informatics mathematics

IN PARTNERSHIP WITH:
Ecole nationale supérieure des techniques avancées

Activity Report 2013

Project-Team FLOWERS

Flowing Epigenetic Robots and Systems

RESEARCH CENTER
Bordeaux - Sud-Ouest

THEME
Robotics and Smart environments

Table of contents

Project-Team FLOWERS

Keywords: Developmental Robotics, Human-robot Interaction, Machine Learning, Language Acquisition, Models Of Human Cognitive Development

Creation of the Team: 2008 April 01*, updated into Project-Team:* 2011 January 01.

1. Members

Research Scientists
Pierre-Yves Oudeyer [Team leader, Inria, Senior Researcher, HdR]
Manuel Lopes [Inria, Researcher]

Faculty Members
David Filliat [ENSTA, Professor, HdR]
Alexander Gepperth [ENSTA, Associate Professor]
Freek Stulp [ENSTA, Associate Professor]

Engineers
Benjamin Clement [Inria, FP7 ERC/09-EXPLORERS project, from Feb 2013 until Aug 2013]
Paul Fudal [Inria, FP7 ERC/09-EXPLORERS project, until Sep 2013]
Antoine Hoarau [Inria]
Yoan Mollard [Inria, from Dec 2013]
Pierre Rouanet [Inria, FP7 ERC/09-EXPLORERS project, from Jan 2013]
Didier Roy [Inria, from Sept 2013]
Jérome Béchu [ENSTA, from Apr 2013]

PhD Students
Fabien Benureau [Univ. Bordeaux I]
Jan Thomas Cederborg [Inria, FP7 ERC/09-EXPLORERS project, finished Dec 2013]
Mai Nguyen [Inria, finished Nov 2013]
Louis-Charles Caron [ENSTA]
Alexandre Armand [ENSTA, CIFRE Renault]
Duceux Guillaume [ENSTA]
Jonathan Grizou [Inria, FP7 ERC/09-EXPLORERS project]
Natalia Lyubova [ENSTA, ANR MACSi, until Oct 2013, PostDoc from Nov 2013]
Thibaut Munzer [Inria, FP7 EU-STREP 3rdHand, from Nov 2013]
Ievgen Perederieiev [Inria, Conseil Régional d'Aquitaine]
Olivier Mangin [Inria, FP7 ERC/09-EXPLORERS project, until Aug 2013]
Matthieu Lapeyre [Inria, FP7 ERC/09-EXPLORERS project, until Aug 2013]
Yuxin Chen [ENSTA, started oct 2013]

Post-Doctoral Fellows
Mathieu Lefort [Inria, from Nov 2013]
Clément Moulin-Frier [Inria, FP7 ERC/09-EXPLORERS project, until Nov 2014]
Thomas Degris [Inria, ANR /10- MACSi project, until Oct 2013]

Administrative Assistant
Nicolas Jahier [Inria]

Others
Jules Brochard [ENS Cachan, Master Student, from May 2013 until Aug 2013]
Axel Davy [ENS Paris, Master Student, from Jun 2013 until Aug 2013]
Julie Golliot [Inria, Master Student, from May 2013 until Aug 2013]
Brice Miard [Univ. Bordeaux II, Master Student, from Nov 2013 until Dec 2013]

Chloé Rozenbaum [Inria, Master Student, from May 2013 until Jul 2013]
Caio Tomazelli Da Silva Oliveira [Inria, Master Student, from May 2013 until Aug 2013]

2. Overall Objectives

2.1. Introduction

Can a robot learn like a child? Can it learn new skills and new knowledge in an unknown and changing environment? How can it discover its body and its relationships with the physical and social environment? How can its cognitive capacities continuously develop without the intervention of an engineer? What can it learn through natural social interactions with humans?

These are the questions that are being investigated in the FLOWERS research team at Inria Bordeaux Sud-Ouest. Rather than trying to imitate the intelligence of adult humans like in the field of Artificial Intelligence, we believe that trying to reconstruct the processes of development of the child's mind will allow for more adaptive, more robust and more versatile machines. This approach is called developmental robotics, or epigenetic robotics, and imports concepts and theories from developmental psychology. As most of these theories are not formalized, this implies a crucial computational modeling activity, which in return provides means to assess the internal coherence of theories and sketch new hypothesis about the development of the human child's sensorimotor and cognitive abilities.

Our team focuses in particular on the study of developmental constraints that allow for efficient open-ended learning of novel sensorimotor and interaction skills in embodied systems. In particular, we study constraints that guide exploration in large sensorimotor spaces:

- Mechanisms of intrinsically motivated exploration and active learning, including artificial curiosity, allowing in particular to self-organize developmental trajectories and collect efficiently learning data;
- Mechanisms of adequately constrained optimization and statistical inference for sensorimotor skill acquisition (e.g. for optimizing motor policies in real robots);
- Mechanisms for social learning, e.g. learning by imitation or demonstration, which implies both issues related to machine learning and human-robot interaction;
- Constraints related to embodiment, in particular through the concept of morphological computation, as well as the structure of motor primitives/muscle synergies that can leverage the properties of morphology and physics for simplifying motor control and perception;
- Maturational constraints which, coupled with the other constraints, can allow the progressive release of novel sensorimotor degrees of freedom to be explored;

We also study how these constraints on exploration can allow a robot to bootstrap multimodal perceptual abstractions associated to motor skills, in particular in the context of modelling language acquisition as a developmental process grounded in action.

Among the developmental principles that characterize human infants and can be used in developmental robots, FLOWERS focuses on the following three principles:

- **Exploration is progressive.** The space of skills that can be learnt in real world sensorimotor spaces is so large and complicated that not everything can be learnt at the same time. Simple skills are learnt first, and only when they are mastered, new skills of progressively increasing difficulty become the behavioural focus;
- **Internal representations are (partially) not innate but learnt and adaptive.** For example, the body map, the distinction self/non-self and the concept of "object" are discovered through experience with initially uninterpreted sensors and actuators, guided by experience, the overall pre-determined connection structure of the brain, as well as a small set of simple innate values or preferences.
- **Exploration can be self-guided and/or socially guided.** On the one hand, internal and intrinsic motivation systems regulate and organize spontaneous exploration; on the other hand, exploration can be guided through social learning and interaction with caretakers.

2.1.1. Research axis

The work of FLOWERS is organized around the following axis:

- **Curiosity-driven exploration and sensorimotor learning:** intrinsic motivation are mechanisms that have been identified by developmental psychologists to explain important forms of spontaneous exploration and curiosity. In FLOWERS, we try to develop computational intrinsic motivation systems, and test them on robots, allowing to regulate the growth of complexity in exploratory behaviours. These mechanisms are studied as active learning mechanisms, allowing to learn efficiently in large inhomogeneous sensorimotor spaces;
- **Cumulative learning of sensorimotor skills:** FLOWERS develops machine learning algorithms that can allow embodied machines to acquire cumulatively sensorimotor skills. In particular, we develop optimization and reinforcement learning systems which allow robots to discover and learn dictionaries of motor primitives, and then combine them to form higher-level sensorimotor skills.
- **Natural and intuitive social learning:** FLOWERS develops interaction frameworks and learning mechanisms allowing non-engineer humans to teach a robot naturally. This involves two sub-themes: 1) techniques allowing for natural and intuitive human-robot interaction, including simple ergonomic interfaces for establishing joint attention; 2) learning mechanisms that allow the robot to use the guidance hints provided by the human to teach new skills;
- **Discovering and abstracting the structure of sets of uninterpreted sensors and motors:** FLOWERS studies mechanisms that allow a robot to infer structural information out of sets of sensorimotor channels whose semantics is unknown, for example the topology of the body and the sensorimotor contingencies (propriocetive, visual and acoustic). This process is meant to be open-ended, progressing in continuous operation from initially simple representations to abstract concepts and categories similar to those used by humans.
- Body design and role of the body in sensorimotor and social development We study how the physical properties of the body (geometry, materials, distribution of mass, growth, ...) can impact the acquisition of sensorimotor and interaction skills. This requires to consider the body as an experimental variable, and for this we develop special methodologies for designing and evaluating rapidly new morphologies, especially using rapid prototyping techniques like 3D printing.
- **Intelligent Tutoring Systems:**FLOWERS develops methods for online personalization of teaching sequences for educational software and MOOCs. This work builds on top of online optimization methods and motivational research previously developed.

2.2. Highlights of the Year

In April 2013 at the International Conference on Robotics and Automation in Karlsruhe, Freek Stulp received the “King-Sun Fu Best Paper Award of the IEEE Transactions on Robotics”. As T-RO has on of the highest impact factors, this is considered to be the highest paper prize in robotics. It is the first time this prize has been awarded to an article on machine learning.

The team has announced in october 2013 the open-source release of the Poppy humanoid robot. Poppy is to our knowledge the first humanoid robot in the world to be at the same time open-source (hardware and software) and based on 3D printing techniques. It is based on robust, flexible, easy-to-use hardware and software. Its development aims at providing an affordable and hackable humanoid robot for science, education, art and geeks. Poppy was initially made for our research project about understanding the role of morphology in biped locomotion, and full-body physical and social interaction in robots and humans. The robot has generated a huge enthusiasm from geeks, academic laboratories, and educational institutions, and within the first two months already 200 beta-testers registered to rebuild their own copy of the robot. Dozens of articles appeared on the internet and printed press, and the Poppy videos was viewed nearly 40k times. Web site: http://www.poppy-project.org.

The Flowers team made major achievements in diffusing science and technology towards the general public. Pierre-Yves Oudeyer published a popular science book entitled "Aux sources de la parole" at Odile Jacob, and was invited to talk about our research on major radio channels (e.g. France Inter, France Culture, France Info). http://www.pyoudeyer.com/AuxSourcesDeLaParole.htm

The team also initiated the development of educational activities in "écoles primaires" and "collèges" to have kids discover robotics and programming, as well as ran experiments in "école primaires" in Aquitaine to test novel educational software to help children learn mathematics, and developped within the KidLearn ADT project. This was achieved thanks to the arrival of Didier Roy, former math teacher in college, in the team.

The Flowers team is now coordinating the European project "Semi-autonomous 3rdHand" (coord. Manuel Lopes). The goal is to develop a semi-autonomous robot assistant that acts as a third hand of a human worker in factories, which may be a transformative technology for industry in the coming years. It aims to elaborate techniques allowing to instruct even by an untrained layman worker, allow for efficient knowledge transfer between tasks and enable a effective collaboration between a human worker with a robot third hand. http://3rdhandrobot.eu

The Flowers team started the work on Intelligent Tutoring Systems. The project Kidlearn is a research project studying how machine learning can be applied to intelligent tutoring systems. It aims at developing methodologies and software which adaptively personalize sequences of learning activities to the particularities of each individual student. First experiments were realized in elementary schools of Région Aquitaine, where 6-7 year old kids learnt elements of mathematics with our software. https://flowers.inria.fr/research/kidlearn/

An associated team, called Neurocuriosity, was created between Flowers and the Cognitive Neuroscience lab of Jacqueline Gottlieb at Univ. Columbia, NY. The goal of this associated team is to investigate mechanisms of spontaneous exploration and learning in humans by setting up experiments allowing to confirm or falsify predictions made by computational models previously developped by the team. This constitutes a crucial collaboration between developmental robotics and cognitive neuroscience. This joint work already led to a major publication on curiosity and information seeking, in the prestigious Trends in Cognitive Science journal (impact factor: 16.5).[10]

Thomas Cederborgs PhD thesis "A Formal Approach to Social Learning: Exploring Language Acquisition Through Imitation" won the ”ThesAqt” prize, awarded by Region Aquitaine who gives this awards to excellent theses in the region.

3. Research Program

3.1. Research Program

Research in artificial intelligence, machine learning and pattern recognition has produced a tremendous amount of results and concepts in the last decades. A blooming number of learning paradigms - supervised, unsupervised, reinforcement, active, associative, symbolic, connectionist, situated, hybrid, distributed learning... - nourished the elaboration of highly sophisticated algorithms for tasks such as visual object recognition, speech recognition, robot walking, grasping or navigation, the prediction of stock prices, the evaluation of risk for insurances, adaptive data routing on the internet, etc... Yet, we are still very far from being able to build machines capable of adapting to the physical and social environment with the flexibility, robustness, and versatility of a one-year-old human child.

Indeed, one striking characteristic of human children is the nearly open-ended diversity of the skills they learn. They not only can improve existing skills, but also continuously learn new ones. If evolution certainly provided them with specific pre-wiring for certain activities such as feeding or visual object tracking, evidence shows that there are also numerous skills that they learn smoothly but could not be “anticipated” by biological evolution, for example learning to drive a tricycle, using an electronic piano toy or using a video game joystick. On the contrary, existing learning machines, and robots in particular, are typically only able to learn a single pre-specified task or a single kind of skill. Once this task is learnt, for example walking with two legs, learning is over. If one wants the robot to learn a second task, for example grasping objects in its visual field, then an engineer needs to re-program manually its learning structures: traditional approaches to task-specific machine/robot learning typically include engineer choices of the relevant sensorimotor channels, specific design of the reward function, choices about when learning begins and ends, and what learning algorithms and associated parameters shall be optimized.

As can be seen, this requires a lot of important choices from the engineer, and one could hardly use the term “autonomous” learning. On the contrary, human children do not learn following anything looking like that process, at least during their very first years. Babies develop and explore the world by themselves, focusing their interest on various activities driven both by internal motives and social guidance from adults who only have a folk understanding of their brains. Adults provide learning opportunities and scaffolding, but eventually young babies always decide for themselves what activity to practice or not. Specific tasks are rarely imposed to them. Yet, they steadily discover and learn how to use their body as well as its relationships with the physical and social environment. Also, the spectrum of skills that they learn continuously expands in an organized manner: they undergo a developmental trajectory in which simple skills are learnt first, and skills of progressively increasing complexity are subsequently learnt.

A link can be made to educational systems where research in several domains have tried to study how to provide a good learning experience to learners. This includes the experiences that allow better learning, and in which sequence they must be experienced. This problem is complementary to that of the learner that tries to learn efficiently, and the teacher here has to use as efficienlty the limited time and motivational resources of the learner. Several results from psychology [70] and neuroscience [10] have argued that the human brain feels intrinsic pleasure in practicing activities of optimal difficulty or challenge. A teacher must exploit such activities to create positive psychological states of flow [76].

A grand challenge is thus to be able to build robotic machines that possess this capability to discover, adapt and develop continuously new know-how and new knowledge in unknown and changing environments, like human children. In 1950, Turing wrote that the child’s brain would show us the way to intelligence: “Instead of trying to produce a program to simulate the adult mind, why not rather try to produce one which simulates the child’s” [133]. Maybe, in opposition to work in the field of Artificial Intelligence who has focused on mechanisms trying to match the capabilities of “intelligent” human adults such as chess playing or natural language dialogue [91], it is time to take the advice of Turing seriously. This is what a new field, called developmental (or epigenetic) robotics, is trying to achieve [100] [135]. The approach of developmental robotics consists in importing and implementing concepts and mechanisms from developmental psychology [107], cognitive linguistics [75], and developmental cognitive neuroscience [95] where there has been a considerable amount of research and theories to understand and explain how children learn and develop. A number of general principles are underlying this research agenda: embodiment [72] [119], grounding [89], situatedness [66], self-organization [131] [120], enaction [134], and incremental learning [73].

Among the many issues and challenges of developmental robotics, two of them are of paramount importance: exploration mechanisms and mechanisms for abstracting and making sense of initially unknown sensorimotor channels. Indeed, the typical space of sensorimotor skills that can be encountered and learnt by a developmental robot, as those encountered by human infants, is immensely vast and inhomogeneous. With a sufficiently rich environment and multimodal set of sensors and effectors, the space of possible sensorimotor activities is simply too large to be explored exhaustively in any robot’s life time: it is impossible to learn all possible skills and represent all conceivable sensory percepts. Moreover, some skills are very basic to learn, some other very complicated, and many of them require the mastery of others in order to be learnt. For example, learning to manipulate a piano toy requires first to know how to move one’s hand to reach the piano and how to touch specific parts of the toy with the fingers. And knowing how to move the hand might require to know how to track it visually.

Exploring such a space of skills randomly is bound to fail or result at best on very inefficient learning [15]. Thus, exploration needs to be organized and guided. The approach of epigenetic robotics is to take inspiration from the mechanisms that allow human infants to be progressively guided, i.e. to develop. There are two broad classes of guiding mechanisms which control exploration:

1. **internal guiding mechanisms,** and in particular intrinsic motivation, responsible of spontaneous exploration and curiosity in humans, which is one of the central mechanisms investigated in FLOWERS, and technically amounts to achieve online active self-regulation of the growth of complexity in learning situations;
2. **social learning and guidance,** a learning mechanisms that exploits the knowledge of other agents in the environment and/or that is guided by those same agents. These mechanisms exist in many different forms like emotional reinforcement, stimulus enhancement, social motivation, guidance, feedback or imitation, some of which being also investigated in FLOWERS;

3.1.1. Internal guiding mechanisms

In infant development, one observes a progressive increase of the complexity of activities with an associated progressive increase of capabilities [107], children do not learn everything at one time: for example, they first learn to roll over, then to crawl and sit, and only when these skills are operational, they begin to learn how to stand. The perceptual system also gradually develops, increasing children perceptual capabilities other time while they engage in activities like throwing or manipulating objects. This make it possible to learn to identify objects in more and more complex situations and to learn more and more of their physical characteristics.

Development is therefore progressive and incremental, and this might be a crucial feature explaining the efficiency with which children explore and learn so fast. Taking inspiration from these observations, some roboticists and researchers in machine learning have argued that learning a given task could be made much easier for a robot if it followed a developmental sequence and "started simple" [68] [80]. However, in these experiments, the developmental sequence was crafted by hand: roboticists manually build simpler versions of a complex task and put the robot successively in versions of the task of increasing complexity. And when they wanted the robot to learn a new task, they had to design a novel reward function.

Thus, there is a need for mechanisms that allow the autonomous control and generation of the developmental trajectory. Psychologists have proposed that intrinsic motivations play a crucial role. Intrinsic motivations are mechanisms that push humans to explore activities or situations that have intermediate/optimal levels of novelty, cognitive dissonance, or challenge [70] [76] [79]. The role and structure of intrinsic motivation in humans have been made more precise thanks to recent discoveries in neuroscience showing the implication of dopaminergic circuits and in exploration behaviors and curiosity [78] [92] [125]. Based on this, a number of researchers have began in the past few years to build computational implementation of intrinsic motivation [15] [117] [123] [69] [93] [105] [124]. While initial models were developed for simple simulated worlds, a current challenge is to manage to build intrinsic motivation systems that can efficiently drive exploratory behaviour in high-dimensional unprepared real world robotic sensorimotor spaces [117][15] [118] [122]. Specific and complex problems are posed by real sensorimotor spaces, in particular due to the fact that they are both high-dimensional as well as (usually) deeply inhomogeneous. As an example for the latter issue, some regions of real sensorimotor spaces are often unlearnable due to inherent stochasticity or difficulty, in which case heuristics based on the incentive to explore zones of maximal unpredictability or uncertainty, which are often used in the field of active learning [74] [90] typically lead to catastrophic results. The issue of high dimensionality does not only concern motor spaces, but also sensory spaces, leading to the problem of correctly identifying, among typically thousands of quantities, those latent variables that have links to behavioral choices. In FLOWERS, we aim at developing intrinsically motivated exploration mechanisms that scale in those spaces, by studying suitable abstraction processes in conjunction with exploration strategies.

3.1.2. Socially Guided and Interactive Learning

Social guidance is as important as intrinsic motivation in the cognitive development of human babies [107]. There is a vast literature on learning by demonstration in robots where the actions of humans in the environment are recognized and transferred to robots [67]. Most such approaches are completely passive: the human executes actions and the robot learns from the acquired data. Recently, the notion of interactive learning has been introduced in [132], [71], motivated by the various mechanisms that allow humans to socially guide a robot [121]. In an interactive context the steps of self-exploration and social guidances are not separated and a robot learns by self exploration and by receiving extra feedback from the social context [132], [96] [106].

Social guidance is also particularly important for learning to segment and categorize the perceptual space. Indeed, parents interact a lot with infants, for example teaching them to recognize and name objects or characteristics of these objects. Their role is particularly important in directing the infant attention towards objects of interest that will make it possible to simplify at first the perceptual space by pointing out a segment of the environment that can be isolated, named and acted upon. These interactions will then be complemented by the children own experiments on the objects chosen according to intrinsic motivation in order to improve the knowledge of the object, its physical properties and the actions that could be performed with it.

In FLOWERS, we are aiming at including intrinsic motivation system in the self-exploration part thus combining efficient self-learning with social guidance [109], [110]. We also work on developing perceptual capabilities by gradually segmenting the perceptual space and identifying objects and their characteristics through interaction with the user [102] and robots experiments [94]. Another challenge is to allow for more flexible interaction protocols with the user in terms of what type of feedback is provided and how it is provided [98].

4. Application Domains

4.1. Applications

- **Personal robotics.** Many indicators show that the arrival of personal robots in homes and everyday life will be a major fact of the 21st century. These robots will range from purely entertainment or educative applications to social companions that many argue will be of crucial help in our aging society. For example, UNECE evaluates that the industry of entertainment, personal and service robotics will grow from $5.4Bn to $17.1Bn over 2008-2010. Yet, to realize this vision, important obstacles need to be overcome: these robots will have to evolve in unpredictable homes and learn new skills while interacting with non-engineer humans after they left factories, which is out of reach of current technology. In this context, the refoundation of intelligent systems that developmental robotics is exploring opens potentially novel horizons to solve these problems.
- **Human-Robot Collaboration.** Robots play a vital role for industry and ensure the efficient and competitive production of a wide range of goods. They replace humans in many tasks which otherwise would be too difficult, too dangerous, or too expensive to perform. However, the new needs and desires of the society call for manufacturing system centered around personalized products and small series productions. Human-robot collaboration could widen the use of robot in this new situations if robots become cheaper, easier to program and safe to interact with. The most relevant systems for such applications would follow an expert worker and works with (some) autonomy, but being always under supervision of the human and acts based on its task models.
- **Video games.** In conjunction with entertainment robotics, a new kind of video games are developing in which the player must either take care of a digital creature (e.g. Neopets), or tame it (e.g. Nintendogs), or raise/accompany them (e.g. Sims). The challenges entailed by programming these creatures share many features with programming personal/entertainment robots. Hence, the video game industry is also a natural field of application for FLOWERS.
- **Environment perception in intelligent vehicles.** When working in simulated traffic environments, elements of FLOWERS research can be applied to the autonomous acquisition of increasingly abstract representations of both traffic objects and traffic scenes. In particular, the object classes of vehicles and pedestrians are if interest when considering detection tasks in safety systems, as well as scene categories ("scene context") that have a strong impact on the occurrence of these object classes. As already indicated by several investigations in the field, results from present-day simulation technology can be transferred to the real world with little impact on performance. Therefore, applications of FLOWERS research that is suitably verified by real-world benchmarks has direct applicability in safety-system products for intelligent vehicles.

- **Automated Tutoring Systems.** Optimal teaching and efficient teaching/learning environments can be applied to aid teaching in schools aiming both at increase the achievement levels and the reduce time needed. From a practical perspective, improved models could be saving millions of hours of students' time (and effort) in learning. These models should also predict the achievement levels of students in order to influence teaching practices.

5. Software and Platforms

5.1. Perception Tools

Participants: David Filliat [correspondant], Natalia Lyubova, Louis-Charles Caron, Alexander Gepperth.

5.1.1. Perception Abstraction Engine

Participants: David Filliat [correspondant], Natalia Lyubova.

PAE (Perception Abstraction Engine) is a C++ library developed to provide a uniform interface to existing visual feature detector such as SIFT, SURF, MSER, superpixels, etc... Its main goal is to be able to use these various feature detectors in a "bag of feature" approach for applications such as robot localisation and object recognition. Several approach are also implemented for the visual vocabularies, in particular the fast incremental vocabularies developed in the team.

The library provide common C++ interfaces to feature detectors, visual features and visual vocabularies. A factory approach make it possible to change the feature detectors and visual vocabularies types and parameters through configuration strings, without the need to recompile. Some applications are also included in the library, in particular topological robot localization (room recognition) and visual object recognition. An Urbi interface is also provided for these modules.

5.1.2. Incremental object discovery

Participants: Natalia Lyubova [correspondant], David Filliat.

This software makes it possible to detect, model and recognize objects in a scenario of interaction between a humanoid robot and a human teacher. It is based either on standard images, or on the kinect camera to take advantage of the depth information. The software is written in C++ and relies mainly on PAE and OpenCV.

The software implements several modules: candidate object segmentation based on motion information, keypoint-based object tracking, incremental object model construction integrating multiple features (keypoints + superpixels) and object categorisation based on mutual information with robot motors (making it possible to segment robot parts, objects and humans). Based on all these modules, it is possible for the robot to learn objects shown by a human partner and to improve the objects models by manipulating them when they are put in front of the robot.

5.1.3. Object recognition from a 3-D point cloud

Participants: Louis-Charles Caron [correspondant], Alexander Gepperth, David Filliat.

This software scans the 3-D point cloud of a scene to find objects and match them against a database of known objects. The process consists in 3 stages. The segmentation step finds the objects in the point cloud, the feature extraction computes discriminating properties to be used in the classification stage for object recognition.

The segmentation is based on simple assumptions about the geometry of an indoor scene. Successive RANSACs are used to find large planes, which correspond to the floor, ceiling and walls. The cloud is stripped from the points belonging to these planes. The remaining points are clustered, meaning that close-by points are considered to form a single object.

Objects are characterized by their shape and color. Color histograms and SIFT features are computed, using the PAE library, to capture the visual appearance of the objects. Their shape is encoded by computing thousands of randomly chosen SURFLET features to construct a relative frequency histrogram.

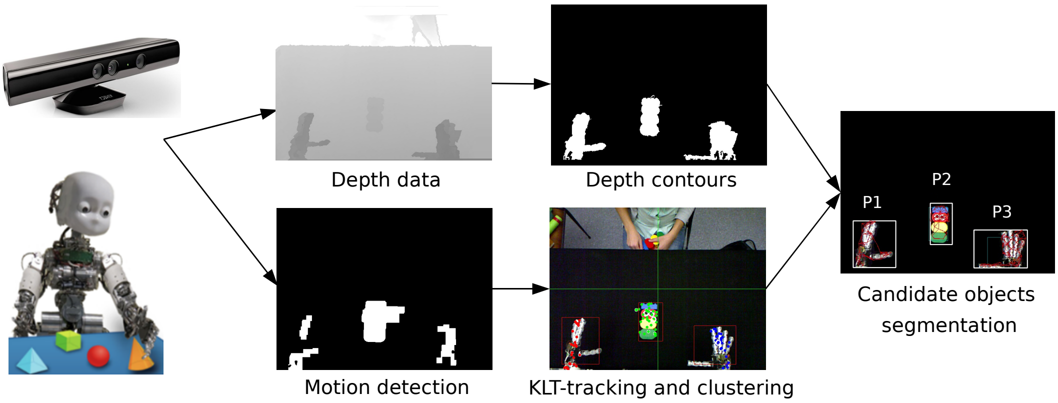

Figure 1. System Overview of the Incremental object discovery Software.

An early classification is done using each of the 3 features separately. For the color features a bag of words approach (from PAE) is used. For the shape feature, the minimum squared distance between the object's histogram and that of all objects in the database is calculated. Classification scores are then fused by a feed-forward neural network to get the final result [81].

5.1.4. PEDDETECT: GPU-accelerated person detection demo

Participant: Alexander Gepperth [correspondant].

PEDDETECT implements real-time person detection in indoor or outdoor environments. It can grab image data directly from one or several USB cameras, as well as from pre-recorded video streams. It detects mulitple persons in 800x600 color images at frame rates of >15Hz, depending on available GPU power. In addition, it also classifies the pose of detected persons in one of the four categories "seen from the front", "seen from the back", "facing left" and "facing right". The software makes use of advanced feature computation and nonlinear SVM techniques which are accelerated using the CUDA interface to GPU programming to achieve high frame rates. It was developed in the context of an ongoing collaboration with Honda Research Institute USA, Inc.

5.1.5. A Python OptiTrack client

Participant: Pierre Rouanet [correspondant].

This python library allows you to connect to an OptiTrack from NaturalPoint (http://www.naturalpoint.com/optitrack/). This camera permits the tracking of 3D markers efficiently and robustly. With this library, you can connect to the Motive software used by the OptiTrack and retrieve the 3D position and orientation of all your tracked markers directly from python.

5.2. Datasets

5.2.1. Choreography dataset 1 and 2

Participants: Olivier Mangin [correspondant], Haylee Fogg.

These databases contain choreography motions recorded through a kinect device. In the first dataset, these motions have a combinatorial structure: from a given set of primitive dance motions, choreographies are constructed as simultaneous execution of some of these primitive motions. Primitive dance motions are chosen from a total set of 48 motions and are spanned over one or two limbs, either the legs (e.g. walk, squat), left or right arm (e.g. wave hand, punch) or both arms (e.g. clap in hands, paddle). Complex choreographies are produced as the simultaneous demonstration of two or three of these primitive motion: either one for legs and one for both arm, or one for legs and one for each arm. The dataset has been used in the experiments from [104] for studying learning techniques allowing to identify dictionaries of motion primitives, and is publicly available at https://flowers.inria.fr/choreography_database.html.

The second dataset only contains choreographies composed of a single motion. It contains 110 records of each gesture from a set of 10 simple gestures and was used in the experiments from [53]. The dataset is publicly available at https://flowers.inria.fr/choreo2.

5.2.2. Development on NoFish Platform

Participants: Paul Fudal [correspondant], Sao Mai Nguyen.

NoFish (fig. 2) platform is a setup used by Mai Nguyen to perform severals experiments following her PhD work on social learning and intrinsic motivation. The setup consists to an ErgoRobot (fig. 3) with a fishing rod attached to the tip and a cap made with a red juggling ball. The robot is plugged on an ethernet power-switch used to turn it off if something wrong happens with the robot (for example : the robot try to go in a position he cannot reach and its motors are forcing too much). Tracking the cap is made with Full HD camera on the ceiling. At last, a video-projector prints informations on the floor (fig. 4) which helps for interactions between humans and the robots.

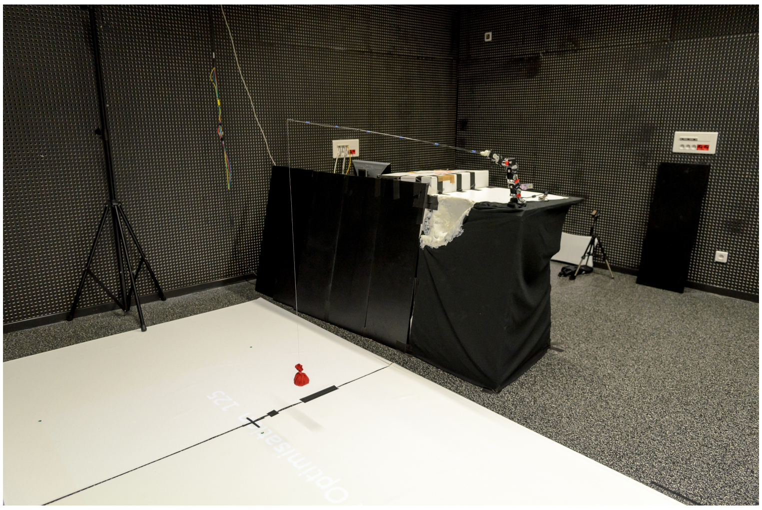

Figure 2. Illustration of the NoFishPlatform

Controlling the robot is made using URBI [1] scripts allowing different control level from single motor control to pre-programmed primitive (for example : reseting the robot or make it going to its starting position). URBI is also used to perform action through the power-switch if the robot must be turned off.

Tracking the cap is made by a program written using OpenCV and keep up to date the cap's coordinates by sending them to URBI through a network socket.

[1] http://www.gostai.com/products/jazz/urbi/

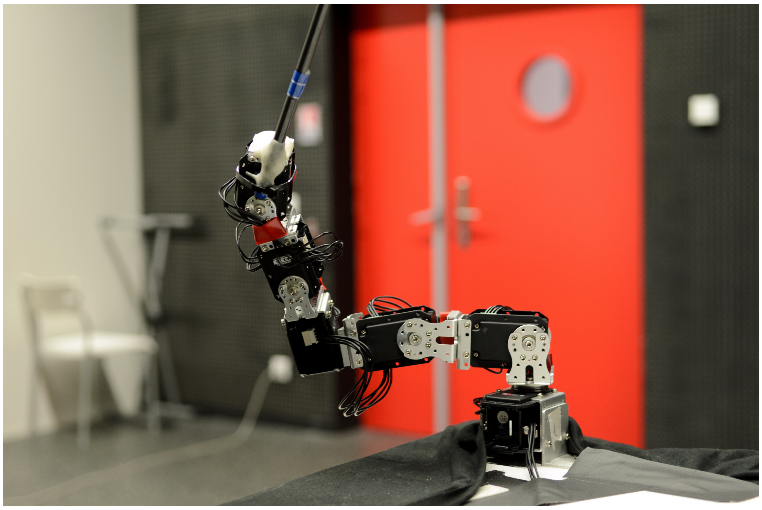

Figure 3. Illustration of NoFish ErgoRobot

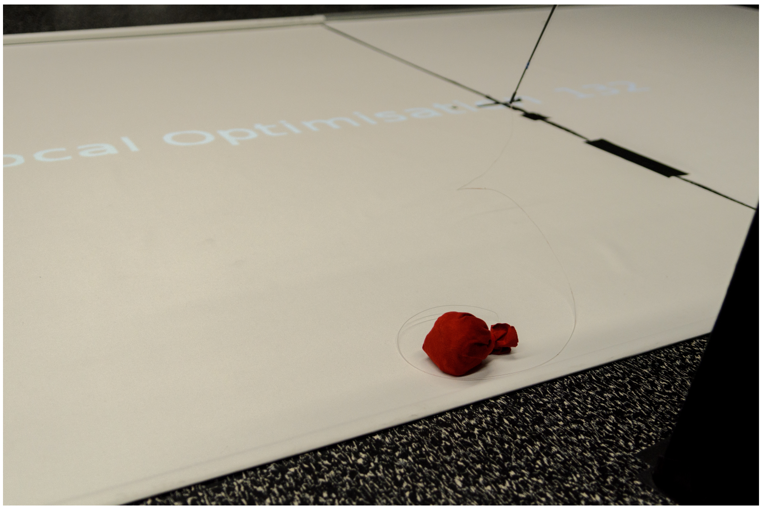

Figure 4. NoFish cap and informations printed on the floor

An other program written in JAVA and Processing [2] allows to print informations on the floor; it consists of a server receiving through a network socket texts and shapes informations to print on the floor and a JAVA object which can be used with compatible software. This component where really useful for giving simple and direct information during an experiment and also during interaction between a human and the robot for socially guided experiments.

The main program is written in Matlab and includes all the explained previous components and softwares. It also includes a forward kinematic calculation module used to ensure a movement sent can be safely be played by the robot to avoid it breaking himself; this module gives step by step informations of what will happen when the robot will play the movement like, for example, if the fishing cap will touch the floor at the end or not which permits to keep the robot safe and speed up experiments by ignoring dangerous moves for the robot and useless ones for the deployed algorithm.

This setup were used by Sao Mai Nguyen to run experiments during her thesis on social learning and intrinsic motivation

5.3. Learning algorithms

5.3.1. KidLearn

Participants: Manuel Lopes [correspondant], Benjamin Clement, Pierre-Yves Oudeyer, Didier Roy.

The KidLearn software provides and Intelligent Tutoring System that optimizes teaching sequences based on the estimated level of each particular student [65]. We implemented a Game of Money that allows students, ages 7-8, to learn how to use money. It includes 3 main components: i) a webserver that handles the requests and stores the experiments in a databased; ii) a GUI that provides the interface for the game; and iii) the optimization software.

Graphical interfaces in ITS can have unwanted side effects. For this reason, the interface was entirely designed with the help of a didactician, with several specific design choices motivated by pedagogical, motivational and attention requirements. For example, the interface, shown in Figure 5. is such that:

- display is as clear and simple as possible;
- there is no chronometer, so that students are not put under time pressure;
- coins and banknotes have realistic visual appearance, and their relative sizes are respected;
- display of prices use visual encodings commonly used in shops;
- the zone for receiving money is automatically cleared in case of error after the student submits it;
- automatic snapping of money and tokens icons in the reception zone, and automatic visual arrangement;
- text quantity is kept to minimum;

Four principal regions are defined in the graphical interface, as shown in Figure 5. The first is the wallet location where users can pick and drag the money to drop them on the repository location to make for the correct price. The object and the price are present in the object location, where the price can the written and/or spoken depending on the parameterization of the activity. The information location is using to display information for the learners such as extra clues when they make a mistake (for which they have to press the light bulb) and feedback. In order to improve the pedagogical success of the activity, the correct solution is presented automatically to the students if they fail to compose the correct price after 3 trials.

5.3.2. RLPark - Reinforcement Learning Algorithms in JAVA

Participant: Thomas Degris [correspondant].

[2] http://processing.org

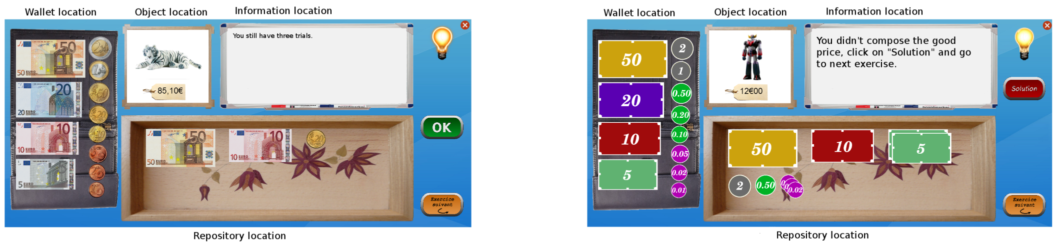

Figure 5. Example interface with four principal regions of interaction: wallet location, repository, object and price, and feedback location.

RLPark is a reinforcement learning framework in Java. RLPark includes learning algorithms, state representations, reinforcement learning architectures, standard benchmark problems, communication interfaces for three robots, a framework for running experiments on clusters, and real-time visualization using Zephyr. More precisely, RLPark includes:

- Online Learning Algorithms: Sarsa, Expected Sarsa, Q-Learning, On-policy and off-policy Actor-Critic with normal distribution (continuous actions) and Boltzmann distribution (discrete action), average reward actor-critic, TD, TD(λ), GTD(λ), GQ(λ), TDC
- State Representations: tile coding (with no hashing, hashing and hashing with mumur2), Linear Threshold Unit, observation history, feature normalization, radial basis functions
- Interface with Robots: the Critterbot, iRobot Create, Nao, Puppy, Dynamixel motors
- Benchmark Problems: mountain car, swing-up pendulum, random walk, continuous grid world

An example of RLpark running an online learning experiment on a reinforcement learning benchmark problem is shown in Figure 6.

RLPark was started in spring 2009 in the RLAI group at the university of Alberta (Canada) when Thomas Degris was a postdoc in this group. RLPark is still actively used by RLAI. Collaborators and users include Adam White, Joseph Modayil and Patrick Pilarski (testing) from the University of Alberta.

RLPark has been used by Richard Sutton, a professor and iCORE chair in the department of computing science at the University of Alberta, for a demo in his invited talk *Learning About Sensorimotor Data* at the Neural Information Processing Systems (NIPS) 2011 [3]. Patrick Pilarski used RLPark for live demos on television (Breakfast Television Edmonton, CityTV, June 5th, 2012) and at TEDx Edmonton on Intelligent Artificial Limbs[4]. So far, RLPark has been used in more than a dozens of publications (see http://rlpark.github.com/publications.html for a list).

RLPark has been ported to C++ by Saminda Abeyruwan, a student of the University of Miami (United States of America). The Horde architecture in RLPark has been optimized for GPU by Clément Gehring, a student of the McGill University in Montreal (Canada).

Future developments include the implementation of additional algorithms (the Dyna architecture, back propagation in neural networks, ...). A paper is under review for the JMLR Machine Learning Open Source Software. Documentation and tutorials are included on the RLPark web site [5]. RLPark is licensed under the open source Eclipse Public License.

[3] http://webdocs.cs.ualberta.ca/~sutton/Talks/Talks.html#sensorimotor
[4] http://www.youtube.com/watch?v=YPc-Ae7zqSo
[5] http://rlpark.github.com

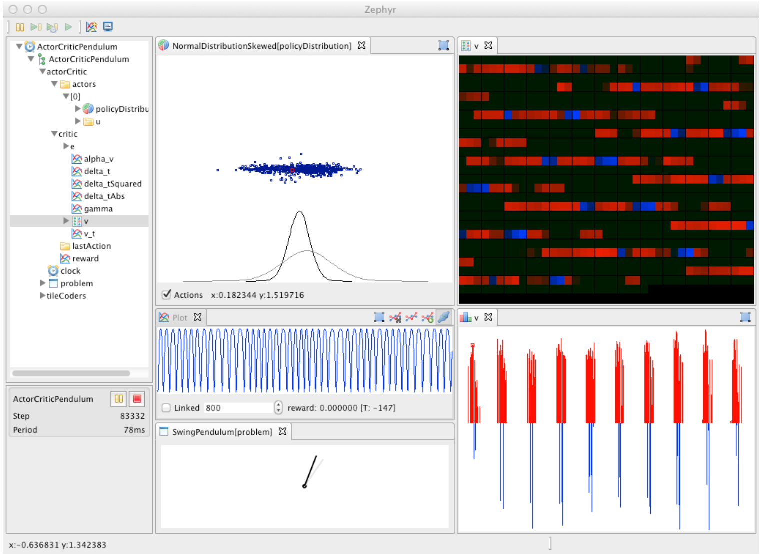

Figure 6. An example of an experiment in RLPark. Zephyr displays two views of a learned weight vector, an animation of the problem, the current policy distribution learned by the algorithm and the reward obtained by the algorithm. Videos are available at: http://rlpark.github.com.

5.3.3. DMP-BBO Matlab library

Participant: Freek Stulp [correspondant].

The dmp_bbo (Black-Box Optimization for Dynamic Movement Primitives) Matlab library is a direct consequence of the insight that black-box optimization outperforms reinforcement learning when using policies represented as Dynamic Movement Primitives. It implements several variants of the PI^{BB} algorithm for direct policy search. It is currently being used and extended by several FLOWERS members (Manuel Lopes, Clément Moulin-Frier) and external collaborators (Jonas Buchli, Hwangbo Jemin of ETH Zurich). This code was used for the following publications: [130], [127], [128].

5.3.4. Self-calibration BCI - Matlab library

Participants: Jonathan Grizou [correspondant], Iñaki Iturrate, Luis Montesano, Manuel Lopes, Pierre-Yves Oudeyer.

The Matlab software implements the algorithms described in [45]. It allows a robot to be instructed a new task by a human using communicative signals initially totally unknown to the robot. It is currently extended and improved in the context of EEG-based brain-machine interfaces (BMIs) [44].

It results in a BCI based control of sequential tasks with feedback signals that do not require any calibration process. As a by-product, the method provides an unsupervised way to train a decoder with the same performance of state-of-the-art supervised classifiers, while keeping the system operational and solving, with a lower performance during the first steps, the unknown task. The algorithm has been tested with online experiments (fig. 7), showing that the users were able to guide from scratch an agent to a desired position.

To improve the efficiency of the algorithm, we introduced a new planning method that uses the uncertainty in the signal-target estimation. This planner is inspired by exploration methods with exploration bonuses that allow guiding to reduce the uncertainty in an efficient way. We showed that trying to follow the best hypothesis does not explore the space significantly to reduce uncertainty and thus identify the correct task. Only through

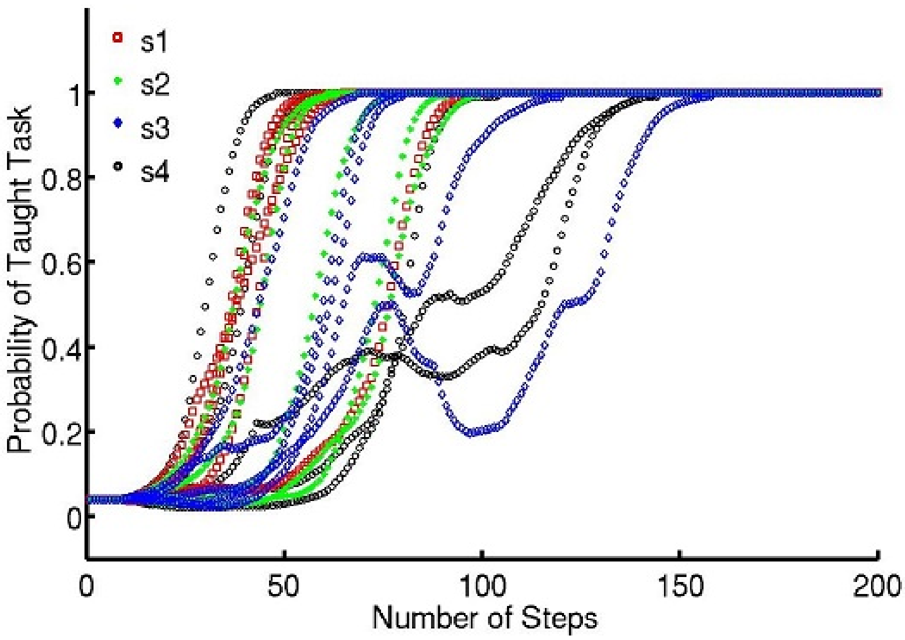

Figure 7. Results from the online BCI experiment for identifying the task. Evolution of the probability of the taught task for each subject and run

an approach that plans how to reduce the uncertainty multiple steps ahead are we sure that the agent will reach states that can only be explained by the correct hypothesis.

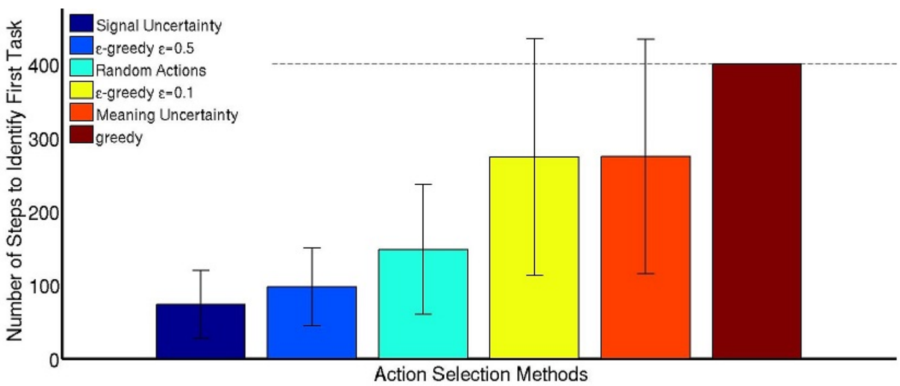

Figure 8. Comparison between different exploration methods. Planning wrt. uncertainty in noth task and signal space is the most efficient method

5.3.5. PROPRE: simulation of developmental concept formation using PYTHON

Participant: Alexander Gepperth [correspondant].

This simulation software implements the algorithms described in [86], [83]. It is available online under the URL www.gepperth.net/downloads.html. The simulation is implemented in PYTHON for easy use, yet the time-critical core functions are written in C.

5.3.6. pyStreamPlayer: synchronized replay of multiple sensor recordings and supplementary data

Participant: Alexander Gepperth [correspondant].

This Python software is intended to facilitate the application of machine learning algorithms by avoiding to work directly with an embodied agent but instead with data recorded in such an agent. Assuming that non-synchronous data from multiple sensors (e.g., camera, Kinect, laser etc.) have been recorded according to a flexible format defined by the pyStreamPlayer architecture, pyStreamPlayer can replay these data while retaining the exact temporal relations between different sensor measurements. As long as the current task does not involve the generation of actions, this software allows to process sensor data as if it was coming from an agent which is usually considerably easier. At the same time, pyStreamPlayer allows to replay arbitrary supplementary information such as, e.g., object information, as if it was coming from a sensor. In this way, supervision information can be stored and accessed together with sensory measurements using an unified interface. pyStreamPlayer has been used to facilitate real-world object recognition tasks, and several of the major databases in this field (CalTech Pedestrian database, HRI RoadTraffic traffic objects database, CVC person database, KITTI traffic objects database) have been converted to the pyStreamPlaer format and now serve as a source of training and test data for learning algorithms.

pyStreamPlayer has been integrated into a ROS node as well, allowing th replay and transmission across networks of distributed processes.

5.3.7. Multimodal: framework around the NMF algorithm for multimodal learning

Participant: Olivier Mangin [correspondant].

The pyhton code provides a minimum set of tools and associated libraries to reproduce the experiments on [53], together with the choreography datasets. The code, publicly available at https://github.com/omangin/multimodal, under the new BSD license, is primarily intended for reproduction of the mulimodal learning experiment mentioned above. It is also expected that the public availability of the code encourages further experimentation by other scientists with data coming from other domains, thus increasing both the impact of the aforementioned publication and the knowledge on the algorithm behaviors. The nonnegative matrix factorization algorithm used in the experiments will also soon be included as a third party project to http://scikit-learn.org. Finally the code is currenlty being used by other members of the team and is expected to play an important role in further collaborations.

5.3.8. Tools for curiosity-driven learning on a robotic arm

Participants: Pierre Rouanet [correspondant], Clément Moulin-Frier.

This library is intended to provide tools for experimenting how curiosity model can facilitate learning of complex tasks such as manipulating objects. First, it provides high-level access to a robotic arm made of dynamixel motors: forward and inverse kinematics, demonstrations recording. Then it wraps the IMLE [77] library which we used for incremental and online learning of the sensorimotor mappings of the robot. Finally, it implements curiosity-driven learning based on the maximization of the learning progress. This modelling is based on recent works by Moulin-Frier and Oudeyer [54], [56], [55] proposing a probabilistic algorithmic architecture unifying various principles of developmental robotics such as motor babbling, goal babbling and curiosity-driven exploration. This architecture has already been successfully applied to model infant speech acquisition in the previously cited papers.

5.4. Software Platforms

5.4.1. Meka robot plateform enhancement and maintenance

Participants: Antoine Hoarau [ADT Engineer Since Nov. 2012], Freek Stulp [Supervisor], David Filliat [Supervisor].

Autonomous human-centered robots, for instance robots that assist people with disabilities, must be able to physically manipulate their environment. There is therefore a strong interest within the FLOWERS team to apply the developmental approach to robotics in particular to the acquisition of sophisticated skills for manipulation and perception. ENSTA-ParisTech has recently acquired a Meka (cf. 9) humanoid robot dedicated to human-robot interaction, and which is perfectly fitted to this research. The goal of this project is to install state-of-the-art software architecture and libraries for perception and control on the Meka robot, so that this robot can be jointly used by FLOWERS and ENSTA. In particular, we want to provide the robot with an initial set of manipulation skills.
The goal is to develop a set of demos, which demonstrate the capabilities of the Meka, and provide a basis on which researchers can start their experiments.

The platform is evolving as the sofware (Ubuntu, ROS, our code) is constantly updated and requires some maintenance so less is needed for later. A few demos were added, as the hand shaking demo, in which the robot detects people via kinect and initiates a hand shake with facial expressions. This demo has been used to setup a bigger human robot interaction experiment, currently tested on subjects at Ensta (cf. 10). The stacking cups demo is more of a manipulation and vision demo : the robot detects its cups by their shape and color, and tries to make a tower with it (cf. 12). This demo has required to manually update the old pr2 tabletop object detector to the new ROS version, create a tool to semi-automatically calibrate the extrinsicts parameters of the kinect, and different set of tools to catch objects from different angles (waiting for moveit to be fully integrated). Finally, we've seen that the robot itself also needs some maintenance; some components broke (a finger tendon), a welding got cold (in the arm) and a few cables experienced fatigue (led matrix and cameras) (cf. 11).
I've also given a talk at Humanoid 2013 at Altanta and the University of Texas at Austin for my participation on [61].

5.4.2. Experiment platform for multiparameters simulations

Participants: Fabien Benureau, Paul Fudal.

Simulations in robotics have many shortcommings. At the same time, they offer high customizability, rapidity of deployment, abscence of failure, consistency across time and scalability. In the context of the PhD work of Fabien Benureau, it was decided to investigate hypothesis first in simulation before moving to real hardware. In order to be able to test a high number of different hypothesis, we developped a software platform that would scale to the computing ressource available.

We designed simple continuous simulations around a of-the-shelf 2D physic engine and wrote a highly modular platform that would automatically deploy experiments on cluster environments, with proper handling of dependencies; our work investigate transfer learning, and some experiments's input data is dependent of the results of another.

So far, this platform and the university cluster has allowed to conduct thousands of simulations in parallel, totaling more than 10 years of simulation time. It has led us to present many diverse experiments in our published work [40], each repeated numerous times. It has allowed us to conduct a multi-parameter analysis on the setup, which led to new insights, which are being presented in a journal article to be submitted in the beginning of this year.

Because of its high modularity, this platform is proving to be highly flexible. We are currently adaptating it to a modified, cluster-ready, version of the V-REP simulator. Those simulations will serve to back ones on similar real-world hardware that are currently setup.

We have released the platform and the complete experiments code when we published the results of [40], allowing to reproduce the results of the paper, and will continue to do so with each published work.

5.4.3. PyPot

Participants: Pierre Rouanet [correspondant], Matthieu Lapeyre.

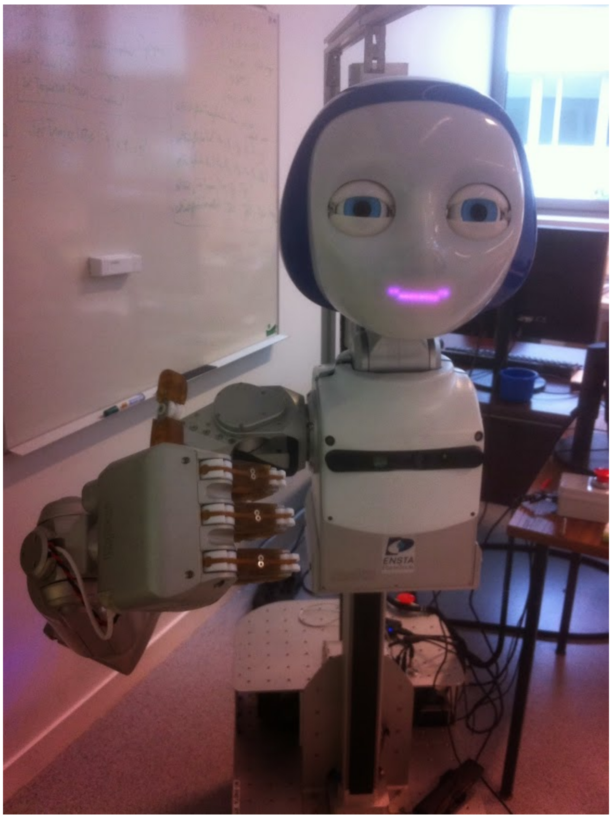

Figure 9. The Meka robot plateform acquired by ENSTA ParisTech

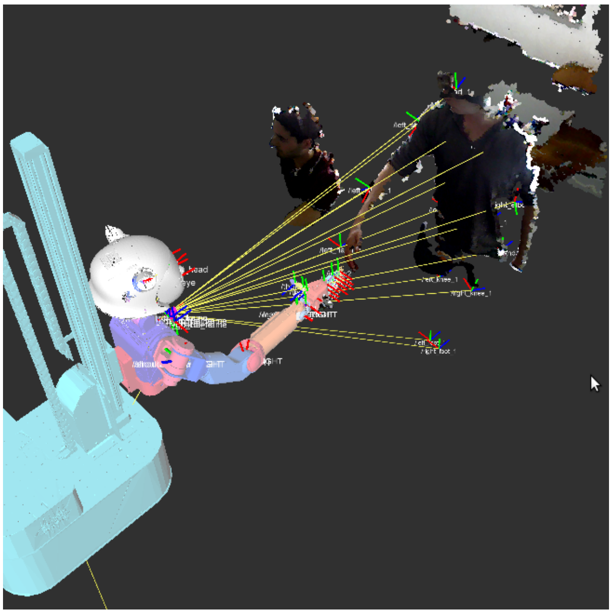

Figure 10. Hand shake demo visualized on Rviz (ROS)

PyPot is a framework developed to make it easy and fast to control custom robots based on dynamixel motors. This framework provides different levels of abstraction corresponding to different types of use. More precisely, you can use PyPot to:

1. directly control robotis motors through a USB2serial device,
2. define the structure of your particular robot and control it through high-level commands,
3. define primitives and easily combine them to create complex behavior.

PyPot has been entirely written in Python to allow for fast development, easy deployment and quick scripting by non-necessary expert developers. It can also benefits from the scientific and machine learning libraries existing in Python. The serial communication is handled through the standard library and thus allows for rather high performance (10ms sensorimotor loop). It is crossed-platform and has been tested on Linux, Windows and Mac OS.

PyPot is part of the Poppy project (http://www.poppy-project.org) and has been released under an open source license GPL V3. More details are available on PyPot website: https://github.com/poppy-project/pypot

5.5. Experimental Setups

5.5.1. Experimental Platform for User Study of Curiosity-driven Exploration

Participants: Pierre Rouanet [correspondant], Jonathan Grizou, Brice Miard, Julie Golliot.

This platform has been developed to investigate curiosity-driven behaviors and more precisely how humans explore new sensori-motor spaces. It consists in several simple games where users control a 2D/3D shape with the movements of their body. They have to discover the mapping between their movements and a shape displayed on the screen and learn how to make the controlled shape match the target one (fig 13).

The software is entirelly written in Python. It includes a Kinect wrapper allowing the access of 3D position of tracked skeleton joints. It provides a framework for creating new games based on the 2D drawing library (pygame). It also includes a web server used to display game instructions, cut-scene videos and questionnaire.

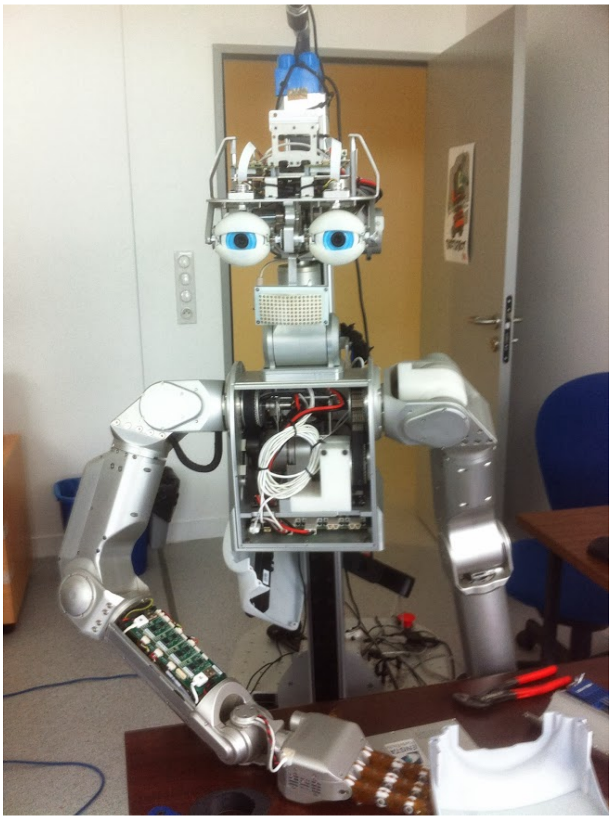

Figure 11. Maintenance is required on the robot

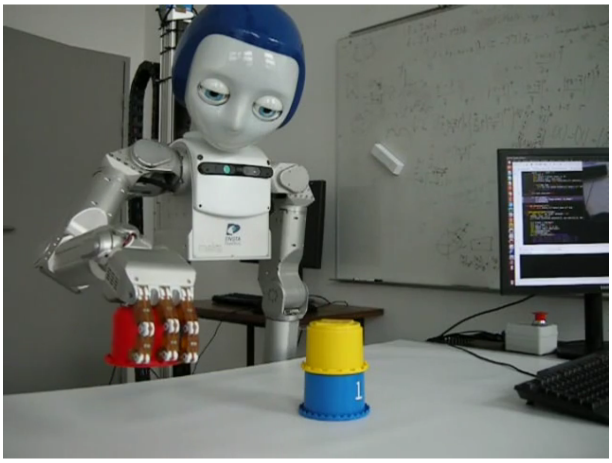

Figure 12. The stacking cup demo (video available at http://cogrob.ensta-paristech.fr)

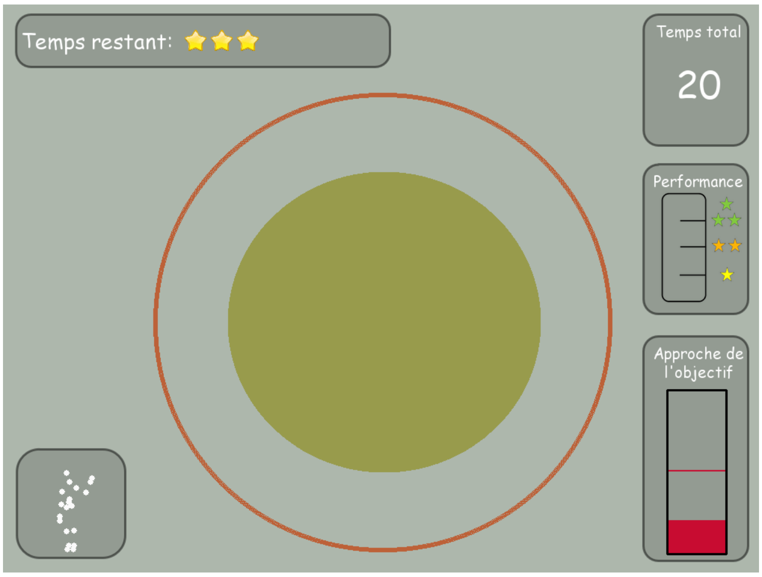

Figure 13. A screenshot representing the game interface as seen by the user.

5.6. Visualization Tools

5.6.1. Zephyr - Realtime Visualization in JAVA

Participant: Thomas Degris [correspondant].

Zephyr is a software to visualize numeric variables and data structure in real time and at different time scale. Zephyr is practical because it requires only minimal changes in the code: it uses Java reflexivity to automatically detect variables in the code to monitor and data structure with an associated dedicated view. Zephyr can easily be extended with new plugins because it is based on the popular Eclipse Rich Client Platform. Consequently, Zephyr takes advantage of an already existing and fully operational Eclipse plugins for many of its functionalities. Finally, Zephyr is distributed with a Java python virtual machine named Jython and a lisp implementation named Clojure. An example of a Zephyr screen is shown in Figure 14.

Zephyr was started in fall 2009 in the RLAI group at the university of Alberta (Canada) when Thomas Degris was a postdoc in this group. Zephyr is still actively used by RLAI. Users include Adam White, Joseph Modayil and Patrick Pilarski from the University of Alberta. Zephyr has been registered on the Eclipse marketplace since October 2011. Documentation about Zephyr is included on its website: http://zephyrplugins.github.com. Zephyr is licensed under the open source Eclipse Public License.

5.6.2. Experimental Setups for User Study of Alignment in Asymmetric Interactions

Participants: Jonathan Grizou [correspondant], Chloé Rozenbaum, Manuel Lopes, Katharina Rohlfing, Pierre-Yves Oudeyer.

This platform has been developed to investigate alignement in asymmetric interactions. We consider a remote construction task, where one user (user A) knows what to build but do not have access to the construction site while its partner (user B) is at the site but do not know what to do. By constraining the communicative channel between the two partners, we study how, and if, they will agree on a similar set of signals to convey information and what type of information they tend to produce.

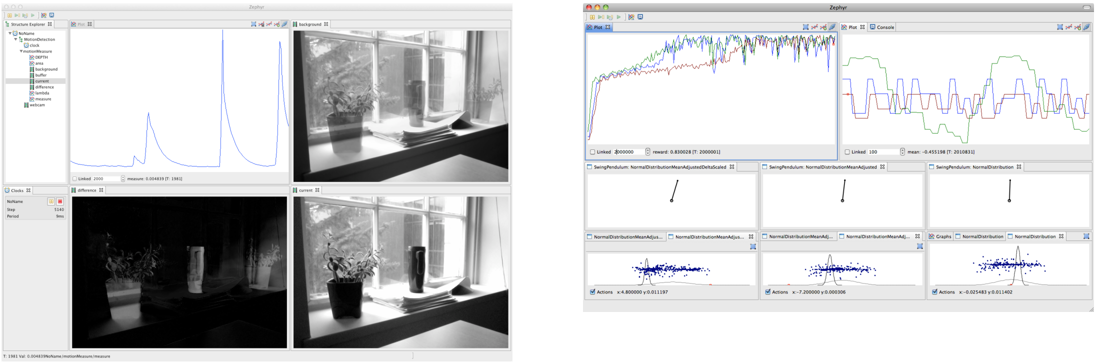

Figure 14. Left: Zephyr showing the different steps of a video processing pipeline in real-time. Right: Zephyr showing different data structure and variables of a reinforcement learning agent at different time scale. A video is available at: http://zephyrplugins.github.com.

The experimental setup consist of box with button, a video recording system and two screens. User A can send signals to user B by pressing buttons (fig. 15). Signals are displayed on a screen (fig. 15) at user B side. User A is not aware of what is displayed on user B screen, neither user B is aware of the relation between button presses and screen events. The video of user B construction scene is streamed to a screen at user B side.

The task consist of bulding arbitrary construction (fig. 15) using colored toy bricks (fig. 15).

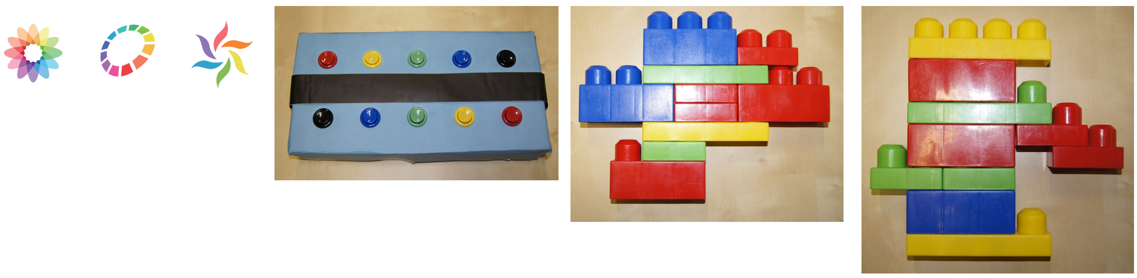

Figure 15. Three examples of sign displayed on the learner screen; The box and the button use as an interface for the teacher to communicate with the learner; Examples of construction presented to the teacher.

5.7. Hardware

5.7.1. Poppy Platform

Participants: Matthieu Lapeyre [correspondant], Pierre Rouanet, Jonathan Grizou, Pierre-Yves Oudeyer [supervisor].

5.7.1.1. Main goals :

No current platform (Nao [87], Darwin Op [88], Nimbro Op [126], HRP-2, ...) does offer both a adapted morphology in the sense of allowing physical interaction (safe, compliant, playful) and optimized for walking. So to explore these challenges we have decided to build a new bio-inspired humanoid robotic platform, called

Poppy, which provides some of the software and hardware features needed to explore both social interaction and biped locomotion for personal robot. It presents the following main features to make it an interesting platform to study how the combination of morphology and social interaction can help the learning:

- Design inspired from the study of the anatomy of the human body and its bio-mechanic
- Dynamic and reactive: we try to keep the weight of the robot as low as possible (geometry of the pieces and smaller motors)
- Social interaction: screen for communication and permits physical interaction thanks to compliance
- Study of the morphology of the leg to improve the biped walking
- Practical platform: low cost, ease of use and easy to reproduce

5.7.1.2. Overview :

Poppy platform (Figure 16) is a humanoid, it is 84cm tall for 3 kg. It has a large sensor motors space including 25 dynamical motors (MX-28 and AX-12), force sensors under its feet and some extra sensors in the head: 2 HD-wide angle-cameras, stereo-micros and an inertial central unit (IMU 9DoF) plus a large LCD Screen (4 inch) for visual communication (e.g. emotions, instructions or debug). The mechanical parts were designed and optimized to be as light as possible while maintaining the necessary strength. For this, the choice of a lattice beam structure manufactured with 3Dprinting polyamide was used.

The poppy morphology is designed based on the actual human body. We have deeply studied the biomechanics of the human body and have extracted some interesting features for humanoid robotics. This inspiration is expressed in the whole structure (e.g. the limb proportions) and in particular in the trunk and legs.

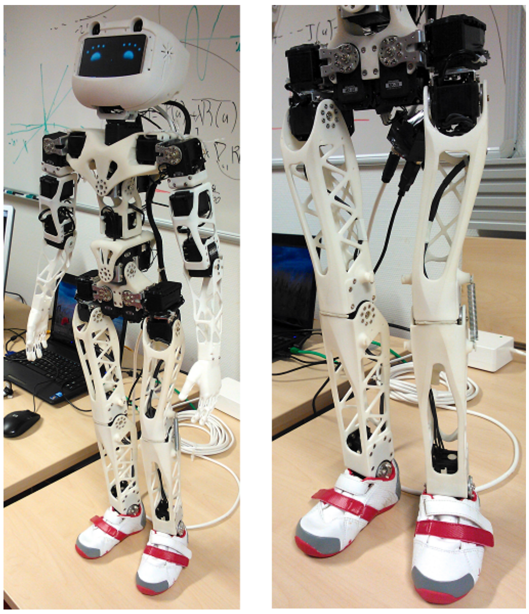

Figure 16. a. Global view of the Poppy platform. b. Zoom on legs design

Poppy uses the bio-inspired trunk system introduced by Acroban [101]. These five motors allow it to reproduce the main changes brought by the human spine. This feature allows the integration of more natural and fluid motion while improving the user experience during physical interactions. In addition, the spine plays a fundamental role in bipedal walking and postural balance by actively participating in the balancing of the robot.

The legs were designed to increase the stability and agility of the robot during the biped walking by combining bio-inspired, semi-passive, lightweight and mechanical-computation features. We will now describe two examples of this approach:

The architecture of the hips and thighs of Poppy uses biomechanical principles existing in humans. The human femur is actually slightly bent at an angle of about 6 degrees. In addition, the implantation of the femoral head in the hip is on the side. This results in a reduction of the lateral hip movement needed to move the center of gravity from one foot to another and a decrease in the lateral falling speed. In the case of Poppy, the inclination of its thighs by an angle of 6 degrees causes a gain of performance of more than 30% for the two above mentioned points.

Another example is Poppy's feet. Poppy has the particularity of having small feet compared to standard humanoids. It has humanly proportioned feet (ie about 15% of its total size).It is also equipped with compliant toes joints (see Figure 17.a). We believe that this feet involve two keys features to obtain a human-like and efficient walking gait. However, that raises problems regarding balance because the support polygon is reduced. We decided to add pressure sensors under each foot in order to get accurate feedback of the current state of the robot (see Figure 17.b).

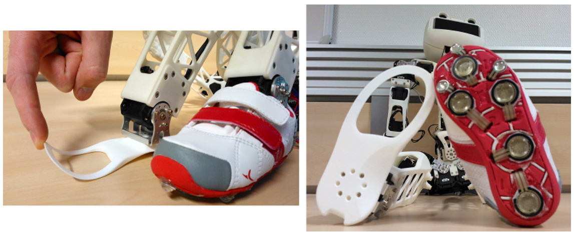

Figure 17. Poppy feet use actual children shoes combine with a compliant feet, toes (a.) and pressure sensors (b.)

5.7.1.3. Open source release :

To allow the distribution in the robotic community, we have decided to make Poppy an open platform. So the software and the hardware are open source. They have each a repository. The PyPot library is under GPLV3 license and is available on a bitbucket repository (https://github.com/poppy-project/pypot). The hardware, Solidworks files and STL needed to print the robot are available under a Creative Commons BY+SA+NC license on a private GrabCAD Workbench repository. People can request access to the hardware repository on the Poppy website (http://www.poppy-project.org/open-platform/).

The platform is currently under beta-testing meaning that we let the community grows little by little to ensure a good support of interesting projects.

For now, there are about 200 people on the GrabCAD project. They have access to all files needed to print the robot. There is also about 60 beta testers. They have access to a private section on the website with documentation and a forum for support.

Several of them are already doing a great work both by reporting bugs and managing to build the robot outside the lab. We are trying to work closely with them as they are a great source of feedback to improve the platform before a more wide distribution.

5.7.1.4. Impact in the community

Poppy has been release open source the 15/10/2013. The announcement has been done by Pierre Yves Oudeyer during the Lift 2013 conference. To prepare this event, the website and a overview video were made. The video, accessible here (http://vimeo.com/76917854) has reached about 40K views until now.

A part of the audience are technology-enthusiast people (about 50% on grabCAD), interested by the fact Poppy is 3D printable and so, highly customizable. More interesting, we received a large number of beta request from various applications domains (see 18) around the world (see 19).

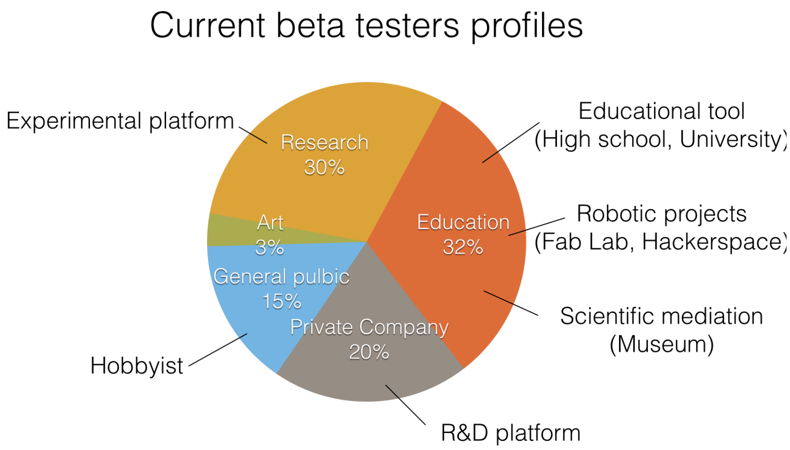

Figure 18. The 60 selected beta testers profiles chart.

5.7.1.4.1. Education:

We receive a lot of request from Fablab around the world (France, USA, New-Zeland, South Africa,...). All have great educational projects for teaching robotics, mechanics and computer science to children. Some of them are close beta testers and we have weekly interactions. In the same topic, several schools, engineering schools and universities showed interest to use Poppy as experimental support. One of our close beta tester is a Bordeaux high school. They are planning to use Poppy as support for mechanics, computer sciences but also architecture or philosophy.

5.7.1.4.2. Art:

A current art project is under construction. A residence with an artist, a dancer and us will take place on the 24/02/2014 to the 05/03/2014. It comes to artistically invest humanoid robotics and thereby examine the relationship of the body to the digital world. The encounter between art and science generates potential new ideas for both disciplines who find themselves at the crossroads of questions relating to the gesture, movement and body. The Poppy Project is focused on morphological and motor aspects in a context human-robot interface. The look of an artist and the movements of a dancer are testing this interface with an unprecedented and direct manere.

5.7.1.4.3. Research:

A large number of researchers showed interest in the platform. Most of them are interested to use it as a experimental tool. They want to address challenges such as balance and walking control, use of force controlled motors, explore human-robot interactions. On this last topic, the Bristol robotic lab target to use tele-operation to investigate what factors are important in terms of appearance and behaviour for credible and trustworthy interaction. For this purpose, they will develop functional hands for gesture and grasping task.

5.7.1.5. Next step:

We are currently working with beta testers on several improvements to make Poppy more accessible, more easy to use and more polyvalent. Two internship students will arrive on the project to work on the embedded electronic and on the feet design. We are targeting to release the final version during the summer 2014. Also we are thinking about the creation of a "Poppy pack" including all necessary components and tools to easily build the robot.

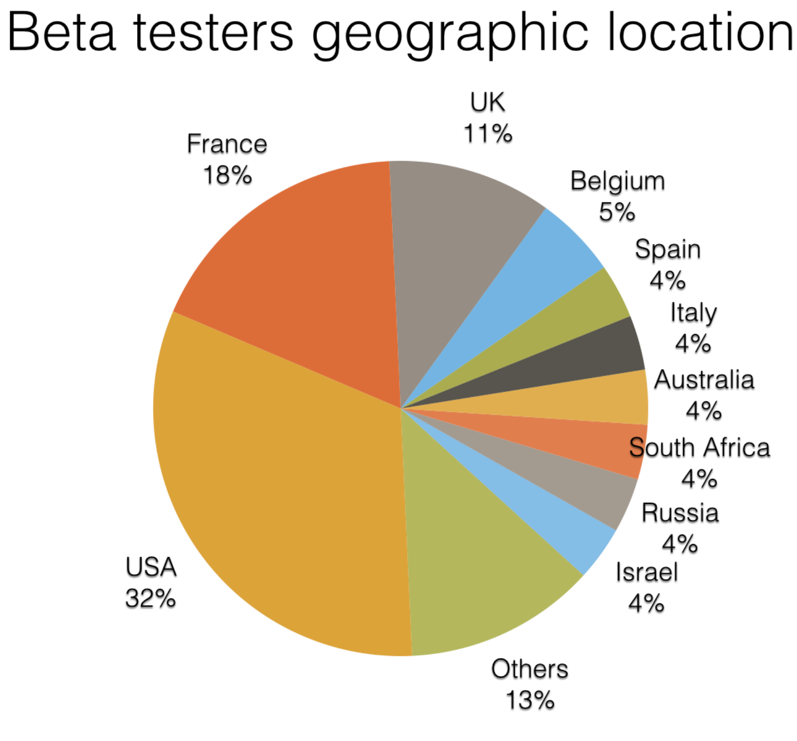

Figure 19. Where are located the beta testers?

Another very important point is the community management. We are currently adding new collaborative tools on the Poppy website. These tools are tested with beta testers. The challenge is to offer the good tools to provide an efficient support to future users and to encourage people to contribute. This work is done in collaboration with Stephane Ribas (D2T inria Grenoble).

6. New Results

6.1. Robotic and Computational Models of Human Development

6.1.1. Computational models of information-seeking, curiosity and attention

Participants: Pierre-Yves Oudeyer, Manuel Lopes.

An associated team, called Neurocuriosity, was created between Flowers and the Cognitive Neuroscience lab of Jacqueline Gottlieb at Univ. Columbia, NY. The goal of this associated team is to investigate mechanisms of spontaneous exploration and learning in humans by setting up experiments allowing to confirm or falsify predictions made by computational models previously developped by the team. This constitutes a crucial collaboration between developmental robotics and cognitive neuroscience. This joint work already led to a major publication on curiosity and information seeking, in the prestigious Trends in Cognitive Science journal (impact factor: 16.5). [27]
Abstract: 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.

6.1.1.1. Formalizing Imitation Learning

Participants: Thomas Cederborg, Pierre-Yves Oudeyer.

An original formalization of imitation learning was elaborated. Previous attempts to systematize imitation learning has been limited to categorizing different types of demonstrator goals (for example defining success in terms of the sequential joint positions of a dance, or in terms of environmental end states), and/or been limited to a smaller subset of imitation (such as learning from tele-operated demonstrations). The formalism proposed attempts to describe a large number of different types of learning algorithms using the same notation. Any type of algorithm that modifies a policy based on observations of a human, is treated as an interpretation hypothesis of this behavior. One example would be an update algorithm that updates a policy, partially based on the hypothesis that the demonstrator succeeds at demonstrations with probability 0.8, or an update algorithm that assumes that a scalar value is an accurate evaluation of an action compared to the latest seven actions. The formalism aims to give a principled way of updating these hypotheses, either rejecting some of a set of hypotheses regarding the same type of behavior, or set of parameters of an hypothesis. Any learning algorithm that modifies policy based on observations of a human that wants an agent to do something or act in some way, is describable as an interpretation hypothesis. If the learning algorithm is static, this simply corresponds to an hypothesis that is not updated based on observations. A journal article [26].

6.1.1.2. Self-Organization of Early Vocal Development in Infants and Machines: The Role of Intrinsic Motivation

Participants: Clément Moulin-Frier, Sao Mai Nguyen, Pierre-Yves Oudeyer.

We bridge the gap between two issues in infant development: vocal development and intrinsic motivation. We propose and experimentally test the hypothesis that general mechanisms of intrinsically motivated spontaneous exploration, also called curiosity-driven learning, can self-organize developmental stages during early vocal learning and explain several aspects observed in infants (Figure 20). We introduce a computational model of intrinsically motivated vocal exploration, which allows the learner to autonomously structure its own vocal experiments, and thus its own learning schedule, through a drive to maximize competence progress. This model relies on a physical model of the vocal tract, the auditory system and the agent's motor control, as well as vocalizations of social peers. We present computational experiments that show how such a mechanism can explain the adaptive transition from vocal self-exploration with little influence from the speech environment, to a later stage where vocal exploration becomes influenced by vocalizations of peers (Figure 21). Within the initial self-exploration phase, we show that a sequence of vocal production stages self-organizes, and shares properties with data from infant developmental psychology: the vocal learner first discovers how to control phonation, then focuses on vocal variations of unarticulated sounds, and finally automatically discovers and focuses on babbling with articulated proto-syllables (Figure 22). As the vocal learner becomes more proficient at producing complex sounds, imitating vocalizations of peers starts to provide high learning progress explaining an automatic shift from self-exploration to vocal imitation.

This work has been recently accepted in the journal *Frontiers in Psychology, Cognitive Science* [30].

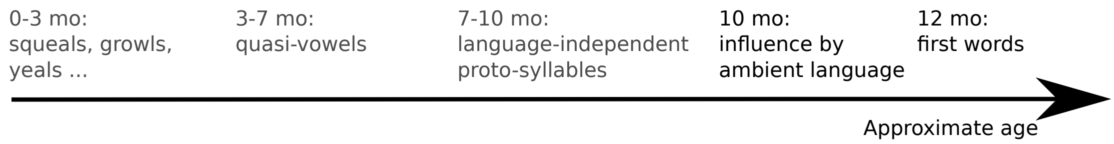

Figure 20. Rapid view of the first year of infant vocal development.

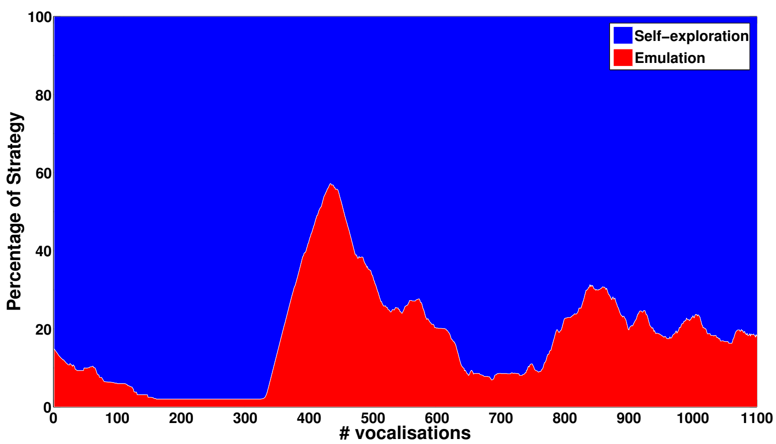

Figure 21. Our model displays an adaptive transition from vocal self-exploration with little influence from the speech environment, to a later stage where vocal exploration becomes influenced by vocalizations of peers.

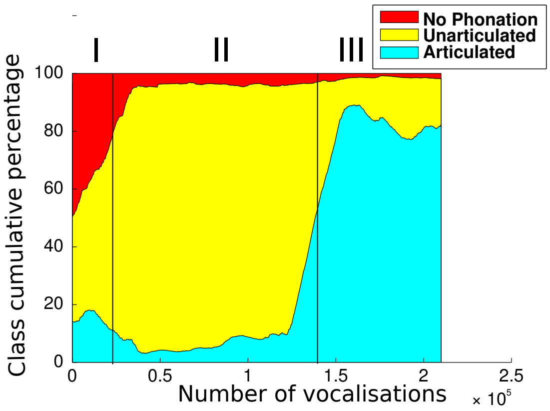

Figure 22. Within the self-exploration phase, our model first discovers how to control phonation, then focuses on vocal variations of unarticulated sounds, and finally automatically discovers and focuses on babbling with articulated proto-syllables.

6.1.1.3. Emergent Proximo-Distal Motor Development through Adaptive Exploration, applied to Reaching and Vocal Learning

Participants: Freek Stulp, Pierre-Yves Oudeyer, Jules Brochard, Clément Moulin-Frier.

Life-long robot learning in the high-dimensional real world requires guided and structured exploration mechanisms. In this developmental context, we have investigated the use of the PI^2-CMAES episodic reinforcement learning algorithm, which is able to learn high-dimensional motor tasks through adaptive control of exploration. By studying PI^2-CMAES in a reaching task on a simulated arm, we observe two developmental properties. First, we show how PI^2-CMAES autonomously and continuously tunes the global exploration/exploitation trade-off, allowing it to re-adapt to changing tasks. Second, we show how PI^2-CMAES spontaneously self-organizes a maturational structure whilst exploring the degrees-of-freedom (DOFs) of the motor space. In particular, it automatically demonstrates the so-called *proximo-distal maturation* observed in humans: after first freezing distal DOFs while exploring predominantly the most proximal DOF, it progressively frees exploration in DOFs along the proximo-distal body axis. These emergent properties suggest the use of PI^2-CMAES as a general tool for studying reinforcement learning of skills in life-long developmental learning contexts. This work was published in the Paladyn Journal of Behavioral Robotics [36].

This model of emergent developmental freezing and unfreezing of degrees of freedom was then applied to infant vocal development. For this aim, we used an articulatory synthesizer which is a computer model of the human vocal tract and the ear. While testing different possibilities, the algorithm eventually creates learning structures, which are more efficient that random motor babbling. Using the algorithm with a vocal synthesizer, we show that it can reproduce a babbling infant's characteristic: the predominance of the jaw over the other articulators, namely the canonical babbling.

This is the first study to our knowledge of emergent maturation in speech. Without presupposing any biological or social constraint, we give a new explanation of the jaw predominance in babbling, based on freezing and freeing the degrees of freedom in an adaptive maturation scheme to improve learning. This provides an original hypothesis regarding the emergence of canonical babbling in infant vocal development.

This last work was performed during the internship of Jules Brochard in 2013 and a journal article is currently being written.

6.1.2. COSMO ("Communicating about Objects using Sensory-Motor Operations"): a Bayesian modeling framework for studying speech communication and the emergence of phonological systems

Participants: Clément Moulin-Frier, Jean-Luc Schwartz, Julien Diard, Pierre Bessiã¨re.

This work began with the PhD thesis of Clement Moulin-Frier at GIPSA-Lab, Grenoble, France, supervised by Jean-Luc Schwartz (GIPSA-Lab, CNRS), Julien Diard (LPNC, CNRS) and Pierre Bessière (College de France, CNRS). A few papers were finalized during his post-doc at FLOWERS in 2012. Firstly, an international journal paper based on the PhD thesis work of Raphael Laurent (GIPSA-Lab), extending Moulin-Frier's model, was published [108], as well as a commentary in *Behavioral and Brain Sciences* [97]. Both these papers provide computational arguments based on a sensory-motor cognitive model to feed the age-old debate of motor vs. auditory theories of speech perception. Secondly, in another journal paper under the submission process, we attempt to derive some properties of phonological systems (the sound systems of human languages) from the mere properties of speech communication. We introduce a model of the cognitive architecture of a communicating agent, called COSMO (for "Communicating about Objects using Sensory-Motor Operations") that allows expressing in a probabilistic way the main theoretical trends found in the speech production and perception literature. This allows a computational comparison of these theoretical trends, helping to identify the conditions that favor the emergence of linguistic codes. We present realistic simulations of phonological system emergence showing that COSMO is able to predict the main regularities in vowel, stop consonant and syllable systems in human languages.

This work is currently under consideration as a target article for a special issue in an international journal. Pierre-Yves Oudeyer joined this process as a member of the editing committee.

6.1.3. Recognizing speech in a novel accent: the Motor Theory of Speech Perception reframed

Participants: Clément Moulin-Frier, Michael Arbib.

Clément Moulin-Frier engaged this work with Michael Arbib during his 6-month visit in 2009 at the USC Brain Project, University of Southern California, Los Angeles, USA, during his PhD thesis at Gipsa-Lab, Grenoble. He continues to write a journal article during his post-doc in the Flowers team in 2012-2013. This paper has been published recently in Biological Cybernetics [29], in which we offer a novel computational model of foreign-accented speech adaptation, together with a thorough analysis of its implications with respect to the motor theory of speech perception.

6.2. Life-Long Robot Learning and Development of Motor and Social Skills

6.2.1. Active Learning and Intrinsic Motivation

6.2.1.1. Active Learning of Inverse Models with Goal Babbling

Participants: Adrien Baranes, Pierre-Yves Oudeyer.

We have continued to elaborate and study our Self-Adaptive Goal Generation - Robust Intelligent Adaptive Curiosity (SAGG-RIAC) architecture as an intrinsically motivated goal exploration mechanism which allows active learning of inverse models in high-dimensional redundant robots. Based on active goal babbling, this allows a robot to efficiently and actively learn distributions of parameterized motor skills/policies that solve a corresponding distribution of parameterized tasks/goals. The architecture makes the robot sample actively novel parameterized tasks in the task space, based on a measure of competence progress, each of which triggers low-level goal-directed learning of the motor policy parameters that allow to solve it. For both learning and generalization, the system leverages regression techniques which allow to infer the motor policy parameters corresponding to a given novel parameterized task, and based on the previously learnt correspondences between policy and task parameters.

We have conducted experiments with high-dimensional continuous sensorimotor spaces in three different robotic setups: 1) learning the inverse kinematics in a highly-redundant robotic arm, 2) learning omnidirectional locomotion with motor primitives in a quadruped robot 2324, 3) an arm learning to control a fishing rod with a flexible wire. We show that 1) exploration in the task space can be a lot faster than exploration in the actuator space for learning inverse models in redundant robots; 2) selecting goals maximizing competence progress creates developmental trajectories driving the robot to progressively focus on tasks of increasing complexity and is statistically significantly more efficient than selecting tasks randomly, as well as more efficient than different standard active motor babbling methods; 3) this architecture allows the robot to actively discover which parts of its task space it can learn to reach and which part it cannot.
This work was published in the journal Robotics and Autonomous Systems [25].

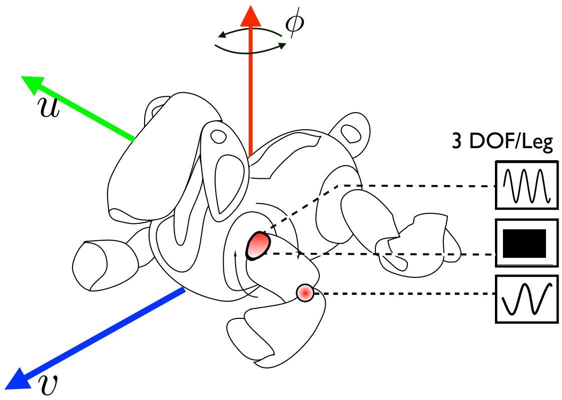

Figure 23. Experimenting SAGG-RIAC for learning an inverse model for omnidirectional locomotion of a quadruped robot. The quadruped robot is controlled using 24 dimensional motor synergies parameterized with 24 continuous values : 12 for the amplitudes and 12 others for the phases of a sinusoid tracked by each motor. Experiments consider a task space u, v, α which corresponds to the 2D position and orientation of the quadruped.

6.2.1.2. Learning Exploration Strategies in Model-based Reinforcement Learning

Participants: Manuel Lopes, Todd Hester, Peter Stone, Pierre-Yves Oudeyer.

We studied how different exploration algorithms can be combine and selected at runtime. Typically the user must hand-tune exploration parameters for each different domain and/or algorithm that they are using. We introduced an algorithm called leo for learning to select among different exploration strategies on-line. This algorithm makes use of bandit-type algorithms to adaptively select exploration strategies based on the rewards received when following them. We show empirically that this method performs well across a set of five domains In contrast, for a given algorithm, no set of parameters is best across all domains. Our results demonstrate that the leo algorithm successfully learns the best exploration strategies on-line, increasing the received reward over static parameterizations of exploration and reducing the need for hand-tuning exploration parameters [46].

6.2.1.3. Active Inverse Reinforcement Learning through Generalized Binary Search

Participants: Manuel Lopes, Francisco Melo.

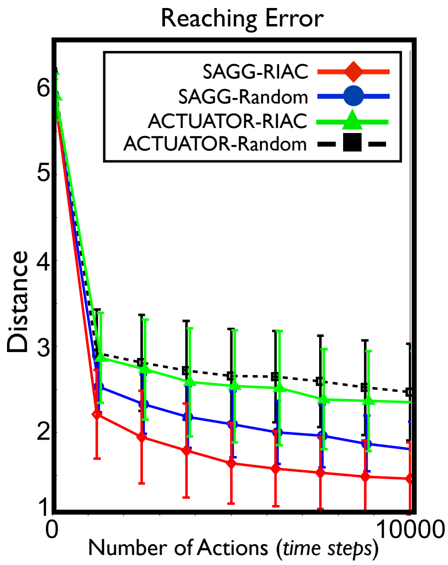

Figure 24. Evolution of the quality of the learnt inverse model for the quadruped robot experiment, depending on various exploration strategies (measured as mean error over a set of uniformly distributed goals generated independantly from learning trials).

We contributed the first aggressive active learning algorithm for nonseparable multi-class classification. We generalize an existing active learning algorithm for binary classification [116] to the multi-class setting, and identify mild conditions under which the proposed method provably retains the main properties of the original algorithm, namely consistency and sample complexity. In particular, we show that, in the binary case, our method reduces to the original algorithm of [116]. We then contribute an extension of our method to multi-label settings, identify its main properties and discuss richer querying strategies. We conclude the paper with two illustrative application examples. The first application features a standard text-classification problem. The second application scenario features a learning from demonstration setting. In both cases we demonstrate the advantage of our active sampling approach against random sampling. We also discuss the performance of the proposed approach in terms of the derived theoretical bounds.

6.2.1.4. Exploration strategies in developmental robotics: a unified probabilistic framework

Participants: Clément Moulin-Frier, Pierre-Yves Oudeyer.

We present a probabilistic framework unifying two important families of exploration mechanisms recently shown to be efficient to learn complex non-linear redundant sensorimotor mappings. These two explorations mechanisms are: 1) goal babbling, 2) active learning driven by the maximization of empirically measured learning progress. We show how this generic framework allows to model several recent algorithmic architectures for autonomous exploration. Then, we propose a particular implementation using Gaussian Mixture Models, which at the same time provides an original empirical measure of the competence progress. Finally, we perform computer simulations on two simulated setups: the control of the end effector of a 7-DoF arm and the control of the formants produced by an articulatory synthesizer. We are able to reproduce previous results from [25] with the advantages of a clean and compact probabilistic framework to efficiently express, implement and compare various exploration strategies on developmental robotics setups.

This work was published in three international conferences [54], [56], [55].

6.2.1.5. Autonomous Reuse of Motor Exploration Trajectories

Participants: Fabien Benureau, Pierre-Yves Oudeyer.

We developped an algorithm for transferring exploration strategies between tasks that share a common motor space in the context of lifelong autonomous learning in robotics. In such context sampling is costly, and exploration can take a long time before finding interesting, learnable data about a task. Our algorithm shows that we can significantly reduce sampling by reusing past data of other learned tasks, with no need of external knowledge or specific task structure. The algorithm does not transfer observations, or make assumptions about how the learning is conducted. Instead, only selected motor commands are transferred between tasks, chosen autonomously according to an empirical measure of learning progress. We show that on a wide variety of variations from a source task, such as changing the object the robot is interacting with or altering the morphology of the robot, this simple and flexible transfer method increases early performance significantly in the new task. We also investigate the limitation of this algorithm on specific situations.

This work has been published at ICDL, in Osaka [40].

6.2.2. Learning and optimization of motor policies

6.2.2.1. Off-Policy Actor-Critic

Participants: Thomas Degris, Martha White, Richard Sutton.

Actor–critic architectures are an interesting candidate for learning with robots: they can represent complex stochastic policies suitable for robots, they can learn online and incrementally and their per-time-step complexity scales linearly with the number of learned weights. Moreover, interesting connections have been identified in the existing literature with neuroscience. Until recently, however, practical actor–critic methods have been restricted to the on-policy setting, in which the agent learns only about the policy it is executing.

In an off-policy setting, on the other hand, an agent learns about a policy or policies different from the one it is executing. Off-policy methods have a wider range of applications and learning possibilities. Unlike on-policy methods, off-policy methods are able to, for example, learn about an optimal policy while executing an exploratory policy, learn from demonstration, and learn multiple tasks in parallel from a single sensory-motor interaction with an environment. Because of this generality, off-policy methods are of great interest in many application domains.

We have presented the first actor-critic algorithm for off-policy reinforcement learning. Our algorithm is online and incremental, and its per-time-step complexity scales linearly with the number of learned weights. We have derived an incremental, linear time and space complexity algorithm that includes eligibility traces and empirically show better or comparable performance to existing algorithms on standard reinforcement-learning benchmark problems. This work was was reproduced independently by Saminda Abeyruwan from the University of Miami.

6.2.2.2. Auto-Actor Critic

Participant: Thomas Degris.

As mentioned above, actor–critic architectures are an interesting candidate for robots to learn new skills in unknown and changing environments. However, existing actor–critic architectures, as many machine learning algorithms, require manual tuning of different parameters to work in the real world. To be able to systematize and scale-up skill learning on a robot, learning algorithms need to be robust to their parameters. The Flowers team has been working on making existing actor–critic algorithms more robust to make them suitable to a robotic setting. Results on standard reinforcement learning benchmarks are encouraging. This work will be submitted to international conference related with reinforcement learning. Interestingly, the methods developed in this work also offer a new formalism to think about different existing themes of Flowers research such as curiosity and maturational constraints.

6.2.2.3. Deterministic Policy Gradient Algorithms

Thomas Degris and colleagues from UCL and Deepming have consider deterministic policy gradient algorithms for reinforcement learning with continuous actions. The deterministic pol- icy gradient has a particularly appealing form: it is the expected gradient of the action-value func- tion. This simple form means that the deter- ministic policy gradient can be estimated much more efficiently than the usual stochastic pol- icy gradient. To ensure adequate exploration, we introduce an off-policy actor-critic algorithm that learns a deterministic target policy from an exploratory behaviour policy. We demonstrate that deterministic policy gradient algorithms can significantly outperform their stochastic counter- parts in high-dimensional action spaces. [58]

6.2.2.4. Relationship between Black-Box Optimization and Reinforcement Learning

Participant: Freek Stulp.

Policy improvement methods seek to optimize the parameters of a policy with respect to a utility function. There are two main approaches to performing this optimization: reinforcement learning (RL) and black-box optimization (BBO). In recent years, benchmark comparisons between RL and BBO have been made, and there has been several attempts to specify which approach works best for which types of problem classes.

We have made several contributions to this line of research by: 1) Defining four algorithmic properties that further clarify the relationship between RL and BBO. 2) Showing how the derivation of ever more powerful RL algorithms displays a trend towards BBO. 3) Continuing this trend by applying two modifications to the state-of-the-art PI^2 algorithm, which yields an algorithm we denote PI^{BB}. We show that PI^{BB} is a BBO algorithm, and, more specifically, that it is a special case of the state-of-the-art CMAES algorithm. 4) Demonstrating that the simpler PI^{BB} achieves similar or better performance than PI^2 on several evaluation tasks. 5) Analyzing why BBO outperforms RL on these tasks. These contributions have been published on HAL [129], and in Paladyn: Journal of Behavioral Robotics [36].

6.2.2.5. Probabilistic optimal control: a quasimetric approach

Participants: Steve N'Guyen, Clément Moulin-Frier, Jacques Droulez.

During his previous post-doc at the Laboratoire de Physiologie de la Perception et de l'Action (College de France, Paris), Clément Moulin-Frier joined Jacques Droulez and Steve N'Guyen to work on an alternative and original approach of probabilistic optimal control called *quasimetric*. A journal paper was published in PLoS ONE in December 2013 [31], where the authors propose a new approach dealing with decision making under uncertainty.

6.2.3. Social learning and intrinsic motivation

6.2.3.1. Socially Guided Intrinsic Motivation for Skill Learning

Participants: Sao Mai Nguyen, Pierre-Yves Oudeyer.

We have explored how social interaction can bootstrap the learning of a robot for motor learning. We first studied how simple demonstrations by teachers could have a bootstrapping effect on autonomous exploration with intrinsic motivation by building a learner who uses both imitation learning and SAGG-RIAC algorithm [25], and thus designed the SGIM-D (Socially Guided Intrinsic Motivation by Demonstration) algorithm [32], [24], [32] [114], [111]. We then investigated on the reasons of this bootstrapping effect [113], to show that demonstrations by teachers can both enhance more tasks to be explored, as well as favor more easily generalized actions to be used. This analysis is generalizable for all algorithms using social guidance and goal-oriented exploration. We then proposed to build a strategic learner who can learn multiple tasks and with multiple strategies. An overview and theoretical study of multi-task, multi-strategy Strategic Learning is presented in [99]. We also forsook to build a learning algorithm for more natural interaction with the human users. We first designed the SGIM-IM algorithm so that it can determine itself when it should ask for help from the teacher while trying to explore autonomously as long as possible so as to use as little of the teacher's time as possible [112]. After tackling with the problem of how and when to learn, we also investigated an active learner who can determine who to ask for help: in the case of two teachers available, SGIM-IM can determine which strategy to adopt between autonomous exploration and learning by demonstration, and which teacher enhances most learning progress for the learner [115], and ask him for help.

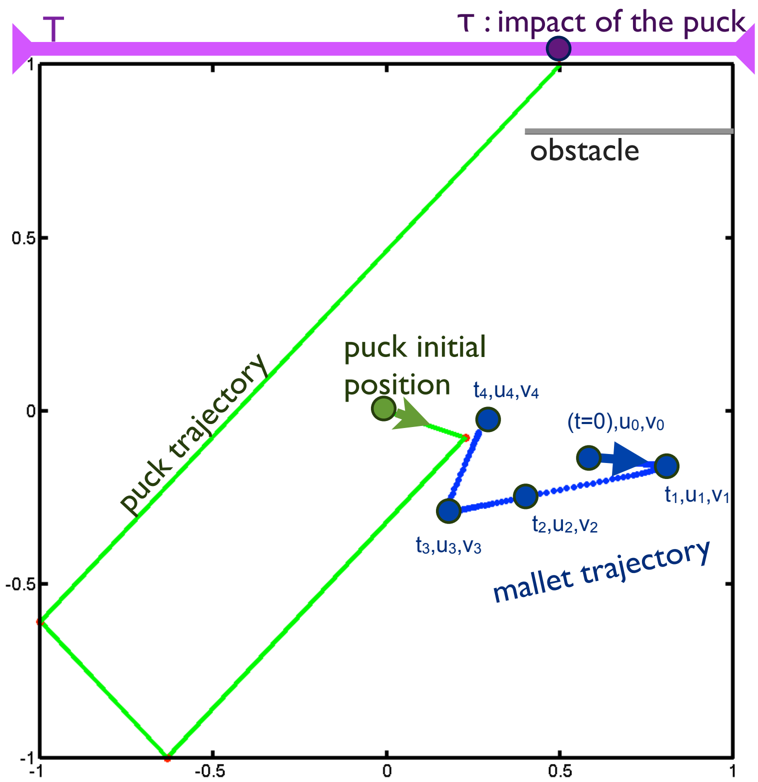

Figure 25. Illustration of SGIM-D and SGIM-IM algorithms

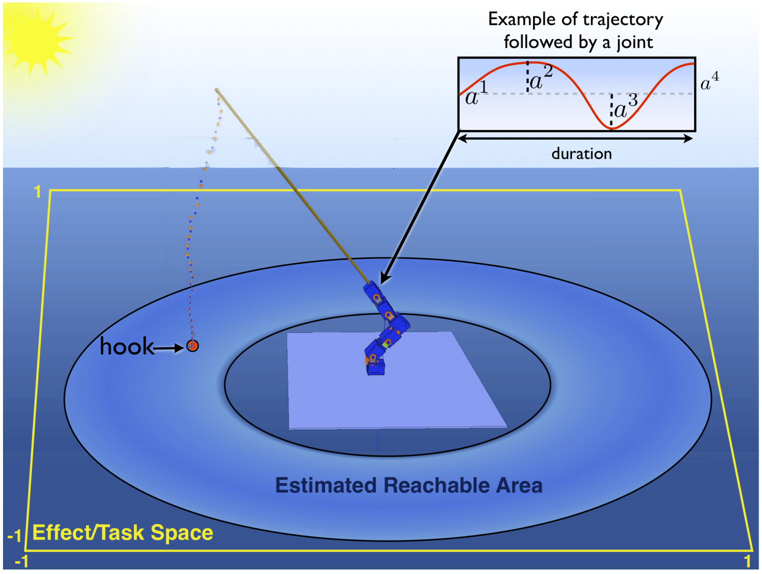

Figure 26. Illustration of SGIM-D and SGIM-IM algorithms

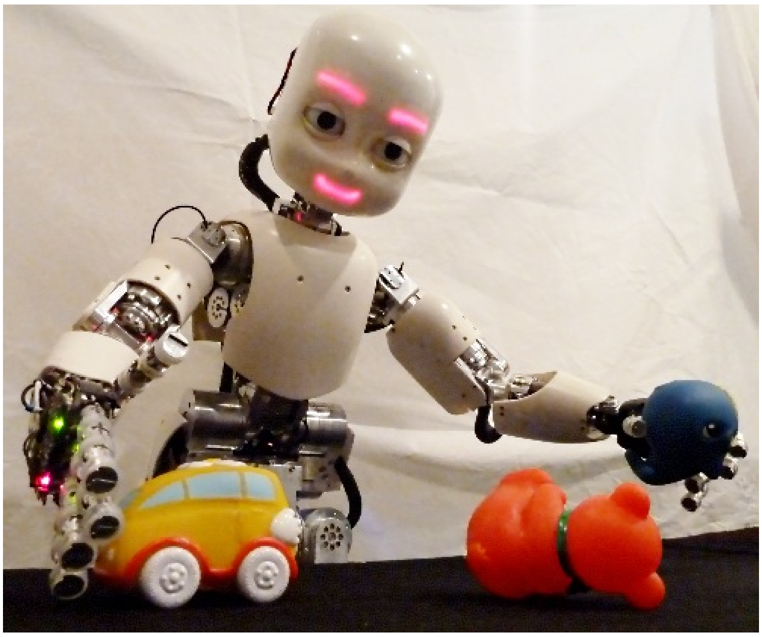

Figure 27. Illustration of SGIM-D and SGIM-IM algorithms

While the above results have been shown in simulation environments: of a simple deterministic air hockey game (fig. 25), and a stochastic fishing experiment with a real-time physical simulator (fig. 26), we are now building the experimental setup of the fishing experiment in order to carry out the experiments with naive users.

6.2.3.2. Adaptive task execution for implicit human-robot coordination

Participants: Ievgen Perederieiev, Manuel Lopes, Freek Stulp.

We began a project which goal is to study how computational models of multi-agent systems can be applied in situations where one agent is a human. We aim at applications where robots collaborate with humans for achieving complex tasks..

A very important capability for efficient collaborative work is the mutual agreement of a task and the ability to predict the behavior of others. We address such aspect by studying methods that increase the predictability of the robot actions. An efficient motor execution becomes the one that not just optimize speed and minimizes energy but also the one that improves the reliability of the team behavior. We are studying policy gradient methods and working on policy improvement algorithms (PI^2, CEM and $CMAES$). A feasibility study will consider a simple task between a robot and a person where the goal is to coordinate the way a set of three colored buttons is pressed.

6.2.4. Unsupervised learning of motor primitives

6.2.4.1. Clustering activities

Participants: Manuel Lopes, Luis Montesano, Javier Almingol.

Learning behaviors from data has applications in surveillance and monitoring systems, virtual agents and robotics among others. In our approach, ww assume that in a given unlabeled dataset of multiple behaviors, it is possible to find a latent representation in a controller space that allows to generate the different behaviors. Therefore, a natural way to group these behaviors is to search a common control system that generate them accurately.

Clustering behaviors in a latent controller space has two major challenges. First, it is necessary to select the control space that generate behaviors. This space will be parameterized by a set of features that will change for different behaviors. Usually, each controller will minimize a cost function with respect to several task features. The latent representation is in turn defined by the selected features and their corresponding weight. Second, an unknown number of such controllers is required to generate different behaviors and the grouping must be based on the ability of the controller to generate the demonstrations using a compact set of controllers.

We propose a Dirichlet Process based algorithm to cluster behaviors in a latent controller space which encodes the dynamical system generating the observed trajectories. The controller uses a potential function generated as a linear combination of features. To enforce sparsity and automatically select features for each cluster independently, we impose a conditional Laplace prior over the controller parameters. Based on this models, we derive a sparse Dirichlet Process Mixture Model (DPMM) algorithm that estimates the number of behaviors and a sparse latent controller for each of them based on a large set of features [38].

6.2.4.2. Learning the Combinatorial Structure of Demonstrated Behaviors with Inverse Reinforcement Control

Participants: Olivier Mangin, Pierre-Yves Oudeyer.

We have elaborated and illustrated a novel approach to learning motor skills from demonstration. This approach combines ideas from inverse reinforcement learning, in which actions are assumed to solve a task, and dictionary learning. In this work we introduced a new algorithm that is able to learn behaviors by assuming that the observed complex motions can be represented in a smaller dictionary of concurrent tasks. We developed an optimization formalism and show how we can learn simultaneously the dictionary and the mixture coefficients that represent each demonstration. We presented results on a toy problem and shown that our algorithm finds an efficient set of primitive tasks where naive approaches such as PCA and using a dictionary built from random examples fail to achieve the same tasks. These results that where presented as [60], extend the ones from [103].

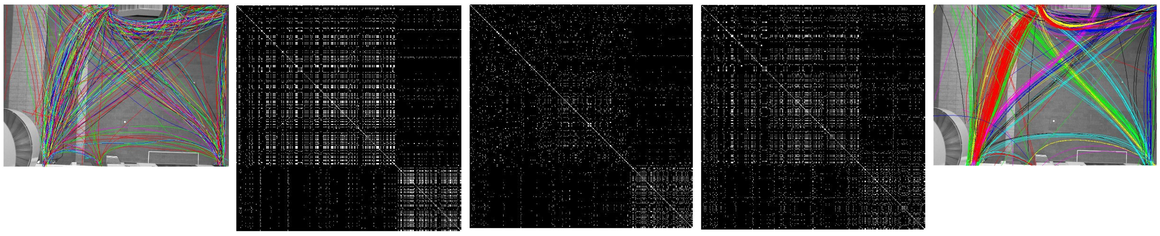

(a) (b) (c) (d) (e)

Figure 28. EIFPD dataset. (a) Trajectories of the EIFPD to be clustered (color is non-informative). (b-d) correspondence matrix for the 474 trajectories for the labeled ground truth, the KMeans in measurement space and the DPMM, respectively. (e) Reconstructed trajectories from the initial point using the estimated parameters of the DPMM algorithm. Due to the large number of clusters (37), colors are repeated for different clusters.

6.2.4.3. Interaction of Maturation and Intrinsic Motivation for Developmental Learning of Motor Skills in Robots

Participants: Adrien Baranes, Pierre-Yves Oudeyer.

We have introduced an algorithmic architecture that couples adaptively models of intrinsic motivation and physiological maturation for autonomous robot learning of new motor skills. Intrinsic motivation, also called curiosity-driven learning, is a mechanism for driving exploration in active learning. Maturation denotes here mechanisms that control the evolution of certain properties of the body during development, such as the number and the spatio-temporal resolution of available sensorimotor channels. We argue that it is useful to introduce and conceptualize complex bidirectional interactions among these two mechanisms, allowing to actively control the growth of complexity in motor development in order to guide efficiently exploration and learning. We introduced a model of maturational processes, taking some functional inspiration from the myelination process in humans, and show how it can be coupled in an original and adaptive manner with the intrinsic motivation architecture SAGG-RIAC (Self-Adaptive Goal Generation - Robust Intelligent Adaptive Curiosity algorithm), creating a new system, called McSAGG-RIAC. We then conducted experiments to evaluate both qualitative and quantitative properties of these systems when applied to learning to control a high-dimensional robotic arm, as well as to learning omnidirectional locomotion in a quadruped robot equipped with motor synergies. We showed that the combination of active and maturational learning can allow to gain orders of magnitude in learning speed as well as reach better generalization performances. A journal article is currently being written.

6.3. Autonomous and Social Perceptual Learning

6.3.1. The Impact of Human-Robot Interfaces on the Learning of Visual Objects

Participants: Pierre Rouanet, Pierre-Yves Oudeyer, Fabien Danieau, David Filliat.

We have continued and finalized a large-scale study of the impact of interfaces allowing non-expert users to efficiently and intuitively teach a robot to recognize new visual objects. We identified challenges that need to be addressed for real-world deployment of robots capable of learning new visual objects in interaction with everyday users. We argue that in addition to robust machine learning and computer vision methods, well-designed interfaces are crucial for learning efficiency. In particular, we argue that interfaces can be key in helping non-expert users to collect good learning examples and thus improve the performance of the overall learning system. Then, we have designed four alternative human-robot interfaces: three are based on the use of a mediating artifact (smartphone, wiimote, wiimote and laser), and one is based on natural human gestures (with a Wizard-of-Oz recognition system). These interfaces mainly vary in the kind of feedback provided to the user, allowing him to understand more or less easily what the robot is perceiving, and thus guide his way of

providing training examples differently. We then evaluated the impact of these interfaces, in terms of learning efficiency, usability and user's experience, through a real world and large scale user study. In this experiment, we asked participants to teach a robot twelve different new visual objects in the context of a robotic game. This game happens in a home-like environment and was designed to motivate and engage users in an interaction where using the system was meaningful. We then analyzed results that show significant differences among interfaces. In particular, we showed that interfaces such as the smartphone interface allows non-expert users to intuitively provide much better training examples to the robot, almost as good as expert users who are trained for this task and aware of the different visual perception and machine learning issues. We also showed that artifact-mediated teaching is significantly more efficient for robot learning, and equally good in terms of usability and user's experience, than teaching thanks to a gesture-based human-like interaction.
This work was published in the IEEE Transactions on Robotics [34].

(a) draw the attention toward an object (b) trigger the capture (c) encircle the area of the object (d) the new learning example

Figure 29. Smartphone Interface. To make the robot collect a new learning example, users have to first draw the robot's attention toward the object they want to teach through simple gestures. Once the robot sees the object, they touch the head of the robot to trigger the capture. Then, they directly encircle the area of the image that represents the object on the screen. The selected area is then used as the new learning example. The combination of the video stream and the gestures facilitate the achievement of joint attention.

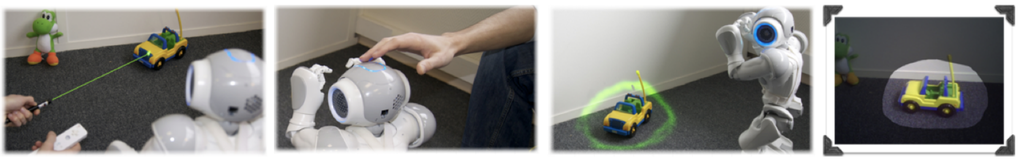

(a) draw the attention toward an object (b) trigger the capture (c) encircle the area of the object (d) the new learning example

Figure 30. Wiimote + laser pointer interface. With this interface users can draw the robot's attention with a laser pointer toward an object. The laser spot is automatically tracked by the robot. They can ensure that the robot detects the spot thanks to haptic feedback on the Wiimote. Then, they can touch the head of the robot to trigger the capture of a new learning example. Finally, they encircle the object with the laser pointer to delimit its area which will be defined as the new learning example.

6.3.2. Developmental object learning through manipulation and human demonstration

Participants: Natalia Lyubova, David Filliat.

The goal of this work is to design a visual system for a humanoid robot. We used a developmental approach that allows a humanoid robot to continuously and incrementally learn entities through interaction with a human partner in a first stage before categorizing these entities into objects, humans or robot parts and using this knowledge to improve objects models by manipulation in a second stage. This approach does not require

Figure 31. The real world environment designed to reproduce a typical living room. Many objects were added in the scene in order to make the environment cluttered.

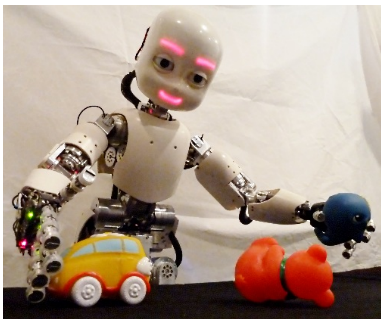

Figure 32. iCub performing curiosity-driven exploration and active recognition of visual objects in 3D

prior knowledge about the appearance of the robot, the human or the objects. The proposed perceptual system segments the visual space into proto-objects, analyses their appearance, and associates them with physical entities. Entities are then classified based on the mutual information with proprioception and on motion statistics. The ability to discriminate between the robot's parts and a manipulated object then allows to update the object model with newly observed object views during manipulation. We evaluate our system on an iCub robot, showing the independence of the self-identification method on the robot's hands appearances by wearing different colored gloves. The interactive object learning using self-identification shows an improvement in the objects recognition accuracy with respect to learning through observation only [52], [51].

6.3.3. A Comparison of Geometric and Energy-Based Point Cloud Semantic Segmentation Methods

Participants: Mathieu Dubois, Alexander Gepperth, David Filliat.

The software we developped for object segmentation and recognition rely on a geometric segmentation of the space. We tested alternative methods for this semantic segmentation task in which the goal is to find some relevant classes for navigation such as wall, ground, objects, etc. Several effective solutions have been proposed, mainly based on the recursive decomposition of the point cloud into planes. We compare such a solution to a non-associative MRF method inspired by some recent work in computer vision.

The results [42] shows that the geometric method gives superior results for the task of semantic segmentation in particular for the object class. This can be explained by the fact that it incorporates a lot of domain knowledge (namely that indoor environments are made of planes and that objects lie on top of them). However, MRF segmentation gives interesting results and has several advantages. First most of it's components can be used for other purpose or in other, less constrained, environments where domain knowledge is not available. For instance we could try to recognize more precisely the objects. Second it requires less tuning since most parameters are learned from the database. Third, it uses the appearance information which could help to identify different types of ground or wall (this was one of the goal in the CAROTTE challenge). Last but not least, as it gives a probabilistic output, it allows the robot to draw hypothesis on the environment and adapt its behavior. Therefore we think it is interesting to investigate improvements to improve the exploitation of the structure of the point clouds.

6.3.4. Efficient online bootstrapping of sensory representations

Participant: Alexander Gepperth.

This work [86] is a simulation-based investigation exploring a novel approach to the open-ended formation of multimodal representations in autonomous agents. In particular, we addressed here the issue of transferring (bootstrapping) features selectivities between two modalities, from a previously learned or innate reference representation to a new induced representation. We demonstrated the potential of this algorithm by several experiments with synthetic inputs modeled after a robotics scenario where multimodal object representations are bootstrapped from a (reference) representation of object affordances, focusing particularly on typical challenges in autonomous agents: absence of human supervision, changing environment statistics and limited computing power. We proposed an autonomous and local neural learning algorithm termed PROPRE (projection-prediction) that updates induced representations based on predictability: competitive advantages are given to those feature-sensitive elements that are inferable from activities in the reference representation, the key ingredient being an efficient online measure of predictability controlling learning. We verified that the proposed method is computationally efficient and stable, and that the multimodal transfer of feature selectivity is successful and robust under resource constraints. Furthermore, we successfully demonstrated robustness to noisy reference representations, non-stationary input statistics and uninformative inputs.

6.3.5. Simultaneous concept formation driven by predictability

Participants: Alexander Gepperth, Louis-Charles Caron.

This work [83] was conducted in the context of developmental learning in embodied agents who have multiple data sources (sensors) at their disposal. We developed an online learning method that simultaneously discovers meaningful concepts in the associated processing streams, extending methods such as PCA, SOM or sparse coding to the multimodal case. In addition to the avoidance of redundancies in the concepts derived from single modalities, we claim that meaningful concepts are those who have statistical relations across modalities. This is a reasonable claim because measurements by different sensors often have common cause in the external world and therefore carry correlated information. To capture such cross-modal relations while avoiding redundancy of concepts, we propose a set of interacting self-organization processes which are modulated by local predictability. To validate the fundamental applicability of the method, we conducted a plausible simulation experiment with synthetic data and found that those concepts which are predictable from other modalities successively "grow", i.e., become overrepresented, whereas concepts that are not predictable become systematically under-represented. We additionally explored the applicability of the developed method to real-world robotics scenarios.

6.3.6. The contribution of context: a case study of object recognition in an intelligent car

Participants: Alexander Gepperth, Michael Garcia Ortiz.

In this work [84], we explored the potential contribution of multimodal context information to object detection in an "intelligent car". The used car platform incorporates subsystems for the detection of objects from local visual patterns, as well as for the estimation of global scene properties (sometimes denoted scene context or just context) such as the shape of the road area or the 3D position of the ground plane. Annotated data recorded on this platform is publicly available as the a "HRI RoadTraffic" vehicle video dataset, which formed the basis for the investigation. In order to quantify the contribution of context information, we investigated whether it can be used to infer object identity with little or no reference to local patterns of visual appearance. Using a challenging vehicle detection task based on the "HRI RoadTraffic" dataset, we trained selected algorithms (context models) to estimate object identity from context information alone. In the course of our performance evaluations, we also analyzed the effect of typical real-world conditions (noise, high input dimensionality, environmental variation) on context model performance. As a principal result, we showed that the learning of context models is feasible with all tested algorithms, and that object identity can be estimated from context information with similar accuracy as by relying on local pattern recognition methods. We also found that the use of basis function representations [1] (also known as "population codes" allows the simplest (and therefore most efficient) learning methods to perform best in the benchmark, suggesting that the use of context is feasible even in systems operating under strong performance constraints.

6.3.7. Co-training of context models for real-time object detection

Participant: Alexander Gepperth.

In this work [85], we developed a simple way to reduce the amount of required training data in context-based models of real- time object detection and demonstrated the feasibility of our approach in a very challenging vehicle detection scenario comprising multiple weather, environment and light conditions such as rain, snow and darkness (night). The investigation is based on a real-time detection system effectively composed of two trainable components: an exhaustive multiscale object detector (signal-driven detection), as well as a module for generating object-specific visual attention (context models) controlling the signal-driven detection process. Both parts of the system require a significant amount of ground-truth data which need to be generated by human annotation in a time-consuming and costly process. Assuming sufficient training examples for signal-based detection, we showed that a co-training step can eliminate the need for separate ground-truth data to train context models. This is achieved by directly training context models with the results of signal-driven detection. We demonstrated that this process is feasible for different qualities of signal-driven detection, and maintains the performance gains from context models. As it is by now widely accepted that signal-driven object detection can be significantly improved by context models, our method allows to train strongly improved detection systems without additional labor, and above all, cost.

6.4. Robot Multimodal Learning of Language and Action

6.4.1. Learning semantic components from sub-symbolic multi-modal perception

Participants: Olivier Mangin, Caio Tomazelli Da Silva Oliveira, Pierre-Yves Oudeyer.

Perceptual systems often include sensors from several modalities. However, existing robots do not yet sufficiently discover patterns that are spread over the flow of multimodal data they receive. In this work we establish a framework to learns multimodal components from perception. We use a nonnegative matrix factorization algorithm to learn a dictionary of components that represent meaningful elements present in the multimodal perception, without providing the system with a symbolic representation of the semantics. In [53] we illustrate this framework by showing how a learner discovers word-like components from observation of gestures made by a human together with spoken descriptions of the gestures, and how it captures the semantic association between the two. These experiments were further extended during the internship of Caio Tomazelli Da Silva Oliveira. Importantly these experiments provide an example of language grounding into perception, and feature global understanding of a linguistic task without requiring its compositional understanding. The code of the experiments from [53] as well as the motion dataset have been made publicly available to improve the reproducibility of the experiments.

6.4.2. Curiosity-driven exploration and interactive learning of visual objects with the ICub robot

Participants: Mai Nguyen, Natalia Lyubova, Damien Gerardeaux-Viret, David Filliat, Pierre-Yves Oudeyer.

We studied how various mechanisms for cognition and learning, such as curiosity, action selection, imitation, visual learning and interaction monitoring, can be integrated in a single embodied cognitive architecture. We have conducted an experiment with the iCub robot for active recognition of objects in 3D through curiosity-driven exploration, in which the robot can manipulate the robot or ask a human user to manipulate objects to gain information and recognise better objects (fig. 27). For this experiment carried out within the MACSi project, we address the problem of learning to recognise objects in a developmental robotics scenario. In a life-long learning perspective, a humanoid robot should be capable of improving its knowledge of objects with active perception. Our approach stems from the cognitive development of infants, exploiting active curiosity-driven manipulation to improve perceptual learning of objects. These functionalities are implemented as perception, control and active exploration modules as part of the Cognitive Architecture of the MACSi project. We integrated a bottom-up vision system based on swift feature points and motor-primitive based robot control with the SGIM-ACTS algorithm (Socially Guided Intrinsic Motivation with Active Choice of Task and Strategy as the active exploration module. SGIM-ACTS is a strategic learner who actively chooses which task to concentrate on, and which strategy is better according to this task. It thus monitors the learning progress for each strategy on all kinds of tasks, and actively interacts with the human teacher. We obtained an active object recognition approach, which exploits curiosity to guide exploration and manipulation, such that the robot can improve its knowledge of objects in an autonomous and efficient way. Experimental results show the effectiveness of our approach: the humanoid iCub is now capable of deciding autonomously which actions must be performed on objects in order to improve its knowledge, requiring a minimal assistance from its caregiver. This work constitutes the base for forthcoming research in autonomous learning of affordances. This work have been published in a conference [57] and in a journal paper [28].

6.4.3. Imitation Learning and Language

Participants: Thomas Cederborg, Pierre-Yves Oudeyer.

We have studied how context-dependant imitation learning of new skills and language learning could be seen as special cases of the same mechanism. We argue that imitation learning of context-dependent skills implies complex inferences to solve what we call the ”motor Gavagai problem”, which can be viewed as a generalization of the so-called ”language Gavagai problem”. In a full symbolic framework where percepts and actions are continuous, this allows us to articulate that language may be acquired out of generic sensorimotor imitation learning mechanisms primarily dedicated at solving this motor Gavagai problem. Through the use of a computational model, we illustrate how non-linguistic and linguistic skills can be learnt concurrently, seamlessly, and without the need for symbols. We also show that there is no need to actually represent the distinction between linguistic and non-linguistic tasks, which rather appears to be in the eye of the observer of the system. This computational model leverages advanced statistical methods for imitation learning, where closed-loop motor policies are learnt from human demonstrations of behaviours that are dynamical responses to a multimodal context. A novelty here is that the multimodal context, which defines what motor policy to achieve, includes, in addition to physical objects, a human interactant which can produce acoustic waves (speech) or hand gestures (sign language). This was published in [26].

6.4.4. Learning to Interpret the Meaning of Teaching Signals in Socially Guided Robot Learning

Participants: Manuel Lopes, Jonathan Grizou, Thomas Cederborg, Pierre-Yves Oudeyer.

We elaborated an algorithm to bootstrap shared understanding in a human-robot interaction scenario where the user teaches a robot a new task using teaching instructions yet unknown to it. In such cases, the robot needs to estimate simultaneously what the task is and the associated meaning of instructions received from the user. For this work, we consider a scenario where a human teacher uses initially unknown spoken words, whose associated unknown meaning is either a feedback (good/bad) or a guidance (go left, right, ...). We present computational results, within an inverse reinforcement learning framework, showing that a) it is possible to learn the meaning of unknown and noisy teaching instructions, as well as a new task at the same time, b) it is possible to reuse the acquired knowledge about instructions for learning new tasks, and c) even if the robot initially knows some of the instructions' meanings, the use of extra unknown teaching instructions improves learning efficiency. Published articles: [43], [45].

An extension to the use of brain signals has been made [44]. Do we need to explicitly calibrate Brain Machine Interfaces (BMIs)? Can we start controlling a device without telling this device how to interpret brain signals? Can we learn how to communicate with a human user through practical interaction? It sounds like an ill posed problem, how can we control a device if such device does not know what our signals mean? This paper argues and present empirical results showing that, under specific but realistic conditions, this problem can be solved. We show that a signal decoder can be learnt automatically and online by the system under the assumption that both, human and machine, share the same a priori on the possible signals' meanings and the possible tasks the user may want the device to achieve. We present results from online experiments on a Brain Computer Interface (BCI) and a Human Robot Interaction (HRI) scenario.

6.4.5. Active Learning for Teaching a Robot Grounded Relational Symbols

Participants: Johannes Kulick, Tobias Lang, Marc Toussaint, Manuel Lopes.

The present work investigates an interactive teaching scenario, where a human aims to teach the robot symbols that abstract geometric (relational) features of objects. There are multiple motivations for this scenario: First, state-of-the-art methods for relational Reinforcement Learning demonstrated that we can successfully learn abstracting and well-generalizing probabilistic relational models and use them for goal-directed object manipulation. However, these methods rely on given grounded action and state symbols and raise the classical question Where do the symbols come from? Second, existing research on learning from human-robot interaction has focused mostly on the motion level (e.g., imitation learning). However, if the goal of teaching is to enable the robot to autonomously solve sequential manipulation tasks in a goal-directed manner, the human should have the possibility to teach the relevant abstractions to describe the task and let the robot eventually leverage powerful relational RL methods (see Figure 33). We formalize human-robot teaching of grounded symbols as an Active Learning problem, where the robot actively generates geometric situations that maximize his information gain about the symbol to be learnt. We demonstrate that the learned symbols can be used in a relational RL framework for the robot to learn probabilistic relational rules and use them to solve object manipulation tasks in a goal-directed manner. [47].

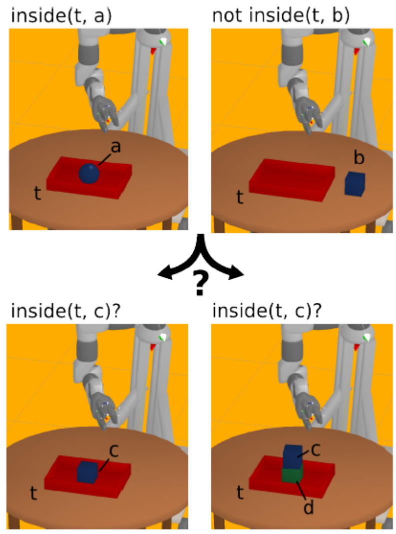

Figure 33. Active learning of symbol descriptions on a real world robot.

6.5. Robot Design and Morphological Computation

6.5.1. The Poppy Humanoid Robot: Leg Design for Biped Locomotion

Participants: Matthieu Lapeyre, Pierre Rouanet, Pierre-Yves Oudeyer.

In this paper introduced for Poppy as a novel humanoid robotic platform designed to jointly address three central goals of humanoid robotics: 1) study the role of morphology in biped locomotion; 2) study full-body compliant physical human-robot interaction; 3) be robust while easy and fast to duplicate to facilitate experimentation. The taken approach relies on functional modeling of certain aspects of human morphology, optimizing materials and geometry, as well as on the use of 3D printing techniques. In this article, we have focused on the presentation of the design of specific morphological parts related to biped locomotion: the hip, the thigh, the limb mesh and the knee. We also presented an initial experiments showing properties of the robot when walking with the physical guidance of a human. [50].

6.5.2. Poppy Humanoid Platform: Experimental Evaluation of the Role of a Bio-inspired Thigh Shape

Participants: Matthieu Lapeyre, Pierre Rouanet, Pierre-Yves Oudeyer.

In this paper, we present an experimental evaluation of the role of the morphology in the Poppy humanoid platform. More precisely, we have investigated the impact of the bio-inspired thigh, bended of 6^{o}, on the balance and biped locomotion. We compare this design with a more traditional straight thigh. We describe both the theoretical model and real experiments showing that the bio-inspired thigh allows the reduction of falling speed by almost 60% (single support phase) and the decrease of the lateral motion needed for the mass transfer from one foot to the other by 30% (double support phase). We also present an experiment where the robot walks on a treadmill thanks to the social and physical guidance of expert users and we show that the bended thigh reduces the upper body motion by about 45% indicating a more stable walk.[48].

6.5.3. Morphological computation and body intelligence

6.5.3.1. Comparative Study of the Role of Trunk in Human and Robot Balance Control

Participants: Matthieu Lapeyre [correspondant], Christophe Halgand, Jean-Renã© Cazalet, Etienne Guillaud, Pierre-Yves Oudeyer.

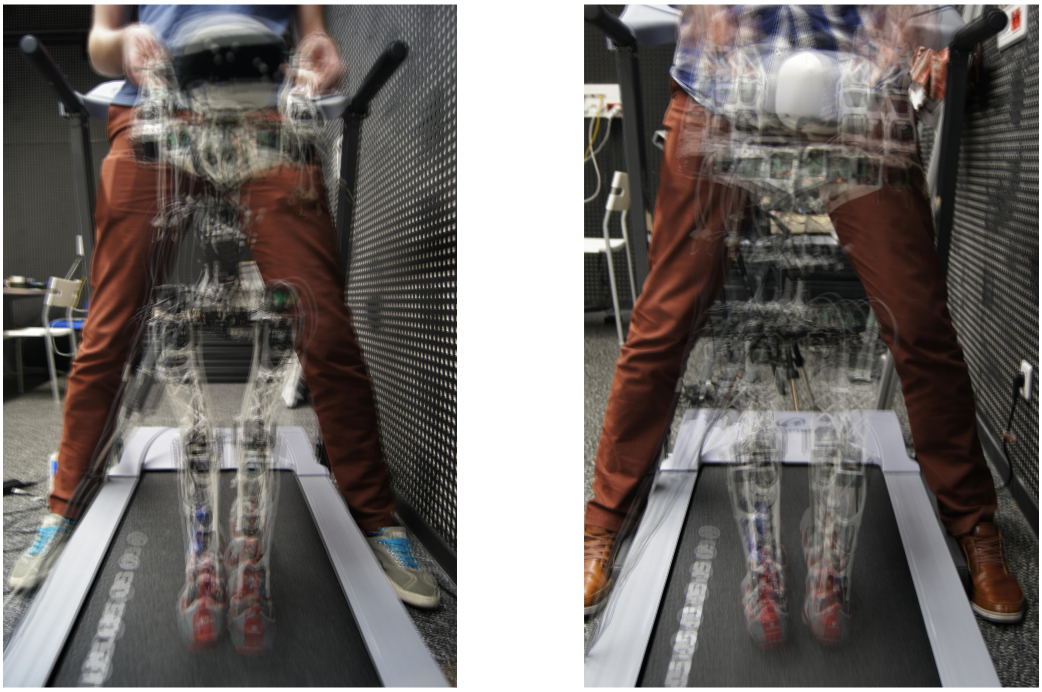

Figure 34. Five pictures have been taken while Poppy was walking and were stacked to obtain a qualitative view of the difference in the walking behavior in function of the morphology of the thigh.

Numerous studies in the field of functional motor rehabilitation were devoted to understanding the functioning of members but few are interested in the coordination of the trunk muscles and the relationship between axial and appendicular motricity which is essential in maintaining balance during travel. Acquiring new knowledge on this subject is a prerequisite in the development of new therapeutic strategies to restore motor function to the overall development of robotic orthosis that would assist the movement. Many robotic orthosis using EMG signals were unfortunately using few joints [82] and a system for controlling a multi articulated spine has not yet been developed. We propose here to use a multidisciplinary approach to define the neuro-mechanical principles where an axial system is operating in synergy with human and robot limbs.

To bring us a theoretical framework, we chose to study the reactions of the Acroban humanoid robot. Including 5 joints in the trunk, Acroban can reproduce in part the fluid movements of the human body [101] and especially to test its behavior when its trunk is held fixed or his arms are no longer used for rebalance. To disrupt postural balance in humans and robots, we have developed a low cost mobile platform (see Figure 35). This platform is made up of a broad stable support $(0.8x5m)$ mounted on a skateboard having a power of 800W. The substitution of the initial order of skate by an embedded microcontroller allows us to generate mono-axial perturbations precise intensity and duration to ensure repeatability of the disturbance. We capture movements (Optitrack 250Hz) and record the acceleration of the platform (accelerometer embedded 2kHz), the center of pressure (WiiBalanceBoard 60Hz), and electromyography (EMG).

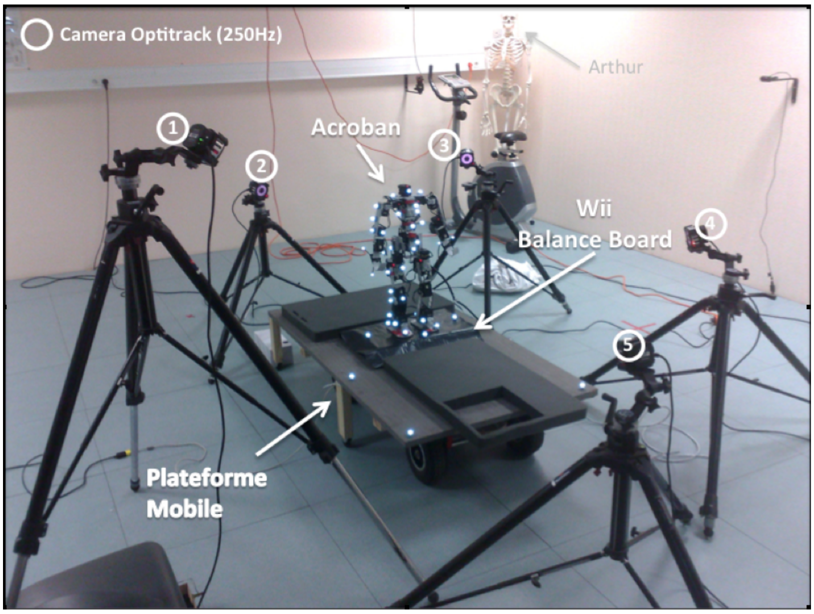

Figure 35. Experimental setup for comparative study of the role of the trunk in human and robot balance control

The experimental device (mobile platform and synchronized recordings) is operational. Preliminary experiments have allowed us to refine the profiles of disturbance on the robot Acroban. The analysis of preliminary results is in progress. Following this study, we hope to improve the modeling of the motor system in humans and robotic simulation as a basis for the development of robotic orthosis axial system. Second, the results provide a basis for improved balancing of Acroban primitives but also the development of future humanoid robots.

6.6. Educational Technologie

6.6.1. KidLearn: Adaptive Personalization of Educational Content with Machine Learning

Kidlearn is a research project studying how machine learning can be applied to intelligent tutoring systems. It aims at developing methodologies and software which adaptively personalize sequences of learning activities to the particularities of each individual student. Our systems aim at proposing to the student the right activity at the right time, maximizing concurrently his learning progress and its motivation. In addition to contributing to the efficiency of learning and motivation, the approach is also made to reduce the time needed to design ITS systems.

Intelligent Tutoring System (ITS) are computer environments designed to guide students in their learning. Through the proposal of different activities, it provides teaching experience, guidance and feedback to improve learning. The FLOWERS team has developed several computational models of artificial curiosity and intrinsic motivation based on research on psychology that might have a great impact for ITS. Results showed that activities with intermediate levels of complexity, neither too easy nor too difficult but just a little more difficult that the current level, provide better teaching experiences. The system is based on the combination of three approaches. First, it leverages Flowers team's recent models of computational models of artificial curiosity and intrinsic motivation based on research in psychology and neuroscience. One overview can be be found in [27]. Second, it uses state-of-the-art Multi-Arm Bandit (MAB) techniques to efficiently manage the exploration/exploitation challenge of this optimization process. Third, it leverages expert knowledge to constrain and bootstrap initial exploration of the MAB, while requiring only coarse guidance information of the expert and allowing the system to deal with didactic gaps in its knowledge. In 2013, we have run a first pilot experiment in elementary schools of Région Aquitaine, where 7-8 year old kids could learn elements of mathematics thanks to an educational software that presented the right exercises at the right time to maximize learning progress. A report is available at: http://arxiv.org/pdf/1310.3174v1.pdf.

6.7. Other applications

6.7.1. Real-time Reaction-Diffusion Simulation: a Machine Learning Technique

Participants: Thomas Degris, Nejib Zemzemi.

Carmen is an Inria team working on modeling the electrical activity of the human heart. Their models are mainly based on reaction-diffusion equations. These methods are expansive in terms of computational costs which limits their use in practice. More specifically, some recent chirurgical intervention techniques on the heart (atrial ablation) requires to identify the source of the electrical wave. Finding such sources requires an optimization procedure. Using classical methods, this procedure is very heavy computationally.

In this project, our goal is to reduce the computational cost using supervised learning techniques. The idea is to replace the incremental resolution of partial differential equations by more suitable data structures for real-time running. Starting from data generated by simulating different excitations scenari on a human atria, this data is afterwords used as a training data set for machine learning algorithms. This approach will allow a faster optimization procedure.

This work is in collaboration with Nejib Zemzemi from the Inria Carmen team. This project is in preliminary steps.

6.7.2. Appearance-based segmentation of indoors/outdoors sequences of spherical views

Participant: David Filliat.

In collaboration with Patrick Rives and Alexandre Chapoulie from the Arobas team at Inria Sophia-Antipolis, we developped a method for environment segmentation based on spherical views [41]. Navigating in large scale, complex and dynamic environments requires reliable representations able to capture metric, topological and semantic aspects of the scene for supporting path planing and real time motion control. In a previous work, we addressed metric and topological representations thanks to a multi-cameras system which allows building of dense visual maps of large scale 3D environments. The map is a set of locally accurate spherical panoramas related by 6dof poses graph. The work presented here is a further step toward a semantic representation. We aim at detecting the changes in the structural properties of the scene during navigation. Structural properties are estimated online using a global descriptor relying on spherical harmonics which are particularly well-fitted to capture properties in spherical views. A change-point detection algorithm based on a statistical Neyman-Pearson test allows us to find optimal transitions between topological places. Results are presented and discussed both for indoors and outdoors experiments.

6.7.3. Modelling Stop Intersection Approaches using Gaussian Processes

Participant: David Filliat.

In collaboration with Javier-Ibanez Guzman and Alexandre Armand from Renault, we developped an approach toward the development of an electronic co-pilot adapted to the driver behavior [39]. Indeed, each driver reacts differently to the same traffic conditions, however, most Advanced Driving Assistant Systems (ADAS) assume that all drivers are the same. This work proposes a method to learn and to model the velocity profile that the driver follows as the vehicle decelerates towards a stop intersection. Gaussian Processes (GP), a machine learning method for non-linear regressions are used to model the velocity profiles. It is shown that GP are well adapted for such an application, using data recorded in real traffic conditions. GP allow the generation of a normally distributed speed, given a position on the road. By comparison with generic velocity profiles, benefits of using individual driver patterns for ADAS issues are presented.

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

7.1.1. Advanced platform for Urban Mobility (PAMU

Participant: David Filliat.

Development of a planning algorithm on a autonomous electric car for Renault SAS. We developped a planning module in order to produce global plans to reach a goal specified in a digital map and to perform local reactive planning to avoid dynamic obstacles. This module is integrated in the PAMU autonomous vallet parking developped by Renault with several academic partners.

7.2. Bilateral Grants with Industry

7.2.1. Development of an Contextual electronic copilot for driving assistance

Participant: David Filliat.

Financing of the CIFRE PhD grant of Alexandre Armand by Renault SAS with the goal of developping an Contextual electronic copilot for driving assistance based on the learning of the behavior of the driver.

8. Partnerships and Cooperations

8.1. Scientific Collaborations (outside consortium projects)

8.1.1. Collaboration and technological transfer with Laboratoire de Physiologie de la Perception et de l'Action (LPPA)

A collaboration is in progress with Jacques Droulez and Steve Nguyen from Laboratoire de Physiologie de la Perception et de l'Action (LPPA), Paris. Poppy represents for them a humanoid platform very interesting because it is relatively flexible and versatile, with more similar proportions to that of humans, which facilitate comparison with the experimental results obtained in humans. The laboratory will evaluate this platform probabilistic methods of control of balance and locomotion.

In the short term the first experimental project with Poppy will test methods of management support, in the case of restoration of balance, in the case of walking to correct or prepare a change of direction. This project will be initiated in the framework of a long internship of master 2 that starts in January. In the future, we would also like to evaluate motor controllers compliant, and learning algorithms. This collaboration involves Matthieu Lapeyre and Pierre-Yves Oudeyer.

8.1.2. Collaborations with Gipsa-Lab, Laboratoire de Psychologie et de Neurocognition (LPNC) and Laboratoire de Physiologie de la Perception et de l'Action (LPPA)

Clément Moulin-Frier is continuing his collaborative work with people he worked with during his PhD thesis at GIPSA-Lab, LPNC and LPPA. See the section entitled "COSMO (Communicating about Objects using Sensory-Motor Operations): a Bayesian modeling framework for studying speech communication and the emergence of phonological systems" for more information. He is also continuing his collaborative work with people he worked with during his post-doc in 2011 at LPPA. See the section entitled "Probabilistic optimal control: a quasimetric approach" for more information.

8.1.3. Collaboration with the Computer Science Department of the University of Zaragoza

A collaboration is in progress with Iñaki Iturrate and Luis Montesano at Zaragoza University, Spain. We aim a develloping a calibration free Brain Computer Interaction system through the use and extension of learning algorithm developped in the team [43], [45], [44]. We focus our effort on error related potentials that occur in the brain while observing or performing a task. They supposedly play a role in human learning as implicit feedback signals that evaluate the correctness or unexpectedness of received stimuli. Our goal is to automatically and reliably detect and classify these signals to provide feedback to artificial systems (e.g. a robot) that learn how to interact and adapt themselves to the user intentions and preferences.

8.2. Regional Initiatives

8.2.1. FUI ROBOT POPULI

This project leaded by Awabot (http://www.awabot.com) funded from 2012 to 2014 aims to investigate, prototype, and test new applications and interactions between the robot and the user to move from niche markets to the general public. This project builds on the theories of Geoffrey Moore (Crossing the Chasm), putting the user at the center of the product design and following the vision of a playful robot and connected to the cloud, where the robot is an interface for advanced interactive entertainment of the future. It brings together partners with complementary expertise to develop and / or adapt and integrate technological bricks missing to fulfill such a vision. Our goal in this project is to develop a robust and low cost navigation system based on RGB-D cameras.

Partners : ARTEFACTS STUDIO, LIRIS, ENSTA ParisTech, GAMAGORA (Université Lyon 2)

8.2.2. PSPC ROMEO 2

This project leaded by Aldebaran Robotics (http://www.aldebaran-robotics.com/) funded from 2012 to 2016 by OSEO aims at developping a humanoid robot for assisting people. The contribution of FLOWERS and ENSTA ParisTech are in the area of human-robot interaction, learning by demonstration, perception and semantic mapping.

Partners : ALL4TEC, Inria, CNRS, VOXLER, SPIROPS, ISIR, UVSQ, CEA LIST, ENSTA ParisTech, STRATE COLLEGE, TELECOM PARISTECH, ASSOCIATION APPROCHE

Web site: http://www.aldebaran-robotics.com/fr/Projets/romeo.html

8.3. National Initiatives

8.3.1. ANR MACSi

An ANR Project (MACSi, ANR Blanc 0216 02), coordinated by ISIR/Univesity Paris VI (Olivier Sigaud), on developmental robotics (motor learning, visual learning, and exploration algorithms on the ICub robot) continued. The MACSi project is a developmental robotics project based on the iCub humanoid robot and the Urbi open source software platform. It is funded an as ANR Blanc project from 2010 to 2013. The project addresses four fundamental challenges, led by four partners:

- How can a robot learn efficient perceptual representations of its body and of external objects given initially only low-level perceptual capabilities? Challenge leader : Inria-ENSTA-ParisTech FLOWERS (Paris).
- How can a robot learn motor representations and use them to build basic affordant reaching and manipulation skills? Challenge leader : ISIR-UPMC-Paris 6 (Paris). ISIR hosts the iCub humanoid robot on which the achievements will be evaluated.
- What guidance heuristics should be used to explore vast sensorimotor spaces in unknown changing bodies and environments? Challenge leader : Inria-ENSTA-ParisTech FLOWERS (Bordeaux).
- How can mechanisms for building efficient representations/abstractions, mechanisms for learning manipulation skills, and guidance mechanisms be integrated in the same experimental robotic architecture and reused for different robots? Challenge leader : GOSTAI company (Paris).

Web site: http://macsi.isir.upmc.fr/

8.4. European Initiatives

8.4.1. FP7 Projects

8.4.1.1. 3rd HAND

Type: COOPERATION

Defi:ICT-2013.2.1 Robotics, Cognitive Systems & Smart Spaces, Symbiotic Interaction

Instrument: Collaborative project

Objectif: Target a) Intelligent robotics systems

Duration: October 2013 - September 2017

Coordinator: Inria, France

Partner: Universitaet Darmstadt, Germany

Partner: Stuttgart University, Germany

Partner: University of Innsbruck, Austria

Inria contact: Manuel Lopes

Abstract: Robots have been essential for keeping industrial manufacturing in Europe. Most factories have large numbers of robots in a fixed setup and few programs that produce the exact same product hundreds of thousands times. The only common interaction between the robot and the human worker has become the so-called ”emergency stop button”. As a result, re-programming robots for new or personalized products has become a key bottleneck for keeping manufacturing jobs in Europe. The core requirement to date has been the production in large numbers or at a high price. Robot-based small series production requires a major breakthrough in robotics: the development of a new class of semi-autonomous robots that can decrease this cost substantially. Such robots need to be aware of the human worker, alleviating him from the monotonous repetitive tasks while keeping him in the loop where his intelligence makes a substantial difference.

In this project, we pursue this breakthrough by developing a semi-autonomous robot assistant that acts as a third hand of a human worker. It will be straightforward to instruct even by an untrained layman worker, allow for efficient knowledge transfer between tasks and enable a effective collaboration between a human worker with a robot third hand. The main contributions of this project will be the scientific principles of semi-autonomous human-robot collaboration, a new semi-autonomous robotic system that is able to: i) learn cooperative tasks from demonstration; ii) learn from instruction; and iii) transfer knowledge between tasks and environments.

8.4.1.2. ERC EXPLORERS

Instrument: ERC Starting Grant

Duration: December 2009 - November 2014

Coordinator: Pierre-Yves Oudeyer, Inria.

Abstract: In spite of considerable and impressive work in artificial intelligence, machine learning, and pattern recognition in the past 50 years, we have no machine capable of adapting to the physical and social environment with the flexibility, robustness and versatility of a 6-months old human child. Instead of trying to simulate directly the adult's intelligence, EXPLORERS proposes to focus on the developmental processes that give rise to intelligence in infants by re-implementing them in machines. Framed in the developmental/epigenetic robotics research agenda, and grounded in research in human developmental psychology, its main target is to build robotic machines capable of autonomously learning and re-using a variety of skills and know-how that were not specified at design time, and with initially limited knowledge of the body and of the environment in which it will operate. This implies several fundamental issues: How can a robot discover its body and its relationships with the physical and social environment? How can it learn new skills without the intervention of an engineer? What internal motivations shall guide its exploration of vast spaces of skills? Can it learn through natural social interactions with humans? How to represent the learnt skills and how can they be re-used? EXPLORERS attacks directly those questions by proposing a series of scientific and technological advances: 1) we will formalize and implement sophisticated systems of intrinsic motivation, responsible of organized spontaneous exploration in humans, for the regulation of the growth of complexity of learning situations; 2) intrinsic motivation systems will be used to drive the learning of forward/anticipative sensorimotor models in high-dimensional multimodal spaces, as well as the building of reusable behavioural macros; 3) intrinsically motivated exploration will be coupled with social guidance from non-engineer humans; 4) an information-theoretic framework will complement intrinsically motivated exploration to allow for the inference of body maps; 5) we will show how learnt basic sensorimotor skills can be re-used to learn the meaning of early concrete words, pushing forward human-robot mutual understanding. Furthermore, we will setup large scale experiments, in order to show how these advances can allow a high-dimensional multimodal robot to learn collections of skills continuously in a weeks-to-months time scale. This project not only addresses fundamental scientific questions, but also relates to important societal issues: personal home robots are bound to become part of everyday life in the 21st century, in particular as helpful social companions in an aging society. EXPLORERS' objectives converge to the challenges implied by this vision: robots will have to be able to adapt and learn new skills in the unknown homes of users who are not engineers.

8.5. International Initiatives

8.5.1. Inria Associate Teams

8.5.1.1. NEUROCURIOSITY

Title: NeuroCuriosity

Inria principal investigator: Manuel Lopes

International Partner (Institution - Laboratory - Researcher):

Columbia Neuroscience (United States) - Jacqueline Goetlieb

Duration: 2013 - 2015

One of the most striking aspects of human behavior is our enormous curiosity, drive for exploration. From a child feverishly examining a new toy with its hands and its eyes, to a tourist exploring a new city, to a scientist studying the brain, humans incessantly want to know. This exuberant curiosity shapes our private and social lives, and is arguably a key cognitive feature that allows our species to understand, control and alter our world. We aim to develop a novel unified biological and computational theory, which explains curiosity in the domain of visual exploration and attention as a deliberate decision motivated by learning progress. This theory will build and improve upon pioneer computational models of intrinsic motivation elaborated in developmental robotics, and be empirically evaluated in the context of visual exploration in monkeys through behavioral and brain imaging techniques. This will be the first attempt at a biological-computational framework of intrinsic motivation and perceptual exploration and their underlying cognitive mechanisms. This collaboration involves Pierre-Yves Oudeyer and Manuel Lopes on the Inria side, and Jacqueline Gottlieb and Adrien Baranes on Univ. Columbia side.

8.6. International Research Visitors

8.6.1. Visits of International Scientists

- Jan Peters, Technische Universitaet Darmstadt
- Marc Toussaint, Stuttgart University
- Justus Piater, University of Innsbruck
- Luis Montesano, University of Zaragoza
- Michael Mistry, Lecturer in Robotics, Intelligent Robotics Lab, University of Birmingham
- Andrej Gams, Post-doc, Biorobotics Laboratory, EPFL
- Adrien Baranes, Columbia University, NY, US
- Katharina Rohlfing, Bielefeld University, Germany
- Yannis Demiris, Imperial College, UK
- Andrew Barto, Univ. Massachussets at Amherst, US

8.6.1.1. Internships

- Jules Brochard, Emergent Proximo-Distal Maturation through Adaptive Exploration
- Axel Davy, Safe exploration in MDPs
- Julie Golliot, Experimental Platform for User Study of Curiosity-driven Exploration
- Brice Miard, Experimental Platform for User Study of Curiosity-driven Exploration
- Chloé Rozenbaum, Learning Simultaneously New Tasks and Feedback Models in Socially Guided Robot Learning
- Caio Tomazelli Da Silva Oliveira, Multimodal learning of speech-action-video primitives

8.6.2. Visits to International Teams

PY Oudeyer visited Gottlieb’s Cognitive Neuroscience lab at Columbia University, NY, US; CITEC at Bielefeld University, Germany.

F. Stulp visited the Max Planck Institute for Intelligent Systems (Stefan Schaal) in Tuebingen, Germany.

In May 2013, Matthieu Lapeyre visited the Bristol Robotic Lab to present the Poppy robot. A close collaboration will begin in 2014, in particular they will hire an engineer to design grasping hand for Poppy.

In May 2013, Jonathan Grizou visited Iñaki Iturrate and Luis Montesano at Zaragoza University, Spain.

In July 2013, Manuel Lopes visit the lab of Andrea Thomaz at Georgia Tech.

In August 2013, Clément Moulin-Frier visited the Honda Research Center in Tokyo as well as Pr. Sawada at Kagawa University, Japan. He gave talks in both labs. He also visited the Intelligent Robotics Laboratory directed by Prof. Hiroshi Ishiguro and Asada Laboratory directed by Prof. Minoru Asada, in Osaka. In October 2013 he visited the Developmental Neuromechanics & Communication Lab at Princeton University, USA.

In August 2013, Fabien Benureau, Olivier Mangin, Mai Nguyen and Jonathan Grizou, visited the Intelligent Robotics Laboratory directed by Prof. Hiroshi Ishiguro, Osaka; the Humanoid Robotics Institute directed by Prof. Atsuo Takanishi, Tokyo; Intelligent Systems and Informatics Laboratory directed by Prof. Yasuo Kuniyoshi, Tokyo; and Asada Laboratory directed by Prof. Minoru Asada, Osaka.

In October 2013, Manuel Lopes, Clément Moulin-Frier and Mai Nguyen visited the Cognitive Neuroscience lab of Jacqueline Gottlieb in New York.

9. Dissemination

9.1. Editorial Activities

9.1.1. Editorial boards

Pierre-Yves Oudeyer has been editor of IEEE CIS Newsletter on AMD, and Fabien Benureau has been associate editor.

Pierre-Yves Oudeyer has been of journals IEEE TAMD, International Journal of Social Robotics, and Frontiers in Neurorobotics.

PY Oudeyer is member of the IEEE CIS Technical Committee on Autonomous Mental Development. (Key instance for the organization of the IEEE ICDL conference series).

M. Lopes is member of the steering committee of the IEEE RAS TC on Robot Learning.

D. Filliat has been associate editor for the IROS 2013 conference.

9.1.2. PC committees and reviewing

- PY Oudeyer has been reviewer for journals (Proceedings of IEEE, Topics in Cognitive Science, Frontiers in Psychology,), member of conference program committees (Humanoids 2013, ICDL-Epirob 2013, IROS 2013, ICRA 2013), and reviewer and expert for several European projects of the ICT Program.
- D Filliat has been member of the evaluation committee for the ANR Apprentissages program, member of the program committee for the REACTS workshop, reviewer for the Autonomous Robots journal and reviewer for conferences (ECMR,ICRA,IROS,IV).
- Jonathan Grizou has been reviewer for ICDL-Epirob 2013 and Robotica 2013 conferences.
- Clément Moulin-Frier has been reviewer for the ICDL-Epirob 2013 conference and the IEEE Transactions on Autonomous Mental Development journal.
- Alexander Gepperth has been reviewer for IROS 2013, ITSC 2013, IJCNN 2013, for the journals Cognitive Computation, Transactions on Intelligent Transportation systems, Neuroco;puting, Neural Networks and Neural Processing Letters. He has been program committe member for IJCNN 2013.
- Freek Stulp was reviewer for journals (IEEE Transactions on Robotics, Autonomous Robots, IEEE Transactions on Autonomous Mental Development, Journal of Intelligent and Robotic Systems), conferences (International Joint Conference on Neural Networks (IJCNN), IEEE International Conference on Intelligent Robots and Systems (IROS), IEEE International Conference on Robotics and Automation (ICRA), IEEE Humanoids), and member of the program committee of ICRA.

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

License: Introduction à Matlab, 21 heures. L3 - ENSTA ParisTech (Alexander Gepperth)

License: Traitement numérique du signal , 21 heures. L3 - ENSTA ParisTech (Alexander Gepperth)

Master: Apprentissage Autonomme, 6 heurs. M2, ENSEIRB - Bordeaux (Pierre-Yves Oudeyer).

Master: Apprentissage Autonomme, 6 heurs. M2, ENSEIRB - Bordeaux (Manuel Lopes).

Master: Robotique de Compagnie, 6 heurs. M2, ENSTA - Paris Tech (Manuel Lopes).

Master: Robotique Mobile, 24 heures. M2, ENSTA - ParisTech (David Filliat).

Master: Vision pour la robotique, 12 heures. M2, University Pierre et Marie Curie (David Filliat).

License: Introduction to Matlab, 21 heures. L3, ENSTA - ParisTech (David Filliat).

Licence 2: Graphe, Langage, Cryptologie, 21 heures. Pole Universitaire Francais de Ho Chi Minh Ville

Master: Option Robotique, Projet Robot Autonome, 32 heures. ENSEIRB, Bordeaux, France.

Licence: Mathématique, 20 heures, niveau (L1), Université de Bordeaux (Olivier Mangin)

Licence: Infomatique, Introduction à la robotique, 30 heures, niveau (L3), Université de Bordeaux (Olivier Mangin et Fabien Bénureau)

9.2.2. Supervision

PhD finished: Mai Nguyen, Bootstrapping Intrinsically Motivated Learning with Human Demonstration, defended nov. 2013 (superv. Pierre-Yves Oudeyer).

PhD finished: Thomas Cederborg, A unified view of context-dependant skill learning and language acquisition, defended dec. 2013 (superv. Pierre-Yves Oudeyer).

PhD finished: Natalia Lyubova, A developmental approach to perception for a humanoid robot [23], defended october 30th, 2013 (superv. David Filliat).

PhD in progress: Louis-Charles Caron, Developmental learning in multimodal sensory-motor loops, started january 2012 (superv. Alexander Gepperth).

PhD in progress: Guillaume Duceux, Navigation and exploration based on RGB-D cameras, started october 2011 (superv. David Filliat).

PhD in progress: Alexandre Armand, Contextual electronic copilot for driving assistance, started feb. 2011 (superv. David Filliat)

PhD in progress: Matthieu Lapeyre, Developmental constraints for biped humanoid walking, started oct. 2010 (superv. Pierre-Yves Oudeyer and Olivier Ly).

PhD in progress: Fabien Benureau, Cumulative, hierarchical and intrinsically motivated learning of robot skills, started oct. 2010 (superv. Pierre-Yves Oudeyer).

PhD in progress: Jonathan Grizou, Fluid simultaneous learning of task and feedback models, started oct. 2011 (superv. Manuel Lopes and Pierre-Yves Oudeyer).

PhD in progress: Olivier Mangin, Learning of sensorimotor primitives with Non-Negative Matrix Factorization, started oct. 2010 (superv. Pierre-Yves Oudeyer).

PhD in progress: Ievgen Perederieiev, Adaptive task execution for implicit human-robot coordination, started oct. 2012 (superv. Manuel Lopes).

PhD in progress: Thibaut Munzer, Learning from Instruction, started oct. 2013 (superv. Manuel Lopes).

PhD in progress: Yuxin Chen, Interactive learning of objects and names on a humanoid robot, started oct. 2013 (superv. David Filliat).

9.2.3. Juries

Edouard Klein (Manuel Lopes, 2013, examinateur) Contributions à l'apprentissage par renforcement inverse, supervised by Yann Guermeur and Matthieu Geist, Université de Lorraine, France

Pedro Sequeira (Manuel Lopes, 2013, examinateur) Socio-Emotional Reward Design for Intrinsically Motivated Learning Agents, supervised by Ana Paiva and Francisco Melo, Instituto Superior Técnico, Portugal

Hemanth Korrapati (03/07/13, David Filliat, Examinateur) : Loop Closure for Topological Mapping and Navigation with Omnidirectional Images. Université Blaise Pascal - Clermont-Ferrand II, Youcef Mezouar (Dir.)

Cédric Meyer (14/06/13, David Filliat, Examinateur) : Théorie de vision dynamique et application à la robotique mobile. Université Pierre et Marie Curie, Ryad Benosman (Dir.)

Srikrishna Bhat (22/01/13, David Filliat, Rapporteur) : Visual words for pose computation. Université de Lorraine (22/01/2013), Marie-Odile Berger (Dir.)

Beata Gryzb (04/13, PY Oudeyer, Rapporteur): Grounding spatial awareness in sensorimotor representations: an interdisciplinary approach, Univ. Madrid, Spain

Xavier Hinault (01/12, PY Oudeyer, Rapporteur): Réseau de neurones récurrent pour le traitement de séquences abstraites et de structures grammaticales, avec une application aux interactions homme-robot, Univ. Lyon 1, France.

Antoine de Rengervé (12/13, PY Oudeyer, Rapporteur): Apprentissage interactif en robotique autonome : vers de nouveaux types d'IHM, Univ. Cergy Pontoise.

Michael Garcia Ortiz (7/2013, A.Gepperth, Rapporteur): Driver behavior prediction, Univ. de Bielefeld (Allemagne)

Sarra Jlassi (28/11/13, Freek Stulp, Examinateur) : Formulation et Etude des Problemes de Commande en Co-manipulation Robotique. Laboratoire des Signaux et Systemes, Supelec (Dir. Yacine Chitour)

9.3. Popularization and Outreach

9.3.1. Popular Science Publications

Pierre-Yves Oudeyer published a popular science book entitled "Aux sources de la parole" at Odile Jacob, http://www.pyoudeyer.com/AuxSourcesDeLaParole.htm

9.3.2. Popular Science Radio Broadcast

PY Oudeyer participated to several popular science radio broadcast:

France Culture (28th october 2013), La parole et l'ordinateur, Interview avec Stéphane Delogeorges, émission Continent Sciences. http://www.franceculture.fr/emission-continent-sciences-la-parole-et-l-ordinateur-2013-10-28

France Inter (27 octobre) Le langage: une auto-organisation ? Interview avec Stéphane Paoli, émission 3D, le journal. http://www.franceinter.fr/emission-3d-le-journal-ou-va-lhumanite-et-langage-une-auto-organisation

RFI (Sept. 2013) Comment s'invente le langage ? (Interview 55 mn), Emission "Autour de la question" de Caroline Lachowsky.http://www.rfi.fr/emission/20130909-1-comment-s-invente-le-langage

9.3.3. Popular Science Talks

Musée des Arts et Métiers (7 nov. 2013) Aux sources de la parole, Parole d'auteurs, rencontre avec PY Oudeyer animée par Daniel Fiévet (France Inter).

9.3.4. Museum Exhibitions, Science Festivals and Public Demonstrations

9.3.4.1. Educatec-Educatice

Kidlearn have been presented on Educatec-Educatice exhibition the 20,21 and 22 November 2013. This exhibition was held in Paris and grouped together 220 education professional and more than 12000 visitors. It was an opportunity to present Kidlearn project and create contact for potential partnership.

9.3.4.2. Cap Sciences exhibition on “Brain and Cognition”

Cap Sciences is an organization in Bordeaux to promote and to communicate about science to the public. Cap Sciences is preparing an exhibit about the brain starting in February 2013. The Flowers team will contribute to this exposition by setting up a booth to explain the complexity of the processing required for intelligent artificial systems (e.g. robots) to transform observations from the environment to actions done in this environment, such processing being done continuously by all living beings, most notably by nervous systems and brains. To explain this idea, the Flowers team is working on a game for the visitors of the exhibit: a player has to drive forward a mobile robot, specifically an iRobot Roomba, while avoiding obstacles. The difficulty for the visitor in this game is that the player is not able to watch the robot in its environment: the player has to control it using only the sensory-information displayed on a computer screen. The player wins when the robot has traveled a given distance in a straight line and in limited time without bumping into an obstacle. This exhibit will start in February 2013 and last for a year. After that, it may move to different locations. More than 100,000 visitors are expected in Cap Sciences, half of them will come from elementary schools. The data generated by the robot and the visitors will be logged and will be available for research on life long learning with robots.

9.3.4.3. NOVAQT

Poppy and Kidlearn have been presented in the NOVAQT exhibition the 5, 6 and 7 December 2013. This exhibition was held in Bordeaux and brought together a selection of 108 innovations created in the region Aquitaine. It was a chance to present our research to 3 different public: industry, education and general public. Feedbacks were very enthusiastic and we have found several potential collaborations with local actors.

9.3.4.4. Fète de la science

On october 13th, the team organised robotics demonstrations for the "Fete de la science" in ENSTA ParisTech. We demontrated physical human-robot interaction using the Meka robot, object segmentation and recognition through a Kinect camera and navigation and mapping with a Pioneer robot. Around 300 people participated in the event.

9.3.5. Education towards the younger

The team also initiated the development of educational activities in "écoles primaires" et "collèges" to have kids discover robotics and programming, as well as ran experiments in "école primaires" in Aquitaine to test novel educational software to help children learn mathematics, and developped within the KidLearn ADT project. This was achieved thanks to the arrival of Didier Roy, former math teacher in college, in the team. Our mediation of robotics activities aim primarily at students, teachers and schools. We focused on three approaches :

- A website, consisting of video presentation of researchers Flowers team, news related to robotics , questions / answers, ready activities for teachers for their students , a directory of links, videos, and pages for “robotics in schools stories " where teachers and students come to post texts, photos and videos to show what happens in robotics in their establishments. A forum is available for researchers to ask questions. We also test robotic kits to be able to provide advices on choices.
- Direct support of schools . For example, we went to Nevers in high school Follereau present what robotics is, interact with students, discuss robotic upcoming events in their school. We subsequently asked to provide technical assistance in the context of robotic operations in the school. We offer activities and we provide technical and educational advice. We move, if needed, in schools or propose meetings by videoconference .
- We formed a working group of a dozen people, most of whom are teachers in primary and secondary and special education (training for disabled adults returning to education) . This group meets regularly to discuss activities, share experiences within schools and consider a robotic curriculum from primary to the end of schooling.

We also participated in Educatec- Educatice event in Paris where education professionals told great interest in robotic activities in education. We participated in the symposium Didacstic didapro in Clermont -Ferrand, about mediation of digital sciences . We worked with cluster Aquitaine Robotics to develop a reflection on the teaching of robotics.

9.3.6. Press, TV and Web coverage

https://flowers.inria.fr/press/

10. Bibliography

Major publications by the team in recent years

[1] A. ANGELI, D. FILLIAT, S. DONCIEUX, J. MEYER. *Fast and incremental method for loop-closure detection using bags of visual words*, in "Robotics, IEEE Transactions on", 2008, vol. 24, n° 5, pp. 1027–1037

[2] A. BARANES, P.-Y. OUDEYER. *RIAC: Robust Intrinsically Motivated Exploration and Active Learning*, in "IEEE Trans. on Auto. Ment. Dev.", 2009, vol. 1, n° 3, pp. 155-169, http://www.pyoudeyer.com/TAMDBaranesOudeyer09.pdf

[3] A. BARANES, P.-Y. OUDEYER. *Active learning of inverse models with intrinsically motivated goal exploration in robots*, in "Robotics and Autonomous Systems", 2013, vol. 61, n° 1, pp. 49 - 73 [*DOI :* 10.1016/J.ROBOT.2012.05.008], http://www.pyoudeyer.com/RAS-SAGG-RIAC-2012.pdf

[4] J. BUCHLI, F. STULP, E. THEODOROU, S. SCHAAL. *Learning Variable Impedance Control*, in "International Journal of Robotics Research", 2011, vol. 30, n° 7, pp. 820-833, http://ijr.sagepub.com/content/early/2011/03/31/0278364911402527

[5] T. DEGRIS, O. SIGAUD, P. WUILLEMIN. *Learning the Structure of Factored Markov Decision Processes in Reinforcement Learning Problems*, in "Proceedings of the 23rd International Conference on Machine learning (ICML)", 2006, pp. 257–264

[6] T. DEGRIS, M. WHITE, R. SUTTON. *Off-Policy Actor-Critic*, in "International Conference on Machine Learning", 2012, http://hal.inria.fr/hal-00764021

[7] D. FILLIAT. *A visual bag of words method for interactive qualitative localization and mapping*, in "Robotics and Automation, 2007 IEEE International Conference on", IEEE, 2007, pp. 3921–3926

[8] A. GEPPERTH. *Efficient online bootstrapping of sensory representations*, in "Neural Networks", December 2012 [*DOI :* 10.1016/J.NEUNET.2012.11.002], http://hal.inria.fr/hal-00763660

[9] A. GEPPERTH, S. REBHAN, S. HASLER, J. FRITSCH. *Biased competition in visual processing hierarchies: a learning approach using multiple cues*, in "Cognitive Computation", March 2011, vol. 3, n° 1, http://hal.archives-ouvertes.fr/hal-00647809/en/

[10] J. GOTTLIEB, P.-Y. OUDEYER, M. LOPES, A. BARANES. *Information-seeking, curiosity, and attention: computational and neural mechanisms.*, in "Trends in Cognitive Sciences", November 2013, vol. 17, n° 11, pp. 585-93 [*DOI :* 10.1016/J.TICS.2013.09.001], http://hal.inria.fr/hal-00913646

[11] M. LOPES, T. LANG, M. TOUSSAINT, P.-Y. OUDEYER. *Exploration in Model-based Reinforcement Learning by Empirically Estimating Learning Progress*, in "Neural Information Processing Systems (NIPS)", Lake Tahoe, United States, December 2012, http://hal.inria.fr/hal-00755248

[12] M. LOPES, F. MELO, L. MONTESANO. *Active learning for reward estimation in inverse reinforcement learning*, in "Machine Learning and Knowledge Discovery in Databases", 2009, pp. 31–46

[13] L. MONTESANO, M. LOPES, A. BERNARDINO, J. SANTOS-VICTOR. *Learning Object Affordances: From Sensory–Motor Coordination to Imitation*, in "Robotics, IEEE Transactions on", 2008, vol. 24, n^{o} 1, pp. 15–26

[14] S. M. NGUYEN, A. BARANES, P.-Y. OUDEYER. *Bootstrapping Intrinsically Motivated Learning with Human Demonstrations*, in "proceedings of the IEEE International Conference on Development and Learning", Frankfurt, Allemagne, 2011, ERC Grant EXPLORERS 240007, http://hal.archives-ouvertes.fr/hal-00645986

[15] P.-Y. OUDEYER, F. KAPLAN, V. HAFNER. *Intrinsic Motivation Systems for Autonomous Mental Development*, in "IEEE Transactions on Evolutionary Computation", 2007, vol. 11, n^{o} 1, pp. 265–286, http://www.pyoudeyer.com/ims.pdf

[16] P.-Y. OUDEYER. , *Self-Organization in the Evolution of Speech*, Studies in the Evolution of Language, Oxford University Press, 2006

[17] P.-Y. OUDEYER. *On the impact of robotics in behavioral and cognitive sciences: from insect navigation to human cognitive development*, in "IEEE Transactions on Autonomous Mental Development", 2010, vol. 2, n^{o} 1, pp. 2–16, http://hal.inria.fr/inria-00541783/en/

[18] P. ROUANET, P.-Y. OUDEYER, F. DANIEAU, D. FILLIAT. *The Impact of Human-Robot Interfaces on the Learning of Visual Objects*, in "IEEE Transactions on Robotics", January 2013, http://hal.inria.fr/hal-00758241

[19] F. STULP, B. BUCHLI, A. ELLMER, M. MISTRY, E. THEODOROU, S. SCHAAL. *Model-free Reinforcement Learning of Impedance Control in Stochastic Force Fields*, in "IEEE Transactions on Autonomous Mental Development", 2012

[20] F. STULP, A. FEDRIZZI, L. MÖSENLECHNER, M. BEETZ. *Learning and Reasoning with Action-Related Places for Robust Mobile Manipulation*, in "Journal of Artificial Intelligence Research (JAIR)", 2012, vol. 43, pp. 1–42

[21] F. STULP, E. THEODOROU, S. SCHAAL. *Reinforcement Learning with Sequences of Motion Primitives for Robust Manipulation*, in "IEEE Transactions on Robotics", 2012, vol. 28, n^{o} 6, pp. 1360-1370

Publications of the year

Doctoral Dissertations and Habilitation Theses

[22] T. CEDERBORG. , *A Formal Approach to Social Learning: Exploring Language Acquisition Through Imitation*, Université Sciences et Technologies - Bordeaux I, December 2013, http://hal.inria.fr/tel-00937615

[23] N. LYUBOVA. , *Approche développementale de la perception pour un robot humanoïde*, Ecole Nationale Supérieure de Techniques Avancées - ENSTA, October 2013, http://hal.inria.fr/tel-00925067

[24] S. M. NGUYEN. , *A Curious Robot Learner for Interactive Goal-Babbling : Strategically Choosing What, How, When and from Whom to Learn*, Université Sciences et Technologies - Bordeaux I, November 2013, http://hal.inria.fr/tel-00936992

Articles in International Peer-Reviewed Journals

[25] A. BARANES, P.-Y. OUDEYER. *Active Learning of Inverse Models with Intrinsically Motivated Goal Exploration in Robots*, in "Robotics and Autonomous Systems", January 2013, vol. 61, n^o 1, pp. 69-73 [*DOI :* 10.1016/J.ROBOT.2012.05.008], http://hal.inria.fr/hal-00788440

[26] T. CEDERBORG, P.-Y. OUDEYER. *From Language to Motor Gavagai: Unified Imitation Learning of Multiple Linguistic and Non-linguistic Sensorimotor Skills*, in "IEEE Transactions on Autonomous Mental Development (TAMD)", 2013, http://hal.inria.fr/hal-00910982

[27] J. GOTTLIEB, P.-Y. OUDEYER, M. LOPES, A. BARANES. *Information-seeking, curiosity, and attention: computational and neural mechanisms*, in "Trends in Cognitive Sciences", November 2013, vol. 17, n^o 11, pp. 585-93 [*DOI :* 10.1016/J.TICS.2013.09.001], http://hal.inria.fr/hal-00913646

[28] S. IVALDI, S. M. NGUYEN, N. LYUBOVA, A. DRONIOU, V. PADOIS, D. FILLIAT, P.-Y. OUDEYER, O. SIGAUD. *Object learning through active exploration*, in "IEEE Transactions on Autonomous Mental Development", 2013, pp. 1-18, http://hal.inria.fr/hal-00919694

[29] C. MOULIN-FRIER, M. A. ARBIB. *Recognizing Speech in a Novel Accent: The Motor Theory of Speech Perception Reframed*, in "Biological Cybernetics", August 2013, vol. 107, n^o 4, pp. 421-447 [*DOI :* 10.1007/S00422-013-0557-3], http://hal.inria.fr/hal-00864213

[30] C. MOULIN-FRIER, S. M. NGUYEN, P.-Y. OUDEYER. *Self-Organization of Early Vocal Development in Infants and Machines: The Role of Intrinsic Motivation*, in "Frontiers in Psychology", 2013, vol. 4, n^o 1006 [*DOI :* 10.3389/FPSYG.2013.01006], http://hal.inria.fr/hal-00927940

[31] S. N'GUYEN, C. MOULIN-FRIER, J. DROULEZ. *Decision Making under Uncertainty: A Quasimetric Approach*, in "PLoS ONE", December 2013, vol. 8, n^o 12 [*DOI :* 10.1371/JOURNAL.PONE.0083411], http://hal.inria.fr/hal-00922767

[32] S. M. NGUYEN, P.-Y. OUDEYER. *Socially Guided Intrinsic Motivation for Robot Learning of Motor Skills*, in "Autonomous Robots", July 2013, 1 p. [*DOI :* 10.1007/S10514-013-9339-Y], http://hal.inria.fr/hal-00936938

[33] P. PETER, K. MRINAL, M. FRANZISKA, F. STULP, B. JONAS, E. THEODOROU, S. STEFAN. *From Dynamic Movement Primitives to Associative Skill Memories*, in "Robotics and Autonomous Systems", 2013, vol. 61, n^o 4, pp. 351-361, http://hal.inria.fr/hal-00789397

[34] P. ROUANET, P.-Y. OUDEYER, F. DANIEAU, D. FILLIAT. *The Impact of Human-Robot Interfaces on the Learning of Visual Objects*, in "IEEE Transactions on Robotics", April 2013, vol. 29, n^o 2, pp. 525-541 [*DOI :* 10.1109/TRO.2012.2228134], http://hal.inria.fr/hal-00758241

[35] O. SIGAUD, F. STULP. *Adaptation de la matrice de covariance pour l'apprentissage par renforcement direct*, in "Revue d'Intelligence Artificielle", 2013, vol. 27, n^o 2, pp. 243-263, http://hal.inria.fr/hal-00922127

[36] F. STULP, P.-Y. OUDEYER. *Exploration through Covariance Matrix Adaptation Enables Developmental Motor Learning*, in "Paladyn. Journal of Behavioral Robotics,", 2013, pp. 128-135, http://hal.inria.fr/hal-00922125

[37] F. STULP, O. SIGAUD. *Robot Skill Learning: From Reinforcement Learning to Evolution Strategies*, in "Paladyn, Journal of Behavioral Robotics", 2013, vol. 4, n^o 1, pp. 49-61, http://hal.inria.fr/hal-00922132

International Conferences with Proceedings

[38] J. ALMINGOL, L. MONTESANO, M. LOPES. *Learning Multiple Behaviors from Unlabeled Demonstrations in a Latent Controller Space*, in "International Conference on Machine Learning (ICML-13)", Atlanta, United States, 2013, pp. 136-144, http://hal.inria.fr/hal-00871852

[39] A. ARMAND, D. FILLIAT, J. IBANEZ-GUZMAN. *Modelling Stop Intersection Approaches using Gaussian Processes*, in "IEEE International Conference on Intelligent Transportation Systems - ITSC", Netherlands, 2013, xx p. , http://hal.inria.fr/hal-00919680

[40] F. BENUREAU, P.-Y. OUDEYER. *Autonomous Reuse of Motor Exploration Trajectories*, in "International Conference on Development and Learning and on Epigenetic Robotics", Osaka, Japan, August 2013, http://hal.inria.fr/hal-00850759

[41] A. CHAPOULIE, P. RIVES, D. FILLIAT. *Appearance-based segmentation of indoors/outdoors sequences of spherical views*, in "IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, IROS'2013", Tokyo, Japan, November 2013, pp. 1946-1951, http://hal.inria.fr/hal-00845450

[42] M. DUBOIS, P. K. ROZO, A. GEPPERTH, F. A. GONZALEZ, D. FILLIAT. *A Comparison of Geometric and Energy-Based Point Cloud Semantic Segmentation Methods*, in "European Conference on Mobile Robotics (ECMR)", Spain, 2013, http://hal.inria.fr/hal-00919627

[43] J. GRIZOU, I. ITURRATE, L. MONTESANO, M. LOPES, P.-Y. OUDEYER. *Interactive Task Estimation From Unlabelled Teaching Signals*, in "International Workshop on Human-Machine Systems, Cyborgs and Enhancing Devices", Manchester, United Kingdom, October 2013, http://hal.inria.fr/hal-00858210

[44] J. GRIZOU, I. ITURRATE, L. MONTESANO, M. LOPES, P.-Y. OUDEYER. *Zero-calibration BMIs for sequential tasks using error-related potentials*, in "IROS 2013 Workshop on Neuroscience and Robotics", Tokyo, Japan, November 2013, http://hal.inria.fr/hal-00872484

[45] J. GRIZOU, M. LOPES, P.-Y. OUDEYER. *Robot Learning Simultaneously a Task and How to Interpret Human Instructions*, in "Joint IEEE International Conference on Development and Learning an on Epigenetic Robotics (ICDL-EpiRob)", Osaka, Japan, 2013, http://hal.inria.fr/hal-00850703

[46] T. HESTER, P. STONE, M. LOPES. *Learning Exploration Strategies in Model-Based Reinforcement Learning*, in "International Conference on Autonomous Agents and Multiagent Systems, AAMAS'13", St. Paul, MN, United States, 2013, http://hal.inria.fr/hal-00871861

[47] J. KULICK, T. LANG, M. TOUSSAINT, M. LOPES. *Active Learning for Teaching a Robot Grounded Relational Symbols*, in "International Joint Conference on Artificial Intelligence", Beijing, China, 2013, http://hal.inria.fr/hal-00871858

[48] M. LAPEYRE, P. ROUANET, P.-Y. OUDEYER. *Poppy Humanoid Platform: Experimental Evaluation of the Role of a Bio-inspired Thigh Shape*, in "Humanoids 2013", Atlanta, United States, October 2013, http://hal.inria.fr/hal-00861110

[49] M. LAPEYRE, P. ROUANET, P.-Y. OUDEYER. *Poppy: a New Bio-Inspired Humanoid Robot Platform for Biped Locomotion and Physical Human-Robot Interaction*, in "Proceedings of the 6th International Symposium on Adaptive Motion in Animals and Machines (AMAM)", Darmstadt, Germany, March 2013, http://hal.inria.fr/hal-00788433

[50] M. LAPEYRE, P. ROUANET, P.-Y. OUDEYER. *The Poppy Humanoid Robot: Leg Design for Biped Locomotion*, in "IEEE/RSJ International Conference on Intelligent Robots and Systems", Tokyo, Japan, November 2013, http://hal.inria.fr/hal-00852858

[51] N. LYUBOVA, D. FILLIAT, S. IVALDI. *Improving object learning through manipulation and robot self-identification*, in "IEEE International Conference on Robotics and Biomimetics", China, 2013, http://hal.inria.fr/hal-00919649

[52] N. LYUBOVA, S. IVALDI, D. FILLIAT. *Developmental object learning through manipulation and human demonstration*, in "ICRA Mobile Manipulation Workshop on Interactive Perception", Germany, 2013, http://hal.inria.fr/hal-00839519

[53] O. MANGIN, P.-Y. OUDEYER. *Learning Semantic Components from Subsymbolic Multimodal Perception*, in "Joint IEEE International Conference on Development and Learning an on Epigenetic Robotics (ICDL-EpiRob)", Osaka, Japan, IEEE, 2013, http://hal.inria.fr/hal-00842453

[54] C. MOULIN-FRIER, P.-Y. OUDEYER. *Exploration strategies in developmental robotics: a unified probabilistic framework*, in "ICDL-Epirob - International Conference on Development and Learning, Epirob", Osaka, Japan, 2013, http://hal.inria.fr/hal-00860641

[55] C. MOULIN-FRIER, P.-Y. OUDEYER. *Learning how to reach various goals by autonomous interaction with the environment: unification and comparison of exploration strategies*, in "1st Multidisciplinary Conference on Reinforcement Learning and Decision Making (RLDM2013), Princeton University, New Jersey", Princeton, United States, 2013, http://hal.inria.fr/hal-00922537

[56] C. MOULIN-FRIER, P.-Y. OUDEYER. *The role of intrinsic motivations in learning sensorimotor vocal mappings: a developmental robotics study*, in "Interspeech", Lyon, France, ISCA, 2013, http://hal.inria.fr/hal-00860655

[57] S. M. NGUYEN, S. IVALDI, N. LYUBOVA, A. DRONIOU, D. GERARDEAUX-VIRET, D. FILLIAT, V. PADOIS, O. SIGAUD, P.-Y. OUDEYER. *Learning to recognize objects through curiosity-driven manipulation with the iCub humanoid robot*, in "IEEE International Conference on Development and Learning and on Epigenetic Robotics - ICDL-EPIROB", Japan, 2013, pp. 1–8, http://hal.inria.fr/hal-00919674

[58] D. SILVER, G. LEVER, N. HEESS, T. DEGRIS, D. WIERSTRA, M. RIEDMILLER. *Deterministic Policy Gradient Algorithms*, in "ICML", Beijing, China, June 2014, http://hal.inria.fr/hal-00938992

National Conferences with Proceedings

[59] D. FILLIAT. *Software architecture of the PACOM project*, in "Proceedings of the 8th National Conference on Control Architectures of Robots", France, 2013, http://hal.inria.fr/hal-00839518

Conferences without Proceedings

[60] O. MANGIN, P.-Y. OUDEYER. *Inverse Reinforcement Learning of a Dictionary of Primitive Tasks*, in "Workshop on Hierarchical and Structured Learning for Robotics", Berlin, Germany, Gerhard Neumann and George Konidaris and Freek Stulp and Jan Peters, June 2013, http://hal.inria.fr/hal-00931398

[61] F. STULP, G. RAIOLA, A. HOARAU, S. IVALDI, O. SIGAUD. *Learning Compact Parameterized Skills with a Single Regression*, in "IEEE-RAS International Conference on Humanoid Robots", United States, 2013, http://hal.inria.fr/hal-00922135

[62] F. STULP, O. SIGAUD. *Policy Improvement: Between Black-Box Optimization and Episodic Reinforcement Learning*, in "Journées Francophones Planification, Décision, et Apprentissage pour la conduite de systèmes", France, 2013, http://hal.inria.fr/hal-00922133

Scientific Books (or Scientific Book chapters)

[63] P.-Y. OUDEYER, A. BARANES, F. KAPLAN. *Intrinsically Motivated Learning of Real World Sensorimotor Skills with Developmental Constraints*, in "Intrinsically Motivated Learning in Natural and Artificial Systems", G. BALDASSARRE, M. MIROLLI (editors), Springer, February 2013, http://hal.inria.fr/hal-00788611

[64] P. OUDEYER. *Self-Organization: Complex Dynamical Systems in the Evolution of Speech*, in "The Language Phenomenon", P. BINDER, K. SMITH (editors), Springer, 2013, http://hal.inria.fr/hal-00788509

Other Publications

[65] M. LOPES, B. CLEMENT, D. ROY, P.-Y. OUDEYER. , *Multi-Armed Bandits for Intelligent Tutoring Systems*, 2013, http://hal.inria.fr/hal-00913669

References in notes

[66] L. STEELS, R. BROOKS (editors). , *The Artificial Life Route to Artificial Intelligence: Building Embodied, Situated Agents*, L. Erlbaum Associates Inc.Hillsdale, NJ, USA, 1995

[67] B. ARGALL, S. CHERNOVA, M. VELOSO. *A Survey of Robot Learning from Demonstration*, in "Robotics and Autonomous Systems", 2009, vol. 57, n^{o} 5, pp. 469–483

[68] M. ASADA, S. NODA, S. TAWARATSUMIDA, K. HOSODA. *Purposive Behavior Acquisition On A Real Robot By Vision-Based Reinforcement Learning*, in "Machine Learning", 1996, vol. 23, pp. 279-303

[69] A. BARTO, S. SINGH, N. CHENTANEZ. *Intrinsically Motivated Learning of Hierarchical Collections of Skills*, in "Proceedings of the 3rd International Conference on Development and Learning (ICDL 2004)", Salk Institute, San Diego, 2004

[70] D. BERLYNE. , *Conflict, Arousal and Curiosity*, McGraw-Hill, 1960

[71] C. BREAZEAL. , *Designing sociable robots*, The MIT Press, 2004

[72] R. BROOKS, C. BREAZEAL, R. IRIE, C. C. KEMP, B. SCASSELLATI, M. WILLIAMSON. *Alternative essences of intelligence*, in "Proceedings of 15th National Conference on Artificial Intelligence (AAAI-98", AAAI Press, 1998, pp. 961–968

[73] A. CLARK. , *Mindware: An Introduction to the Philosophy of Cognitive Science*, Oxford University Press, 2001

[74] D. COHN, Z. GHAHRAMANI, M. JORDAN. *Active learning with statistical models*, in "Journal of artificial intelligence research", 1996, vol. 4, pp. 129–145

[75] W. CROFT, D. CRUSE. , *Cognitive Linguistics*, Cambridge Textbooks in Linguistics, Cambridge University Press, 2004

[76] M. CSIKSZENTHMIHALYI. , *Flow-the psychology of optimal experience*, Harper Perennial, 1991

[77] B. DAMAS, J. SANTOS-VICTOR. *Online Learning of Single- and Multivalued Functions with an Infinite Mixture of Linear Experts*, in "Neural Computation", 2013, vol. 25, n^{o} 11, pp. 3044–3091

[78] P. DAYAN, W. BELLEINE. *Reward, motivation and reinforcement learning*, in "Neuron", 2002, vol. 36, pp. 285–298

[79] E. DECI, R. RYAN. , *Intrinsic Motivation and Self-Determination in Human Behavior*, Plenum Press, 1985

[80] J. ELMAN. *Learning and development in neural networks: The importance of starting small*, in "Cognition", 1993, vol. 48, pp. 71–99

[81] D. FILLIAT, E. BATTESTI, S. BAZEILLE, G. DUCEUX, A. GEPPERTH, L. HARRATH, I. JEBARI, R. PEREIRA, A. TAPUS, C. MEYER, S.-H. IENG, R. BENOSMAN, E. CIZERON, J.-C. MAMANNA, B. POTHIER. *RGBD object recognition and visual texture classification for indoor semantic mapping*, in "Technologies for Practical Robot Applications (TePRA), 2012 IEEE International Conference on", United States, 2012, pp. 127 - 132, http://hal.inria.fr/hal-00755295

[82] C. FLEISCHER, A. WEGE, K. KONDAK, G. HOMMEL. *Application of EMG signals for controlling exoskeleton robots*, in "Biomedizinische Technik", 2006, vol. 51, n^{o} 5/6, pp. 314–319

[83] A. GEPPERTH, L.-C. CARON. *Simultaneous concept formation driven by predictability*, in "International conference on development and learning", San Diego, États-Unis, October 2012

[84] A. GEPPERTH, B. DITTES, M. GARCIA ORTIZ. *The contribution of context information: a case study of object recognition in an intelligent car*, in "Neurocomputing", February 2012, vol. 94 [*DOI :* 10.1016/J.NEUCOM.2012.03.008], http://hal.inria.fr/hal-00763650

[85] A. GEPPERTH. *Co-training of context models for real-time object detection*, in "IEEE Symposium on Intelligent Vehicles", Madrid, Espagne, June 2012, http://hal.inria.fr/hal-00763676

[86] A. GEPPERTH. *Efficient online bootstrapping of sensory representations*, in "Neural Networks", December 2012 [*DOI :* 10.1016/J.NEUNET.2012.11.002], http://hal.inria.fr/hal-00763660

[87] D. GOUAILLIER, V. HUGEL, P. BLAZEVIC, C. KILNER, J. MONCEAUX, P. LAFOURCADE, B. MARNIER, J. SERRE, B. MAISONNIER. *The nao humanoid: a combination of performance and affordability*, in "CoRR, vol. abs/0807.3223", 2008

[88] I. HA, Y. TAMURA, H. ASAMA, J. HAN, D. HONG. *Development of open humanoid platform DARwIn-OP*, in "SICE Annual Conference (SICE), 2011 Proceedings of", IEEE, 2011, pp. 2178–2181

[89] S. HARNAD. *The symbol grounding problem*, in "Physica D", 1990, vol. 40, pp. 335–346

[90] M. HASENJAGER, H. RITTER. , *Active learning in neural networks*, Physica-Verlag GmbHHeidelberg, Germany, Germany, 2002, pp. 137–169

[91] J. HAUGELAND. , *Artificial Intelligence: the very idea*, The MIT PressCambridge, MA, USA, 1985

[92] J.-C. HORVITZ. *Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events*, in "Neuroscience", 2000, vol. 96, n^{o} 4, pp. 651-656

[93] X. HUANG, J. WENG. *Novelty and reinforcement learning in the value system of developmental robots*, in "Proceedings of the 2nd international workshop on Epigenetic Robotics : Modeling cognitive development in robotic systems", C. PRINCE, Y. DEMIRIS, Y. MAROM, H. KOZIMA, C. BALKENIUS (editors), Lund University Cognitive Studies 94, 2002, pp. 47–55

[94] S. IVALDI, N. LYUBOVA, D. GÉRARDEAUX-VIRET, A. DRONIOU, S. ANZALONE, M. CHETOUANI, D. FILLIAT, O. SIGAUD. *Perception and human interaction for developmental learning of objects and affordances*, in "Proc. of the 12th IEEE-RAS International Conference on Humanoid Robots - HUMANOIDS", Japan, 2012, to appear, http://hal.inria.fr/hal-00755297

[95] M. JOHNSON. , *Developmental Cognitive Neuroscience*, 2nd, Blackwell publishing, 2005

[96] W. B. KNOX, P. STONE. *Combining manual feedback with subsequent MDP reward signals for reinforcement learning*, in "Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'10)", 2010, pp. 5–12

[97] R. LAURENT, C. MOULIN-FRIER, P. BESSIÈRE, J.-L. SCHWARTZ, J. DIARD. , *Integrate, yes, but what and how? A computational approach of perceptuo-motor fusion in speech perception.*, 2012, Accepted commentary in Behavioral and Brain Sciences, in Press., http://hal.inria.fr/hal-00765973

[98] M. LOPES, T. CEDERBORG, P.-Y. OUDEYER. *Simultaneous Acquisition of Task and Feedback Models*, in "Development and Learning (ICDL), 2011 IEEE International Conference on", Germany, 2011, pp. 1 - 7 [*DOI :* 10.1109/DEVLRN.2011.6037359], http://hal.inria.fr/hal-00636166/en

[99] M. LOPES, P.-Y. OUDEYER. *The Strategic Student Approach for Life-Long Exploration and Learning*, in "IEEE Conference on Development and Learning / EpiRob 2012", San Diego, United States, November 2012, http://hal.inria.fr/hal-00755216

[100] M. LUNGARELLA, G. METTA, R. PFEIFER, G. SANDINI. *Developmental Robotics: A Survey*, in "Connection Science", 2003, vol. 15, n^{o} 4, pp. 151-190

[101] O. LY, M. LAPEYRE, P. OUDEYER. *Bio-inspired vertebral column, compliance and semi-passive dynamics in a lightweight humanoid robot*, in "Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on", IEEE, 2011, pp. 1465–1472

[102] N. LYUBOVA, D. FILLIAT. *Developmental Approach for Interactive Object Discovery*, in "Neural Networks (IJCNN), The 2012 International Joint Conference on", Australia, 2012, pp. 1-7, http://hal.inria.fr/hal-00755298

[103] O. MANGIN, P.-Y. OUDEYER. *Learning the Combinatorial Structure of Demonstrated Behaviors with Inverse Feedback Control*, in "Human Behavior Understanding, Lecture notes in computer science", Springer, October 2012, vol. 7559, http://hal.inria.fr/hal-00764448

[104] O. MANGIN, P.-Y. OUDEYER. *Learning to recognize parallel combinations of human motion primitives with linguistic descriptions using non-negative matrix factorization*, in "Proceedings of 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems", IEEE, October 2012, http://hal.inria.fr/hal-00764353

[105] J. MARSHALL, D. BLANK, L. MEEDEN. *An Emergent Framework for Self-Motivation in Developmental Robotics*, in "Proceedings of the 3rd International Conference on Development and Learning (ICDL 2004)", Salk Institute, San Diego, 2004

[106] M. MASON, M. LOPES. *Robot Self-Initiative and Personalization by Learning through Repeated Interactions*, in "6th ACM/IEEE International Conference on Human-Robot", Switzerland, 2011 [*DOI :* 10.1145/1957656.1957814], http://hal.inria.fr/hal-00636164/en

[107] P. MILLER. , *Theories of developmental psychology*, 4th, New York: Worth, 2001

[108] C. MOULIN-FRIER, R. LAURENT, P. BESSIÈRE, J.-L. SCHWARTZ, J. DIARD. *Adverse conditions improve distinguishability of auditory, motor and percep-tuo-motor theories of speech perception: an exploratory Bayesian modeling study*, in "Language and Cognitive Processes", 2012, vol. 27, n^{o} 7-8 Special Issue: Speech Recognition in Adverse Conditions, pp. 1240-1263 [*DOI :* 10.1080/01690965.2011.645313], http://hal.archives-ouvertes.fr/hal-00642311

[109] S. M. NGUYEN, A. BARANES, P.-Y. OUDEYER. *Bootstrapping Intrinsically Motivated Learning with Human Demonstrations*, in "IEEE International Conference on Development and Learning", Frankfurt, Germany, 2011, http://hal.inria.fr/hal-00645986/en

[110] S. M. NGUYEN, A. BARANES, P.-Y. OUDEYER. *Constraining the Size Growth of the Task Space with Socially Guided Intrinsic Motivation using Demonstrations.*, in "IJCAI Workshop on Agents Learning Interactively from Human Teachers (ALIHT)", Barcelona, Spain, 2011, http://hal.inria.fr/hal-00645995/en

[111] S. M. NGUYEN, P.-Y. OUDEYER. *Active Choice of Teachers, Learning Strategies and Goals for a Socially Guided Intrinsic Motivation Learner*, in "Paladyn Journal of Bejavioral Robotics", September 2012, vol. 3, n^{o} 3, pp. 136-146 [*DOI :* 10.2478/s13230-013-0110-z], http://hal.inria.fr/hal-00936932

[112] S. M. NGUYEN, P.-Y. OUDEYER. *Interactive Learning Gives the Tempo to an Intrinsically Motivated Robot Learner*, in "IEEE-RAS International Conference on Humanoid Robots", Osaka, Japan, December 2012, http://hal.inria.fr/hal-00762753

[113] S. M. NGUYEN, P.-Y. OUDEYER. *Properties for Efficient Demonstrations to a Socially Guided Intrinsically Motivated Learner*, in "21st IEEE International Symposium on Robot and Human Interactive Communication", Paris, France, September 2012, http://hal.inria.fr/hal-00762758

[114] S. M. NGUYEN, P.-Y. OUDEYER. *Socially Guided Intrinsically Motivated Learner*, in "IEEE International Conference on Development and Learning", San Diego, United States, 2012 [*DOI :* 10.1109/DEVLRN.2012.6400809], http://hal.inria.fr/hal-00936960

[115] S. M. NGUYEN, P.-Y. OUDEYER. *Whom Will an Intrinsically Motivated Robot Learner Choose to Imitate from?*, in "Post-Graduate Conference on Robotics and Development of Cognition: RobotDoC-PhD 2012", Lausanne, Switzerland, September 2012, http://hal.inria.fr/hal-00762762

[116] R. NOWAK. *The Geometry of Generalized Binary Search*, in "Information Theory, Transactions on", 2011, vol. 57, n^{o} 12, pp. 7893–7906

[117] P.-Y. OUDEYER, F. KAPLAN. *Intelligent adaptive curiosity: a source of self-development*, in "Proceedings of the 4th International Workshop on Epigenetic Robotics", L. BERTHOUZE, H. KOZIMA, C. PRINCE, G. SANDINI, G. STOJANOV, G. METTA, C. BALKENIUS (editors), Lund University Cognitive Studies, 2004, vol. 117, pp. 127–130

[118] P.-Y. OUDEYER, F. KAPLAN. *What is intrinsic motivation? A typology of computational approaches*, in "Frontiers in Neurorobotics", 2007, vol. 1, n^{o} 1

[119] P.-Y. OUDEYER. *Sur les interactions entre la robotique et les sciences de l'esprit et du comportement*, in "Informatique et Sciences Cognitives : influences ou confluences ?", C. GARBAY, D. KAISER (editors), Presses Universitaires de France, 2009, http://hal.inria.fr/inria-00420309/en/

[120] P.-Y. OUDEYER. *L'auto-organisation dans l'évolution de la parole*, in "Parole et Musique: Aux origines du dialogue humain, Colloque annuel du Collège de France", S. DEHAENE, C. PETIT (editors), Odile Jacob, 2009, pp. 83-112, http://hal.inria.fr/inria-00446908/en/

[121] A. REVEL, J. NADEL. *How to build an imitator?*, in "Imitation and Social Learning in Robots, Humans and Animals: Behavioural, Social and Communicative Dimensions", K. DAUTENHAHN, C. NEHANIV (editors), Cambridge University Press, 2004

[122] T. SCHATZ, P.-Y. OUDEYER. *Learning motor dependent Crutchfield's information distance to anticipate changes in the topology of sensory body maps*, in "IEEE International Conference on Learning and Development", Chine Shangai, 2009, http://hal.inria.fr/inria-00420186/en/

[123] M. SCHEMBRI, M. MIROLLI, G. BALDASSARRE. *Evolving internal reinforcers for an intrinsically motivated reinforcement-learning robot*, in "IEEE 6th International Conference on Development and Learning, 2007. ICDL 2007.", July 2007, pp. 282-287, http://dx.doi.org/10.1109/DEVLRN.2007.4354052

[124] J. SCHMIDHUBER. *Curious Model-Building Control Systems*, in "Proceedings of the International Joint Conference on Neural Networks, Singapore", IEEE press, 1991, vol. 2, pp. 1458–1463

[125] W. SCHULTZ, P. DAYAN, P. MONTAGUE. *A neural substrate of prediction and reward*, in "Science", 1997, vol. 275, pp. 1593-1599

[126] M. SCHWARZ, M. SCHREIBER, S. SCHUELLER, M. MISSURA, S. BEHNKE. *NimbRo-OP Humanoid TeenSize Open Platform*

[127] F. STULP, P.-Y. OUDEYER. *Emergent Proximo-Distal Maturation through Adaptive Exploration*, in "International Conference on Development and Learning (ICDL)", 2012

[128] F. STULP, O. SIGAUD. *Path Integral Policy Improvement with Covariance Matrix Adaptation*, in "Proceedings of the 29th International Conference on Machine Learning (ICML)", 2012

[129] F. STULP, O. SIGAUD. , *Policy Improvement Methods: Between Black-Box Optimization and Episodic Reinforcement Learning*, 2012, 34 p. , http://hal.inria.fr/hal-00738463

[130] F. STULP. *Adaptive Exploration for Continual Reinforcement Learning*, in "International Conference on Intelligent Robots and Systems (IROS)", 2012

[131] E. THELEN, L. B. SMITH. , *A dynamic systems approach to the development of cognition and action*, MIT PressCambridge, MA, 1994

[132] A. L. THOMAZ, C. BREAZEAL. *Teachable robots: Understanding human teaching behavior to build more effective robot learners*, in "Artificial Intelligence Journal", 2008, vol. 172, pp. 716-737

[133] A. TURING. *Computing machinery and intelligence*, in "Mind", 1950, vol. 59, pp. 433-460

[134] F. VARELA, E. THOMPSON, E. ROSCH. , *The embodied mind : Cognitive science and human experience*, MIT PressCambridge, MA, 1991

[135] J. WENG, J. MCCLELLAND, A. PENTLAND, O. SPORNS, I. STOCKMAN, M. SUR, E. THELEN. *Autonomous mental development by robots and animals*, in "Science", 2001, vol. 291, pp. 599-600