Developmental Robotics

Fabricating open-source baby robots

Pierre-Yves Oudeyer
Project-Team INRIA-ENSTA-ParisTech FLOWERS

http://www.pyoudeyer.com
http://flowers.inria.fr

informatics mathematics

erc European Research Council

ENSTA ParisTech

Behavioural and Cognitive ***Development*** in Human Infants

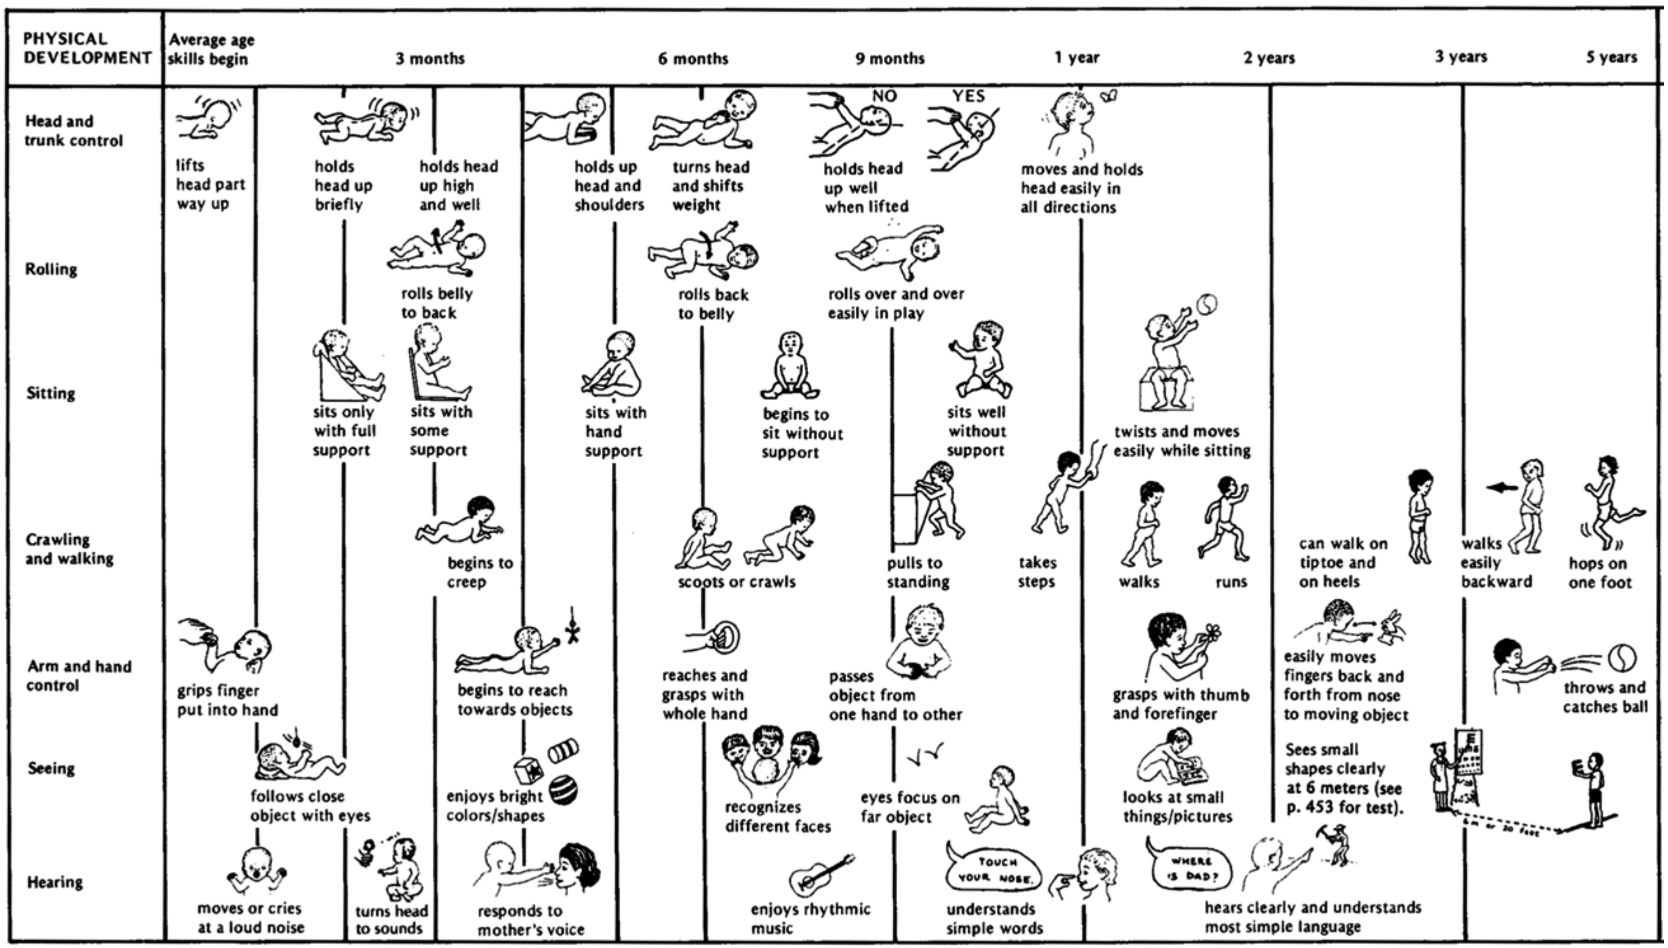

- How do developmental structures form?
- How do developmental structures impact the acquisition of novel skills?

Human
cognitive
development

Cognitive
development
in robots

Developmental robotics

Developmental robotics?

Families of developmental « forces »

Intrinsic motivation, active learning

- Autonomous collection of data
- Efficient learning
- Self-organization of developmental trajectories

Cognitive abstractions:

- Perceptual categories grounded in action
- Active goal babbling, macro-actions, macro-states
- Efficient learning in high-dimensions

Social learning, imitation

- Imitation of trajectories and goals
- Learning combinatorial motor primitives
- Optimal teaching

Body morphology and growth :

- Morphology
- Synergies
- Self-organization of movement structures
- Adaptive maturation driven by intrinsic motivation
- Self-organization of maturational schedule

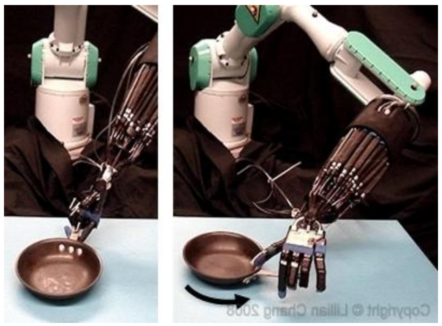

Development of sensorimotor skills

Relation (context, actions <-> effect) and their sequencing/composition

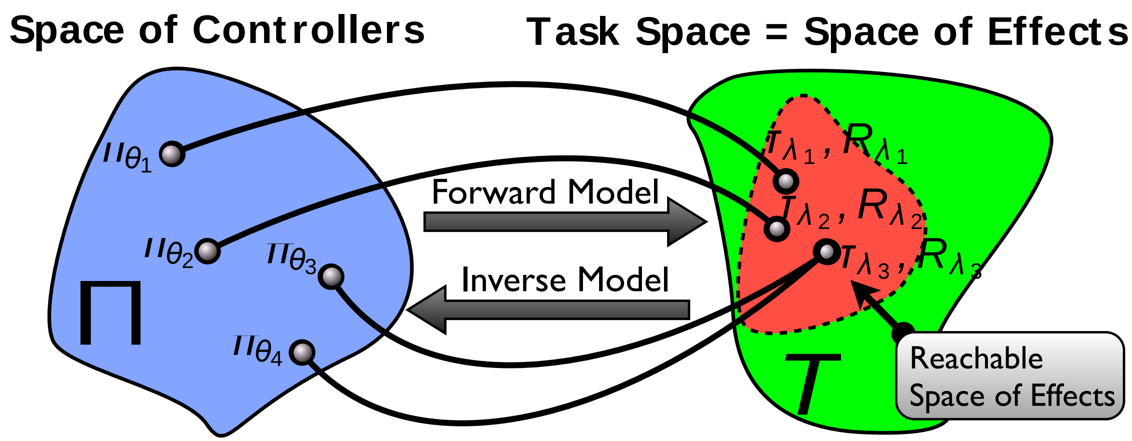

Parameterized by

$$\theta_i \in \mathbb{R}^n$$

Parameterized by

$$\lambda_j \in \mathbb{R}^m$$

- Forward models: Regression algorithms
- Inverse models: (Stochastic) Optimization algorithms
- Sequencing/composition: RL and structure discovery alg.

High-dimensions
Non-linear, redundant
Limited time resources

→ **How to achieve autonomous learning?**

Intrinsic motivation, curiosity and active learning

Hull (1943), White (1959): Basic forms of motivations (e.g. motivation for food and water, for sex, motivation for the maintainance of physical integrity, search for social bonding) can not account for the whole diversity of spontaneous exploratory behaviours of humans.

➔ Intrinsic drive to reduce uncertainty, and to experiencing novelty, surprise, cognitive dissonance, challenge, incongruences, …
➔ Optimal interest = optimal difficulty = neither trivial nor too difficult challenges
Berlyne (1960), White (1960) Csikszentmihalyi (1996)

Review

Cell PRESS

Information-seeking, curiosity, and attention: computational and neural mechanisms

Jacqueline Gottlieb[1,2], Pierre-Yves Oudeyer[3,4], Manuel Lopes[3,4], and Adrien Baranes[1]

[1] Department of Neuroscience, Columbia University, New York, NY, USA
[2] Kavli Institute for Brain Science, Columbia University, New York, NY, USA
[3] Inria, Bordeaux, France
[4] Ensta ParisTech, Paris, France

Intelligent animals devote much time and energy to exploring and obtaining information, but the underlying mechanisms are poorly understood. We review recent developments on this topic that have emerged from the traditionally separate fields of machine learning, eye movements in natural behavior, and studies of curiosity in psychology and neuroscience. These studies show that exploration may be guided by a family of mechanisms that range from automatic biases toward novelty or surprise to systematic searches for learning progress and information gain in curiosity-driven behavior. In addition, eye movements reflect visual information searching in multiple conditions and are amenable for cellular-level investigations. This suggests that the oculomotor system is an excellent model system for understanding information-sampling mechanisms.

Information-seeking in machine learning, psychology and neuroscience

For better or for worse, during our limited existence on earth, humans have altered the face of the world. We invented electricity, submarines, and airplanes, and developed farming and medicine to an extent that has massively changed our lives. There is little doubt that these extraordinary advances are made possible by our cognitive structure, particularly the ability to reason and build causal models of external events. In addition, we would argue that this extraordinary dynamism depends on our high degree curiosity, the burning desire to know and understand. Many animals, especially humans, seem to constantly seek knowledge and information in behaviors ranging from the very small (such as looking at a new storefront) to the very elaborate and sustained (such as reading a novel or carrying out research). Moreover, especially in humans, the search for information seems to be independent of a foreseeable profit, as if learning were reinforcing in and of itself.

Despite the importance of information-seeking for intelligent behavior, our understanding of its mechanisms is in its infancy. In psychology, research on curiosity surged during the 1960s and 1970s and subsequently waned [1] and has shown a moderate revival in neuroscience in recent years [2,3]. Our focus here is on evaluating three lines of investigation that are relevant to this question and have remained largely separate: studies of active learning and exploration in the machine learning and robotics fields, studies of eye movements in natural behavior, and studies of curiosity in psychology and neuroscience.

Glossary

Computational reinforcement learning: defines the problem of how to solve an MDP (or a POMDP) through learning (including trial and error), as well as associated computational methods.
Developmental robotics: research field modeling how embodied agents can acquire novel sensorimotor, cognitive, and social skills in an open-ended fashion over a developmental time span through integration of mechanisms that include maturation, intrinsically and extrinsically motivated learning, and self-organization.
Intrinsic and extrinsic rewards: normative accounts of behavior based on computational reinforcement learning and optimal control theory rely on the concept of a reward to assign value to alternative options, and often distinguish between extrinsic and intrinsic rewards. Extrinsic rewards are associated with classical task-directed learning and encode objectives such as finding food and winning a chess game. By contrast, intrinsic rewards are associated with internal cognitive variables such as esthetic pleasure, information-seeking, and epistemic disclosure. Intrinsic rewards may be based on uncertainty, surprise, and learning progress, and they may be either learnt or innate.
Markov decision process (MDP): defines the problem of selecting the optimal actions at each state to maximize future expected rewards in a Markov process.
Markov process (MP): mathematical model of the evolution of a system in which the prediction of a future state depends only on the current state and on the applied action, and not on the path by which the system reached the current state.
Metacognition: capability of a cognitive system to monitor its own abilities – for example, its knowledge, competence, memory, learning, or thoughts – and act according to the results of this monitoring. An example is a system capable of estimating how much confidence or uncertainty it has or how much learning progress it has achieved, and then using these estimates to select actions.
Optimization: mechanism that is often used in machine learning to search for

Corresponding author: Gottlieb, J. (jg2141@columbia.edu).

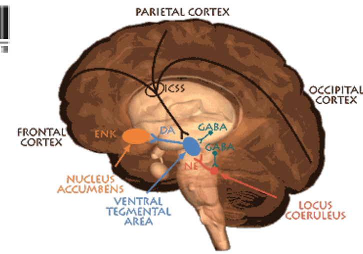

Neurosciences

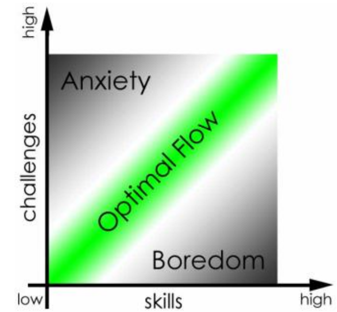

Psychology

Developmental and cognitive robotics

Trends in Cognitive Science, Nov. 2013

Active learning, intrinsic motivation

Intrinsic Motivation

Berlyne (1960), Csikszentmihalyi (1996)
Dayan and Belleine (2002)

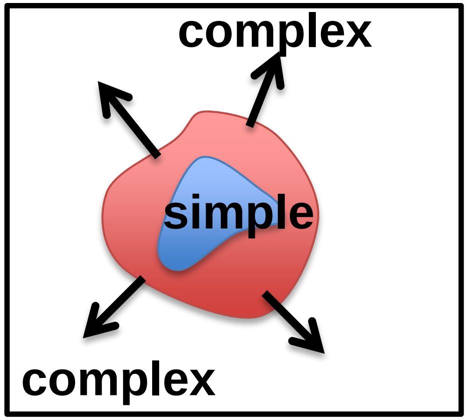

$(S(t), p_q)$

(Schmidhuber, 1991)
(Barto, Singh and Chentanez, 2004)
(Oudeyer, Kaplan, Hafner, 2007)
(Baldassarre, 2011)
…

$$\underset{(SVR, GPR, NN, \dots)}{predict} : (S(t), p_q) \rightarrow \tilde{S}(t + D)$$

$$e(S(t), p_q) = \left| \tilde{S}(t + D) - S(t + D) \right|$$

$$R(S(t), p_q) = -\frac{de}{dt} \text{ in the vicinity of } (S(t), p_q)$$

➔ Non-stationary function, difficult to model

➔ Algorithms for empirical evaluation of de/dt with statistical regression

➔ IAC (2004, 2007), R-IAC (2009), SAGG-RIAC (2010) McSAGG-RIAC (2011), SGIM (2011), Smoothed Beta distribution (2011), SGIM-ACTS (2012)

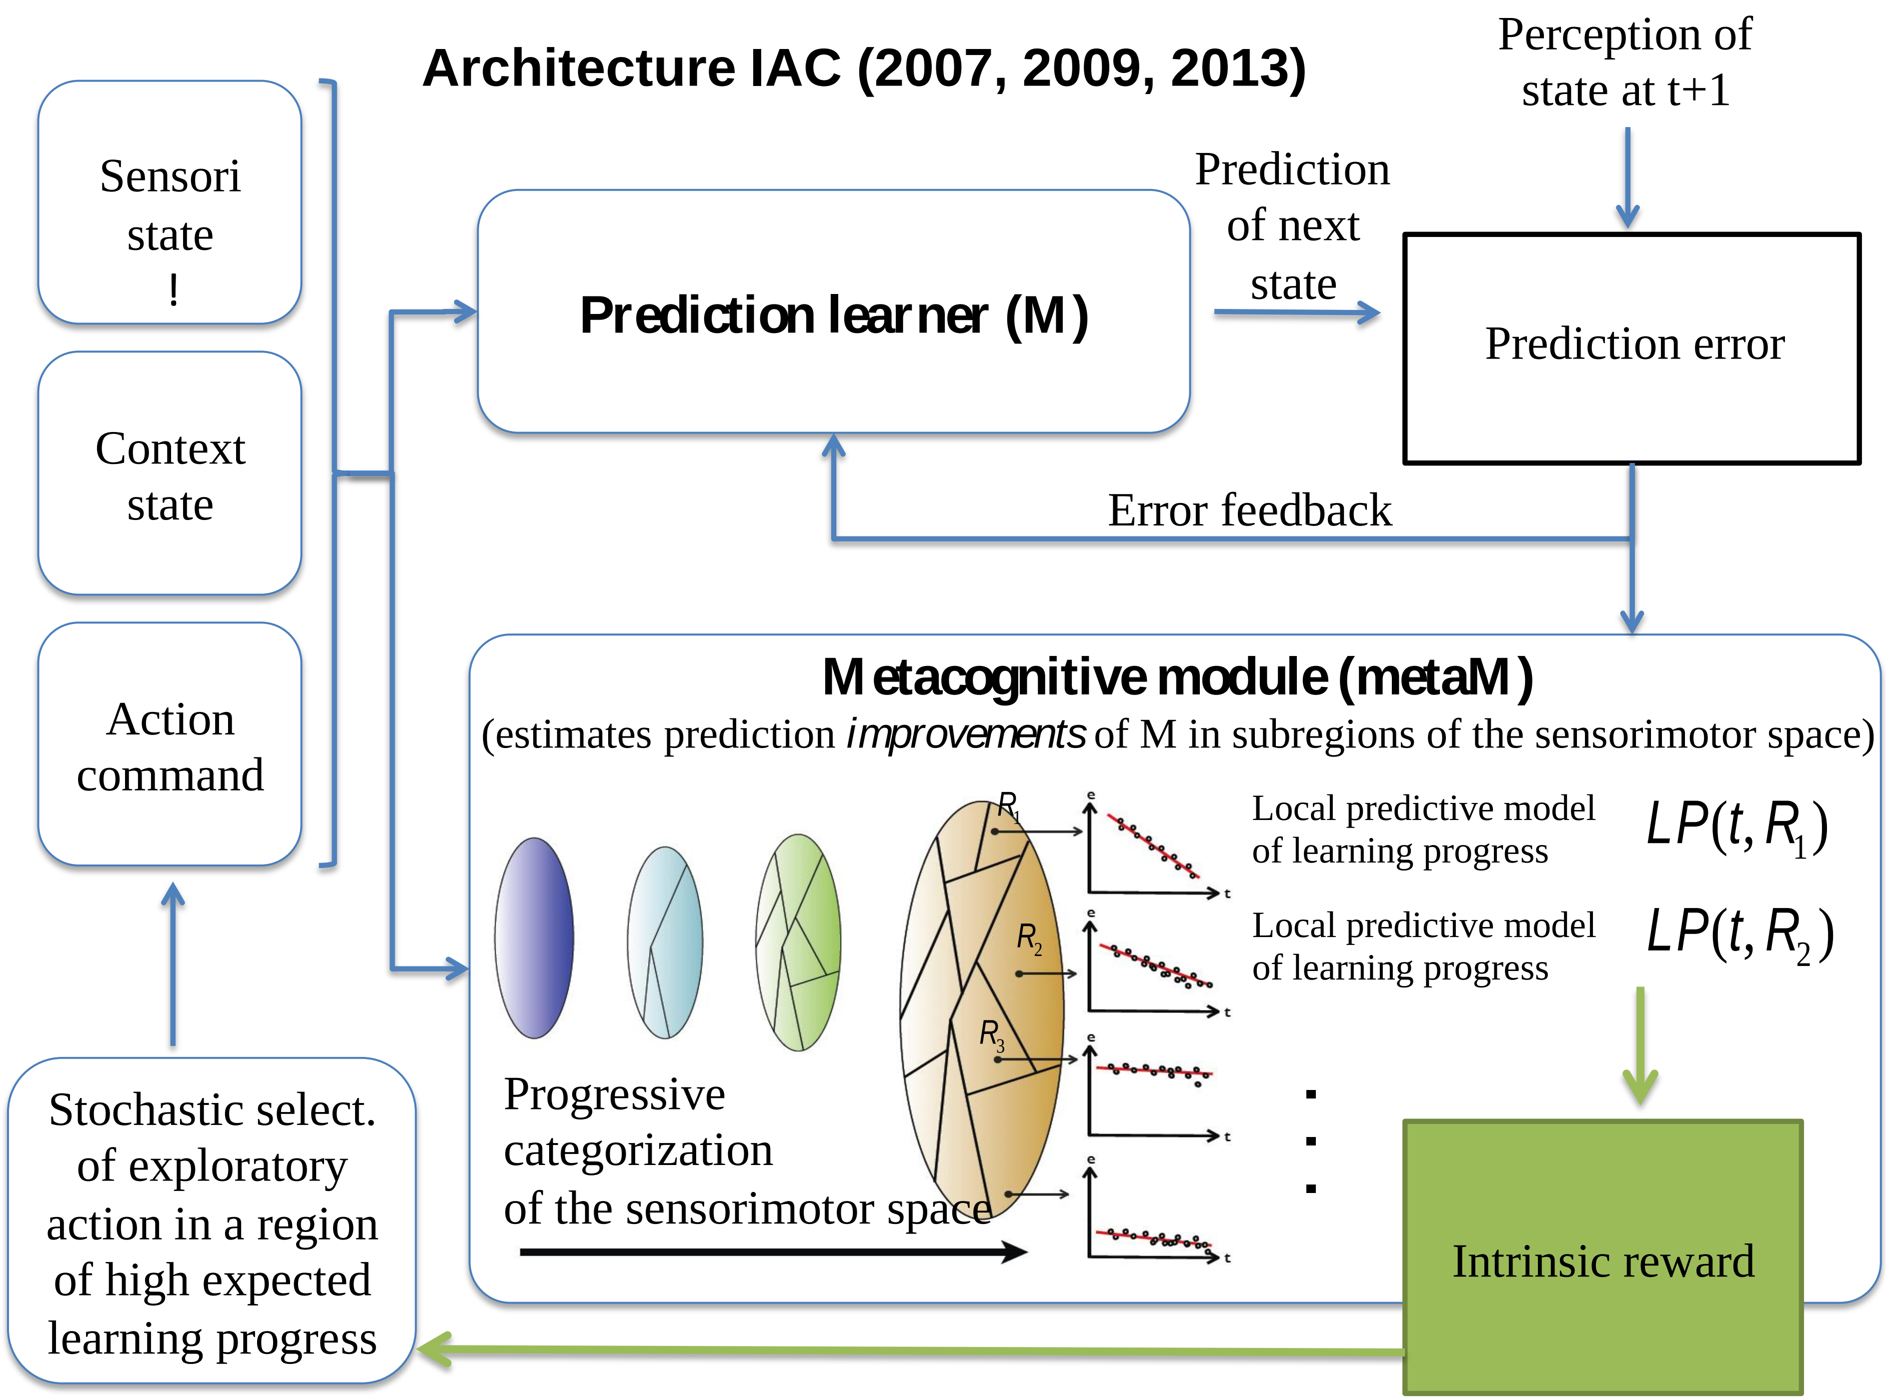
Architecture IAC (2007, 2009, 2013)
Sensori state !
Context state
Action command
Prediction learner (M)
Prediction of next state
Perception of state at t+1
Prediction error
Error feedback
Metacognitive module (metaM)
(estimates prediction improvements of M in subregions of the sensorimotor space)
R1
R2
R3
Local predictive model of learning progress
LP(t, R1)
Local predictive model of learning progress
LP(t, R2)
Progressive categorization of the sensorimotor space
Intrinsic reward
Stochastic select. of exploratory action in a region of high expected learning progress

The Playground Experiments

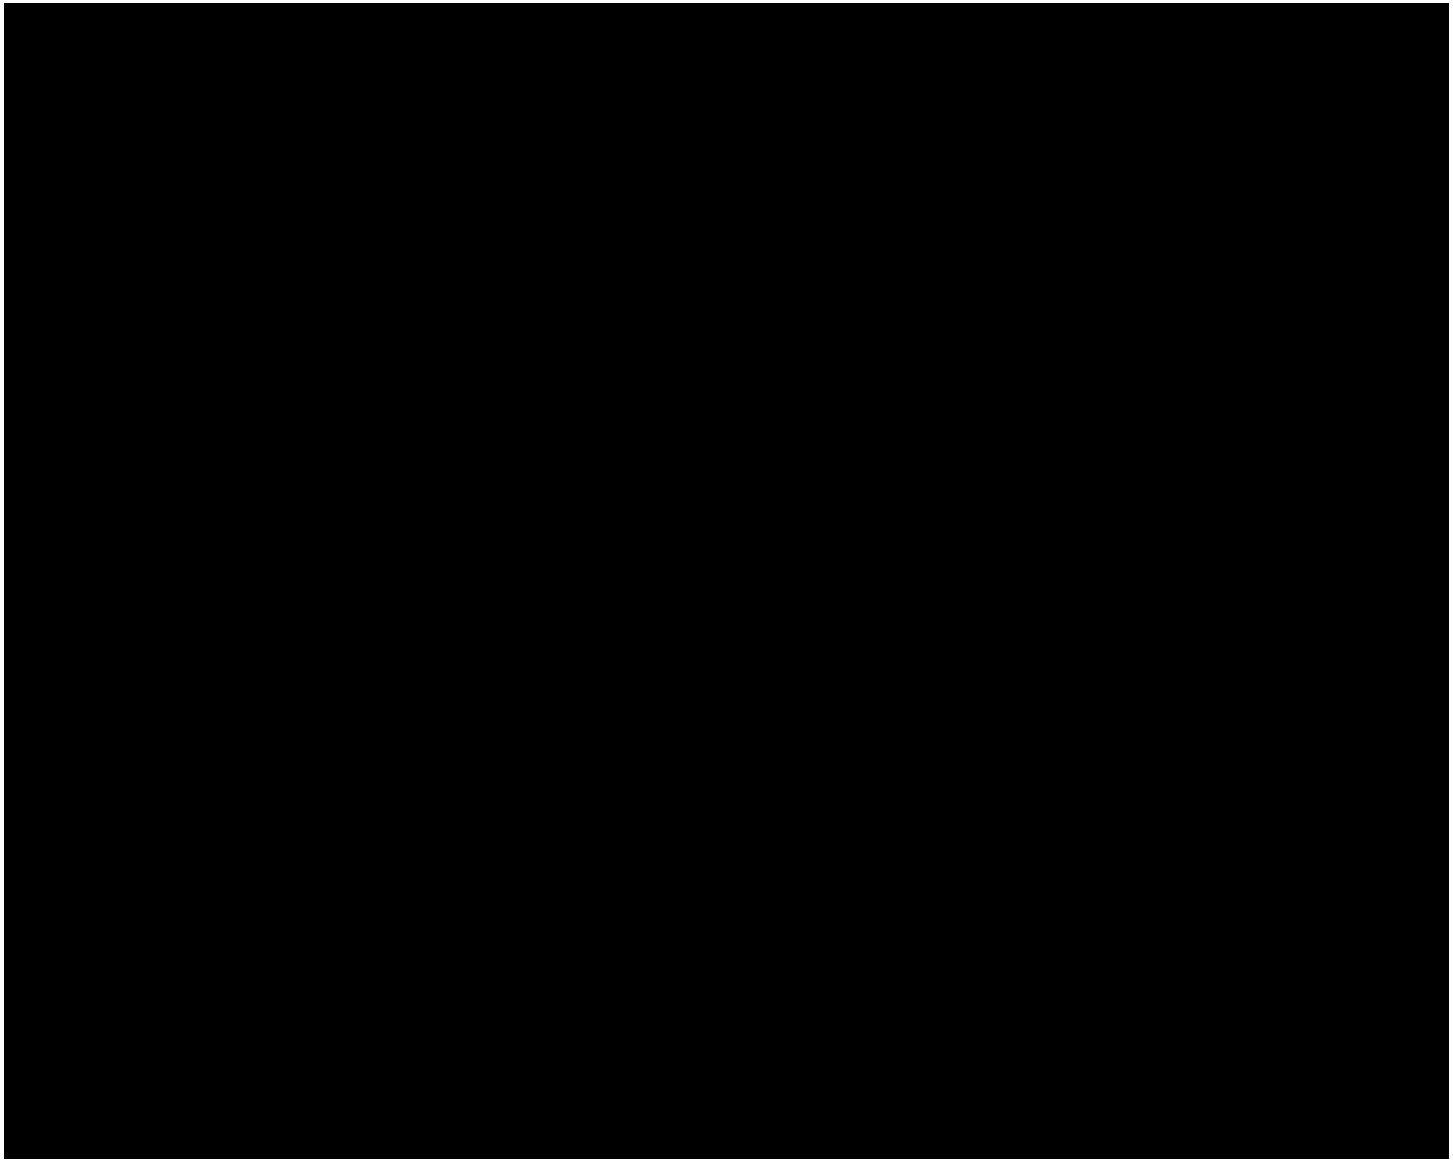

(Oudeyer et al., IEEE Trans. EC 2007; Connection Science 2006)

Curiosity-driven active **Goal Babbling**

(Oudeyer and Kaplan, 2007; Baranes and Oudeyer, 2009, 2010, 2013;
see also Rolf and Steil, 2009, 2010, 2013)

Redundancy of sensorimotor spaces

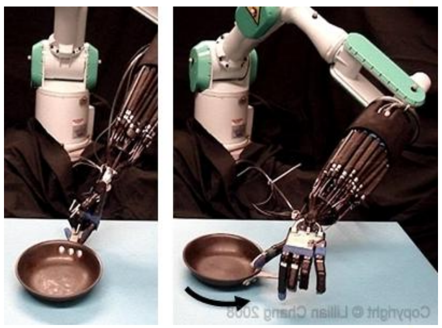

(Context, Movement)
→
Effect

Space of Controllers **Task Space = Space of Effects**

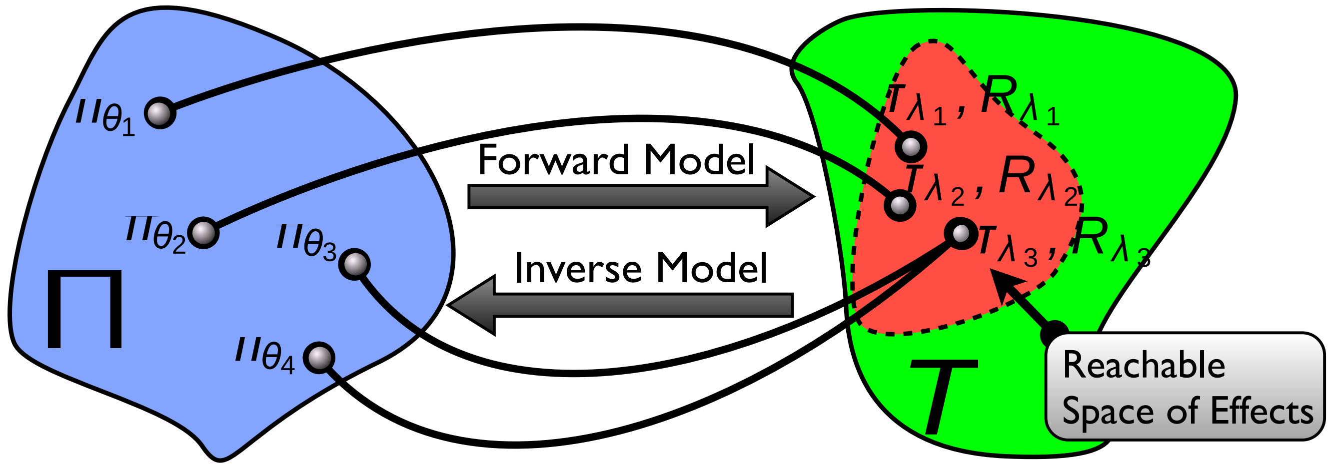

Curiosity-driven goal babbling

Active learning of omnidirectional locomotion

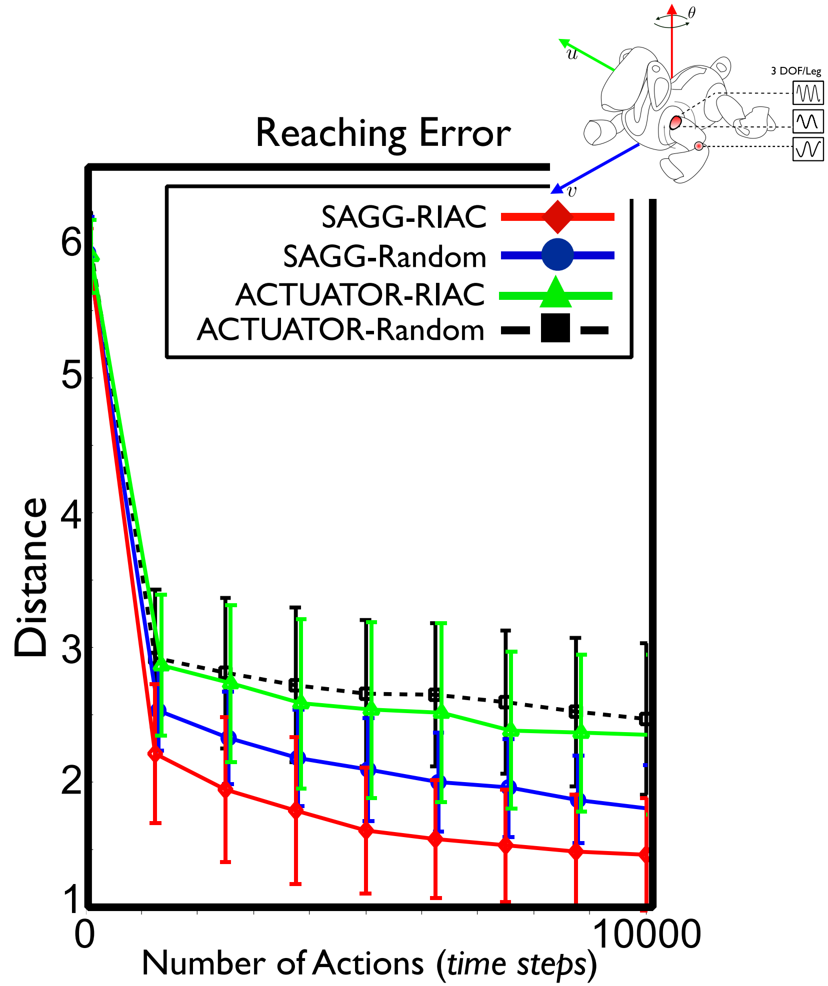

Control Space: $[-1;1]^{24}$ Task Space: $[-1;1]^{3}$

➔ Performance higher than more classical active learning algorithms in real sensorimotor spaces (non-stationary, non homogeneous)
(Baranes and Oudeyer, IEEE TAMD 2009; Robotics and Autonomous Systems; 2013)

Curiosity-driven learning of visual affordances

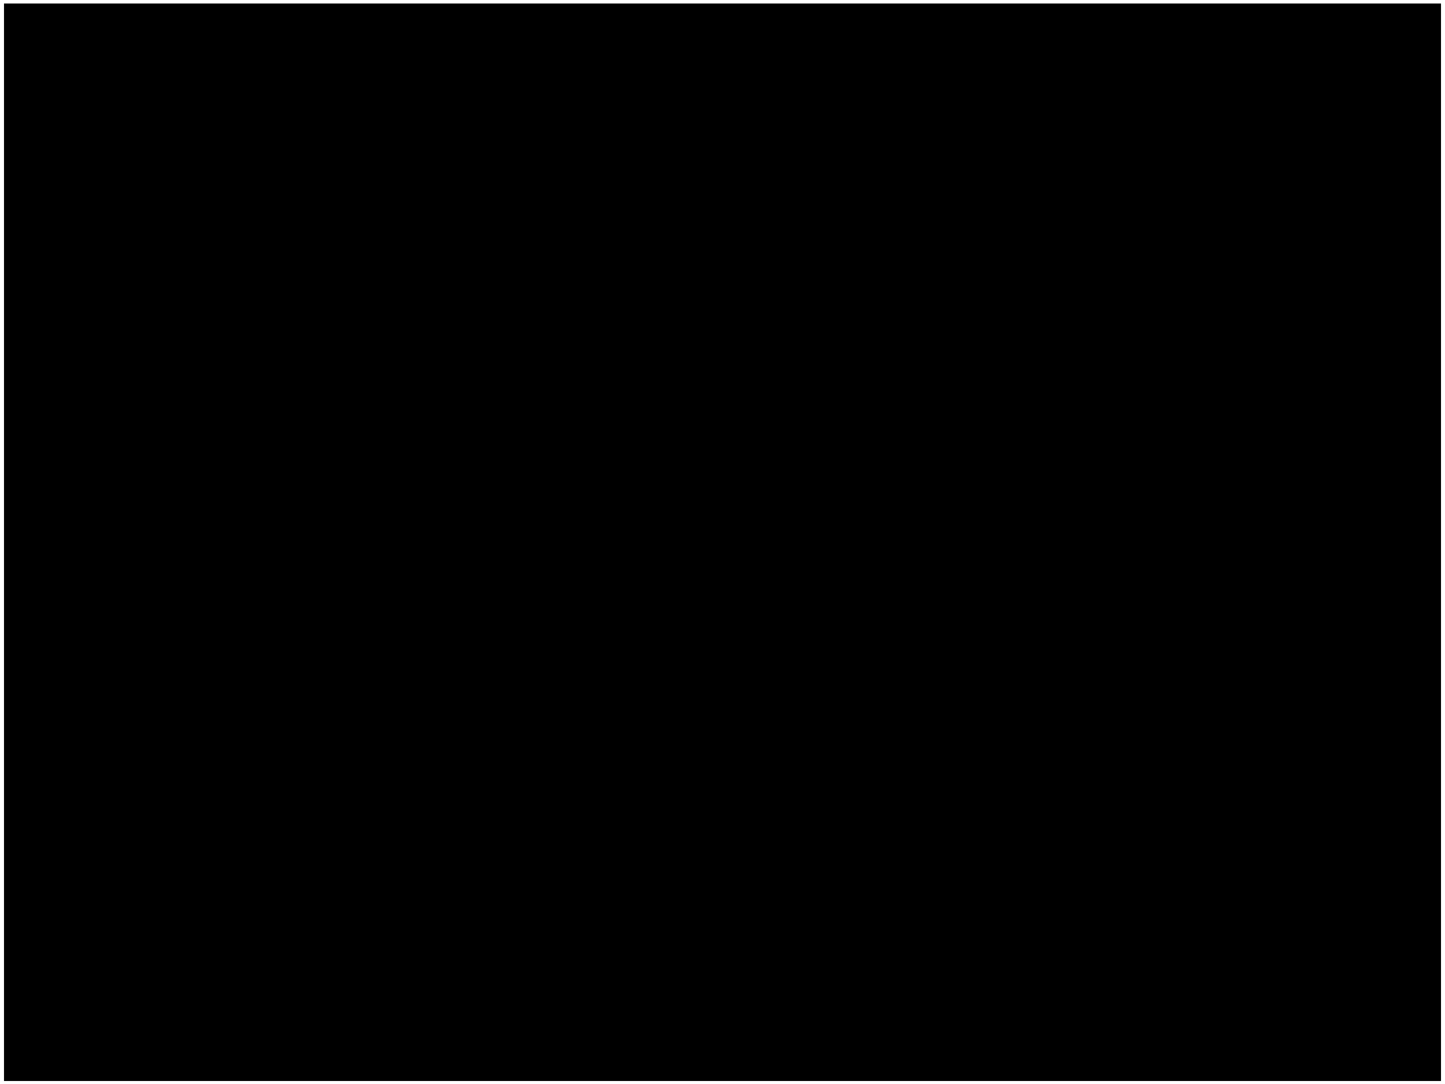

- **Learning to recognize objects based on perceptuo-motor affordances**
- **Collaboration with Univ. Paris VI**
- **Icub awarded through Robocub Open Call**
- **ANR MACSi**

(Nguyen et al.; ICDL-Epirob 2013; Ivaldi et al., IEEE TAMD 2013)

Predictions and experiments about monkey/human spontaneous exploration

- Collaboration with J. Gottlieb, Univ. Columbia, US
- Since jan 2013: **Associated Team Inria-Columbia Neurocuriosity**

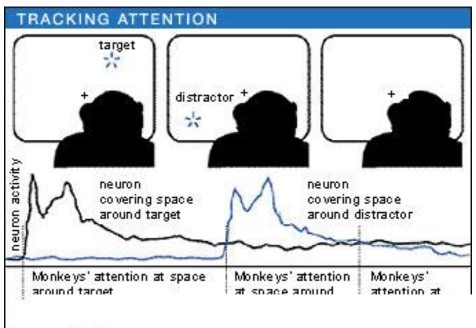

Neuroscience of visual attention and exploration in monkeys

Structure of intrinsically motivated exploration of multiple sensorimotor « games » in humans

➔(Gottlieb, Oudeyer et al. (in press)
« Information seeking, curiosity and attention: computational and neural mechanisms »
Trends in Cognitive Science)

Social learning, imitation

Hierarchical curiosity-driven learning

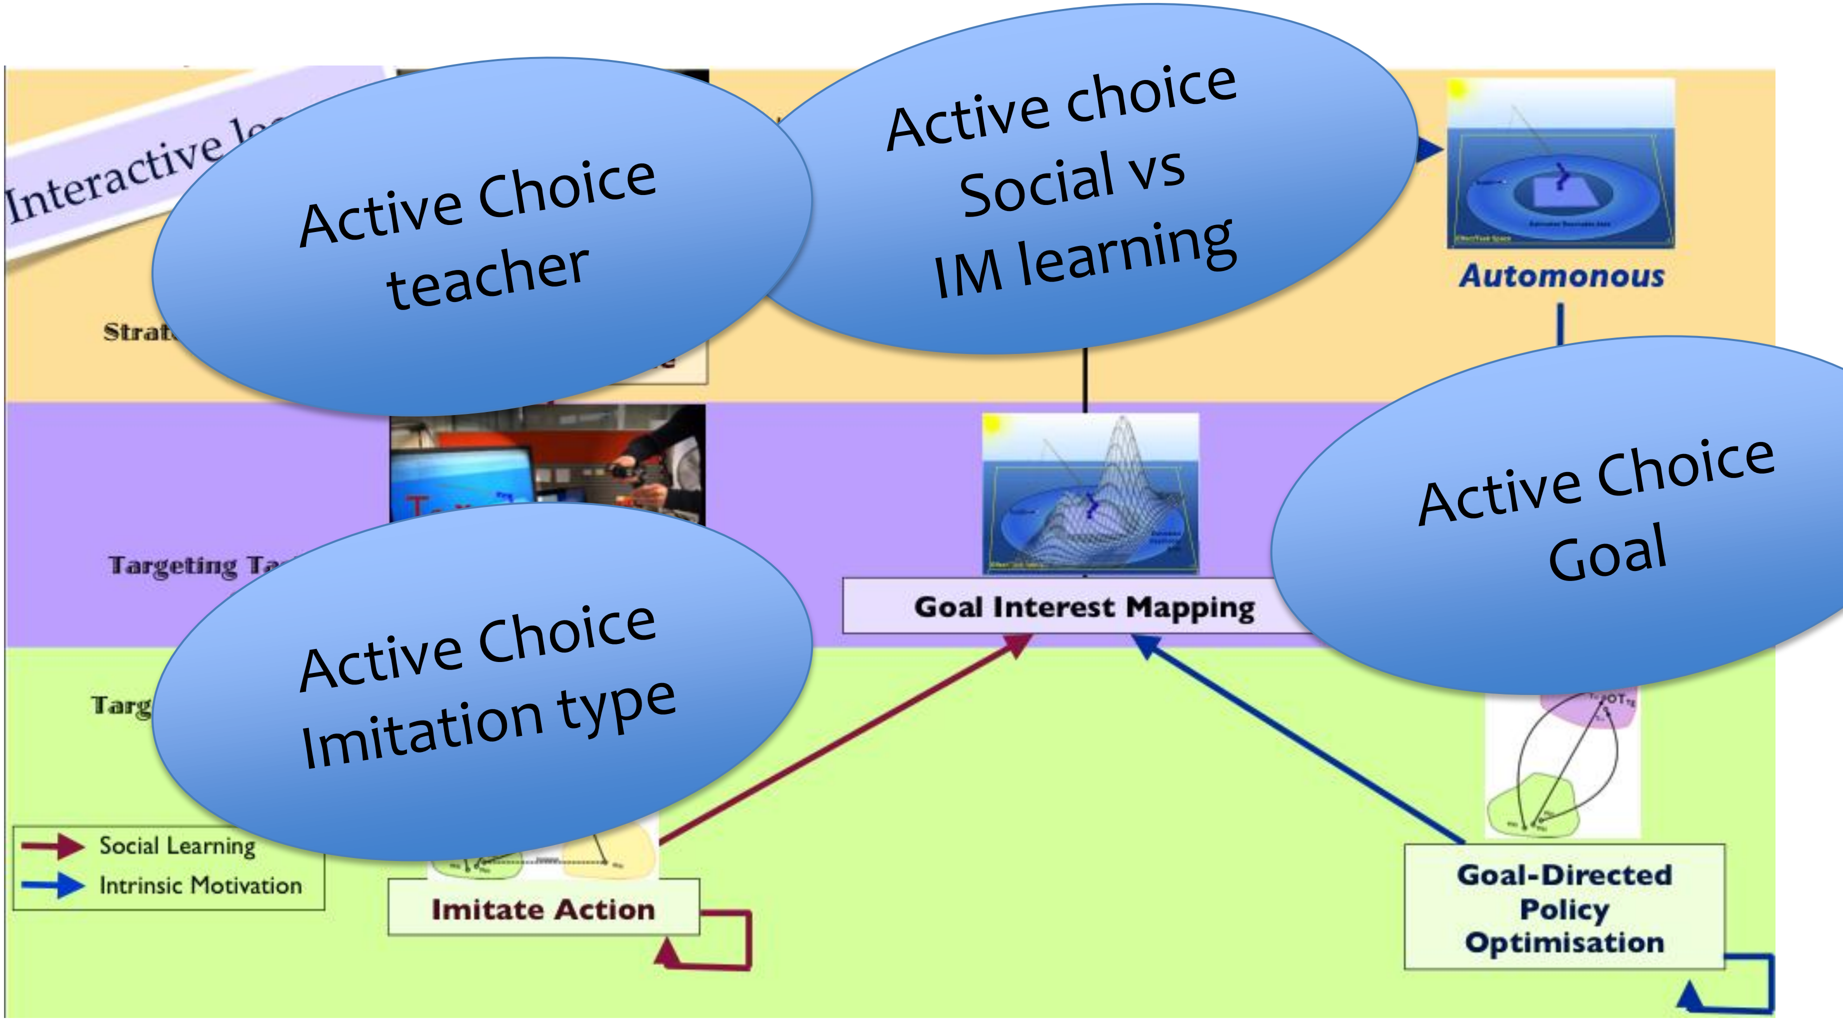

(Nguyen and Oudeyer, Palad. Behav. Rob., 2013; Autonomous Rob. 2013)

Learning to use a fishing rod (essential tool, isnt'it ?)

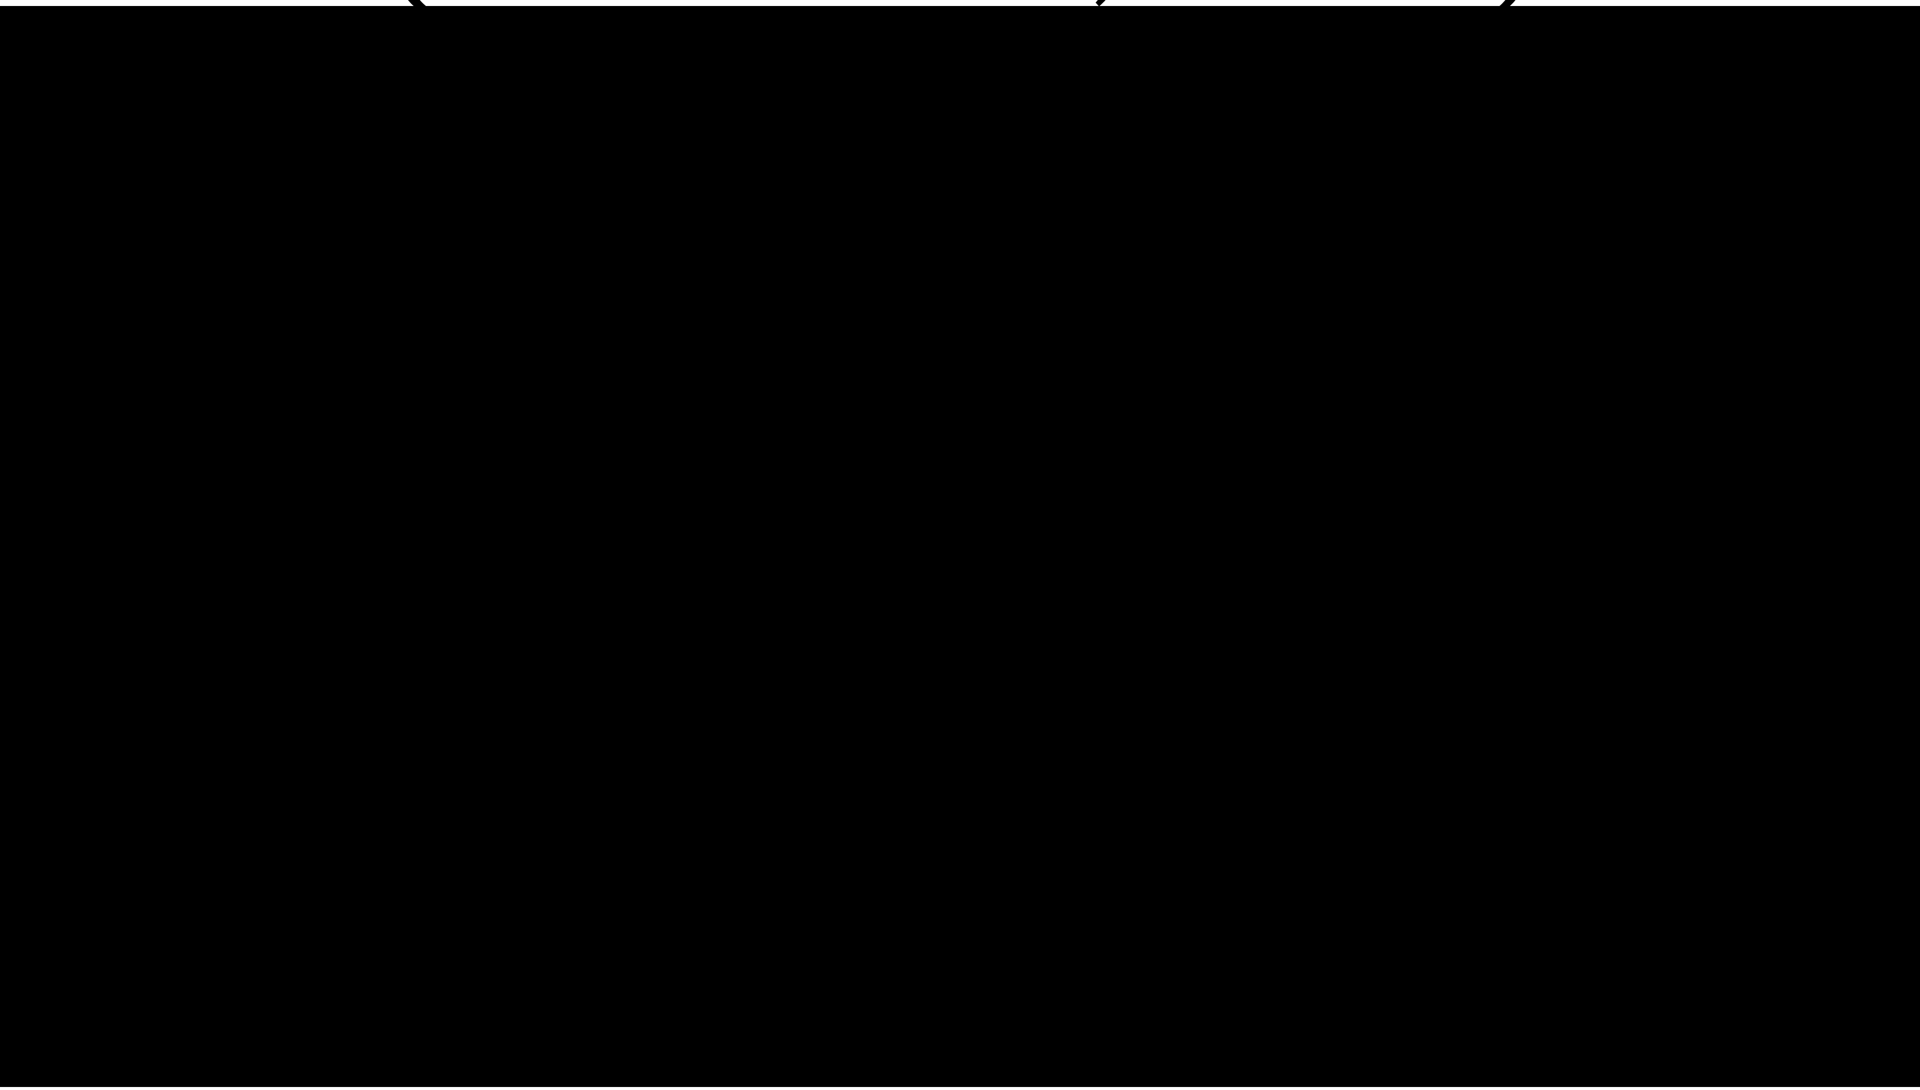

(Nguyen, Oudeyer, Fudal, Béchu, 2010-13)

The role of body morphology

Morphology and self-organization of biped locomotion

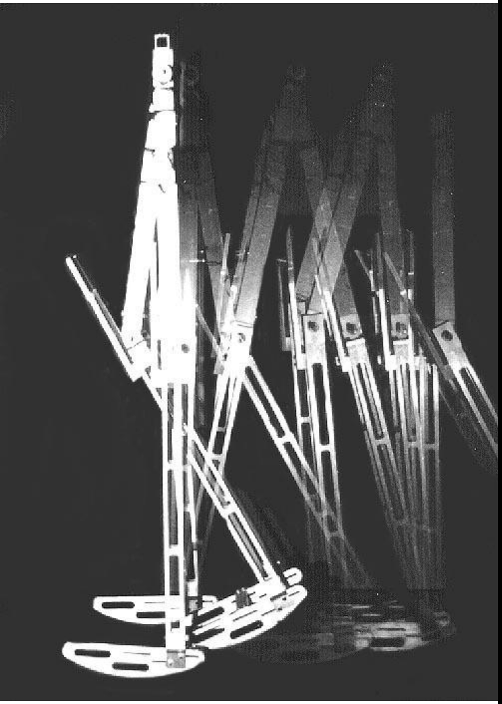

Tad McGeer (McGeer, 1990), Nagoya Univ. (2005)

Morphological computation

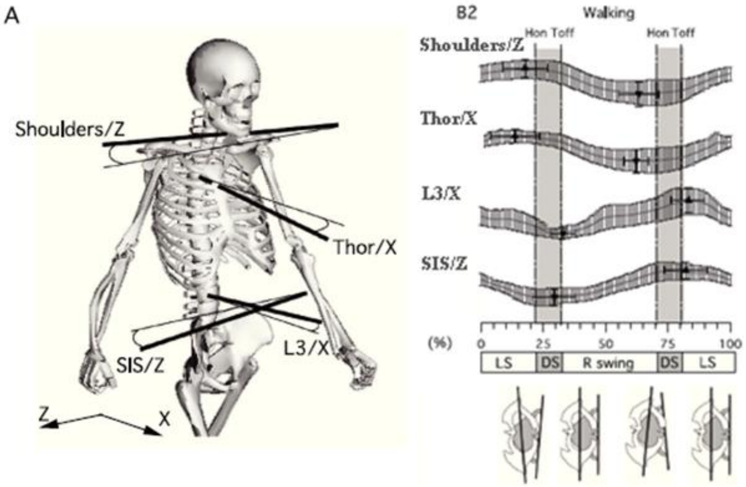

(Ceccato et Cazalets, 2009)

- **Collaboration with Labri/Univ. Bordeaux I**
- **Collaboration with J-R. Cazalets, Integrative Neuroscience Institute, Bordeaux**

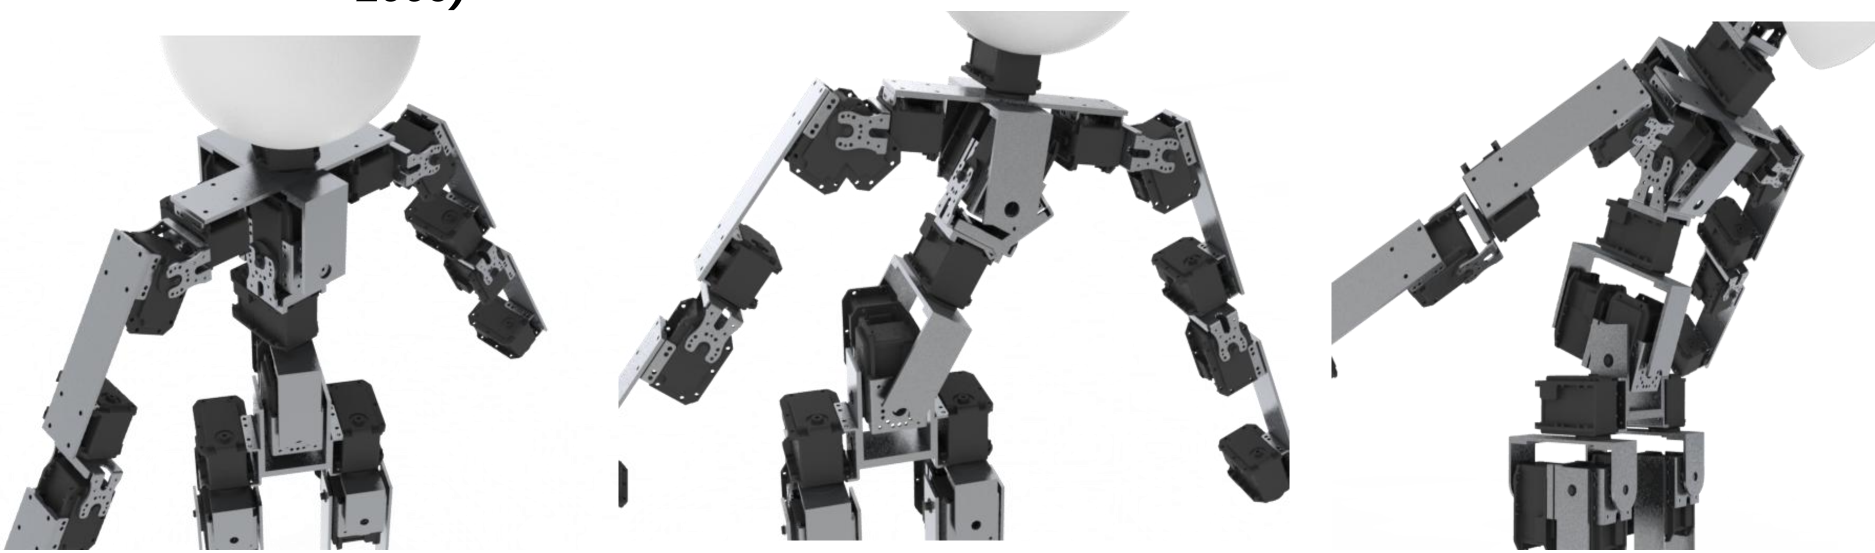

The Acroban humanoid *(Ly, Lapeyre, Oudeyer, 2011, IROS)*

Désimbriquer
corps, cerveau et environnement

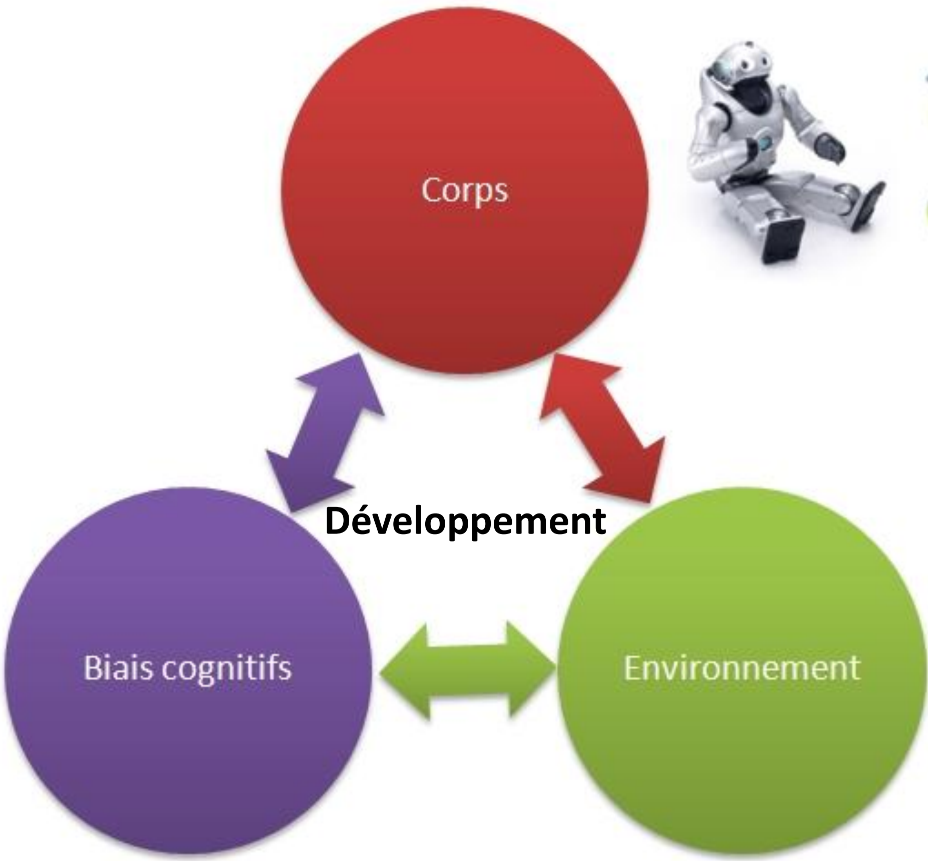

Plateformes « off-the-shelf »

Robots
industriels

- Dangereux
- Rigide
- Non-reconfigurable
- Fixe au labo
- Cher
- Difficile à réparer soi-même

Robots low-cost
off-the-shelf

- Peu précis
- Capacités motrices réduites
- Rigide
- Difficile à réparer soi-même

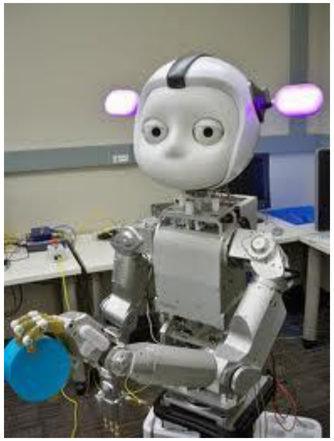

Robots industriels souple

- Dangereux
- Non-reconfigurables
- Fixe au labo
- Difficile à réparer soi-mê
- Cher

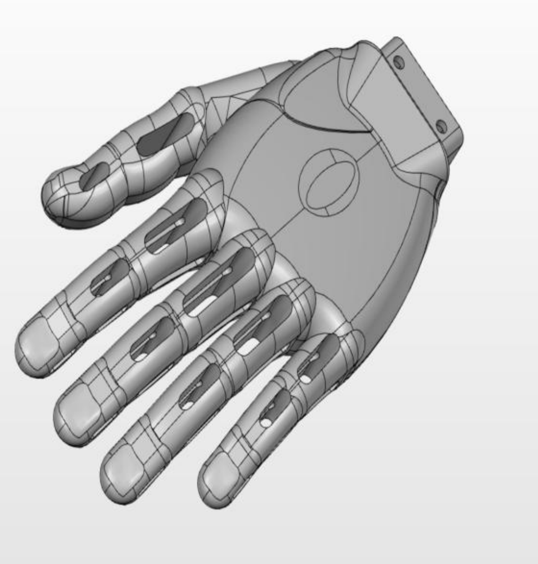

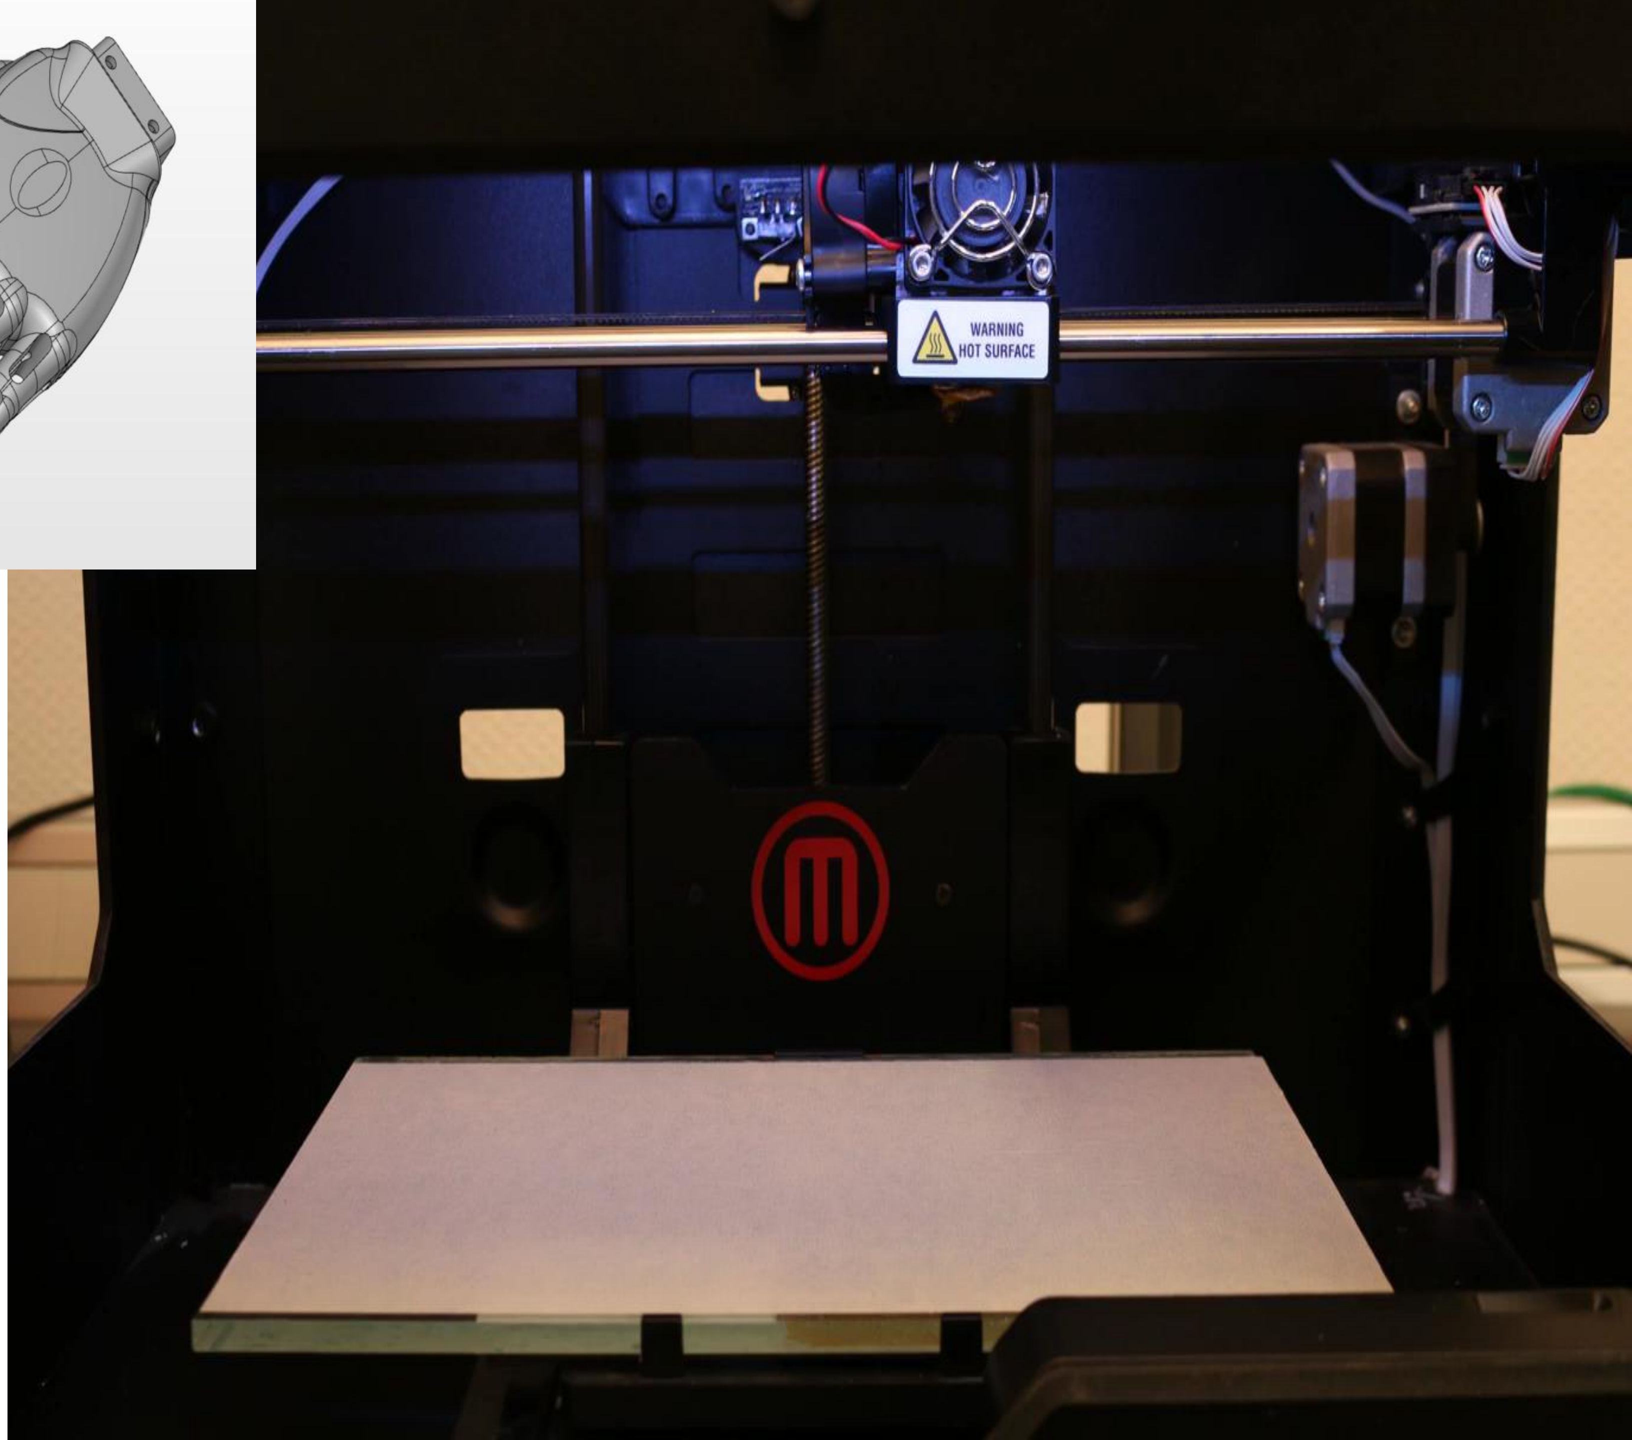

3D printing

Poppy : robot humanoide DIY open-source

[redacted]

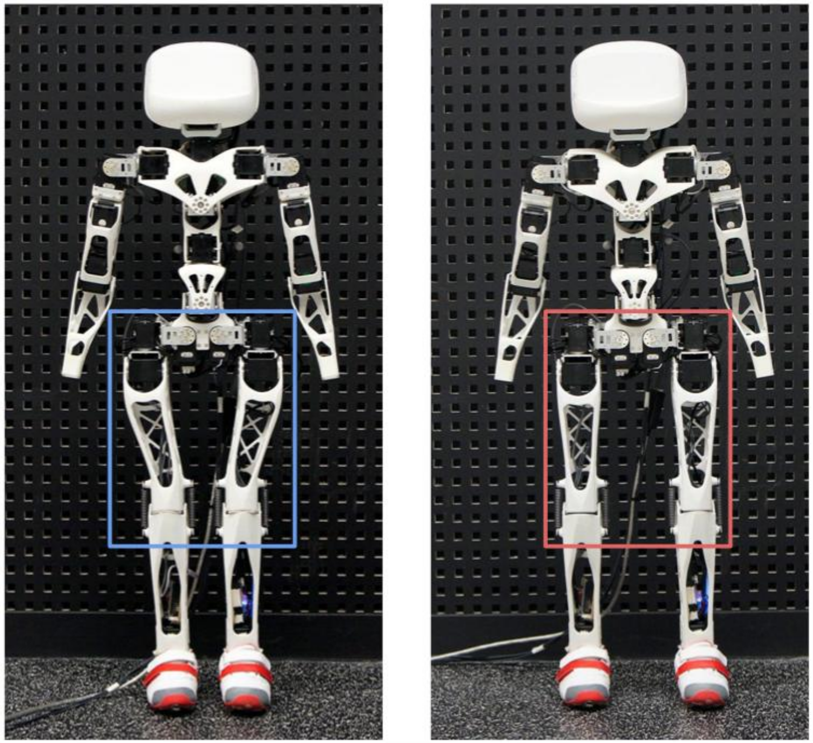

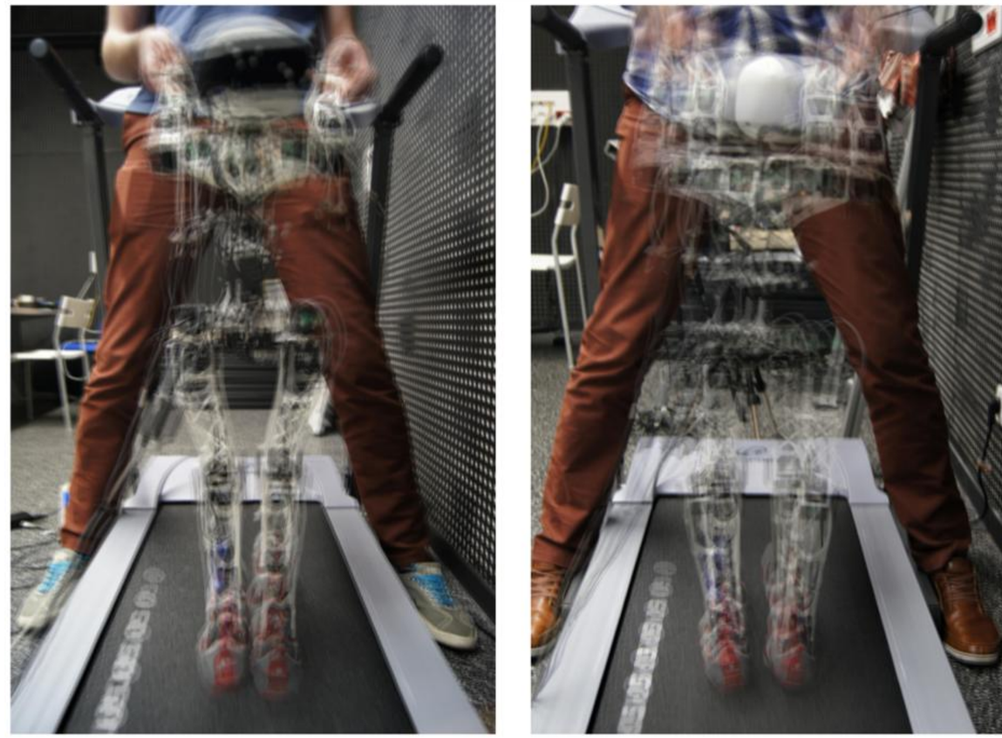

(a) bended thigh

(b) straight thigh

Feets

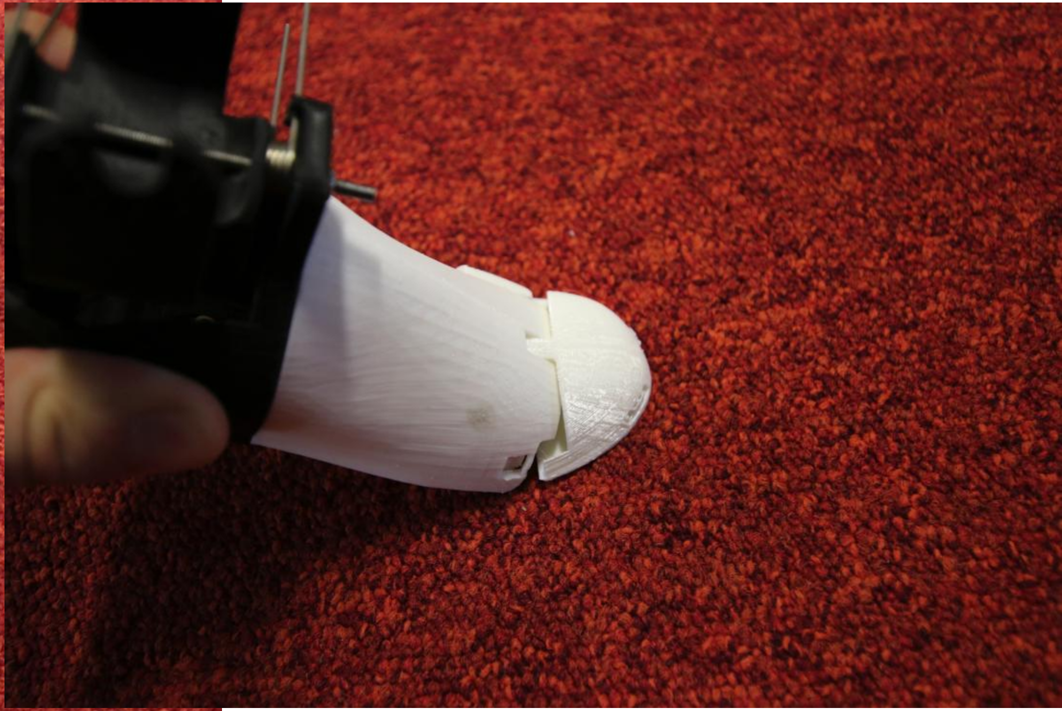

Coupe du monde

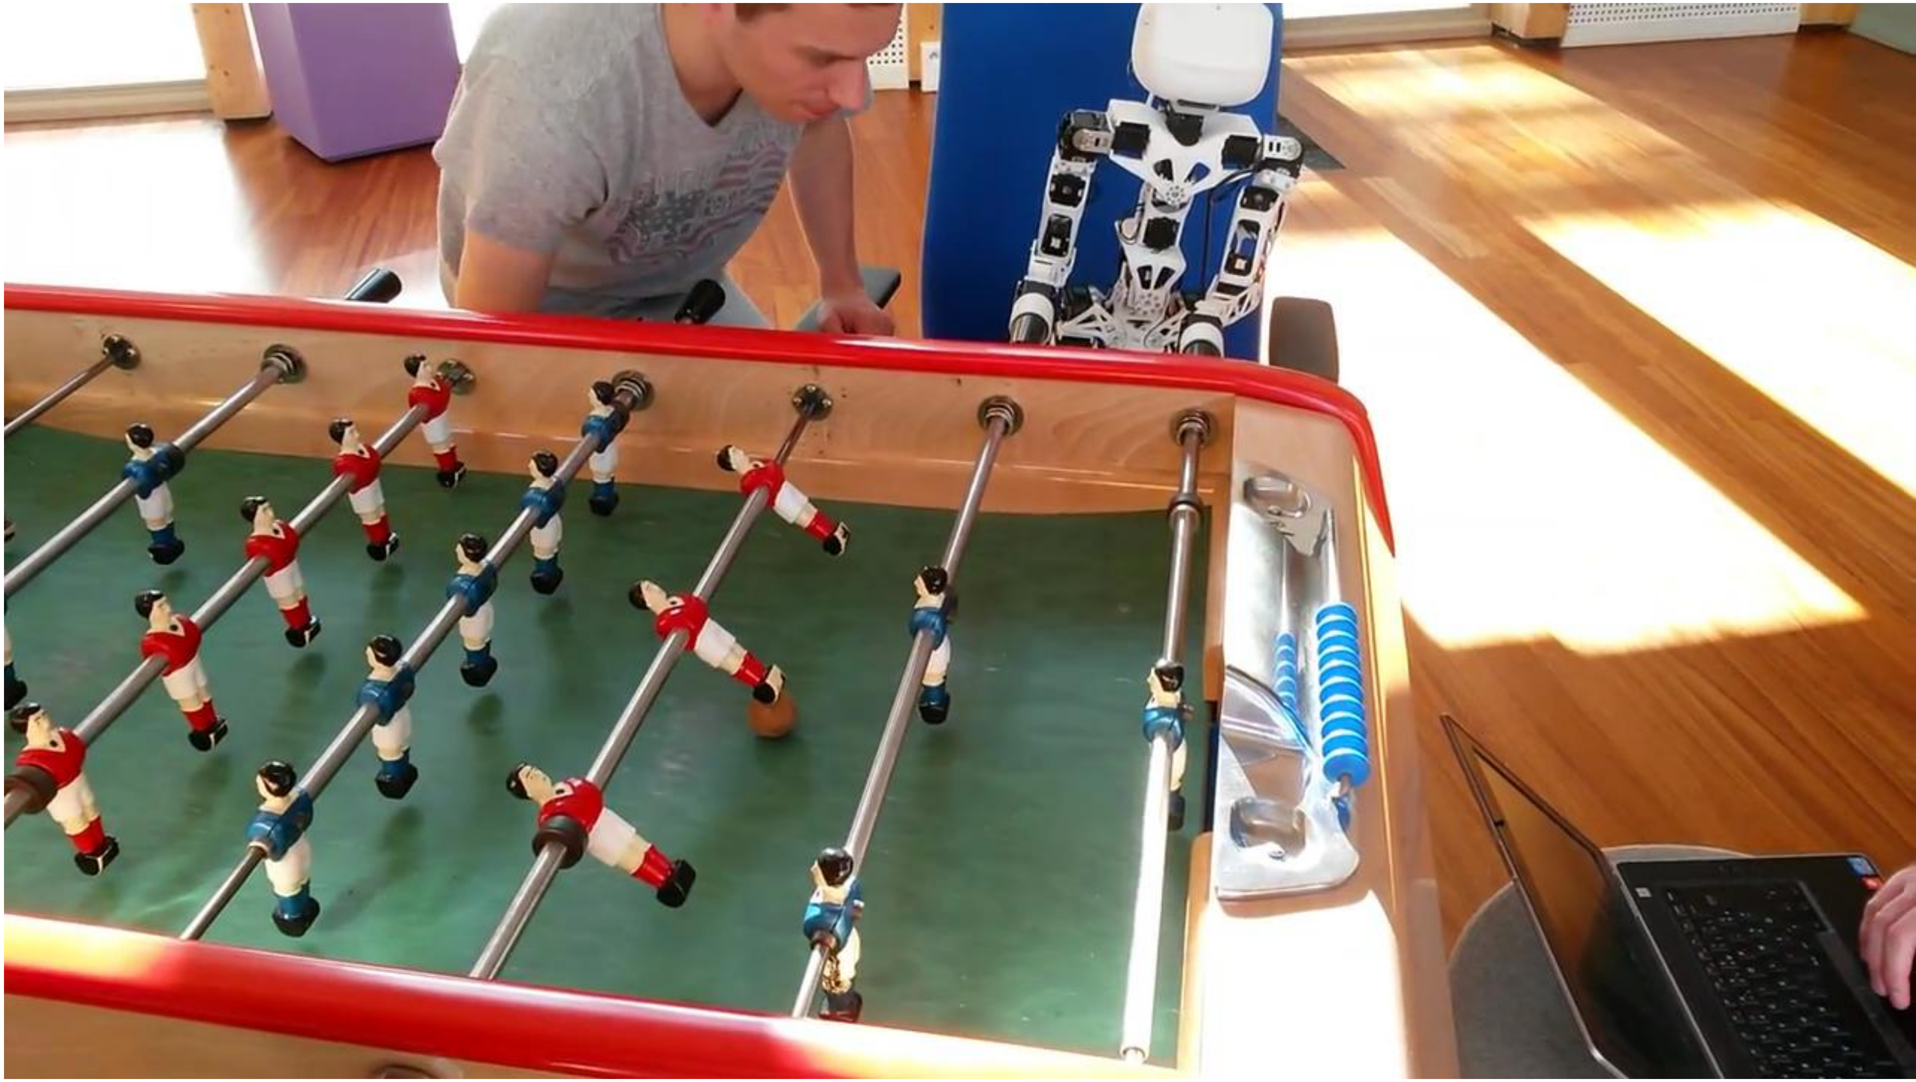

Éducation

Hackathon à la cité des sciences (Paris)

- Outil pédagogique
  - Design
  - Mécanique
  - Informatique
  - Electronique
- Projet de groupes
- Formation à l'impression 3D
- Hacker le robot (morphologie)

POPPYPROJECT

POPPY RESEARCH NEWS ABOUT US

THE POPPY PLATFORM

Poppy is an Open-source humanoid platform based on robust, flexible, easy-to-use hardware and software.

Designed by the Flowers Lab at inria Bordeaux (France), its development aims at providing an affordable humanoid robot for research and education.

Our current research with Poppy focuses on the study of the morphology, the learning of biped locomotion, and physical & social human robot interaction.

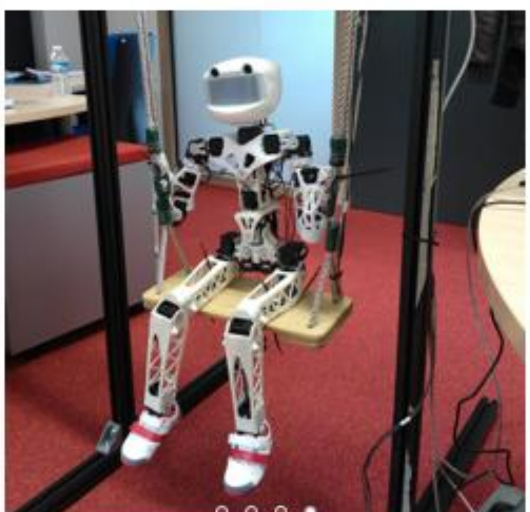

- **OPEN SOURCE**
  Both software and hardware are available under an open source licence for academics.
- **AFFORDABLE**
  The overall materials needed to build your own Poppy robot cost around 7500€ (including motors, electronics and 3D printed parts).
- **OPTIMIZED FOR BIPED LOCOMOTION**
  The morphological optimization is mainly expressed on the locomotive system (legs and trunks) to increase the robot robustness, agility and stability during the walking.
- **SOCIAL AND PHYSICAL HUMAN-ROBOT INTERACTION**
  Physical interaction with full body compliance and an articulated torso. Optionally, social interaction can be improved with cameras, micros and LCD Screen.
- **EASY TO REPAIR AND DUPLICATE**
  Poppy only uses off-the-shelf components (motors and electronics) and limbs that can be printed with regular 3D printing services.

www.poppy-project.org

Open-source hardware and software

For academics and geeks

Twitter: @poppy_project

Poppy assembly

Rencontre **robotique**
artistique
autour du geste