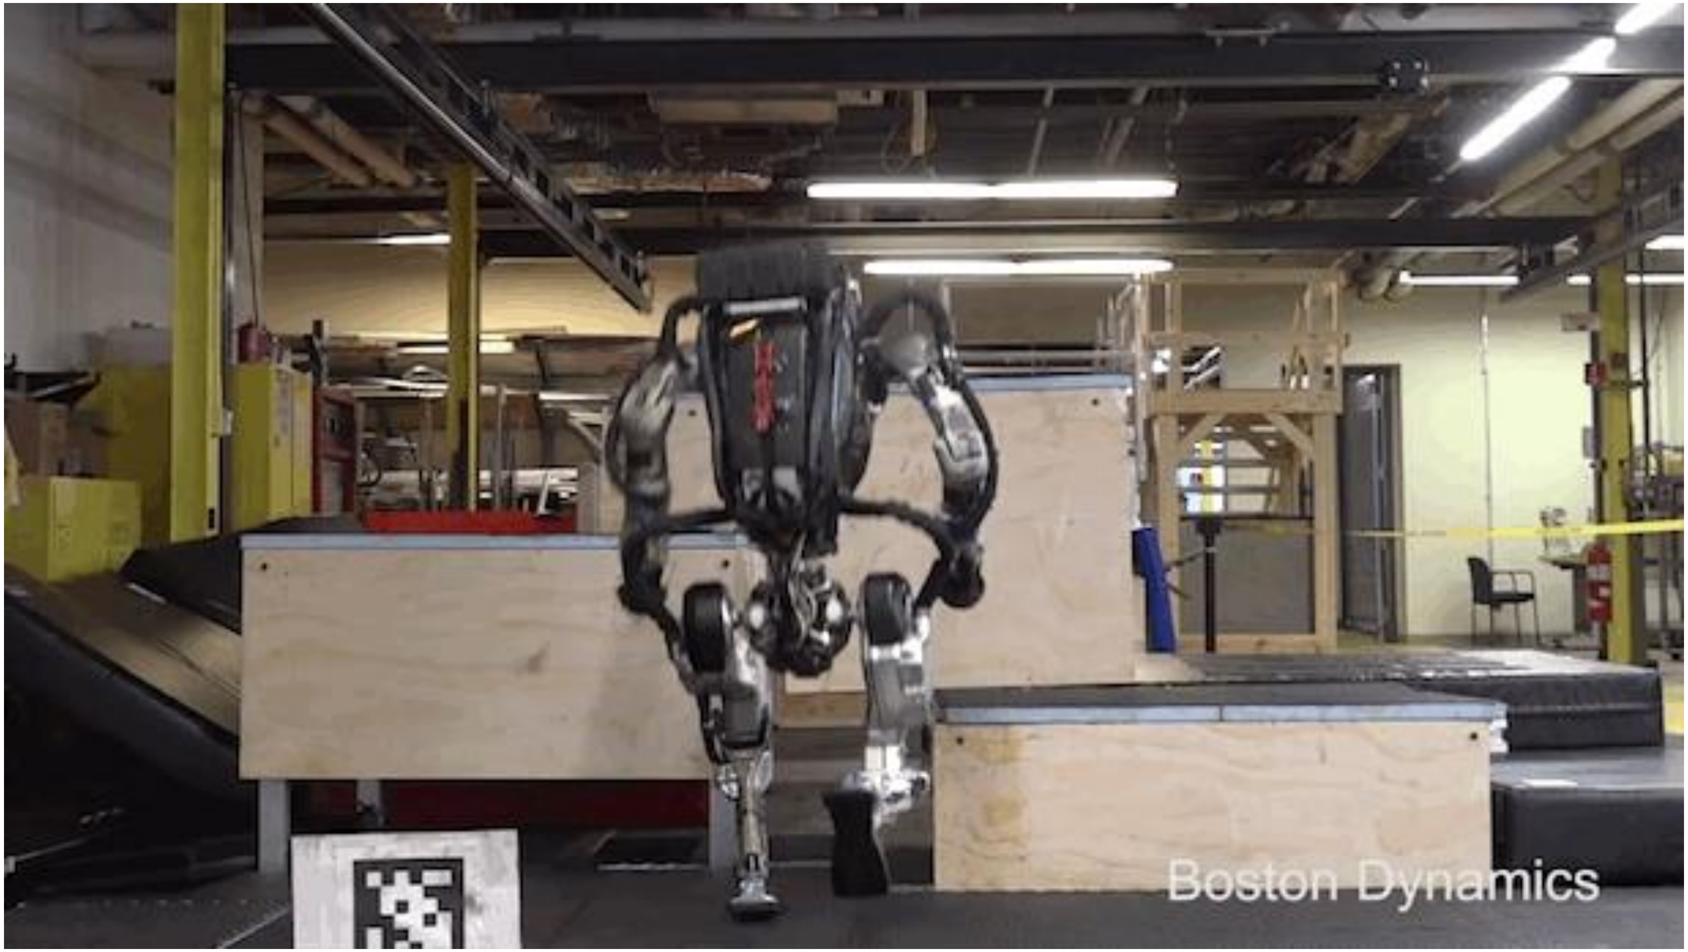

why Actively control our Perception?

A relevant question for several fields


新华


#HazardSpotting

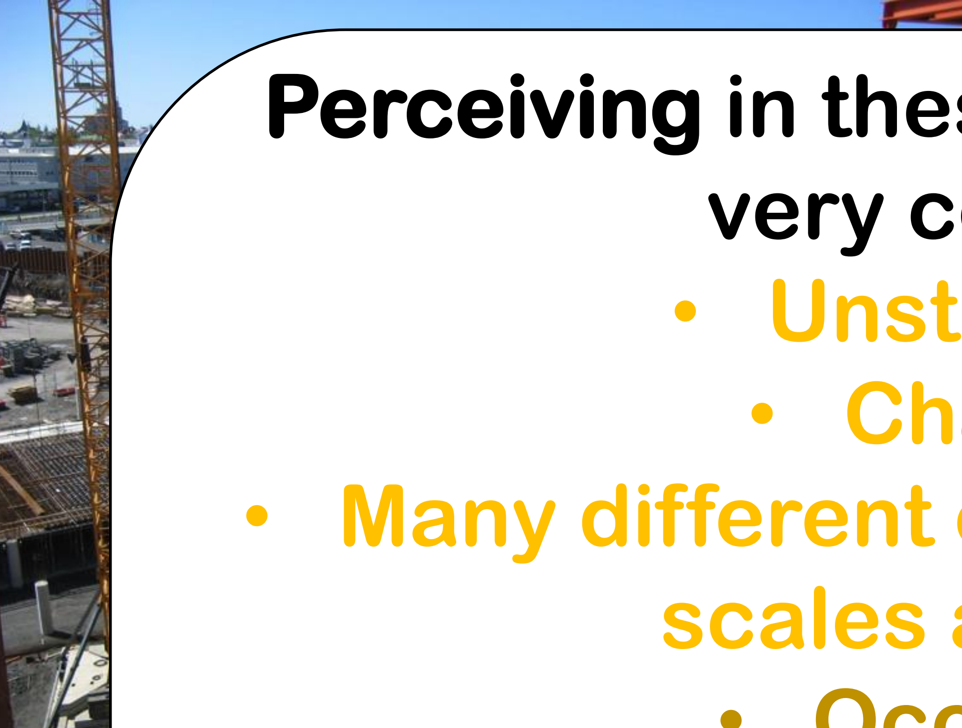

Perceiving in these environments is very complex:

- Unstructured
- Changing
- Many different objects of different scales and shapes
- Occlusions
- Other moving agents to perceive and coordinate with

Currently only humans are able to cope with such level of perceptual complexity…

And humans perceive (pro-)actively…

Occlusions/Interference CV Failures

Rosenfeld, Zemel & Tsotsos. ArXiv 2018

Robotic (not) Success

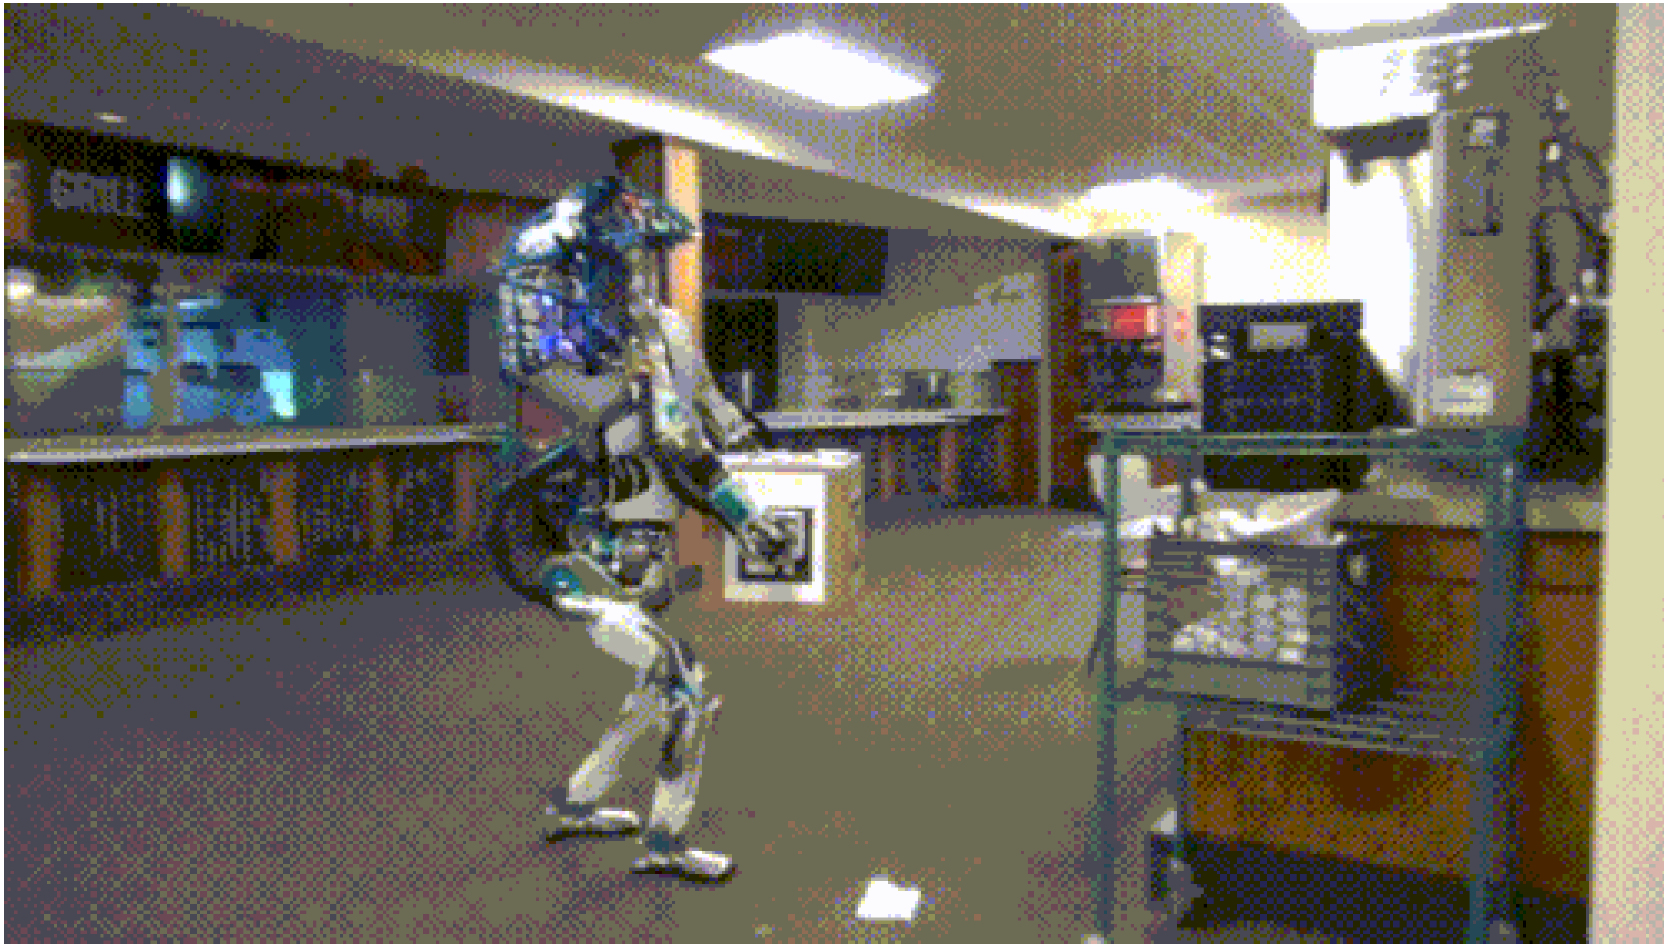

https://www.youtube.com/watch?v=g0TaYhjpOfo

Robotic (not) Success

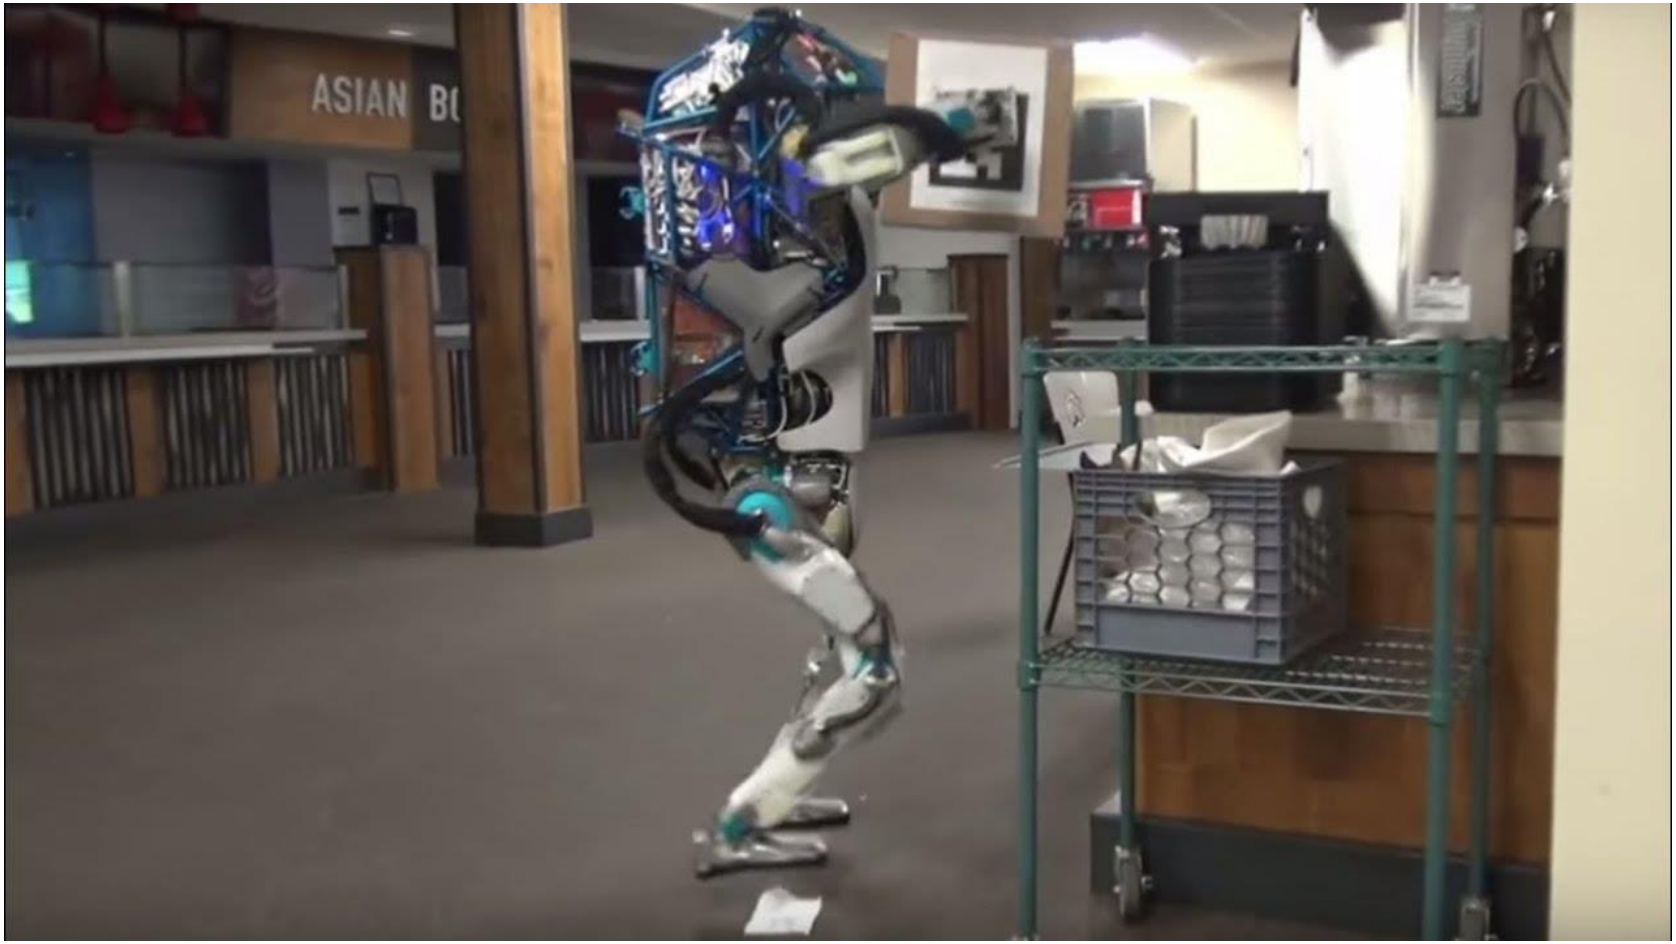

Self occlusion...

https://www.youtube.com/watch?v=g0TaYhjpOfo

Actively Control Sensors to overcome Perceptual Limits

Occlusions, Sensor Resolution, Limited Field of View, Noise and other causes of Aliasing or Partial Observability

Foveal Vision: Human Eye Evolved to be Active and Focus on Relevant Cues

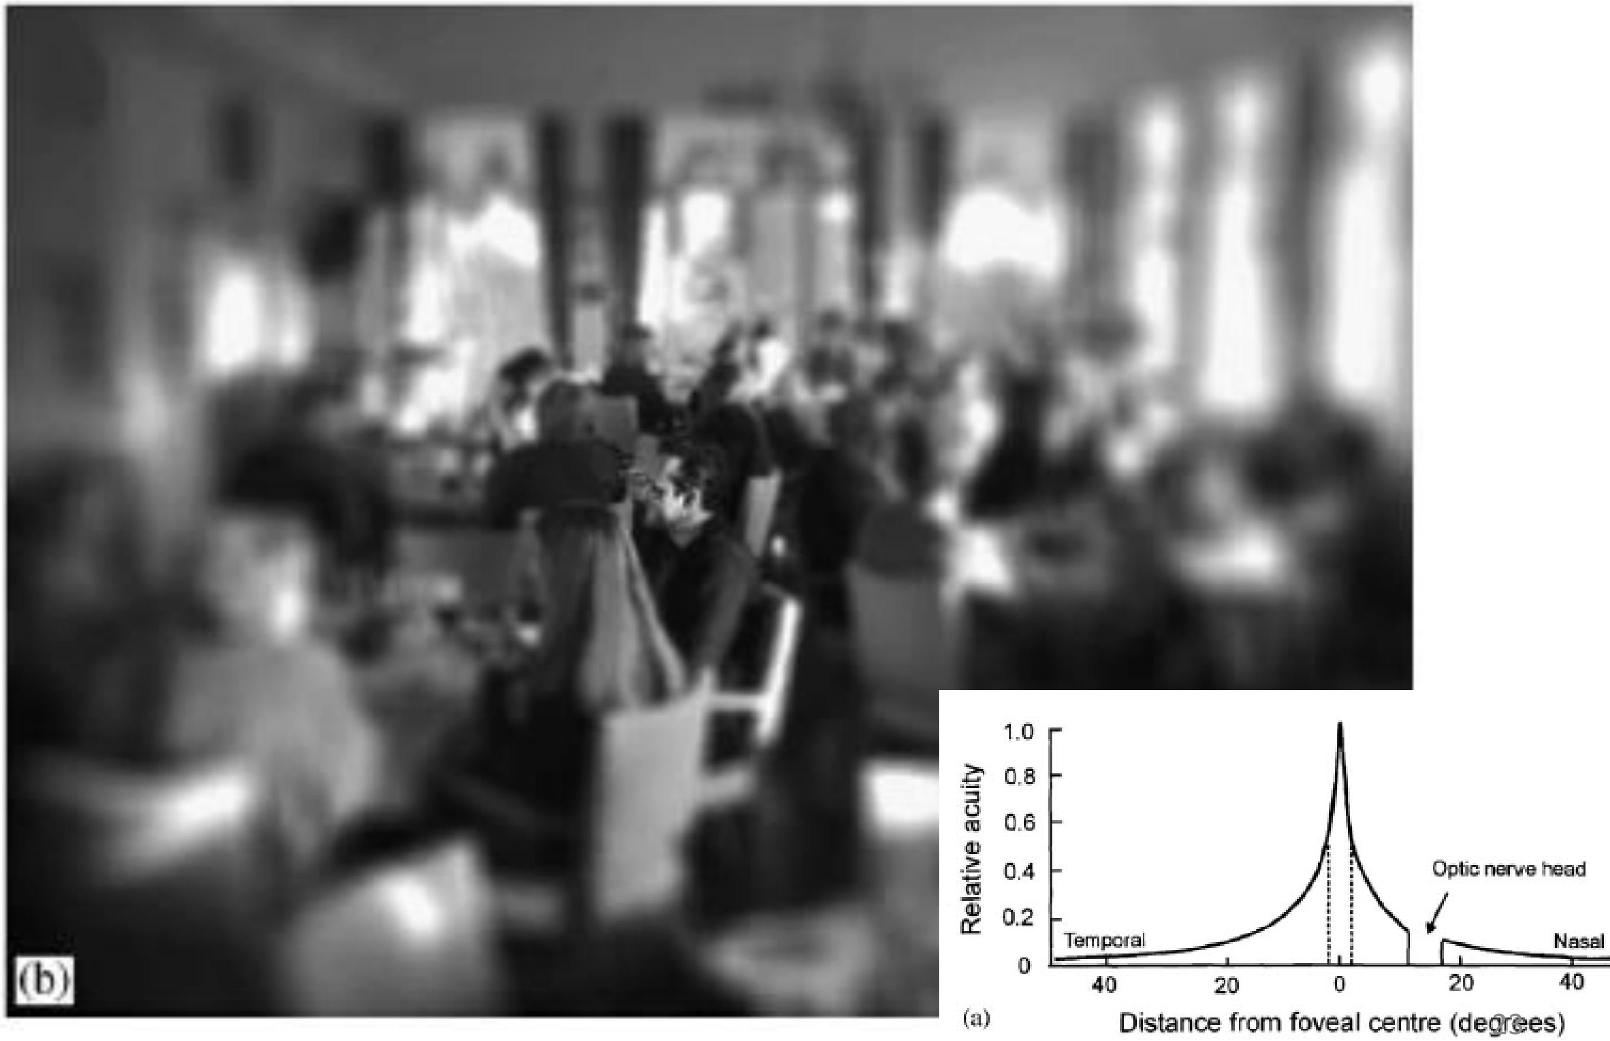

Foveal Vision: Human Eye Evolved to be Active and Focus on Relevant Cues

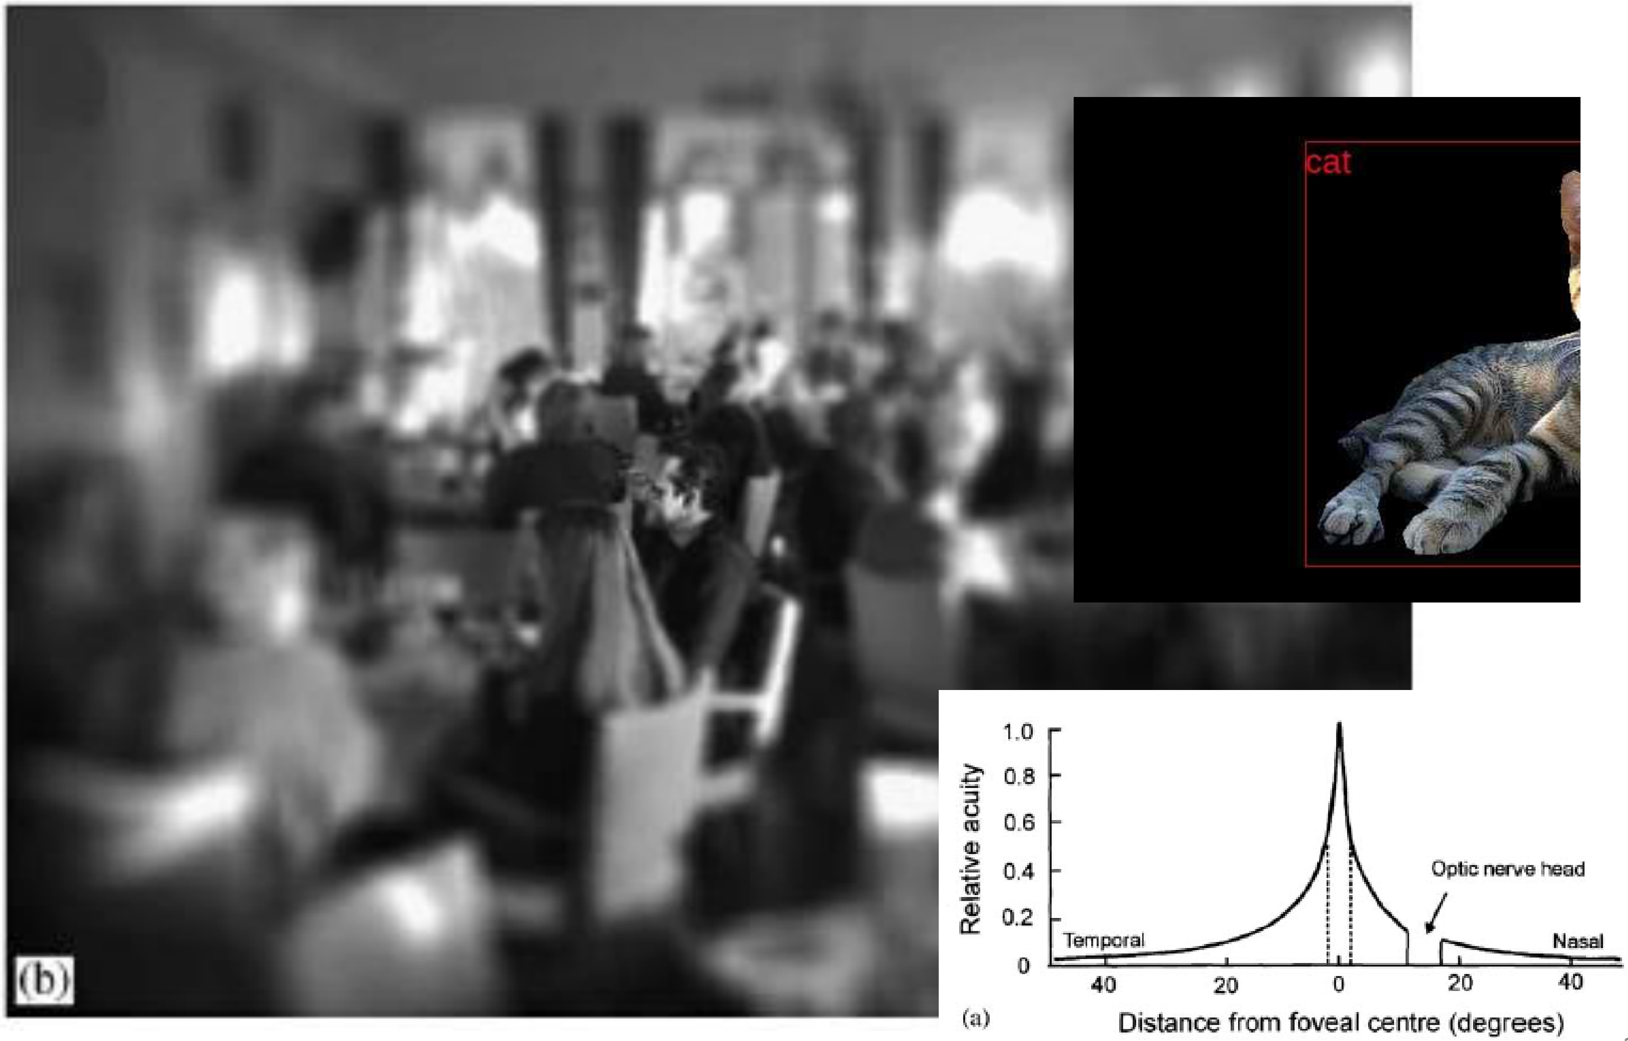

Active Perception **(AP)** Issues

- **Where to look?**
- What to remember?
- When to stop looking and start acting?
  - Enough information?
  - Enough time?
  - Acquired information still valid?

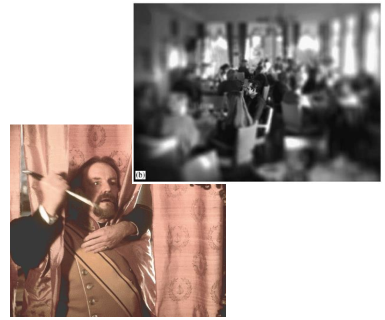

*See also The Frame Problem

Where To look?

Task Based Exploration

(Information On Demand)

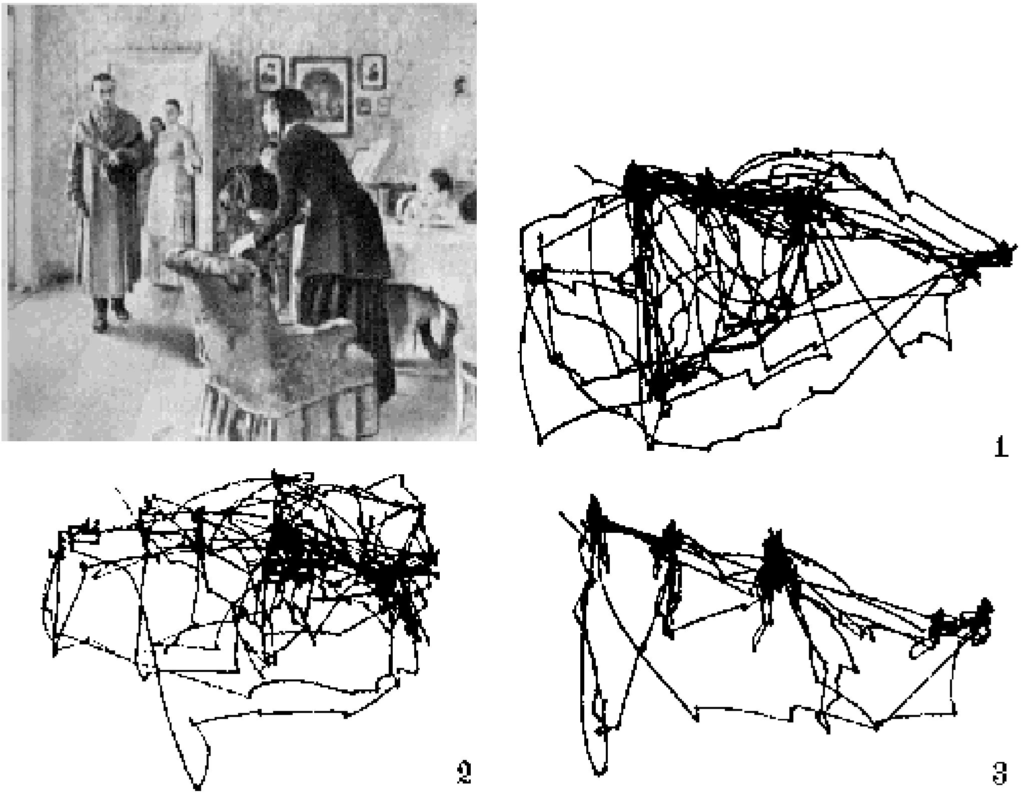

Yarbus 1967

Insufficient Task or Goal Information

- Learning

Constraints

Focus

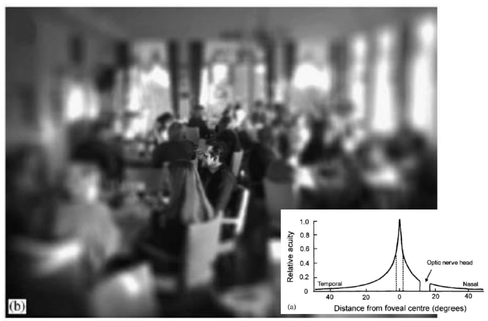

Task

Active Perception and Learning

Active Perception is strongly dependent on the task

Learning a new task may require learning a new Active Perception policy

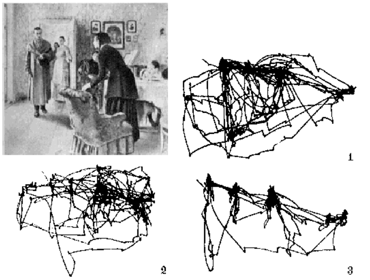

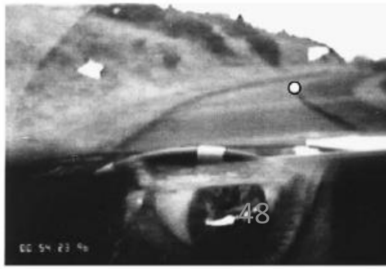

Active Perception and Learning

Active Perception and Learning

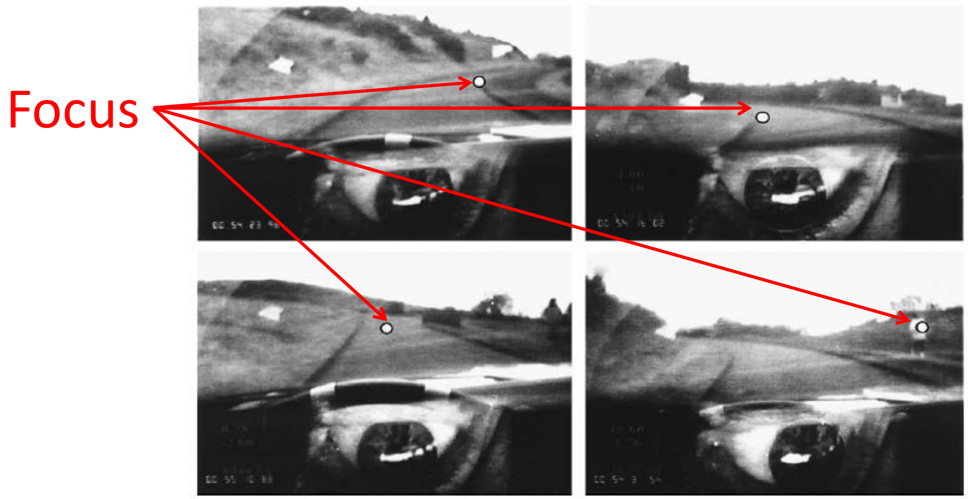

Land 2006

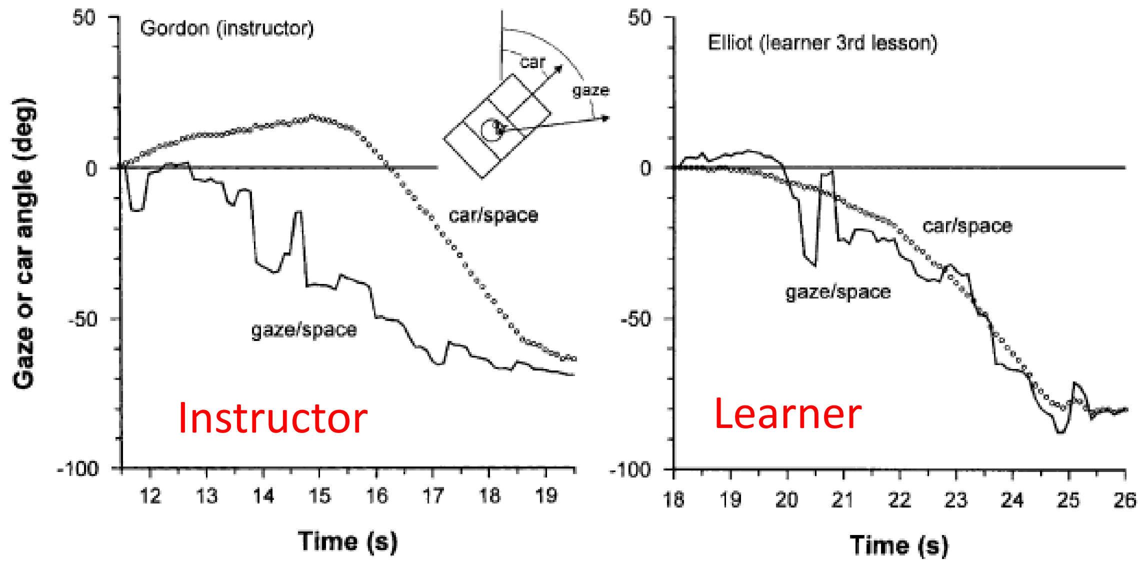

Difficulties of learning perception and action control

- High dimensionality of the observations
- Unknown/Infinite state dimensionality
- Partial Observability
  - check:
    - Ognibene, Volpi et al 2013
    - Whitehead and Lin 1995
- Need to learn reusable skills
- No supervision or immediate feedback
  - e.g. no reward for watching the right object

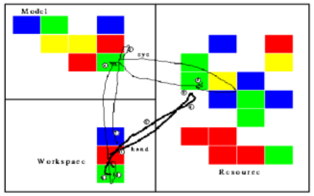

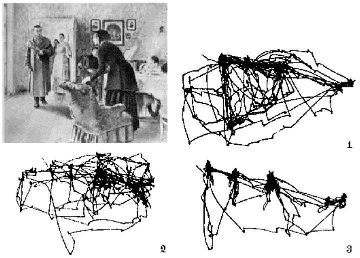

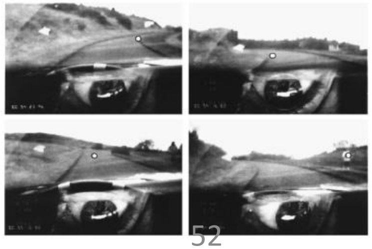

Ognibene & Baldassare 2015

ANN Controlled Camera-Arm Robot

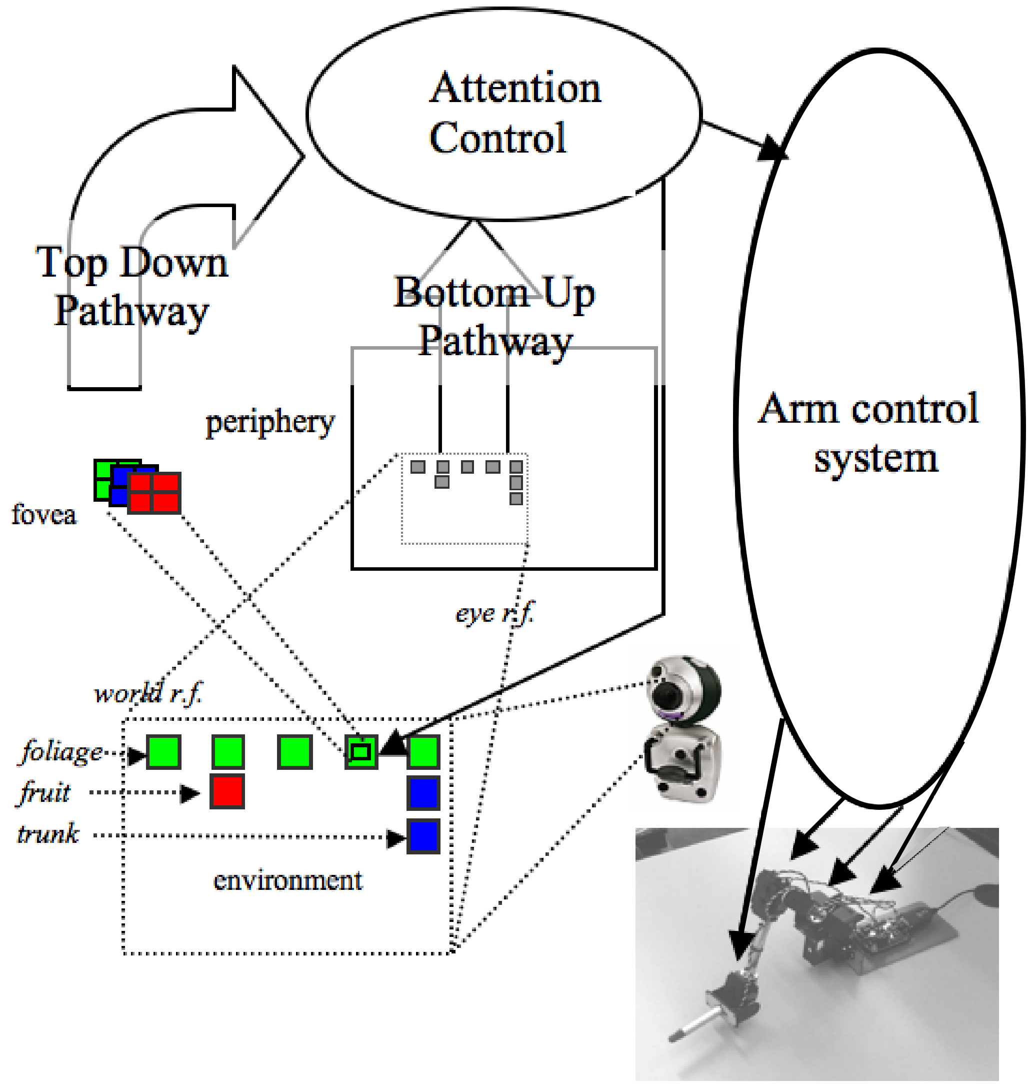

Ognibene & Baldassarre, 2015

Ognibene & Baldassarre, 2015

System Atlas

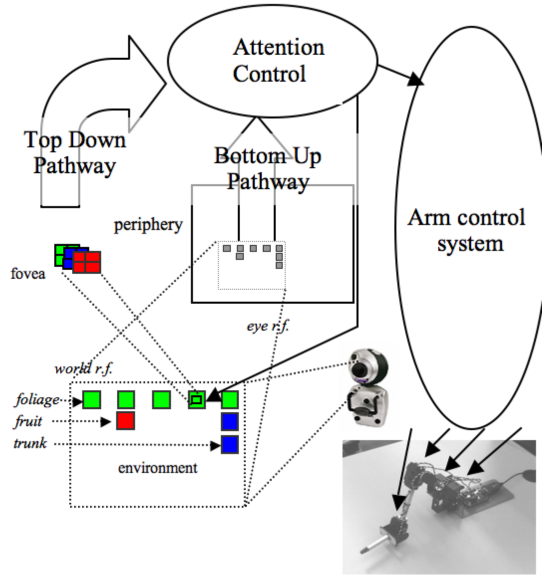

VOTES

PAM

CAMERA

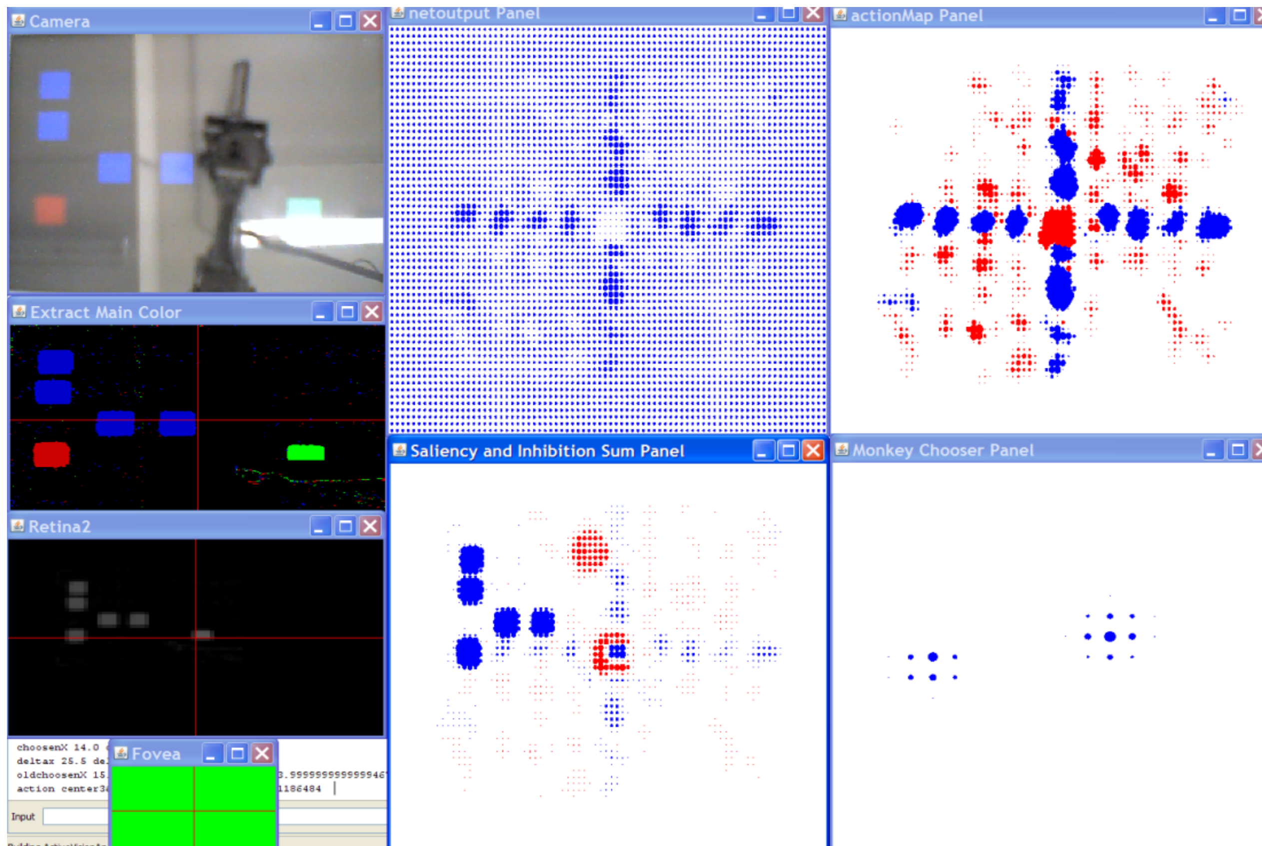

FILTERED IMAGE

BOTTOM-UP ATTENTION

FOVEA

OVERALL SALIENCY

COMPETITION FOR ARM'S ACTION

Base Ecologic Task Behavioural analysis

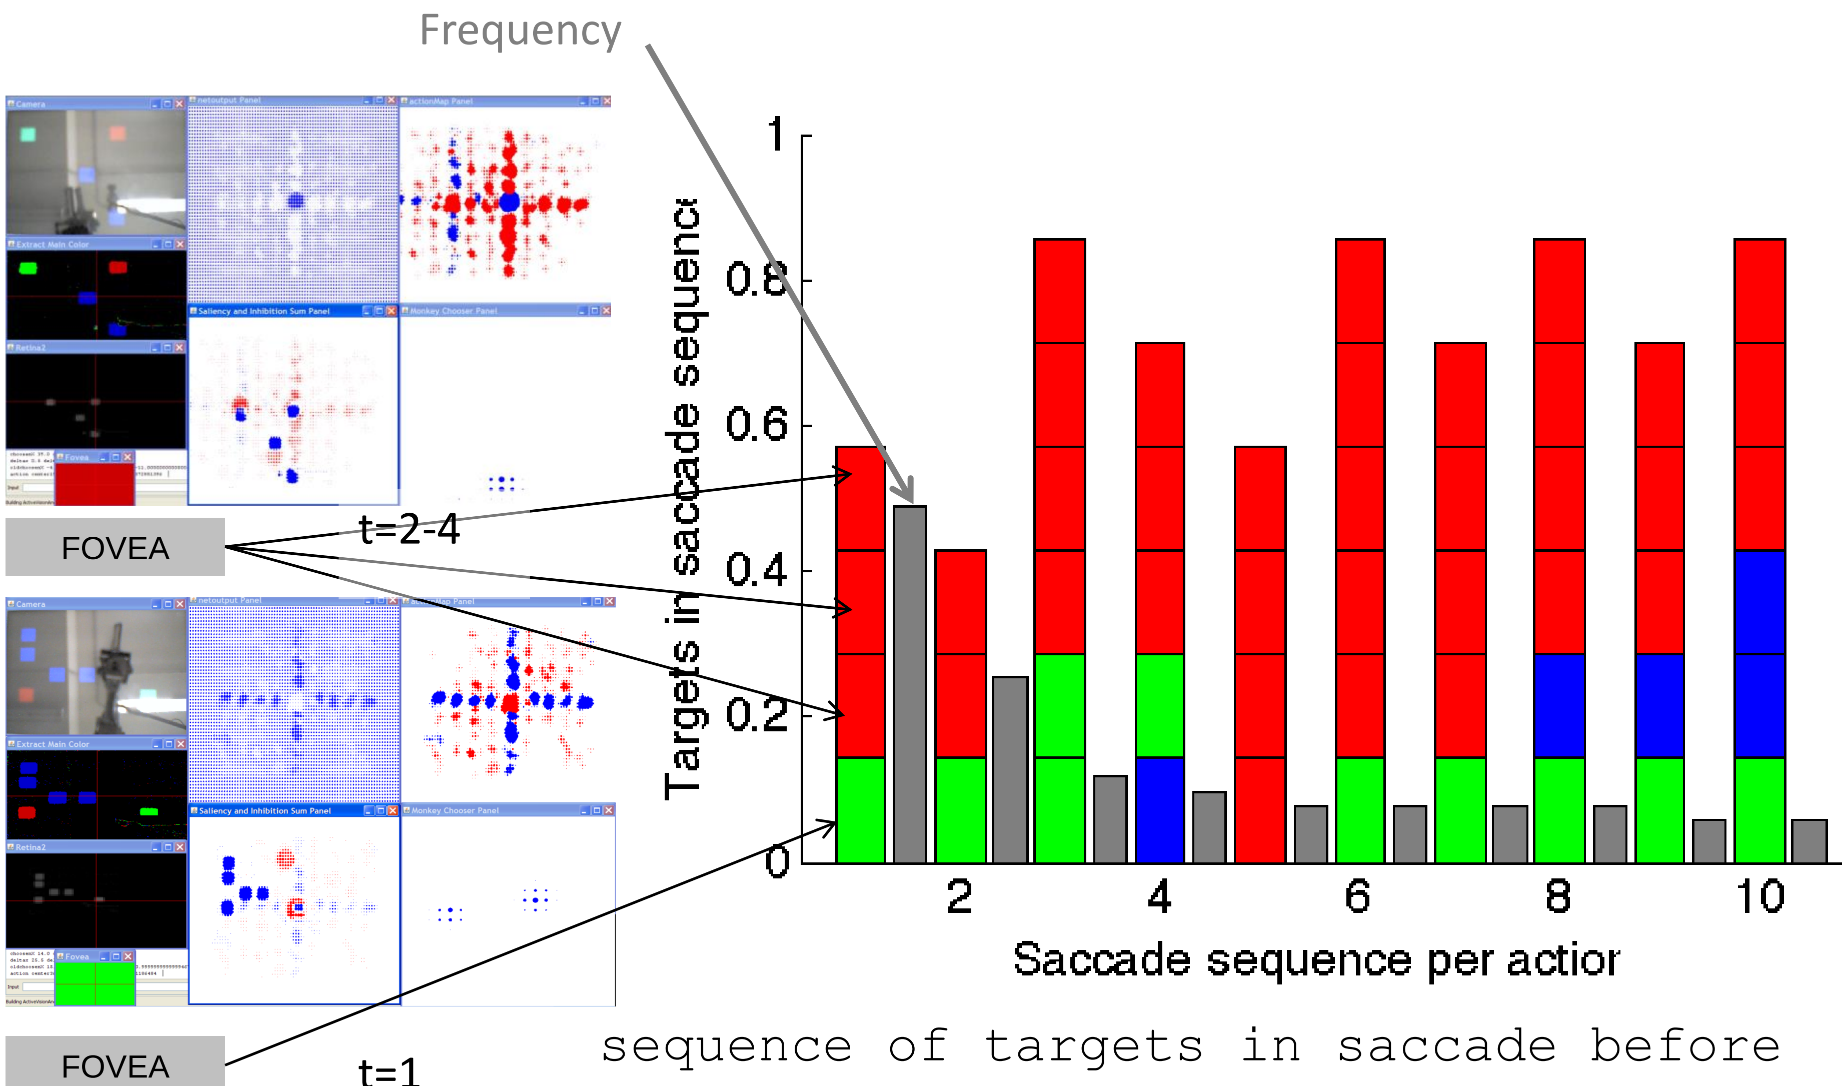

sequence of targets in saccade before reaching action started (after learning)

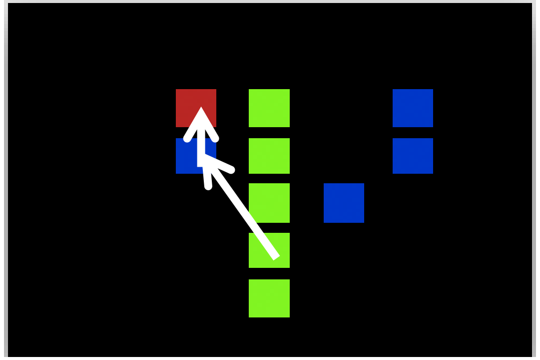

Internal dynamic analysis

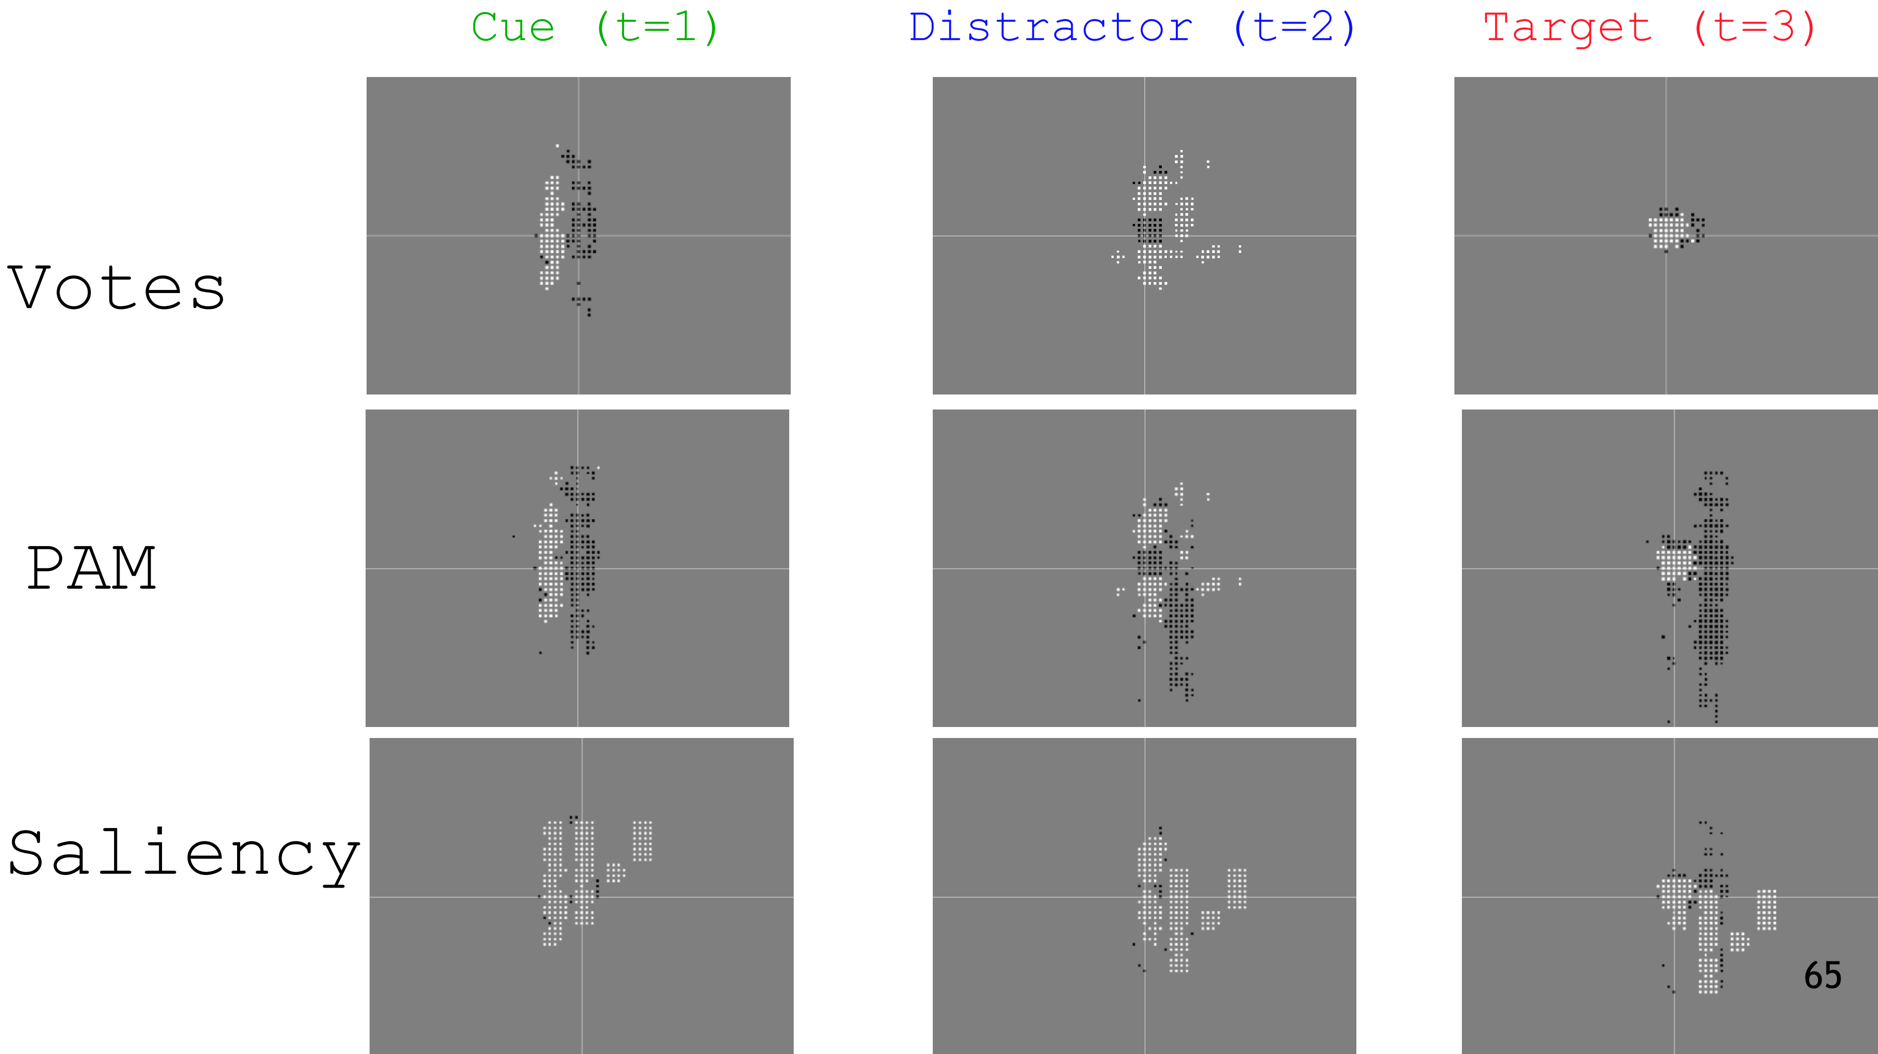

Actual World Structure

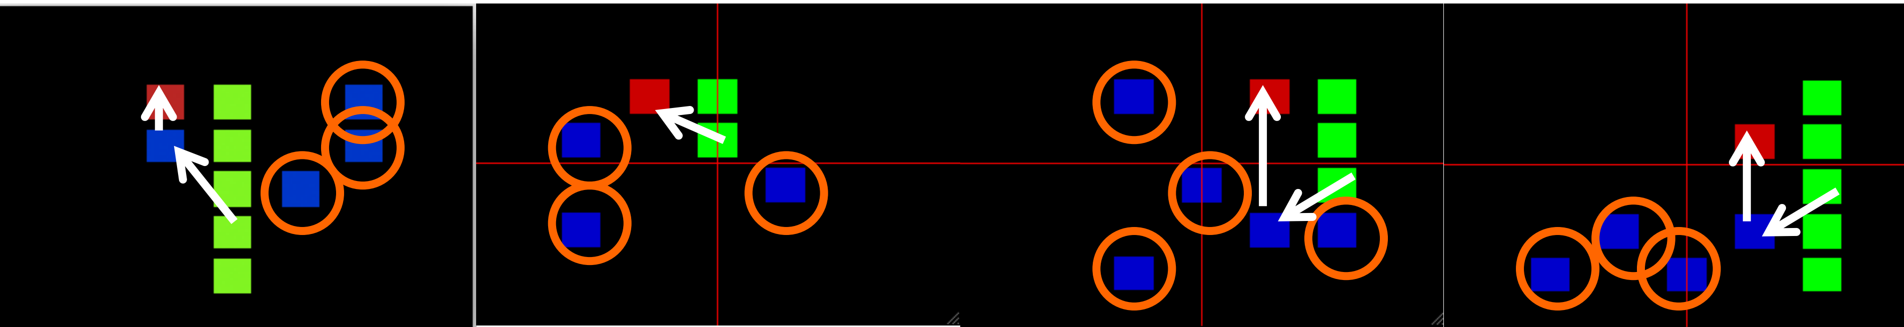

Task: in an environment with salient (bright and big) green object learn to find and touch the red object with no supervision. Agent can see the colour of one object at time.
The blue objects are randomly positioned

Acquired neural representation of blue stimulus describing next gaze movement convenient to perform the task: up,down,near right
It presents order!

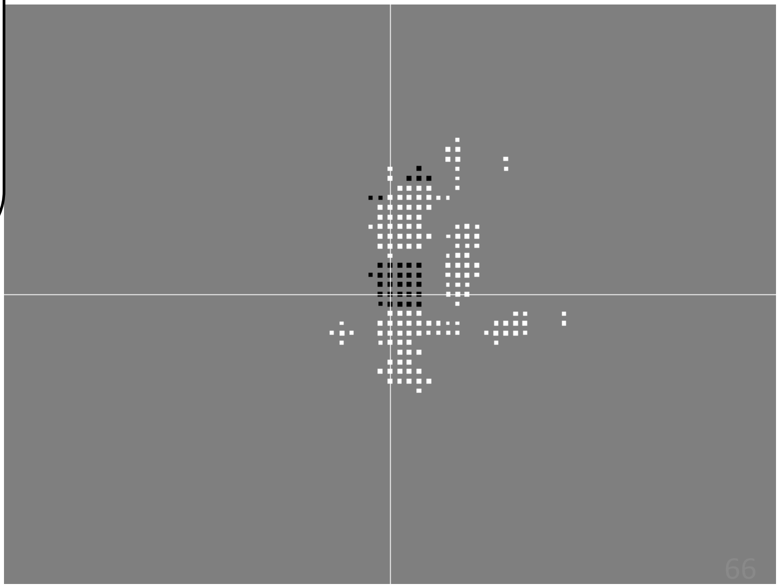

Ognibene & Baldassarre, IEEE TAMD, 2015

Subjective and efficient representations

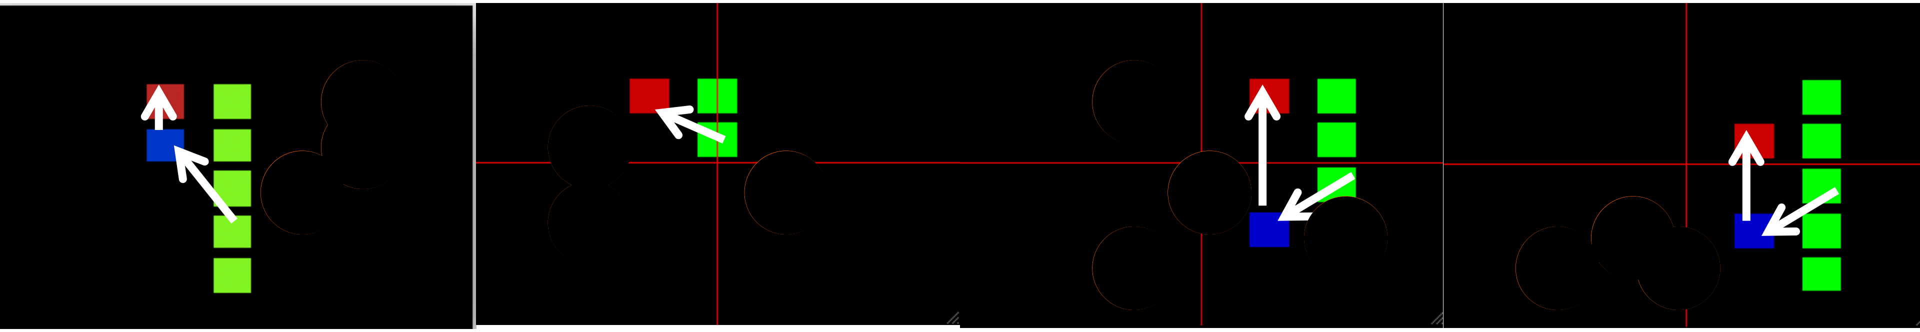

Perceived World biased by Active Perception

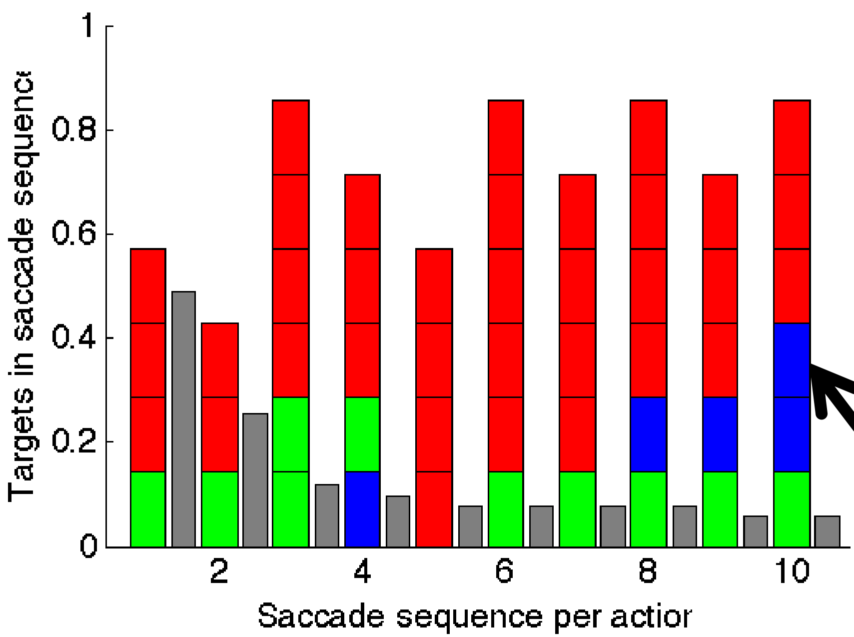

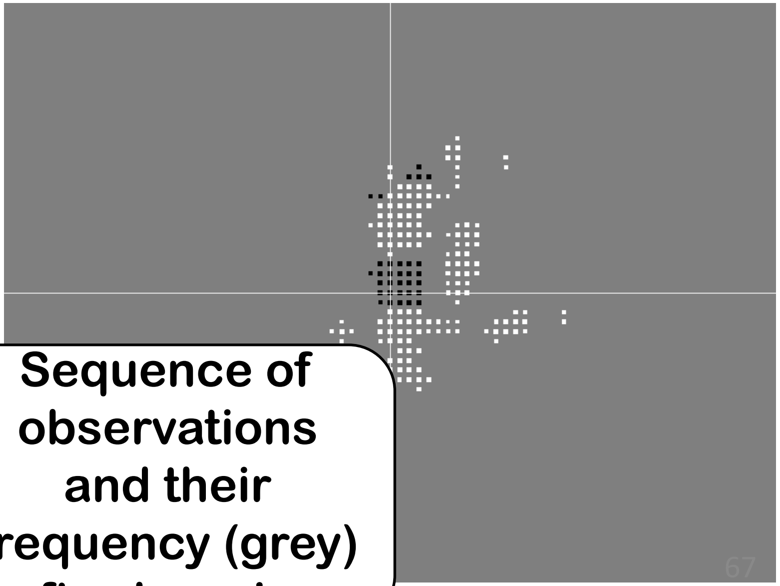

Sequence of observations and their frequency (grey) after learning

Ognibene & Baldassarre, IEEE TAMD, 2014

Experimental Results: Decision Time

- Initially a reaching action starts after 17saccades after scene changed
- Increases to 26 saccades
- Stabilises at about 7 saccades

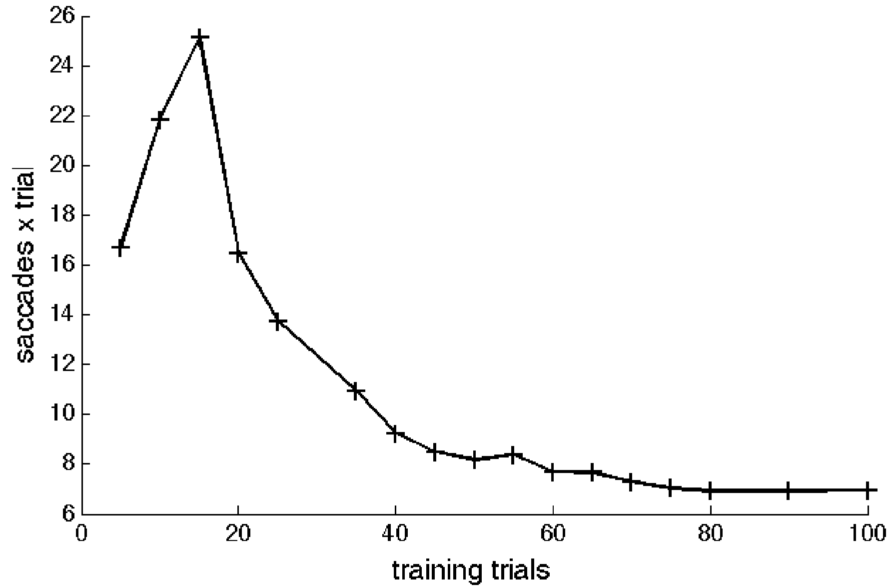

Average number of saccades per reaching action during learning

Saliency Role Test

saliency				number of saccades		
cue	distr.	trials to steady	avg. reward	max	initial avg.	final avg.
0.6	0.6	1	1	14	5.1331	5.0821
0.8	0.6	4	0.9951	60	6.1678	5.8852
1.0	0.6	41	0.9957	55	6.8448	5.9302
0.6	0.8	7	0.9851	42	6.7102	6.1334
0.8	0.8	4	0.9778	35	6.9981	6.4806
1.0	0.8	23	0.9811	66	7.0478	6.2394
0.6	1.0	83	0.9502	145	9.9367	7.8187
0.8	1.0	78	0.9804	387	12.2187	7.4450
1.0	1.0	74	0.9604	1727	16.3086	7.2732

Target saliency 0.8

Only learning is affected by clutter while final performance are minimally affected by cue saliency

The trial to steady is very dependent on the presence of an object that is more salient than the target

Foveal Vision *May* Speed-up Task Learning

Reinforcement Learning Framework for Autonomous Task Learning.
Usually problematic with partial observability (core for AP).

Constraints

Focus

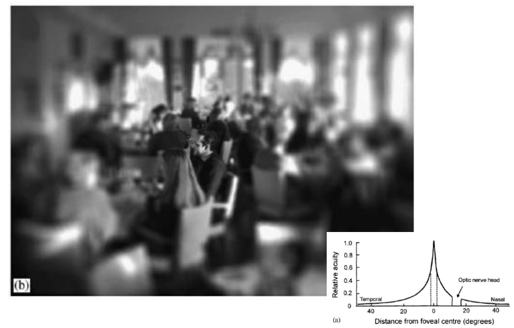

Task

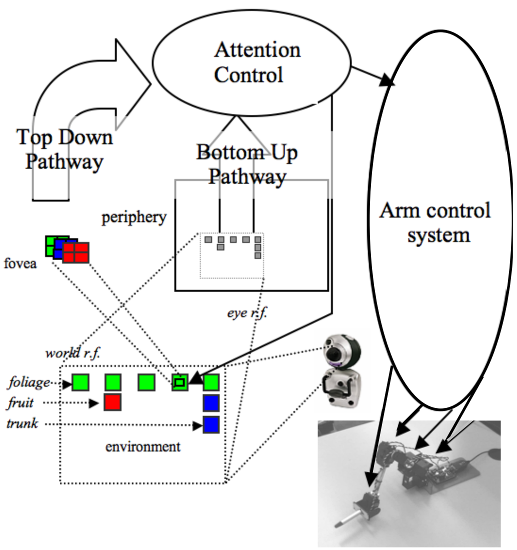

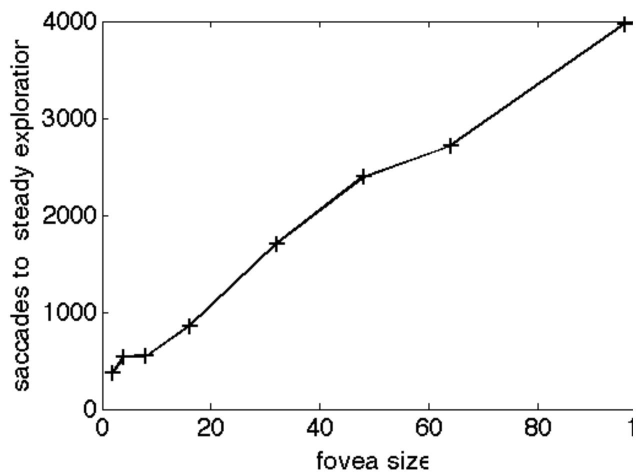

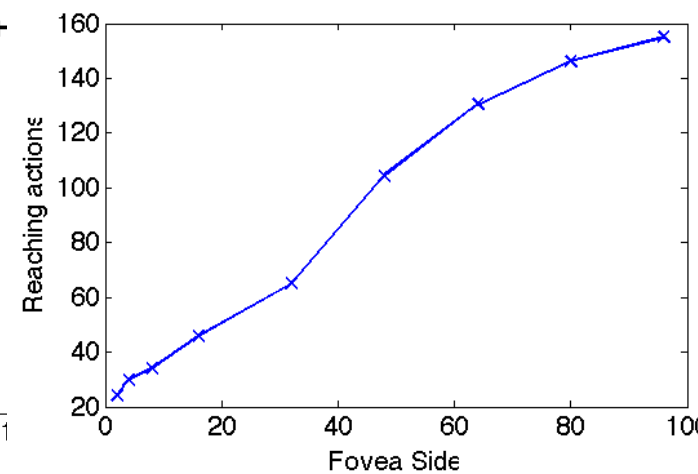

Ognibene & Baldassare IEEE TAMD 2015

Insufficient Task or Goal Information

- Learning
- Deciding what task to execute
- Task/Goal depending on other agents' presence/intentions

Constraints

Focus

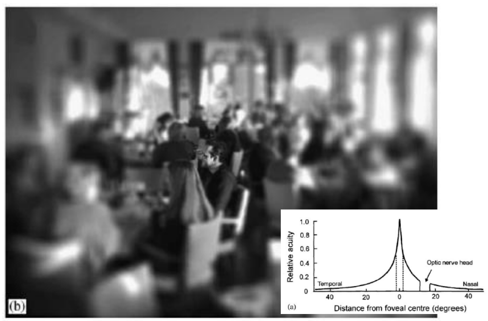

Task

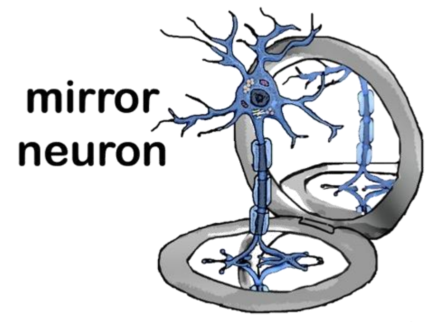

Active Perception and Mirror Neurons

Simulation theory

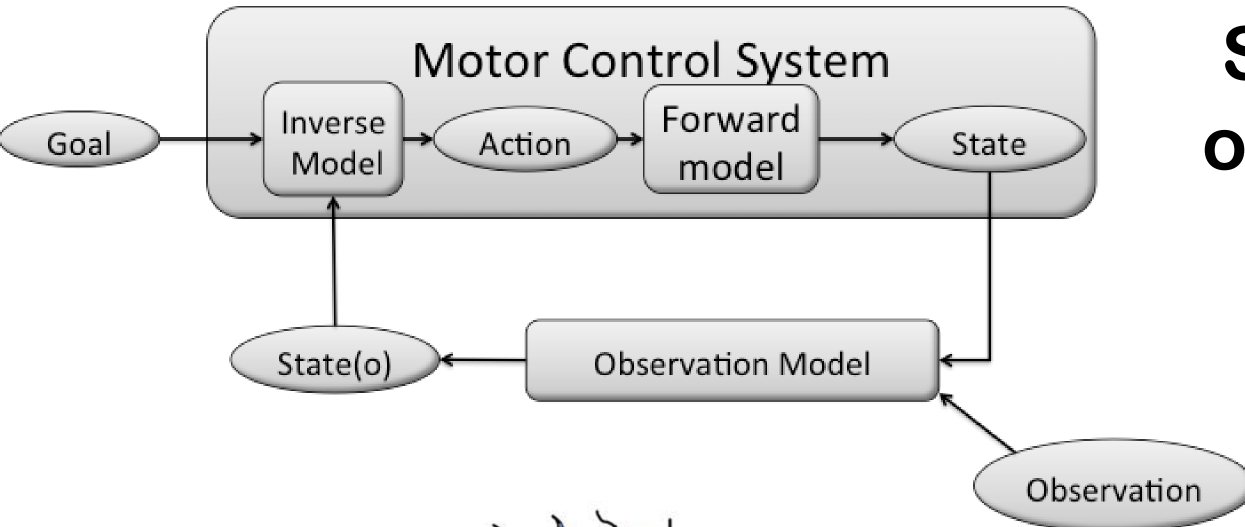

Can Motor Control System predict others' actions?

mirror neuron

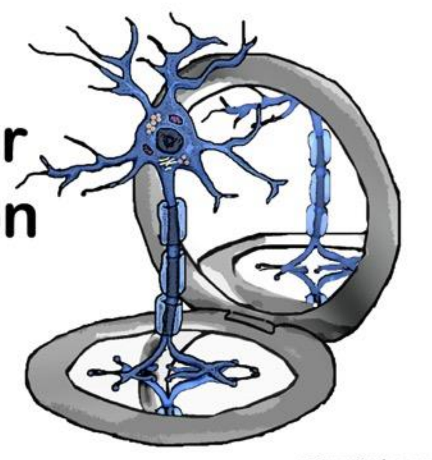

- **Encode action goal**
- **Abstracts trajectory**
- **Needs perceptions**

Active Perception and Mirror Neurons

Simulation theory

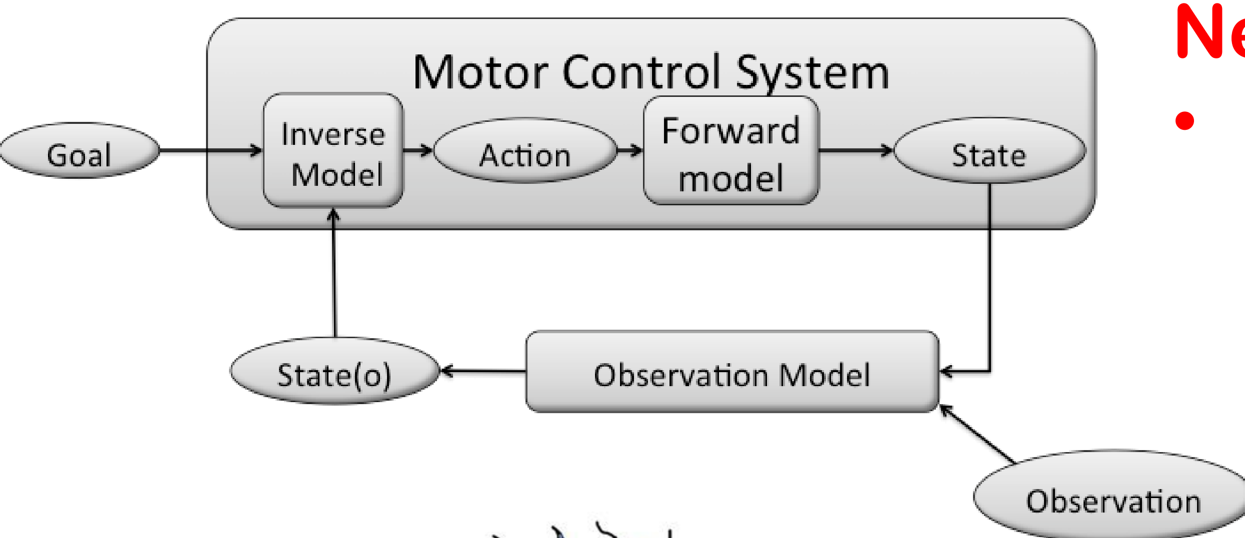

Needs perceptions:

- **Affordance:**
  - **Presence**
  - **Identity**
  - **Position**
  - **Orientation**
- **Effector:**
  - **Timing**
  - **Configuration**
- **Context**

mirror neuron

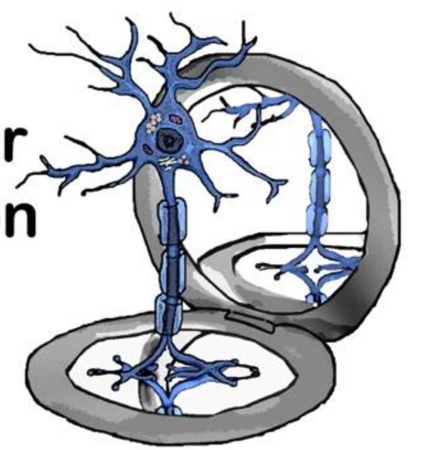

© BumpyBrains.com

Hunting or being hunted? Context in behavior prediction

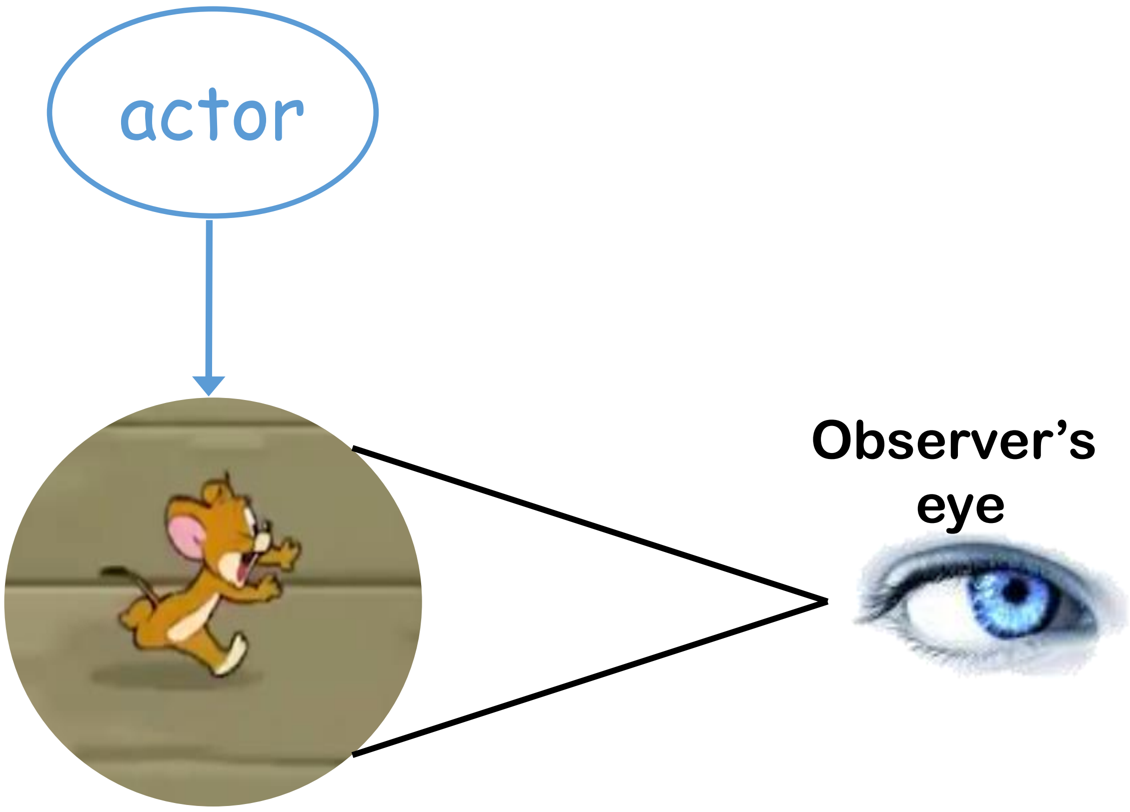

Time constraints and Structured context

Context in behavior prediction

Hunting or hunted?

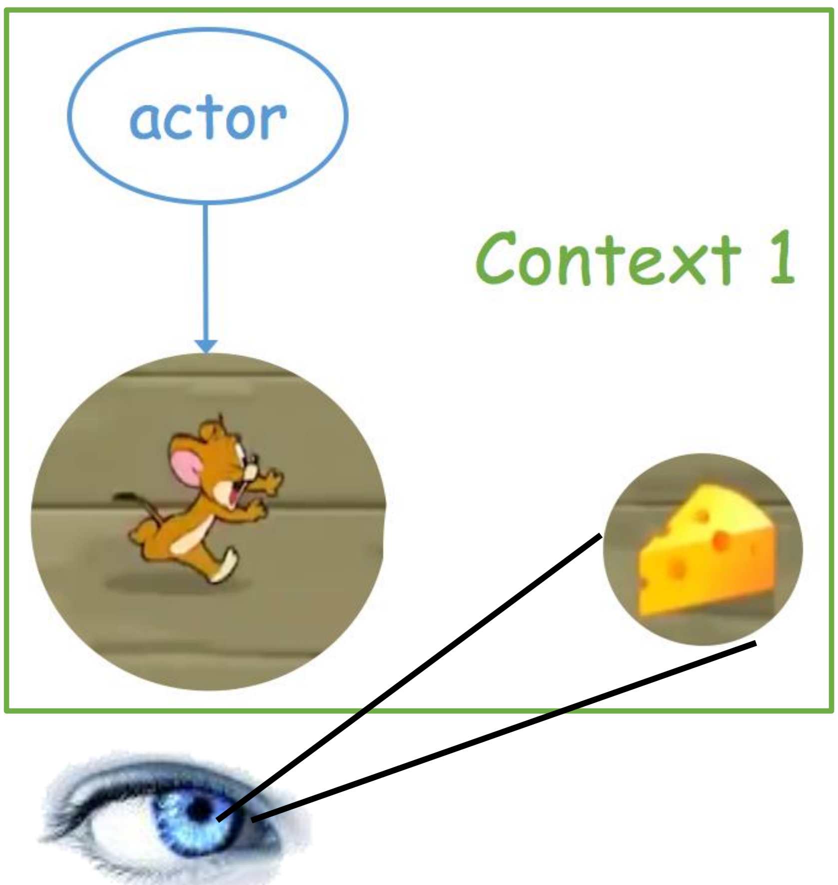

Time constraints and Structured context

Context in behavior prediction

Hunting or hunted?

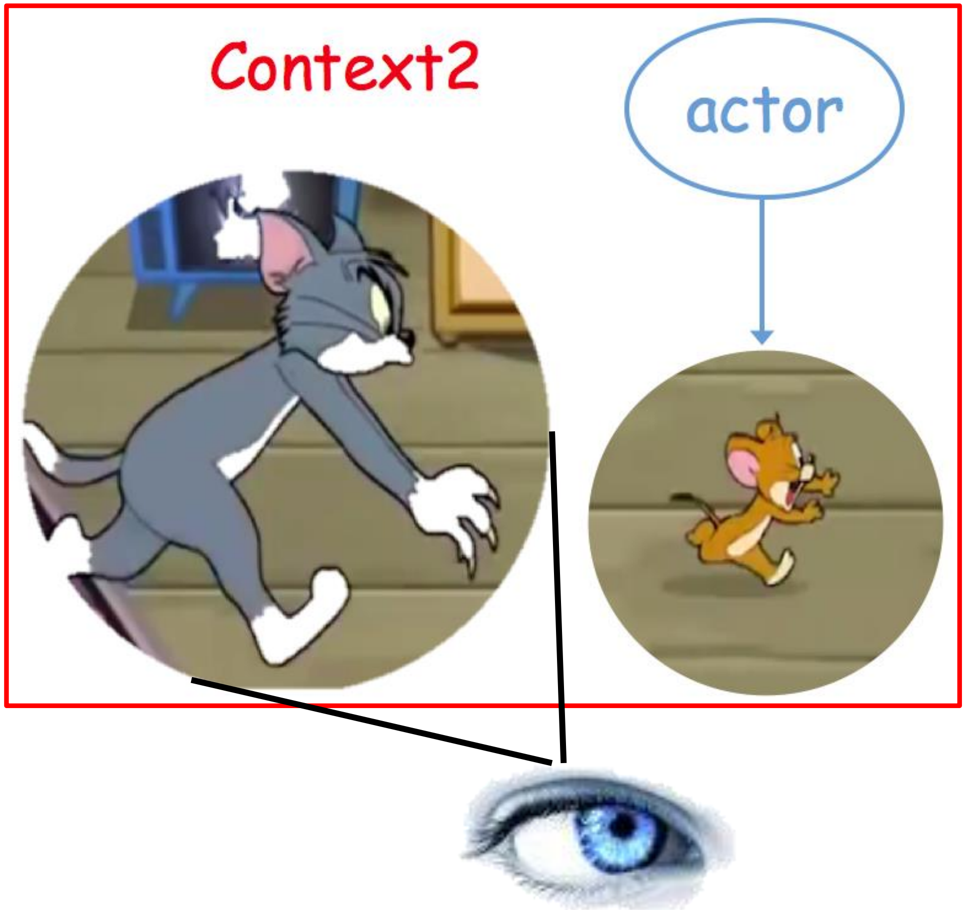

Time constraints and Structured context

Context in behavior prediction

Hunting or hunted?

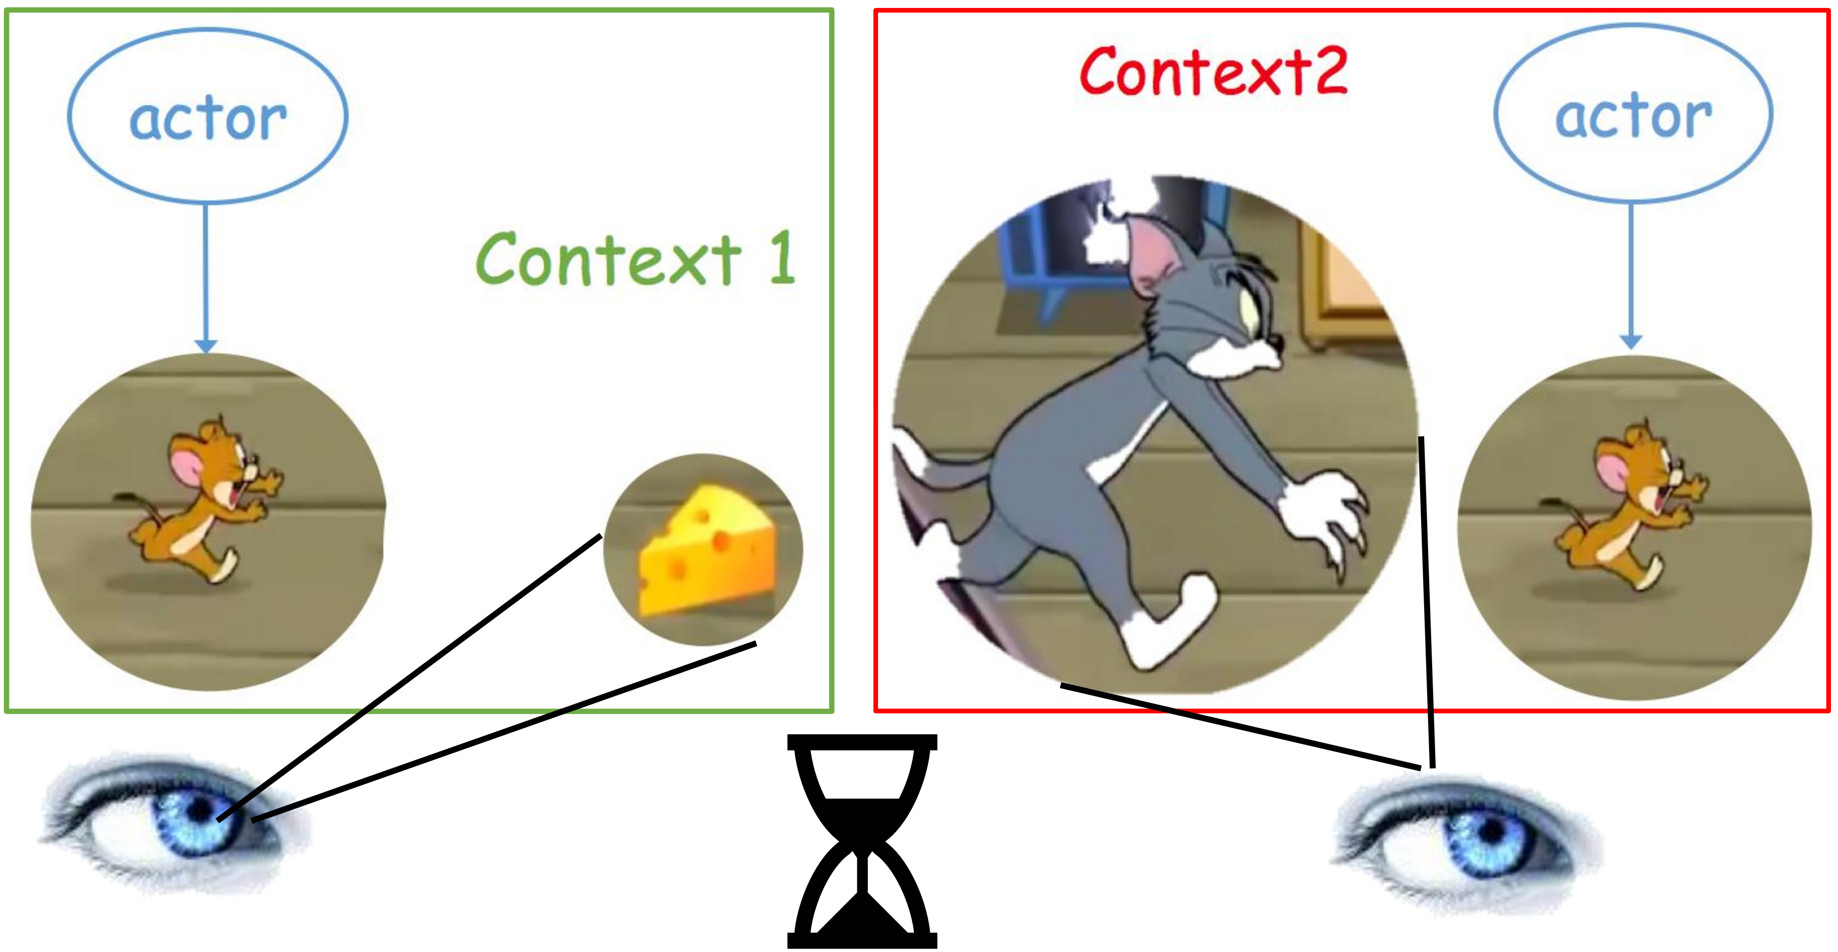

Different hypotheses of target position
Equally probable, not seen

Perceive to reduce uncertainty

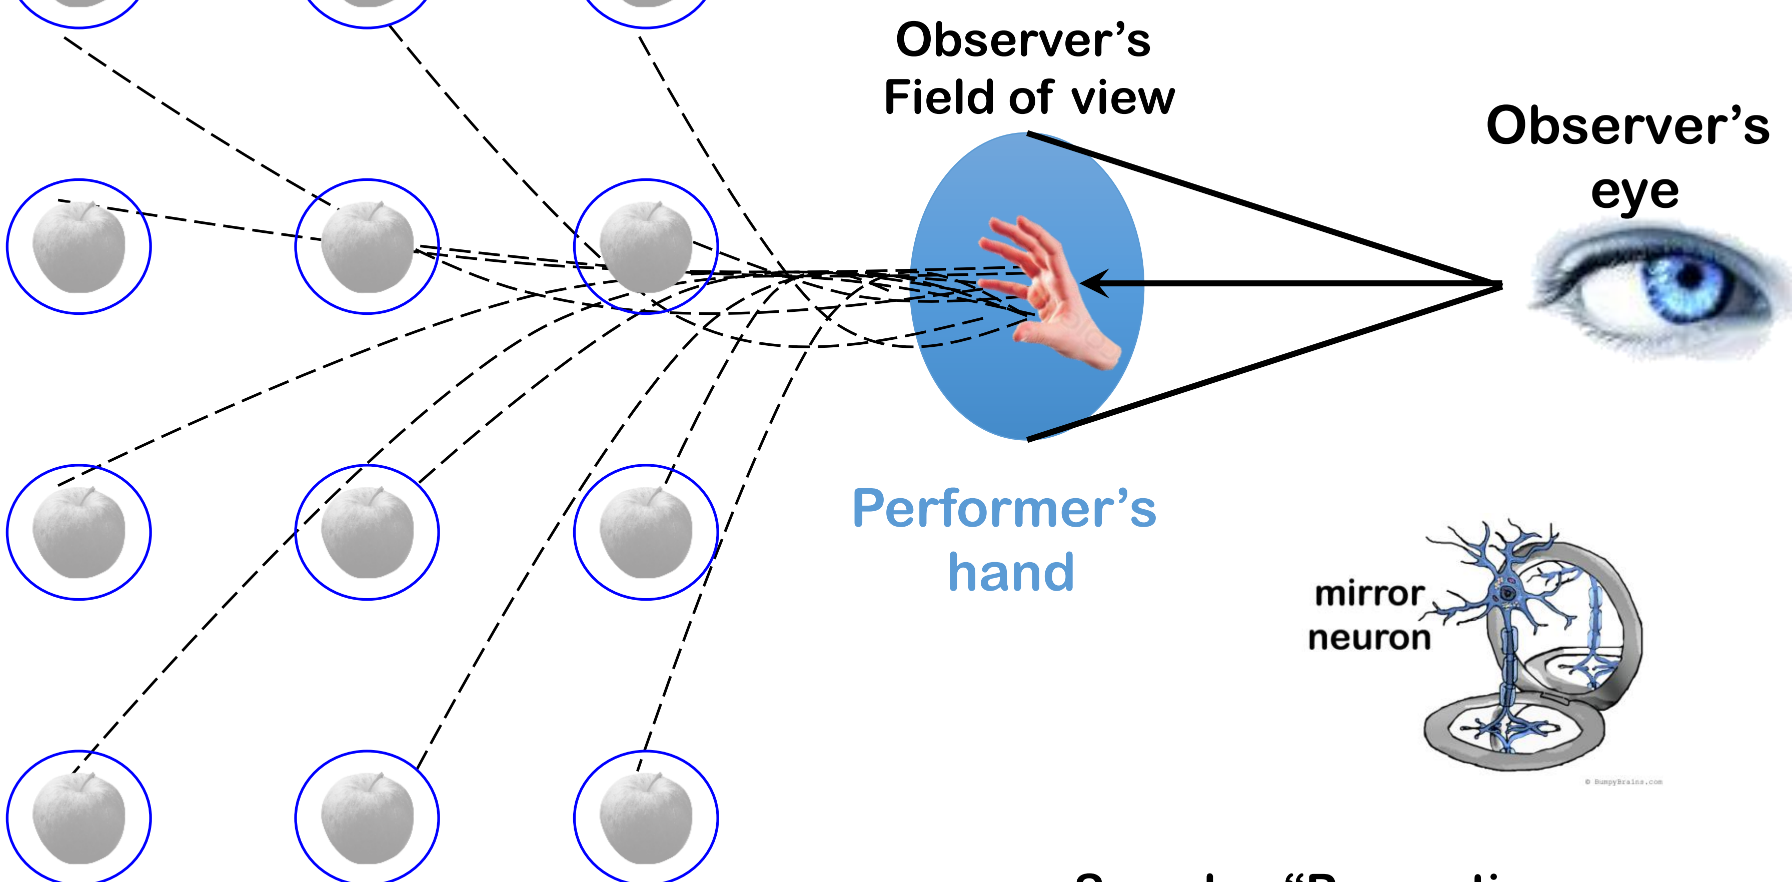

See also "Perceptions as hypotheses: saccades as experiments, Friston et al. 2012"

Hand movement changes distribution on target position

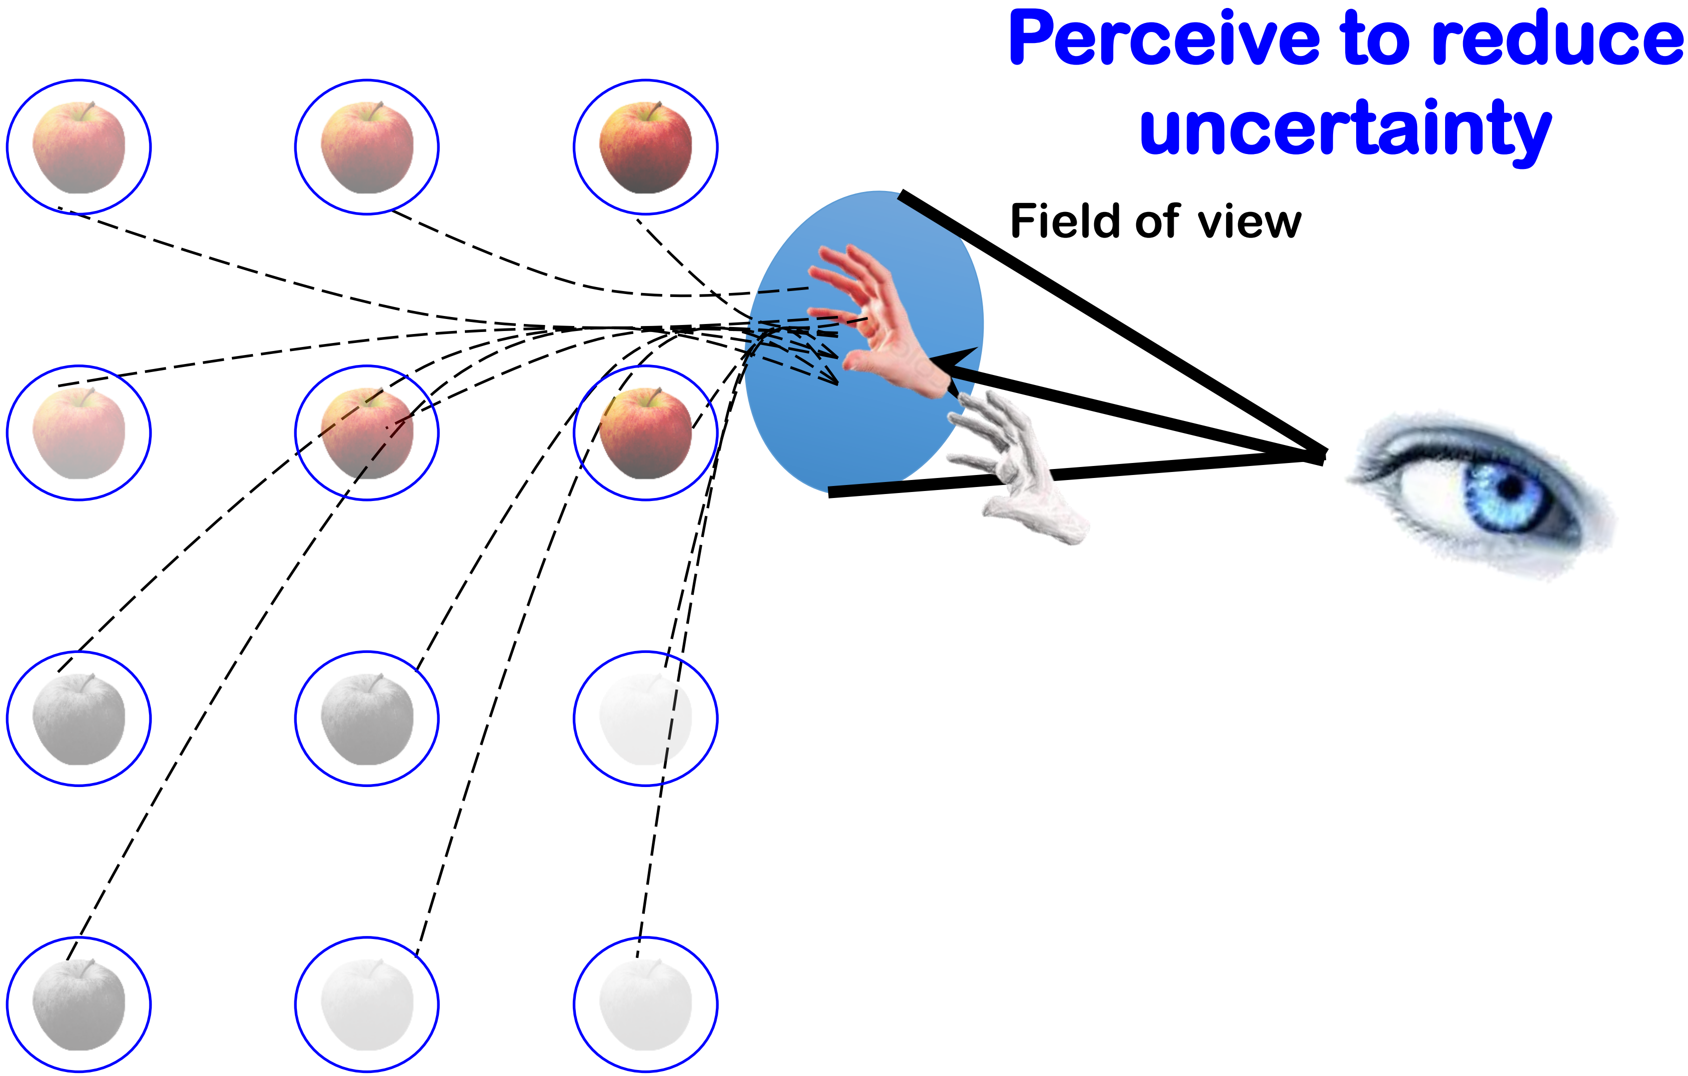

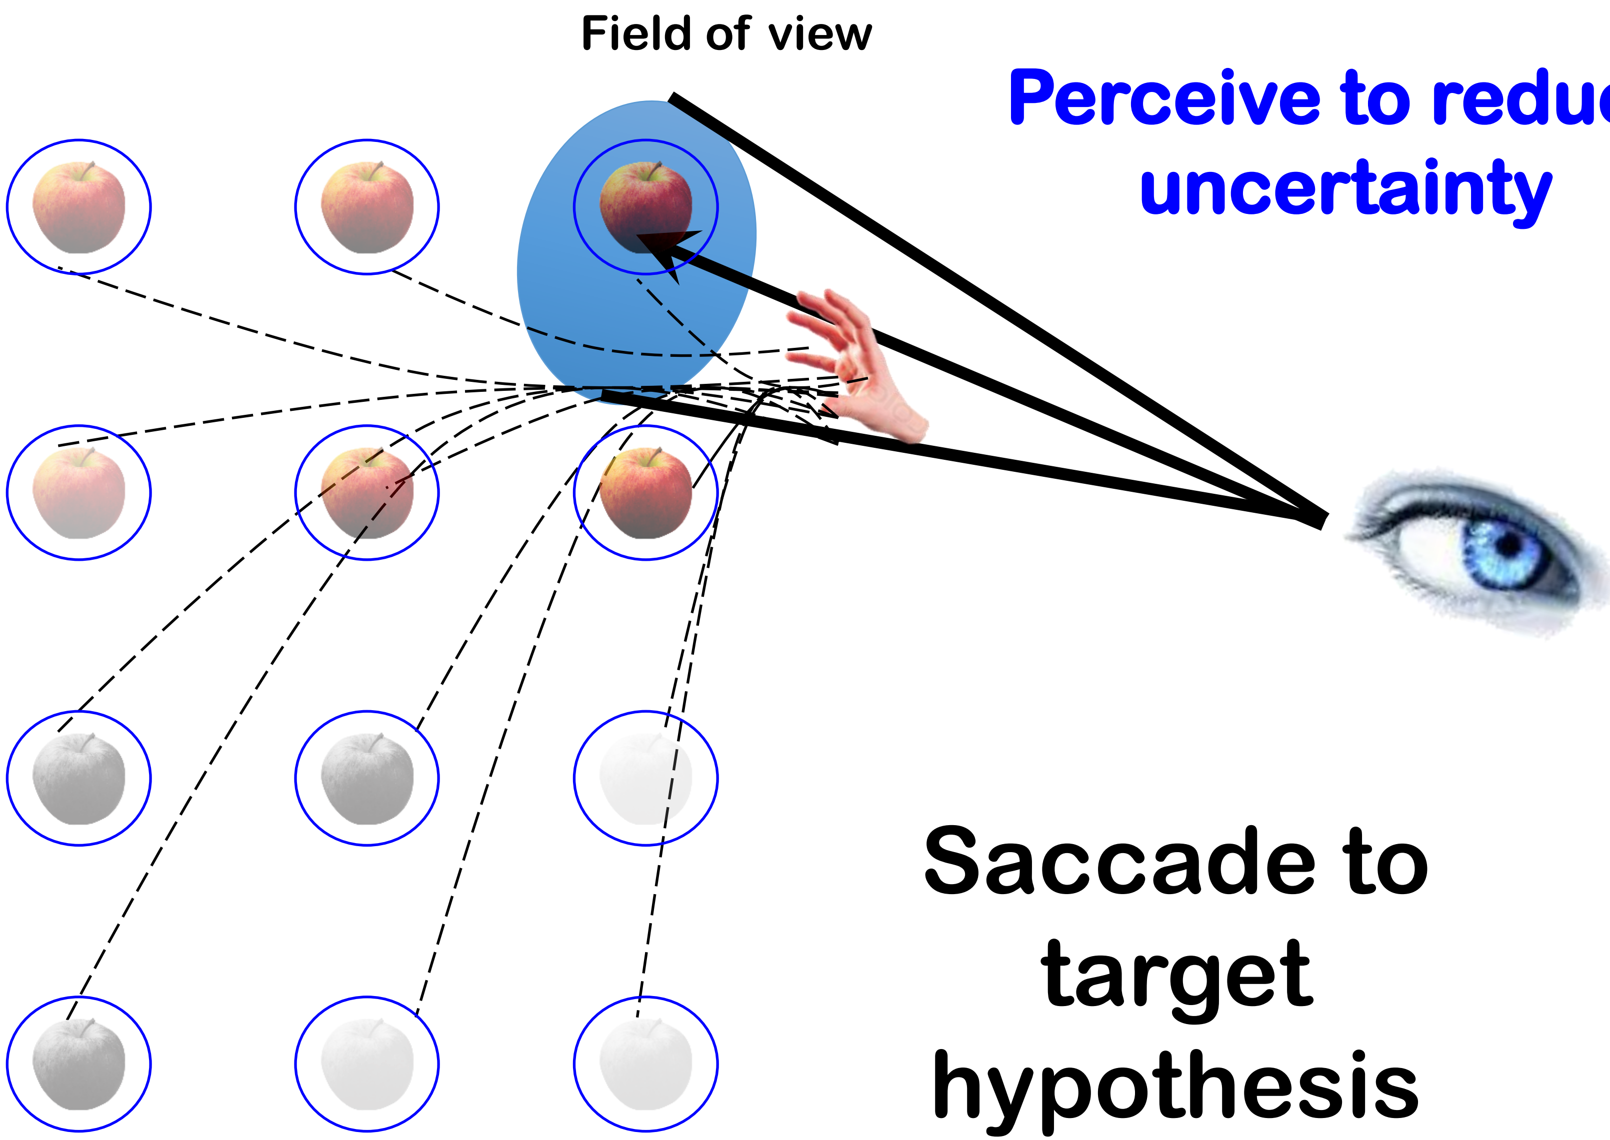

Perceive to reduce uncertainty

Saccade to target hypothesis

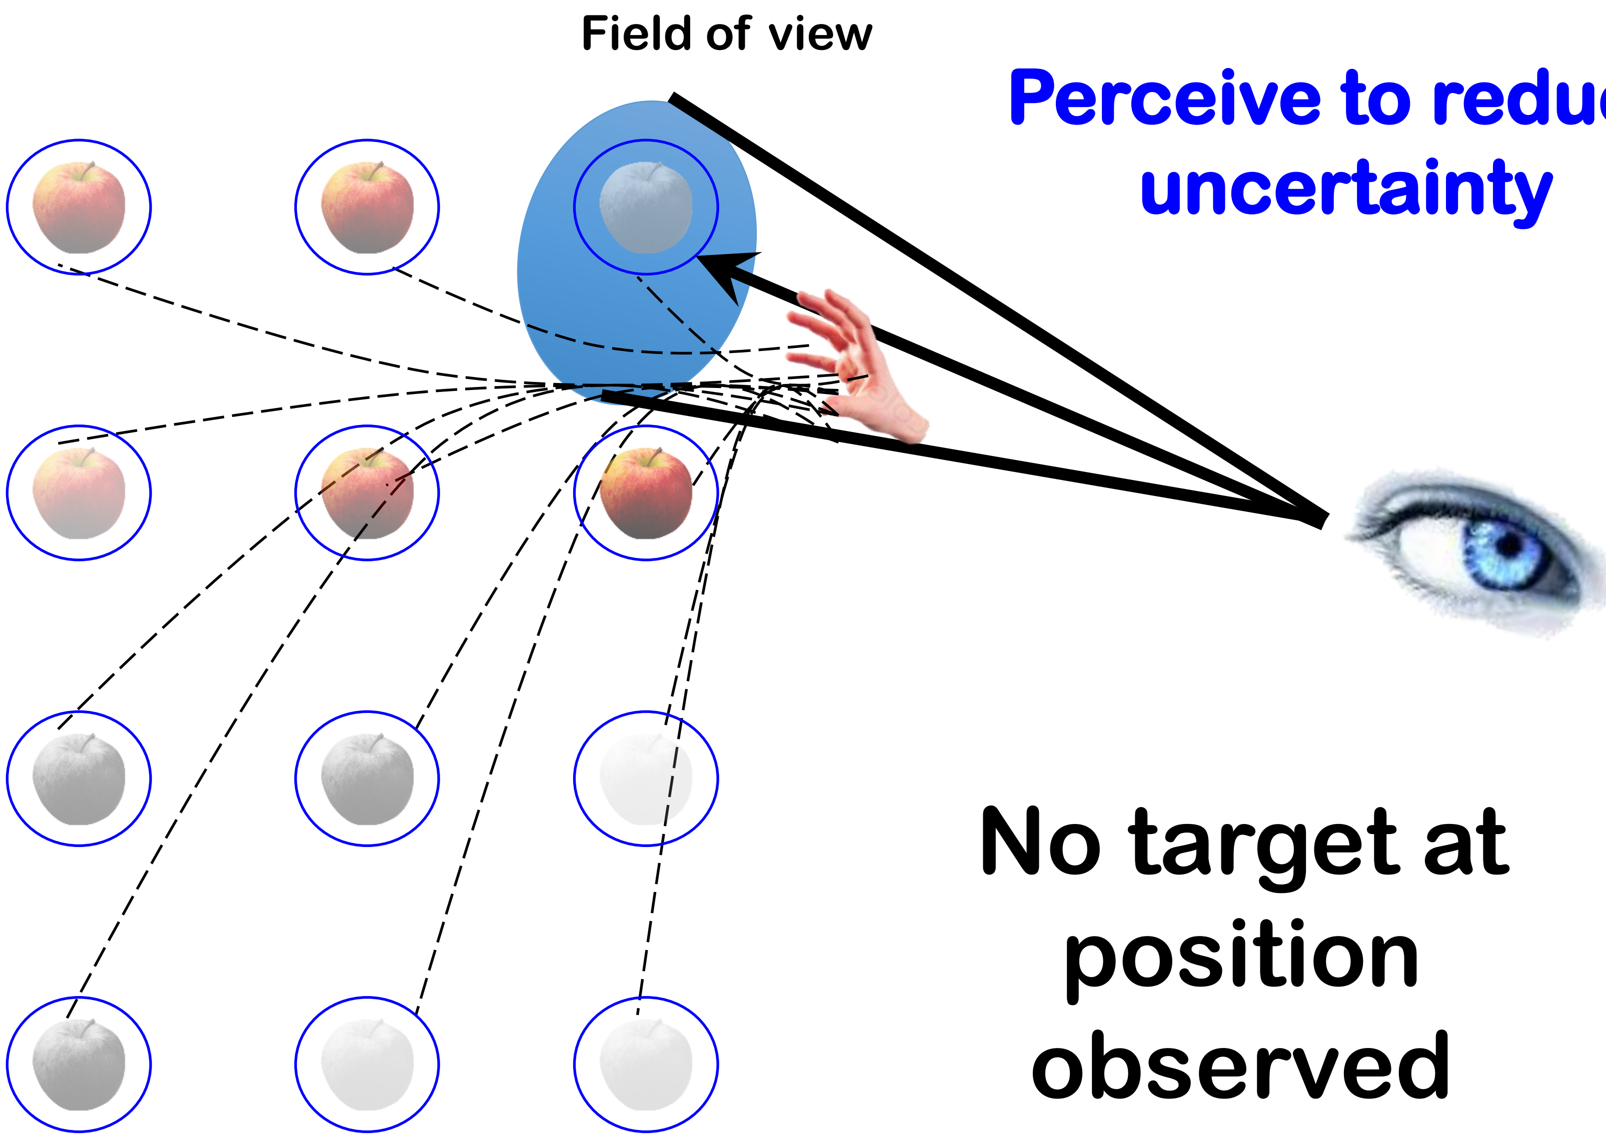
Field of view
Perceive to reduce uncertainty
No target at position observed

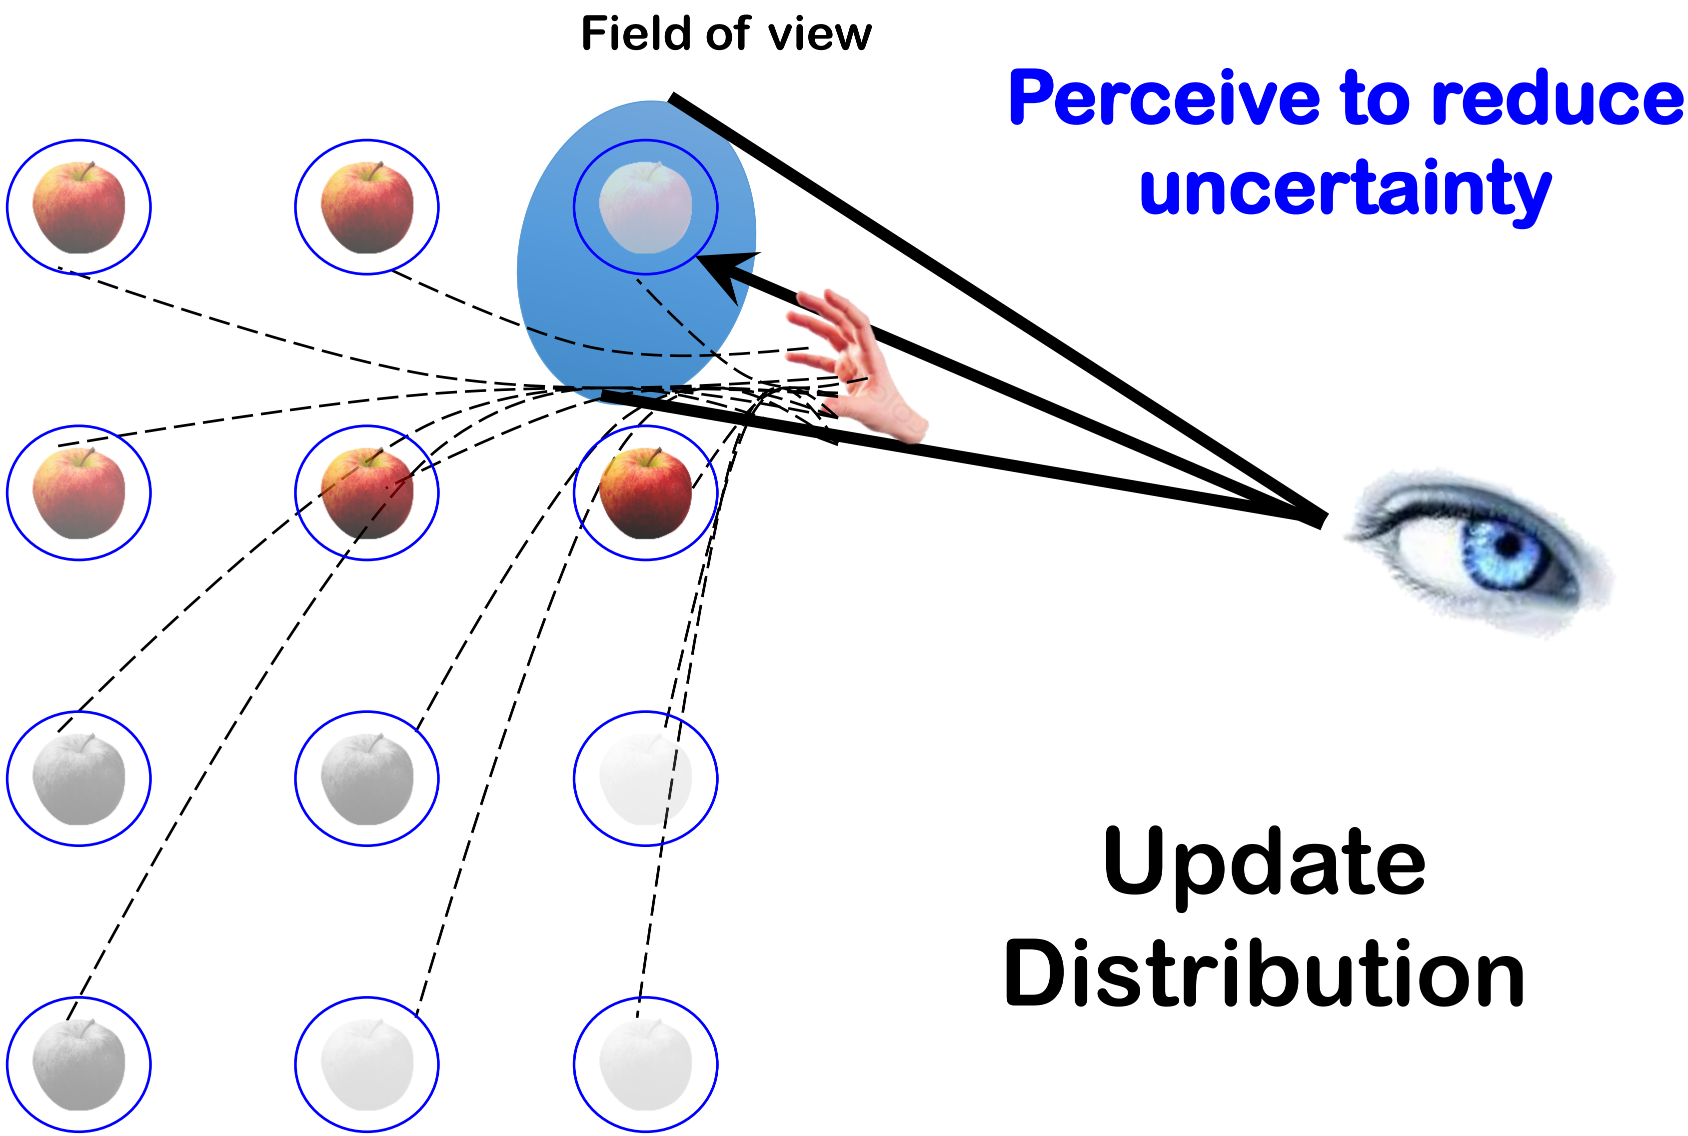
Field of view
Perceive to reduce uncertainty
Update Distribution

Info Gain Perception Control for Action Anticipation

Observe to Minimize Expected Uncertainty of Event (V)

$$\hat{\theta}^t = \underset{\theta^t}{\operatorname{argmin}} \int_{\mathbf{O}} p(\mathbf{o}^t|\mathbf{o}^{0\ldots t-1}, \theta^t) H(V|\mathbf{o}^{o\ldots t}, \theta^{0\ldots t}) \mathrm{d}\mathbf{o}^t$$

Goal Recognition as Inverse *Control* (+ Observer Model)

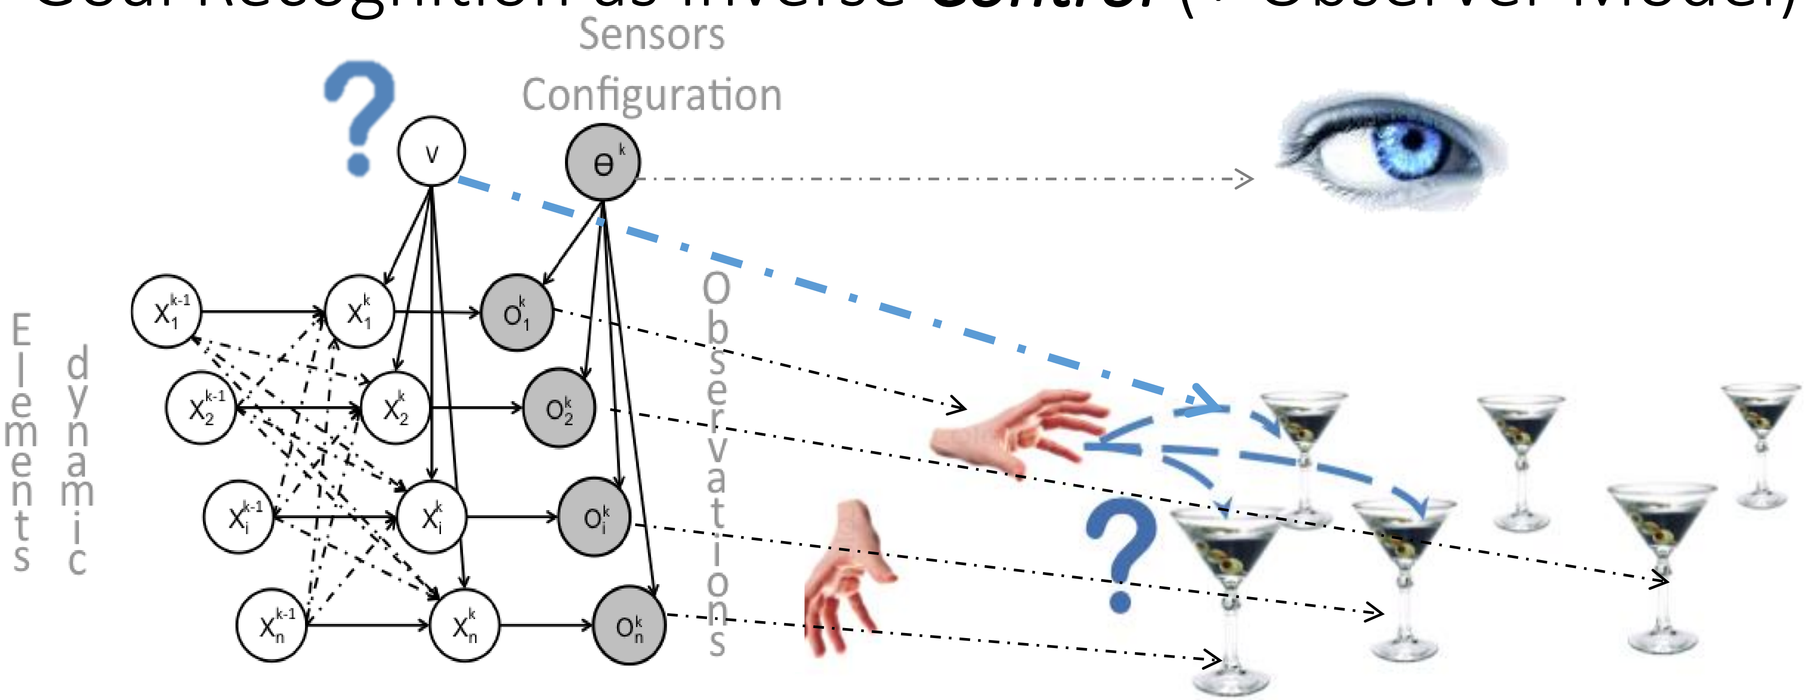

Ognibene & Demiris IJCAI 2013

Info Gain Using Kalman Filters

$$\hat{\theta}^{t+1} \approx \operatorname*{argmin}_{\theta} \sum_{v} P(v) \left[\frac{1}{2} \boxed{\ln \left|\mathbf{S}^{t+1}_{v,\theta^{t+1}}\right|} + \boxed{\ln \sum_{v'}^{V} \left(P(v') \mathcal{N}\left(\bar{\mathbf{o}}_{v,\theta}; \bar{\mathbf{o}}^{t+1}_{v',\theta^{t+1}}, \mathbf{S}^{t+1}_{v',\theta^{t+1}}\right) \right)} \right]$$

Elements Average Expected Entropy

Similarity between predictions for different events

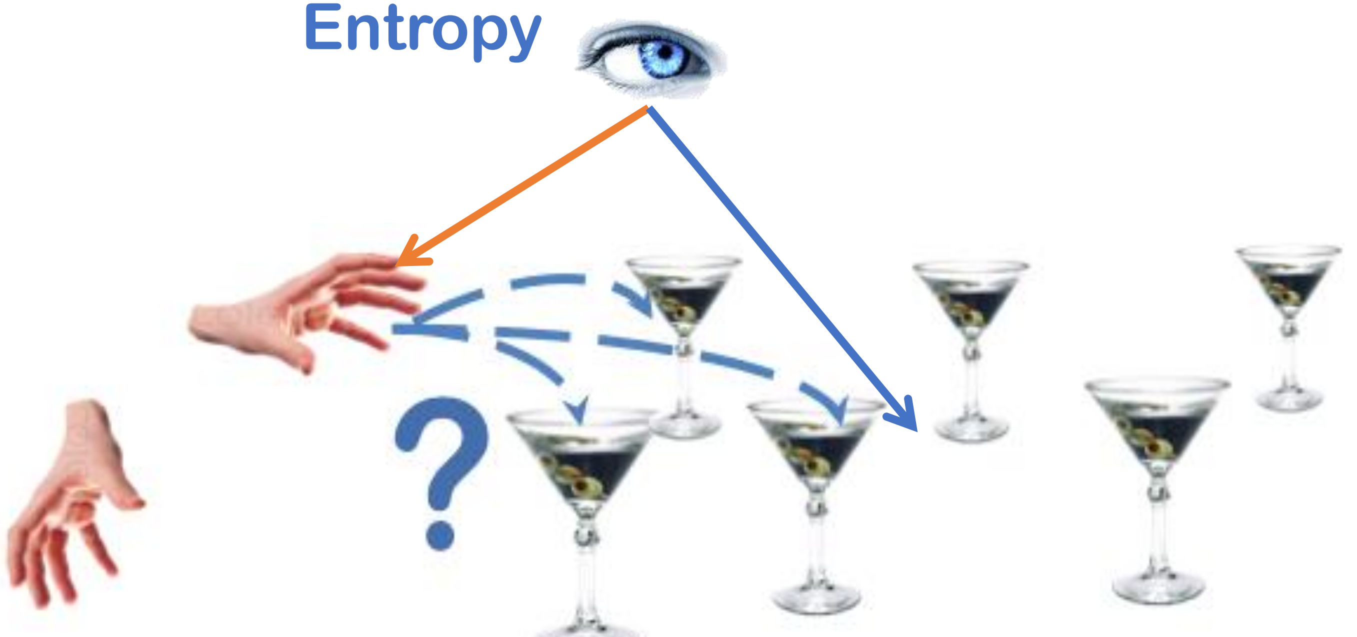

Ognibene & Demiris IJCAI 2013

Results

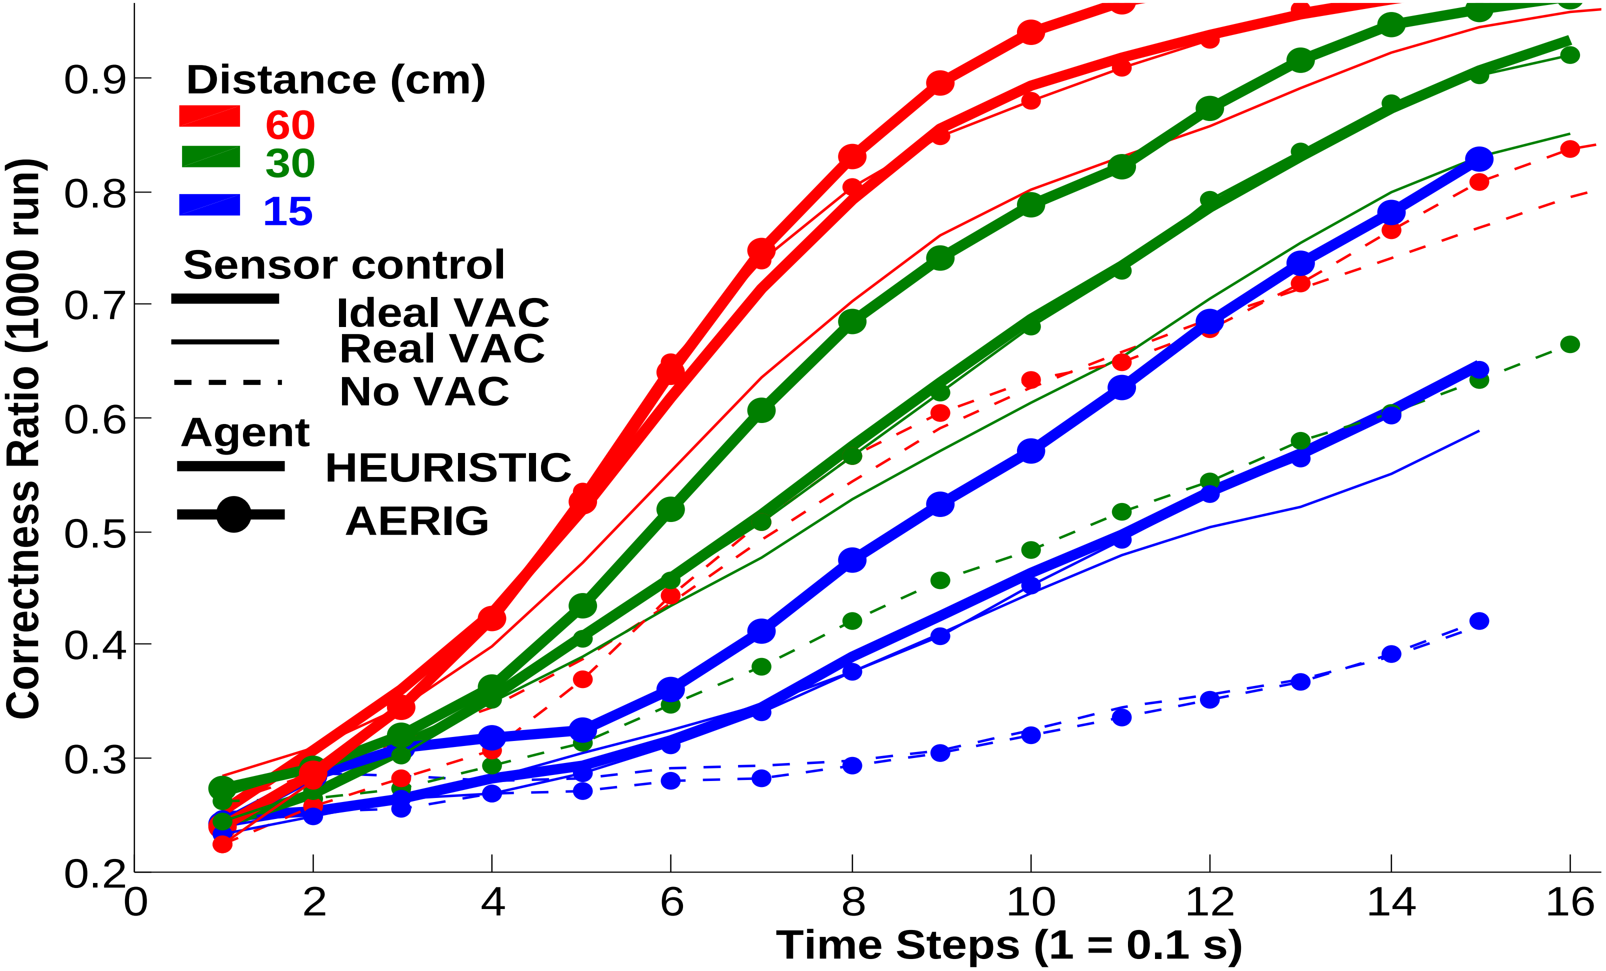

Distance (cm)
60
30
15
Sensor control
Ideal VAC
Real VAC
No VAC
Agent
HEURISTIC
AERIG
Correctness Ratio (1000 run)
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
0
2
4
6
8
10
12
14
16
Time Steps (1 = 0.1 s)

Results: Anticipate Performer Hand

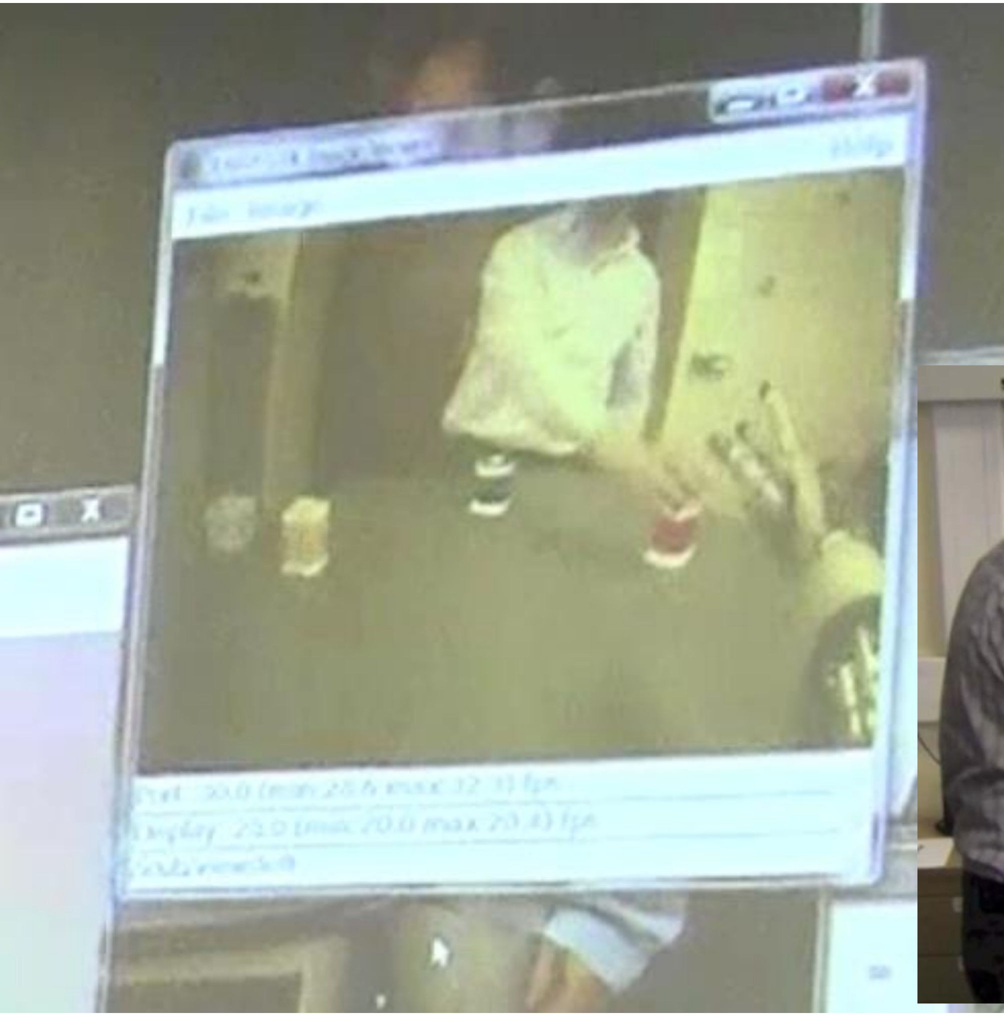

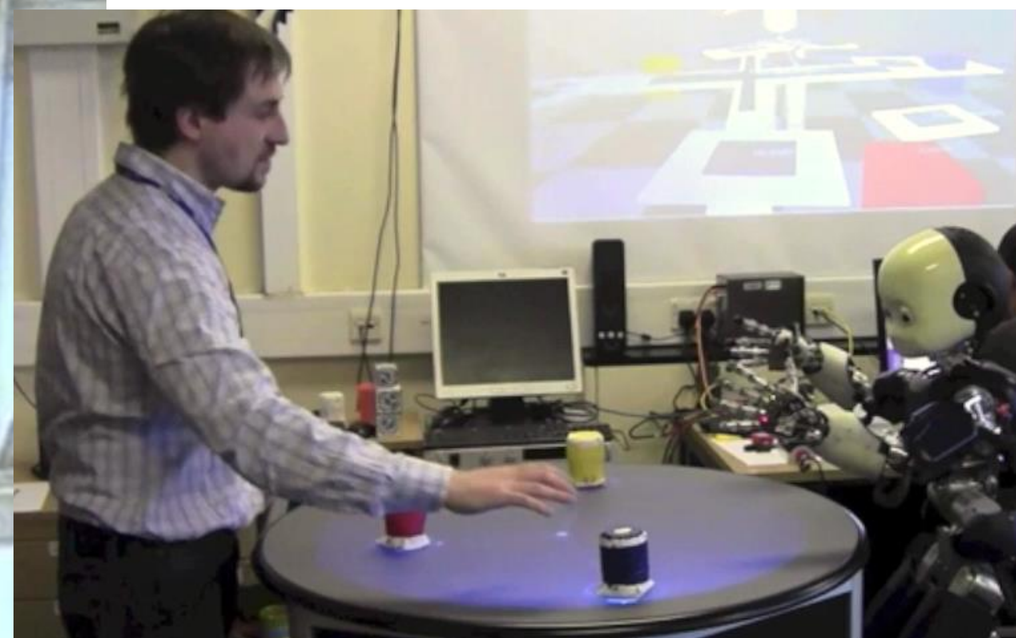

Ognibene & Demiris IJCAI 2013

STARE

Spatio-Temporal Attention Relocation for Multiple Structured Activities Detection

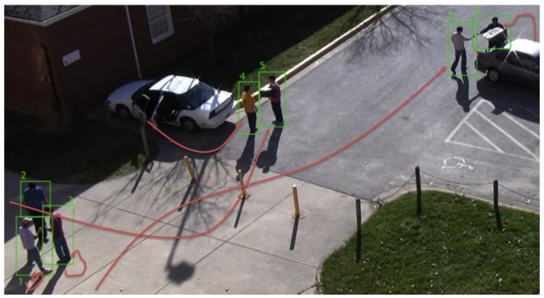

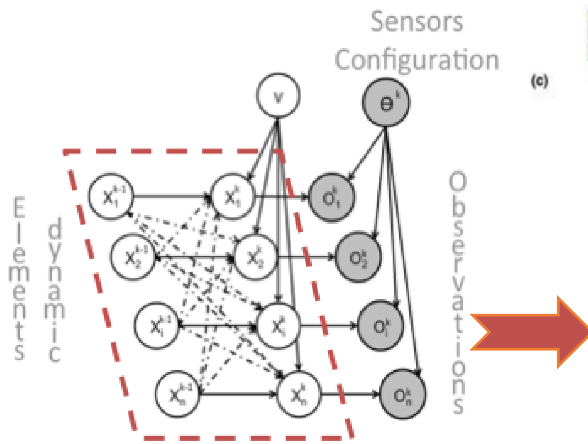

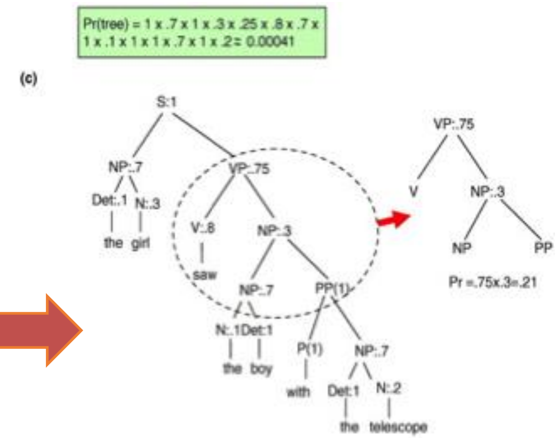

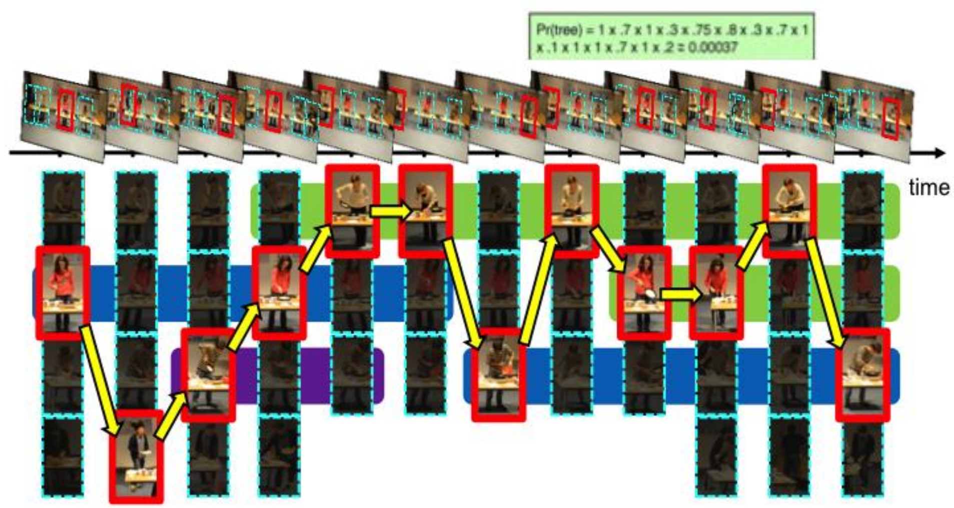

Lee, Ognibene, Chang, Kim, Demiris IEEE Trans Img Proc 2015

Multi-VIEW Image (MVI) in NM

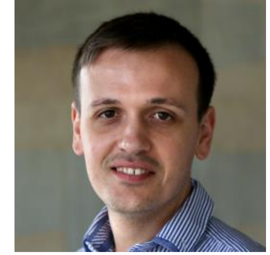

Carmelo Calafiore, Tom Foulsham
University of Essex, Colchester, UK

- 15 views: 5 thetas x 3 phis

An example of Multi-View Video for NN

Carmelo Calafiore, Tom Foulsham
University of Essex, Colchester, UK

View Percentage Frequency in SCM

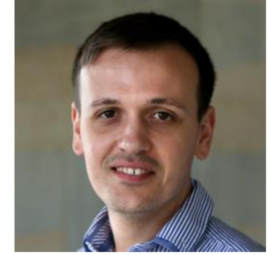


Carmelo Calafiore, Tom Foulsham
University of Essex, Colchester, UK

- Humans in the SCM selected more often the front and middle-height views than the chance.
- They selected less often the bottom, top-back views than the chance;
- Chance is 4.16% (100 % / 24 views);
- The view selection positively correlates with of the NM view accuracy, r=.60, p<.01;
- The view selection negatively correlates with of the NM view RT, r=.62, p<.01;
- People selected views efficiently.

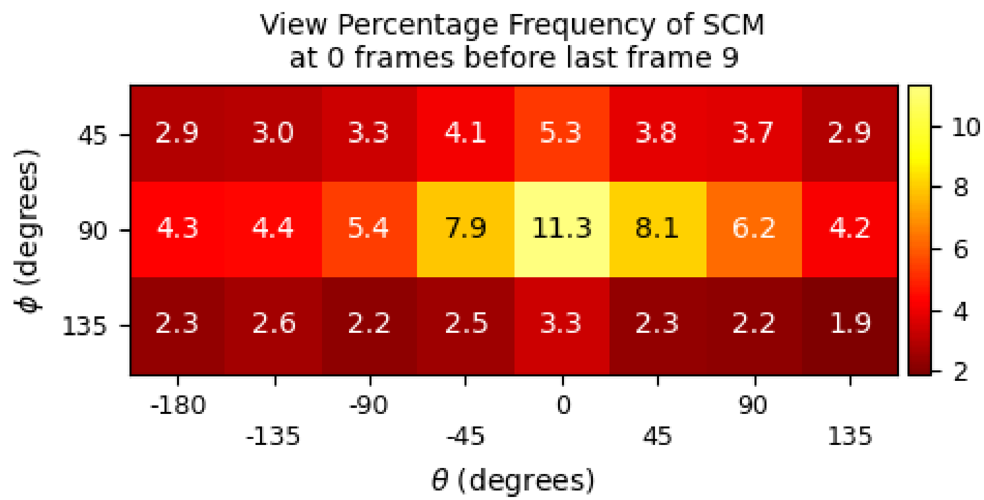

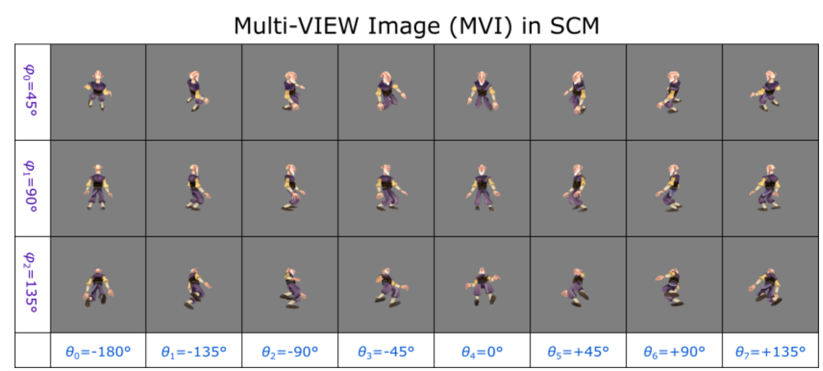

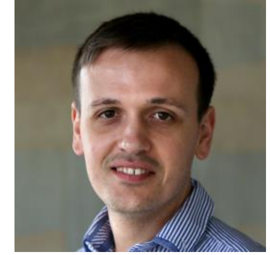

Carmelo Calafiore, Tom Foulsham
University of Essex, Colchester, UK

Multi-VIEW Image (MVI) in IC and ISC

- 40 views: 8 thetas x 5 phis

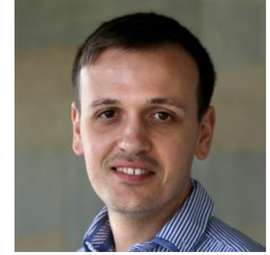 

Carmelo Calafiore, Tom Foulsham
University of Essex, Colchester, UK

Correlation of different DVs of the views

- The view selection does not correlate with the view accuracy of both computer models;
- The view accuracy of the models positively correlates with of the NM view accuracy;
- The view accuracy of the models negatively correlates with of the NM view RT;

Correlation Matrix of the Views

DVs	NM accuracy	NM RT	SCM frequency	SCM N Mov	Resnet Accuracy	LSTM Accuracy
NM accuracy	1.00	-.88	.60	-.53	.37	.38
NM RT	-.88	1.00	-.62	.63	-.59	-.57
SCM frequency	.60	-.62	1.00	-.75	.22	-.07
SCM N Mov	-.53	.63	-.75	1.00	-.30	-.17
Resnet Accuracy	.37	-.59	.22	-.30	1.00	.88
LSTM Accuracy	.38	-.57	-.07	-.17	.88	1.00

P values of the Correlation Matrix of the Views

DVs	NM accuracy	NM RT	SCM frequency	SCM N Mov	Resnet Accuracy	LSTM Accuracy
NM accuracy	.00	.00	.00	.01	.07	.07
NM RT	.00	.00	.00	.00	.00	.00
SCM frequency	.00	.00	.00	.00	.31	.75
SCM N Mov	.01	.00	.00	.00	.16	.44
Resnet Accuracy	.07	.00	.31	.16	.00	.00
LSTM Accuracy	.07	.00	.75	.44	.00	.00

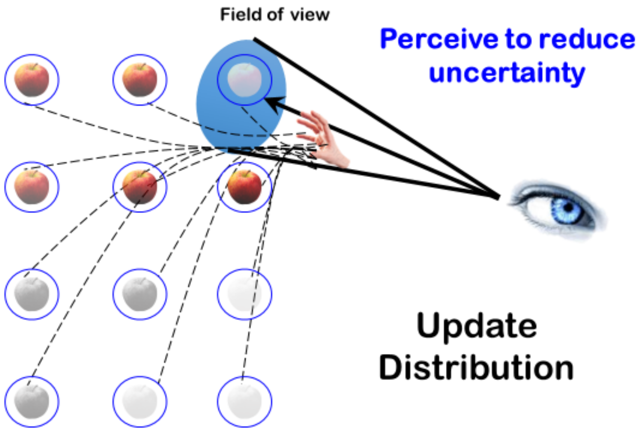

Goal Directed Behaviours

WHAT HAPPENS IF WHO WE OBSERVE IS LOOKING FOR SOMETHING TOO?

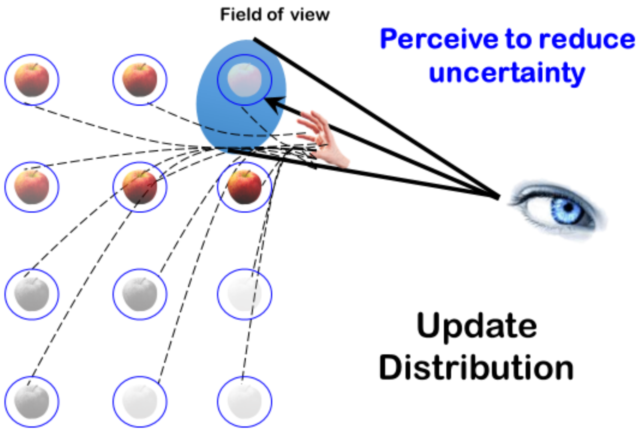

Goal Directed Behaviours

WHAT HAPPENS IF WHO WE OBSERVE IS LOOKING FOR SOMETHING TOO?

How do we learn to predict what partners are missing and **need to see** for any context and task?

Bianco & Ognibene HRI 2020
Bianco & Ognibene ICSR 2019
Bianco & Ognibene CEEC 2019

EMBODIED MODELS FOR EMERGENCE OF THEORY OF MIND

Francesca
Bianco

Theory of mind (ToM), or mentalizing, is the cognitive ability to attribute & represent others' mental states, i.e. intentions, beliefs and desires .

It has advantages for:

- Coordinating and managing false-beliefs
- Proactivity and preparation
- Active perception
- Learning

Ognibene & Demiris IJCAI 2013
Ognibene & Baldassarre AMD 2015
Bianco & Ognibene ICSR 2019
Ognibene et al. ICSR 2019
Bianco & Ognibene HRI 2020

EMBODIED MODELS FOR EMERGENCE OF THEORY OF MIND

Francesca Bianco

Rabinowitz et al NIPS 2018

Learns autonomously to predict "**belief-determined behaviors**" with no explicit information about *other's beliefs*, which are not accessible

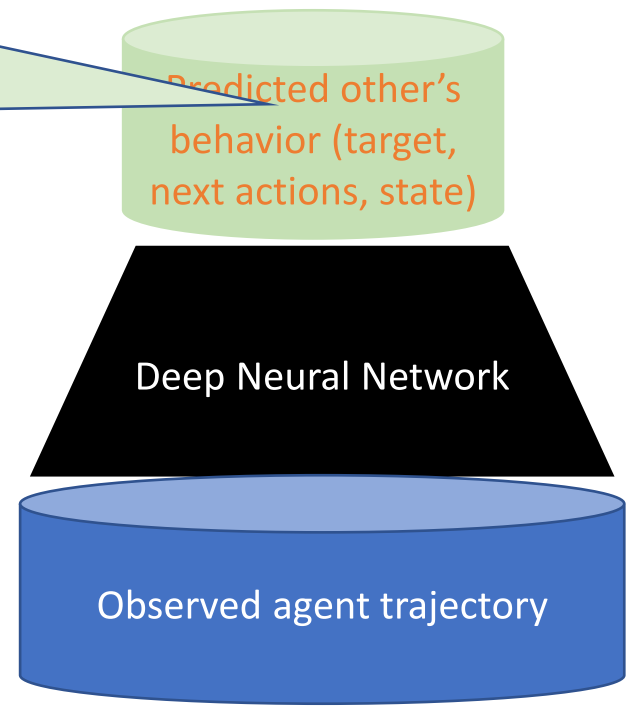

EMBODIED MODELS FOR EMERGENCE OF THEORY OF MIND

Francesca
Bianco

Learns autonomously to predict “**belief-determined behaviors**” with no explicit information about other’s beliefs, which are not accessible

Rabinowitz et al NIPS 2018

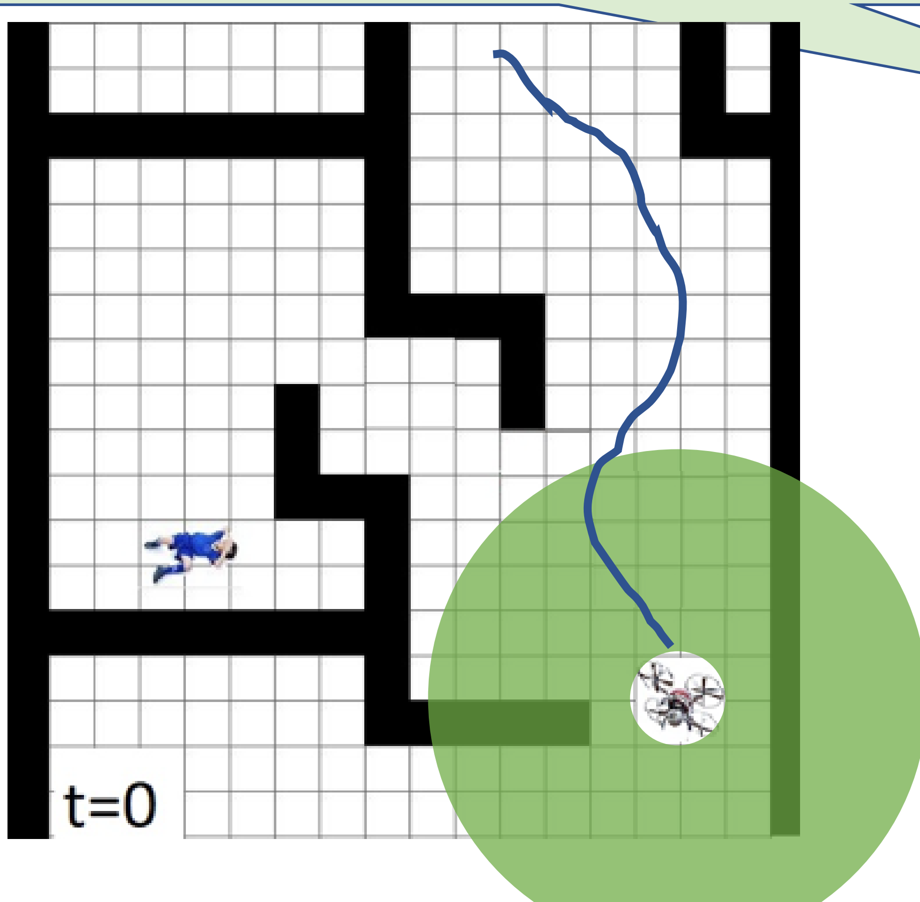

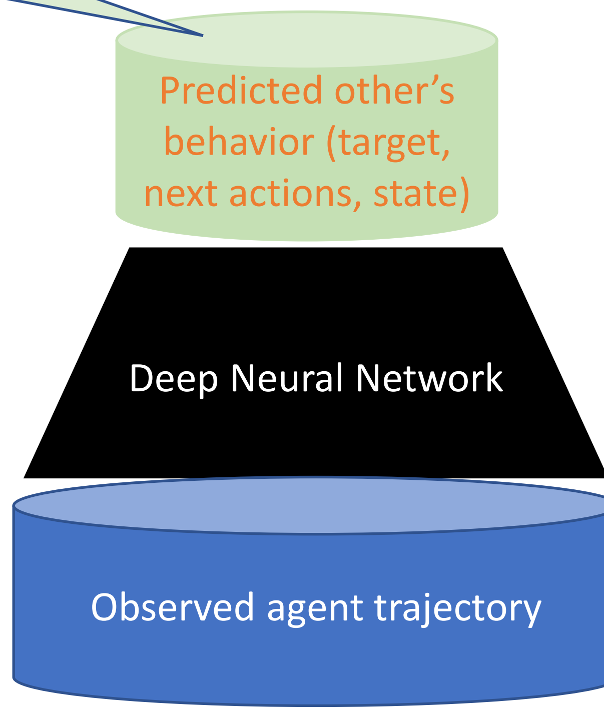

EMBODIED MODELS FOR EMERGENCE OF THEORY OF MIND

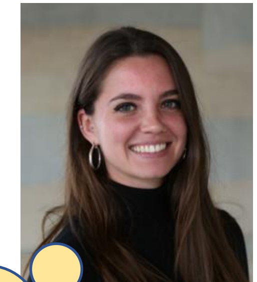

Francesca Bianco

Learns autonomously to predict “**belief-determined behaviors**” with no explicit informati... bout other’s beliefs, w...

...z et al NIPS 2018

- Can we learn a more explicit and «*shared*» representation?

...d other’s ...(target, ..., state)

...eural Network

Observed agent trajectory

t=0

EMBODIED MODELS FOR EMERGENCE OF THEORY OF MIND

Francesca Bianco

In publication

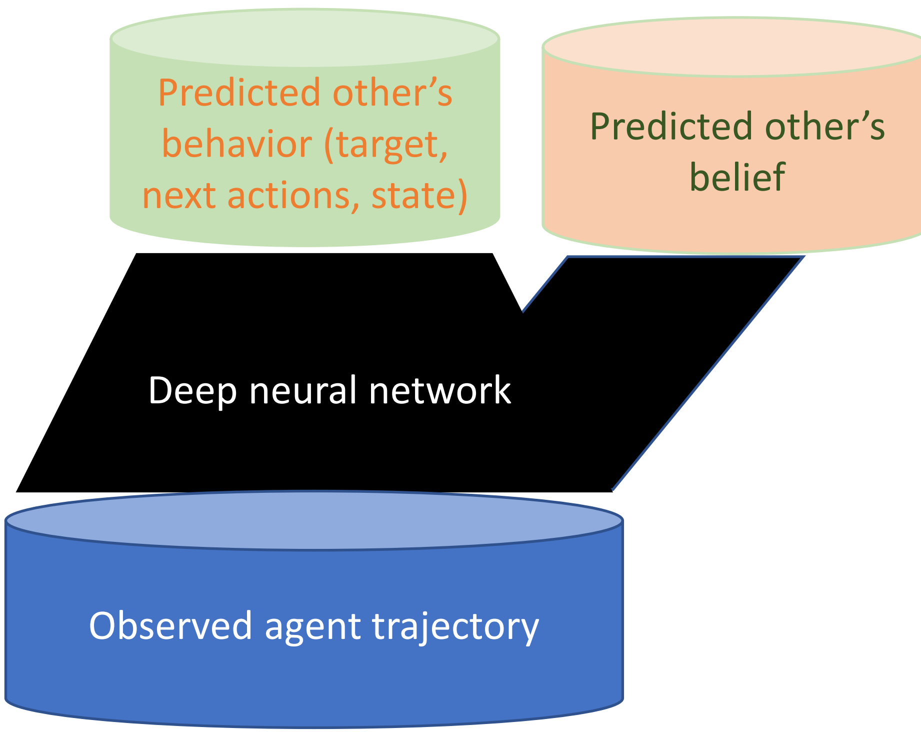

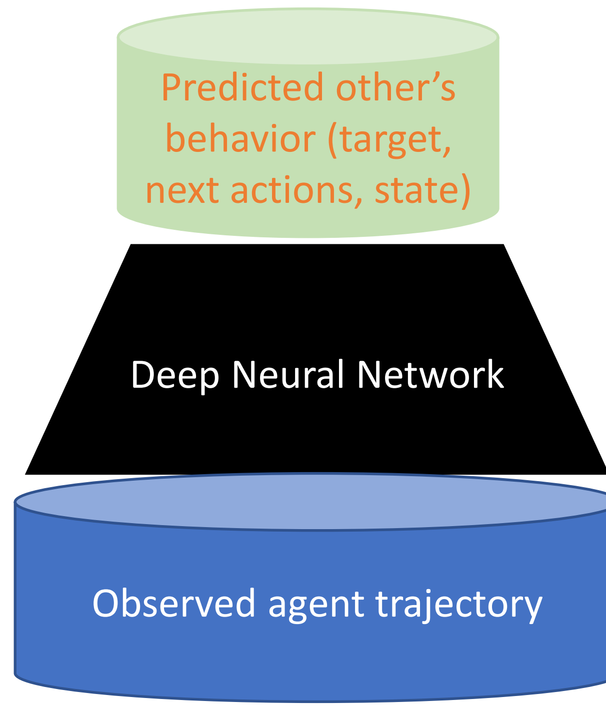

EMBODIED MODELS FOR EMERGENCE OF THEORY OF MIND

Francesca Bianco

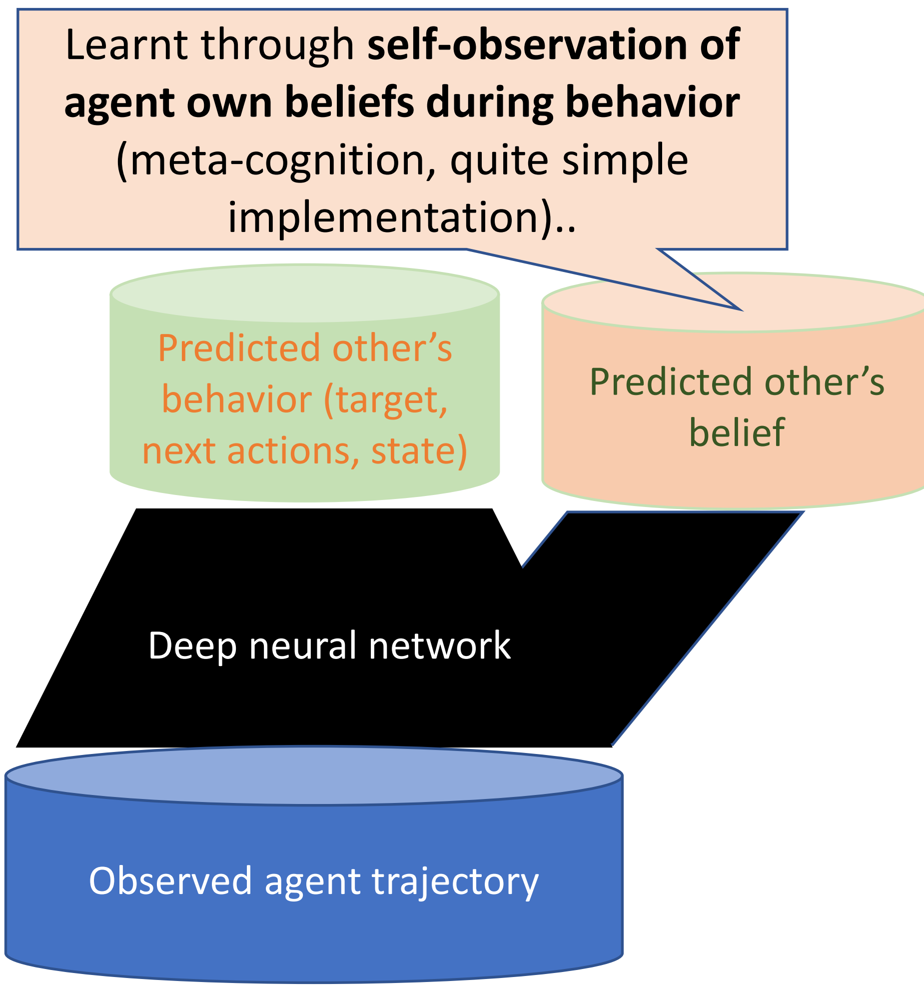

In publication

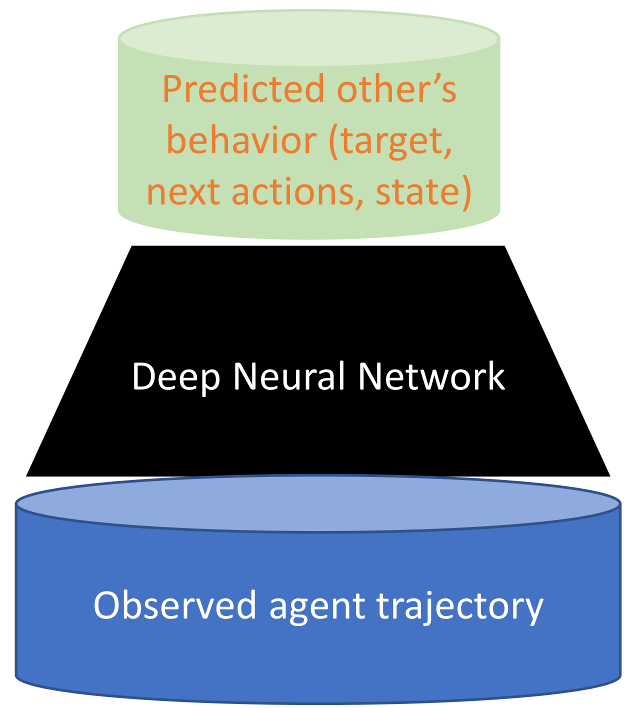

EMBODIED MODELS FOR EMERGENCE OF THEORY OF MIND

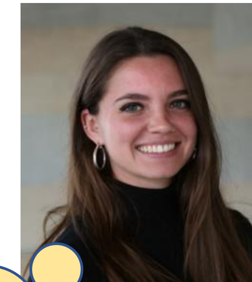

Francesca Bianco

Learnt through **self-observation of agent own beliefs**

- Can we learn a more explicit and «*shared*» representation?
- <u>*Is learning two things at the same time harder?*</u>

other's

target,

state)

Neural Network

Observed agent trajectory

Obser

EMBODIED MODELS FOR EMERGENCE OF THEORY OF MIND

Francesca Bianco

Relative performance gain of explicit belief architecture in generic behavior

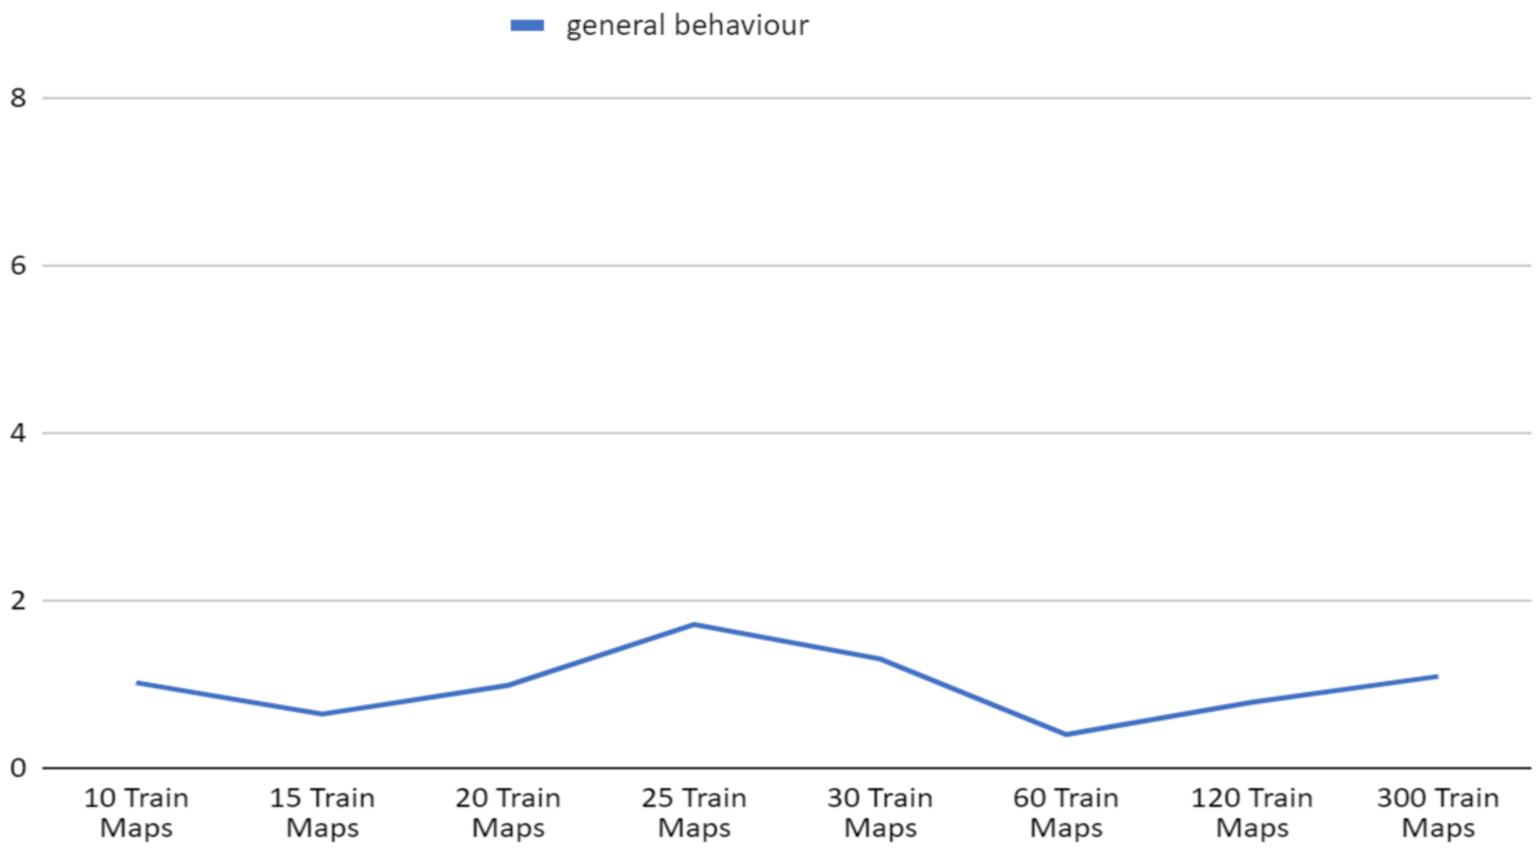

EMBODIED MODELS FOR EMERGENCE OF THEORY OF MIND

Francesca Bianco

Relative performance gain of explicit belief architecture in generic behavior and in the search segment

EMBODIED MODELS FOR EMERGENCE OF THEORY OF MIND

Francesca Bianco

Relative performance gain of explicit belief architecture in generic behavior and in the search segment

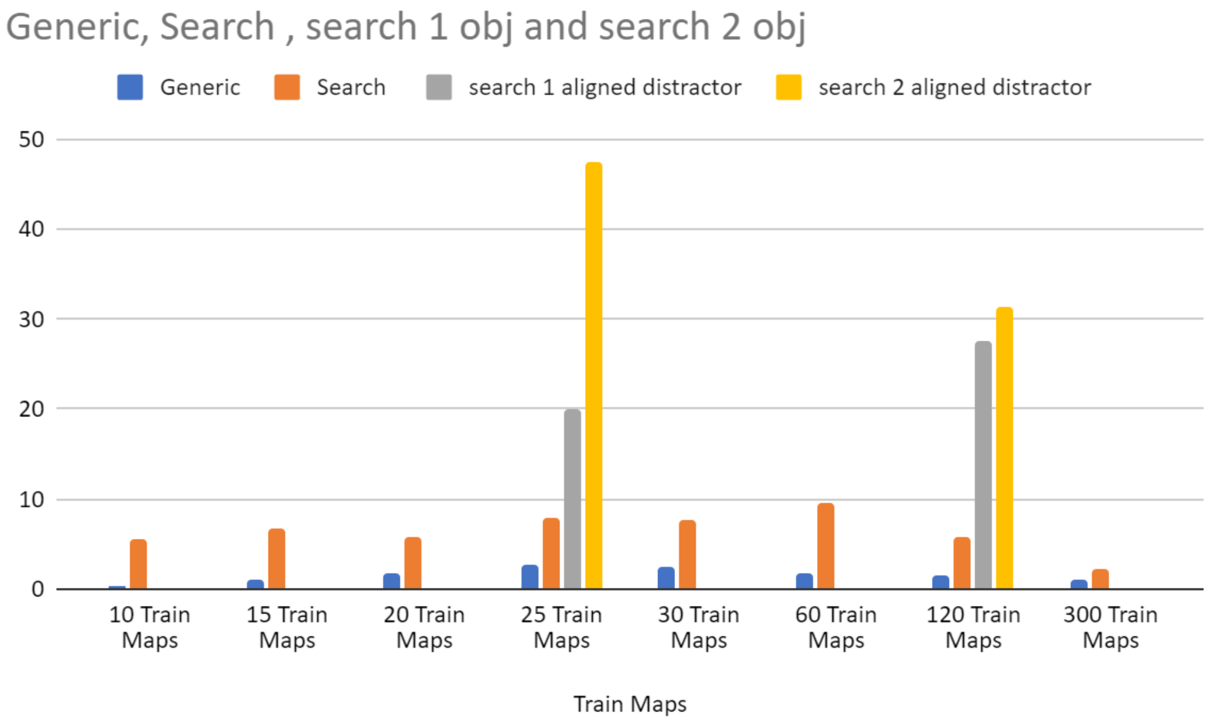

Special thanks to the following colleagues


Thrish Nanayakkara (Imperial)


Karl Friston (UCL)


Hector Geffner (UPF)


Yiannis Demiris (Imperial)

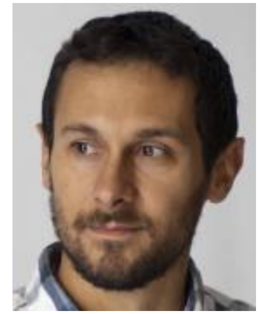
Giovanni Pezzulo (CNR)


Udo Krushwitz (Regensburg)


Tom Foulsham (Essex)


Letizia Marchegiani (Aalborg)


Salvador Soto (UPF)


Davinia Hernandez Leo (UPF)


David Rudrauf (Geneve)


Xiaosi Gu (Icahn, NYC)


Luca Citi (Essex)


Vincenzo G. Fiore (Icahn, NYC)


Giovanni Farinella (Uni CT)

Thanks for your attention!

Special thanks to the following colleagues

Thrish Nanayakkara (KCL)

Karl Friston (UCL)

Hector Geffner (UPF)

Yiannis Demiris (Imperial)

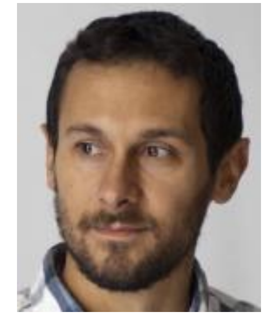

Giovanni Pezzulo (CNR)

Kris De Meyer (KCL)

Salvador Soto (UPF)

Daria Kvasova (UPF)

Vincenzo G. Fiore (UT Dallas)

Xiaosi Gu (UT Dallas)

Kyuhwa Lee (EPFL)

Giuseppe Giglia (Uni Pa)

Giovanni Farinella (Uni CT)

Thanks for your attention