Activity Report 2011

Project-Team TAO

Machine Learning and Optimisation

IN COLLABORATION WITH: Laboratoire de recherche en informatique (LRI)

Table of contents

<div align="center">

Project-Team TAO

</div>

Keywords: Machine Learning, Statistical Learning, Inference, Evolutionary Algorithms

1. Members

Research Scientists

    Marc Schoenauer [Team co-leader, Senior Researcher INRIA, HdR]
    Michèle Sebag [Team co-leader, Senior Researcher CNRS, HdR]
    Anne Auger [Junior Researcher INRIA]
    Cyril Furtlehner [Junior Researcher INRIA]
    Nikolaus Hansen [Junior Researcher INRIA, HdR]
    Balázs Kégl [Junior Researcher CNRS (at LAL), external collaborator, HdR]
    Yann Ollivier [Junior Researcher CNRS, HdR]
    Olivier Teytaud [Junior Researcher INRIA, HdR]

Faculty Members

    Jamal Atif [Associate Professor Université Paris Sud]
    Nicolas Bredèche [Associate Professor Université Paris Sud, HdR]
    Philippe Caillou [Associate Professor Université Paris Sud]
    Florence D'Alche Buc [Professor, Université Evry, délégation INRIA until Aug. 2012, HdR]
    Cécile Germain-Renaud [Professor, Université Paris Sud, HdR]
    Hélène Paugam-Moisy [Professor, Université Lyon 2, délégation INRIA since Sept. 2011, HdR]

Technical Staff

    Jean-Joseph Christophe [since September 2011]
    Hassen Doghmen
    Julien Nauroy
    Raymond Ros [Until May 2011]

PhD Students

    Riad Akrour [SYMBRION, Started November 2010]
    Alejandro Arbelaez [MSR-Adapt, Defended May 2011]
    Ludovic Arnold [Started January 2010]
    Rémi Bardenet [Started September 2009]
    Zyed Bouzarkouna [CIFRE IFP, started December 2008]
    Robert Busa-Fekete [Started October 2010]
    Alexandre Chotard [Started October 2011]
    Adrien Couëtoux [IOMCA, Started September 2010]
    Dawei Feng [Started October 2010]
    Jérémie Decock [ILab METIS, Started October 2011]
    Nicolas Galichet [Started October 2011]
    Moez Hammami [CIFRE EXQUIM, Started May 2011]
    Jean-Baptiste Hoock [MASH, Started November 2009]
    Yoann Isaac [UnsupBrain, Started October 2011]
    Ilya Loshchilov [CSDL, Started October 2009]
    Gaetan Marceau-Caron [CIFRE THALES, Started May 2011]
    Victorin Martin [IMARA-TAO TRAVESTI, Started December 2009]
    Jean-Marc Montanier [Started October 2009]
    Fabien Teytaud [Defended December 2011]
    Weijia Wang [Started October 2010]
    Mouadh Yagoubi [CIFRE PSA, started February 2009]

Post-Doctoral Fellows

David Auger [until August 2011]
Yohei Akimoto [since April 2011]
Matthias Brendel [until October 2011]
Sylvain Chevallier [until February 2011]
Yufei Han [since September 2011]
Verena Heidrich-Meisner [since April 2011]
Mostepha Redouane Khouadjia [since September 2011]
Yusik Kim [until November 2011]
Nadjib Lazaar [since October 2011]
Manuel Loth [since November 2011]
David Meunier [since July 2011]
Anthony Mouraud [until April 2011]
Sébastien Rebecchi
Nataliya Sokolovska [until September 2011]

Administrative Assistant
Marie-Carol Lopes [Tao]

2. Overall Objectives

2.1. Introduction

Data Mining (DM), acknowledged to be one of the main ten challenges of the 21st century[1], aims at building (partial) phenomenological models from the massive amounts of data produced in scientific labs, industrial plants, banks, hospitals or supermarkets. Machine Learning (ML) likewise aims at modeling the complex systems underlying the available data; the main difference between DM and ML disciplines is the emphasis put on the acquisition, storage and management of large-scale data.

DM and ML problems can be set as optimization problems, thus leading to two possible approaches. Note that this alternative has been characterized by H. Simon (1982) as follows. *In complex real-world situations, optimization becomes approximate optimization since the description of the real-world is radically simplified until reduced to a degree of complication that the decision maker can handle. Satisficing seeks simplification in a somewhat different direction, retaining more of the detail of the real-world situation, but settling for a satisfactory, rather than approximate-best, decision.*

The first approach is to simplify the learning problem to make it tractable by standard statistical or optimization methods. The alternative approach is to preserve as much as possible the genuine complexity of the goals (yielding "interesting" models, accounting for prior knowledge): more flexible optimization approaches are therefore required, such as those offered by Evolutionary Computation.

Symmetrically, optimization techniques are increasingly used in all scientific and technological fields, from optimum design to risk assessment. Evolutionary Computation (EC) techniques, mimicking the Darwinian paradigm of natural evolution, are stochastic population-based dynamical systems that are now widely known for their robustness and flexibility, handling complex search spaces (e.g. mixed, structured, constrained representations) and non-standard optimization goals (e.g. multi-modal, multi-objective, context-sensitive), beyond the reach of standard optimization methods.

The price to pay for such properties of robustness and flexibility is twofold. On one hand, EC is tuned, mostly by trials and errors, using quite a few parameters. On the other hand, EC generates massive amounts of intermediate solutions. It is suggested that the principled exploitation of preliminary runs and intermediate solutions, through Machine Learning and Data Mining techniques, can offer sound ways of adjusting the parameters and finding shortcuts in the trajectories in the search space of the dynamical system.

---

[1]MIT Technological Review, fev. 2001.

2.2. Context and overall goal of the project

The overall goals of the project are to model, to predict, to understand, and to control physical or artificial systems. The central claim is that Learning and Optimization approaches must be used, adapted and integrated in a seamless framework, in order to bridge the gap between the system under study on the one hand, and the expert's goal as to the ideal state/functionality of the system on the other hand.

Specifically, our research context involves the following assumptions:

1. The systems under study range from large-scale engineering systems to physical or chemical phenomenons, including robotics and games. Such systems, sometimes referred to as *complex systems*, can hardly be modeled based on first principles due to their size, their heterogeneity and the incomplete information aspects involved in their behavior.

2. Such systems can be observed; indeed selecting the relevant observations and providing a reasonably appropriate description thereof is part of the problem to be solved. A further assumption is that these observations are sufficient to build a reasonably accurate model of the system under study.

3. The available expertise is sufficient to assess the system state, and any modification thereof, with respect to the desired states/functionalities. The assessment function is usually not a well-behaved function (differentiable, convex, defined on a continuous domain, etc), barring the use of standard optimization approaches and making Evolutionary Computation a better suited alternative.

In this context, the objectives of TAO are threefold:

1. using Evolutionary Computation (EC) and more generally Stochastic Optimization to support Machine Learning (ML);

2. using Statistical Machine Learning to support Evolutionary Computation;

3. investigating integrated ML/EC approaches on diversified and real-world applications.

2.3. Highlights

- DAEYAHSP Winner of the Seventh International Planning Competition (Deterministic Temporal Satisficing track) at ICAPS 2011.

- MoGo realized 20 wins out of 20 games in 7x7 Go against 10 different professional players. This is further documented in [24].

- Results in [5] (just accepted) prove that the Nash equilibrium of two-player zero-sum partially observable games is undecidable. This fundamental result notably contradicts published decidability results, which used as a decidability criterion a definition which is not equivalent to optimal play in the Nash sense.

BEST PAPER AWARD :

[19] **Genetic and Evolutionary Computation Conference (GECCO 2011)**. Z. BOUZARKOUNA, A. AUGER, D. Y. DING.

3. Scientific Foundations

3.1. Scientific Foundations

3.1.1. Introduction

This section describes Tao's main research directions, first presented during Tao's evaluation in November 2007. Four strategic issues had been identified at the crossroad of Machine Learning and Evolutionary Computation:

Where	What is the search space and how to search it.
	Representations, Navigation Operators and Trade-offs.
What	What is the goal and how to assess the solutions.
	Optimal Decision under Uncertainty.
How.1	How to bridge the gap between algorithms and computing architectures ?
	Hardware-aware software and Autonomic Computing.
How.2	How to bridge the gap between algorithms and users?
	Crossing the chasm

Six Special Interest Groups (SIGs) have been defined in TAO, investigating the above complementary issues from different perspectives. The comparatively small size of Tao SIGs enables in-depth and lively discussions; the fact that all TAO members belong to several SIGs, on the basis of their personal interests, enforces the strong and informal collaboration of the groups, and the fast information dissemination.

3.1.2. *Representations and Properties*

The choice of the solution space is known to be the crux of both Machine Learning (model selection) and Evolutionary Computation (genotypic-phenotypic mapping).

The first research theme in TAO thus concerns the definition of an adequate representation, or search space $\mathcal{H}$, together with that of adequate navigation operators. $\mathcal{H}$ and its navigation operators must enforce flexible trade-offs between expressiveness and compacity on the one hand, and stability and versatility on the other hand.

Expressiveness/compacity trade-off (static property): $\mathcal{H}$ should simultaneously include sufficiently complex solutions $-$ i.e. good-enough solutions for the problem at hand $-$ and offer a short description for these solutions, thus making it feasible to find them.

Stability/versatility trade-off (dynamic property): while most modifications of a given solution in $\mathcal{H}$ should only marginally modify its behavior (stability), some modifications should lead to radically different behaviors (versatility). Both properties are required for efficient optimization in complex search spaces; stability, also referred to as "strong causality principle" [98] is needed for optimization to do better than random walk; versatility potentially speeds up optimization through creating short-cuts in the search space.

This research direction is investigated in:

- the Complex System SIG (section 6.2) focusing on developmental representations for Design and sequential representations for Temporal Planning;
- the Large and Deep Networks SIG (section 6.6) considering deep or stochastic Neural Network Topologies;
- the Continuous Optimization SIG (section 6.4), concerned with adaptive representations.

3.1.3. *Optimal Decision Under Uncertainty*

Benefiting from the MoGo expertise, TAO investigates several extensions of the Multi-Armed Bandit (MAB) framework and the Monte-Carlo tree search. Some main issues raised by optimal decision under uncertainty are the following:

- Regret minimization and any-time behavior.
  The any-time issue is tightly related to the scalability of Optimal Decision under Uncertainty; typically, MAB was found better suited than standard Reinforcement Learning to large-scale problems as its criterion (the regret minimization) is more amenable to fast approximations.
- Dynamic environments (non stationary reward functions).
  The dynamic environment issue, first investigated in TAO through the On-line Trading of Exploration vs Exploitation Challenge[2], is relevant to e.g. on-line parameter tuning (see section 6.3).
- Use of side information / Multi-variate MAB

The use of side information by MAB is meant to exploit prior knowledge and/or complementary information about the reward. Typically in MoGo, the end of the game can be described at different levels of precision (e.g., win/lose, difference in the number of stones); estimating the local reward estimate depending on the available side information aims at a better robustness.

- Bounded rationality.
  The bounded rationality issue actually regards two settings. The first one considers a number of options which is large relatively to the time horizon, meaning that only a sample of the possible actions can be considered in the imparted time. The second one deals with a finite *unknown* horizon, as is the case for the Feature Selection problem.

- Multi-objective optimization.
  Many applications actually involve antagonistic criteria; for instance autonomous robot controllers might simultaneously want to explore the robot environment, while preserving the robot integrity. The challenge raised by Multi-objective MAB is to find the "Pareto-front" policies for a moderately increased computational cost compared to the standard mono-objective approach.

This research direction is chiefly investigated by the Optimal Decision Making SIG (section 6.5), in interaction with the Complex System and the Crossing the Chasm SIGs (sections 6.2 and 6.3).

3.1.4. Hardware-Software Bridges

Historically, the apparition of parallel architectures only marginally affected the art of programing; the main focus has been on how to rewrite sequential algorithms to make them parallelism-compliant. The use of distributed architectures however calls for a radically different programming style/computational thinking, seamlessly integrating:

- computation: aggregating the local information available with any information provided by other nodes;

- communication: building abstractions of the local node state, to be transmitted to other nodes;

- assessment: modeling other nodes in order to modulate the exploitation (respectively, the abstraction) of the received (resp. emitted) information.

Message passing algorithms such as Page Rank or Affinity Propagation [92] are prototypical examples of distributed algorithms. The analysis is shifted from the static properties (termination and computational complexity) to the dynamic properties (convergence and approximation) of the algorithms, after the guiding principles of complex systems.

Symmetrically, modern computing systems are increasingly viewed as complex systems of their own, due to their ever increasing resources and computational load. The huge need of scalable administration tools, supporting grid monitoring and maintenance of the job running process, paved the way toward Autonomic Computing [94]. Autonomic Computing (AC) Systems are meant to feature self-configuring, self-healing, self-protecting and self-optimizing skills [99]. A key milestone for Autonomic Computing is to provide the system with a phenomenological model of itself (self-aware system), built from the system logs using Machine Learning and Data Mining.

This research direction is investigated in the Complex System SIG (section 6.2) and in the Autonomic Computing SIG (section 6.1).

---

[2]The OTEE Challenge, funded by Touch Clarity Ltd and organized by the PASCAL Network of Excellence, models the selection of news to be displayed by a Web site as a multi-armed bandit, where the user's interests are prone to sudden changes; the OTEE Challenge was won by the TAO team in 2006.

3.1.5. Crossing the chasm

This fourth strategic priority, inspired by Moore's book [97], is motivated by the fact that many outstandingly efficient algorithms never make it out of research labs. One reason for it is the difference between editor's and programmer's view of algorithms. In the perspective of software editors, an algorithm is best viewed as a single "Go" button. The programmer's perspective is radically different: as he/she sees that various functionalities can be ented on the same algorithmic core, the number of options steadily increases (with the consequence that users usually master less than 10% of the available functionalities). Independently, the programmer gradually acquires some idea of the flexibility needed to handle different application domains; this flexibility is most usually achieved through defining parameters and tuning them. Parameter tuning thus becomes a barrier to the efficient use of new algorithms.

This research direction is chiefly investigated by the Crossing the Chasm SIG (section 6.3) and also by the Continuous Optimization SIG (section 6.4).

4. Application Domains

4.1. Application Domains

Since its creation, TAO mainstream applications regard Numerical Engineering, Autonomous Robotics, and Control and Games. Two new fields of applications, due to the arrival of Cécile Germain (Pr UPS, 2005), Philippe Caillou (MdC, 2005), Balázs Kégl (CR CNRS LAL, 2006) and Cyril Furtlehner (CR INRIA, 2007) have been considered: Autonomic Computing and Complex Systems.

Numerical Engineering still is a major source of applications. The successful OMD (Optimization Multi-Disciplinaire) RNTL/ANR project is being resumed by OMD2, started in July 2009. Collaborations with IFP and PSA automobile industry respectively led to Zyed Bouzarkouna's and Mouadh Yagoubi's PhD CIFRE. TAO leads the Work Package "Optimization" in the System@atic CSDL project, responsible for both fundamental research on surrogate models in multi-objective optimization and the setup of a software platform, that lead to Ilya Loshchilov's PhD work. A collaboration with CEA DM2S was conducted as a Digiteo project and lead to Philippe Rolet's PhD on simplified models.

Autonomous Software Robotics is rooted in our participation to the SYMBRION European IP and SyDiN-MaLaS (ANR-JST, coll. University of Kyushu). On this topic, Jean-Marc Montanier started his PhD in Sept. 2009; Vladimir Skortsov did his Post-doc from Sept. 2009 to Sept. 2010; Weijia Wang and Riad Akrour started their PhDs in Sept. 2010. See Section 6.2.

Our activity in **Control and Games** is chiefly visible through Mogo, already mentioned in the Highlights. Another application regards Brain Computer Interfaces: the Digiteo project *Digibrain* (coll. with CEA List and Neurospin), with Cedric Gouy-Pailler's postdoc from October 2009 to October 2010.

Applications related to **Autonomic Computing** became an important part of TAO activities, led by Cécile Germain and Balázs Kégl in tight collaboration with the Laboratoire de l'Accélérateur Linéaire (section 6.1). Applications related to **Social Systems** are led by Philippe Caillou and Cyril Furtlehner, respectively investigating multi-agent models for labor market, and road traffic models (ANR project TRAVESTI, coordinated by C. Furtlehner, started in 2009). Last but not least, the arrival in TAO of Jamal Atif brought a new application fied in image analysis and understanding.

5. Software

5.1. MoGo

Participants: Olivier Teytaud [correspondent], Hassen Doghmen, Jean-Baptiste Hoock.

MoGo and its Franco-Taiwanese counterpart MoGoTW is a Monte-Carlo Tree Search program for the game of Go, which made several milestones of computer-Go in the past (first wins against professional players in 19x19; first win with disadvantageous side in 9x9 Go); MoGo has had new developments as follows:

- A Meta-MCTS module (inspired by the collaboration with Tristan Cazenave in the ANR EXPLO-RA project), which provided both a huge opening book in 9x9 and an approximate solving of 7x7 Go[24].

- Following the "poolRave" modification, introduction of machine learning and statistics into MCTS, such as:

  - Bernstein Races [41] (for offline educating Monte-Carlo simulations).

  These developments have been summarized and compared in [93].

- Variants of Go: we tested variants of Go, in particular blind variants; this suggests that in such frameworks playing theoretically suboptimal moves helps a lot, because such unnatural moves are harder to memorize. A preliminary related publication is [25]; some additional results are to be published. Another interesting variant is random-Go, starting from a randomly generated board; such situations are much harder for humans, and, interestingly, our program was competitive in front of a 6D player (ranked 4th in a world amateur championship and former French champion) on a 19x19 board[40].

MoGo's development team was awarded the 2010 ChessBase award for the best contribution to Computer-Games.

5.2. Covariance Matrix Adaptation Evolution Strategy

Participant: Nikolaus Hansen [correspondent].

The Covariance Matrix Adaptation Evolution Strategy (CMA-ES) is one of the most powerful continuous domain evolutionary algorithms. The CMA-ES is considered state-of-the-art in continuous domain evolutionary computation[3] and has been shown to be highly competitive on different problem classes. The algorithm is widely used in research and industry as witnessed by hundreds of published applications. We provide source code for the CMA-ES in C, Java, Matlab, Octave, Python, and Scilab including the latest variants of the algorithm.

Links: http://www.lri.fr/~hansen/cmaes_inmatlab.html

5.3. COmparing Continuous Optimizers

Participants: Nikolaus Hansen [correspondent], Raymond Ros, Anne Auger, Marc Schoenauer.

COCO (COmparing Continuous Optimizers) is a platform for systematic and sound comparisons of real-parameter global optimizers. COCO provides benchmark function testbeds (noiseless and noisy) and tools for processing and visualizing data generated by one or several optimizers. The code for processing experiments is provided in Matlab and C. The post-processing code is provided in Python. The code has been improved and used for the GECCO 2009 and 2010 workshops on "Black Box Optimization Benchmarking" (BBOB) (see Section 6.4), and will serve as a basis for the test platform in the CSDL project.

Link: http://coco.gforge.inria.fr/doku.php

5.4. GridObservatory

Participants: Michèle Sebag, Cécile Germain-Renaud [correspondent], Julien Nauroy, Yusik Kim.

---

[3]H.-G. Beyer (2007). Evolution Strategies, *Scholarpedia*, page 1965.

The Grid Observatory software suite collects and publishes traces of the EGI (European Grid Initiative) grid usage. With the release and extensions of its portal, the Grid Observatory has made a database of grid usage traces available to the wider computer science community. These data are stored on the grid, and made accessible through a web portal without the need of grid credentials. More than 100 users are currently registered. The GO is supported by an INRIA ADT (Action de Développement Technologique).

In 2011, the suite has been extended to energy consumption. The first barrier to improved energy efficiency of IT systems is the lack of large-scale collections of experimental data. The Green Computing Observatory (GCO), part of the GO initiative monitors a large computing center (Laboratoire de l'Accélérateur Linéaire - LAL) within the EGI grid, and publishes the data through the Grid Observatory. The GCO is supported by the CNRS PEPS program, and by University Paris-Sud through the MRM (Moyens de Recherche Mutualisés) program.

Portal site: http://www.grid-observatory.org

6. New Results

6.1. Autonomous Computing

Participants: Cécile Germain-Renaud, Michèle Sebag, Balázs Kégl, Yusik Kim, Julien Nauroy, Dawei Feng.

Within the classical objectives of Autonomics (self-*), two transversal lines of research have emerged.

Coping with non-stationarity Most existing work on modeling the dynamics of grid behavior assumes a steady-state system and concludes to some form of long-range dependence (slowly decaying correlation) in the associated time-series. But the physical (economic and sociologic) processes governing the grid behavior dispel the stationarity hypothesis. [68] proposes a categorization of the methods integrate non-stationarity into grid modeling. [9] considers a specific class of models: a sequence of stationary processes separated by breakpoints. The model selection question is now defined as identifying the breakpoints and fitting the processes in each segment, together with a validation methodology that empirically addresses the current lack of theoretical results concerning the quality of the estimated model parameters. Even when stationarity is acceptable, the markovian assumption might be too bold. [54] integrate Echo State Network-based regression into a reinforcement learning in continuous state space for fitting the Q function, with application to reactive grid scheduling.

Traces mining In order for an autonomic system to continuously infer knowledge from its monitoring (the so-called MAPE-K, Monitor-Analyze-Plan-Execute-Knowledge) loop, heterogeneous sources of data have to be integrated. [37] exemplifies two use cases of the Grid Observatory data on evaluating the perfomance of the major EGI scheduler, and blackhole detection.

The Green Computing Observatory [38] data include the detailed monitoring of the processors and motherboards, as well as the global site information, such as overall consumption and overall cooling. The data schema for publication is grounded in an ontology of measurements developed in collaboration with the MIS (Modélisation, Information et Systèmes) laboratory of University Picardie Jules Verne.

[42] proposes a new approach for analyzing behavioral traces: as most of them are indeed text documents, state of the art techniques in text mining, including Latent Dirichlet Allocation, can be exploited . The advantages are twofold: providing some level of explanation inferred from the data; and a relatively scalable way to capture the temporal variability of the behavior of interest, while retaining the full dimensionality of the problem at hand. A promising perspective for combining this approach and inferred segmentation has been identified and is currently explored.

6.2. Complex Systems

Participants: Jamal Atif, Nicolas Bredèche, Matthias Brendel, Cyril Furtlehner, Philippe Caillou, Jean-Marc Montanier, Hélène Paugam-Moisy, Marc Schoenauer, Michèle Sebag.

Evolutionary AI Planning: Divide And Evolve (DAE) DAE solves AI-planning problems by using an Evolutionary Algorithm to sequentially divide them into hopefully simpler problems that are handled by some embedded "classical" planner. Within the ANR project DESCARWIN, work has continued in collaboration with Thalès Research & Technology and ONERA Toulouse. A large part of the work this year has been devoted to writing a brand new version of the DAE software, facing difficulties of parallelization [90]. The resulting program entered the 7th International Planning Competition (IPC 2011) at the 21st International Conference on Planning and Scheduling (ICAPS 2011) and won the Gold Medal in the *Temporal Track*. Note that the Silver Medal was won by Vincent Vidal, also member of the DESCARWIN team, using his planner YAHSP2 – the one that won the Gold Medal while embedded in DAE, thus demonstrating one more the added value of the DAE approach. Meanwhile, because DAE has many parameters (like most Evolutionary Algorithms), parameter tuning within DAE remains a difficult task, and an original approach has been proposed to learn the parameters based on some instance features [49], [50], [51]. Note that this method is however relevant of the "Crossing the Chasm" SIG (see Section 6.3), as it can be applied to any optimization algorithm that handles several instances of the same class.

Distributed Autonomous Robotics Resuming work done in 2010, we investigated further the issue of robotic swarm control whenever the environment is partially or completely unknown. This research is at the cross-road of Evolutionary Computation, Machine Learning and Robotics, and a light influence from Evolutionary Ecology, but with a strong focus on engineering (ie. the goal remains to design algorithms). The topic we are interested in is the design of environment-driven self-adaptive distributed algorithms to enable survival at the level of a population of independent robotic units. The population is limited in size, and hardware implementation within real robots has already been achieved [7]. We have also focused our attention on specific aspects of swarm evolutionary dynamics under specific constraints, including the evolution of cooperative and/or altruistic behaviours [53], [52]. This research yielded interesting results, such as the emergence of altruistic behavior under simple, but specific, algorithmic constraint, as well as tuning mechanism to control the level of altruistic behavior in a population of robots. Perspectives of this work is currently under investigation.

The work done in 2010 about the division of labor among asynchronous and decentralized agents, where each agent is modelled from the competition between two spiking neurons, was further analyzed within a spatio-temporal (simulated) frame. The phase transitions between the asynchronous, the aperiodic and periodic synchronous regimes (depending on the sociability and excitability of the agents) was confirmed, with some counter-intuitive results about the overall merits and efficiency of synchronous behaviors [23].

We have also explored objective-driven online learning within real robotic hardware, both for single robot online behavior learning [78] as well as small group of robots for pattern formation learning [43]. Our activity in Evolutionary Robotics has also been strenghtened by the publication of book which gather several contributions from major actors in the field [75], including an introduction paper on current trends and challenges in this domain [76].

From a slightly different perspective, our work on evolving generative and developmental representations has been continued, with an extensive study of robustness within developmental systems [8] and an investigation of the temporal dynamics at work within genetic regulatory networks for design [32]. While not stricly related to robotics, these contributions share the distributed nature of computation and ultimately aim at providing an efficient representation for designing and controlling large scale passive or active assembly of units (e.g. robots with complex morphologies).

Additionally, at the crossroad of Machine Learning and Evolutionary Computation, a new Reinforcement Learning approach based on modelling the user's preferences was proposed [12], [84]; in the so-called *Preference-based Policy Learning*, the robot demonstrates some behaviors, is informed of the user's preferences, builds a model of the user's preferences and self-trains to build a new behavior hopefully more satisfactory according to the conjectured user's preferences.

Statistical Physics Perspective Basic tools from statistical physics (scaling, mean-field techniques and associated distributed algorithms, exactly-solvable models) and probability have been used to model and optimize complex systems, either standalone or combined with MABS approaches. Results are
• In the context of the ANR TRAVESTI project dealing with spatial and temporal modelling of traffic congestion we have studied in [83] some specific properties of the Belief Propagation algorithm used for inference; we have proposed a way to encode dependencies between real variables with a latent binary MRF [89]; we have analyzed macro-states on traffic data and how these relate to belief propagation fixed points [35].
• Also in the ANR Travesti context [36], [88] for modelling congestion at the microscopic level we have proposed a new family of queueing processes where the service rate is coupled stochastically to the number of clients. With this formulation we have been able to relate an asymmetry between acceleration and braking to some condensation mechanism.In this framework we have also proposed a large deviation formulation of the fundamental diagram of traffic flow.
• In the design of multi-objective message passing algorithms we have shown how the state-of-the art MAXSAT solver SP-Y can be used directly by an endogeneous clause elimination procedure, to sample Pareto Front of multi-objective 3-SAT problems [87].

Multi-agent and games • Within the InnovNation serious game project, in collaboration with Paraschool and BlueNove, we developed an e-brainstorming prototype for collaborative ideation. The prototype was tested on 150 students and is being improved to be commercialized. The realist network generator used to analyze the prototype games was also used to study the labor market, and to study the relative importance of friends and colleagues while seeking for a job without ([44] or with [64] variable information transmission speed. We have shown that friends were the most useful when the labor market was at the equilibrium (approximately the same number of jobs and applicants).
• To analyze the logs of multi-agent based simulations (for example for the InnovNation project), we developed a tool to describe homogeneous agent clusters and their evolution ([73]). We use the cluster description to build agent models and generate new simulations with this model to validate the results. • Also in the complex social system analysis context, we applied LSA text mining tools on research projects and patent category descriptions to associate research clusters to their main research fields [70]. This was used to build a classification of the French research clusters based on their context, and especially the adequation of the cluster research specialization with the regional specialization [72].

Image understanding Within the context of image understanding, a new sequential recognition framework has been proposed in [10]. Sequential image understanding refers to the decision making paradigm where objects in an image are successively segmented/recognized following a predefined strategy. Such an approach generally raises some questions about the "best" segmentation sequence to follow and/or how to avoid error propagation. In [10], we propose original approaches to answer these questions in the case where the objects to segment/recognize are represented by a model describing the spatial relations between objects. The process is guided by a criterion derived from visual attention, and more precisely from a saliency map, along with some spatial information to focus the attention. This criterion is used to optimize the segmentation sequence. Spatial knowledge is also used to ensure the consistency of the results and to allow backtracking on the segmentation order if needed. The proposed approach was applied for the segmentation of internal brain structures in magnetic resonance images. The results show the relevance of the optimization criteria and the interest of the backtracking procedure to guarantee good and consistent results. From a logical standpoint, sequential object recognition is formulated as an abduction process in [14], [66]. A scene is viewed as an observation and the task of interpretation is considered as the "best"

explanation considering the prior knowledge about the scene context. Towards this aim, we introduce an algebraic-based framework unifying mathematical morphology, description logics and formal concept analysis. We propose to compute the best explanations of an observation through algebraic erosion over the Concept Lattice of a background theory which is efficiently constructed using tools from Formal Concept Analysis. We show that the defined operators are sound and complete and satisfy important rationality postulates of abductive reasoning.

6.3. Crossing the Chasm

Participants: Alejandro Arbelaez, Anne Auger, Robert Busa-Fekete, Nikolaus Hansen, Balázs Kégl, Manuel Loth, Nadjib Lazaar, Marc Schoenauer, Michèle Sebag.

Due to the departure of both PhD students funded within the Microsoft-INRIA joint lab after their successful defenses (Alvaro Fialho and Alejandro Arbelaez), some of the activities of this SIG have been slightly redefined this year, with the one-month visit of Prof. Th. Runarsson (University of Iceland) in October, and the arrival in November of two new post-docs, also funded by the joint lab (Nadjib Lazaar and Manuel Loth). A new direction of research has appeared, in line with both Adaptive Operator Selection (Alvaro Fialho's PhD) and Continuous Search (Alejandro Arbelaez' PhD).

Bandit-based choice of heuristics in combinatorial optimization This new direction of research deals with heuristic choice within an existing combinatorial solver using bandit-like algorithms, and the very first results deal with scheduling problems and will be published in early 2012 [57].

Adaptive Operator Selection In line with his PhD work, Alvaro Fialho has succesfully used his Adaptive Operator Selection method to the on-line tuning of Differential Evolution in the multi-objective case [96].

Learn and Optimize (LaO), an instance-based parameter-tuning method. Though originally designed for Divide-And-Evolve framework (see Section 6.2), LaO is a generic method that learns the relationship between some instance features and the optimal parameters of the optimizer. The current version [49], [50], [51] uses Neural Network to directly learn the optimal parameters, and average performance increase compared to the default parameter set (that has won the temporal track in the IPC7 competition) is of more than 10%. On-going work uses rankSVM to learn a partial order on the features $\times$ parameter space.

Adaptive Constraint Programming Alejandro Arbelaez defended his PhD on May 31., led under the supervision of Youssef Hamadi and Michèle Sebag [1]. A survey of his PhD work has been published as a book chapter [74] and some of his last work was more concerned with optimizing the collaboration in distributed SAT solving in highly parallel environments [13].

Ranking by calibrated AdaBoost In [22], [21] we describe a learning-to-rank technique based on calibrated multi-class classification. We train a set of multi-class classifiers using AdaBoost.MH, we calibrate them using various techniques to obtain diverse class probability estimates, and, finally, we approximate the Bayes-scoring function (which optimizes the popular Information Retrieval performance measure NDCG), through mixing these estimates into an ultimate scoring function. Our method outperforms many standard ranking algorithms on the LETOR benchmark datasets, most of which are based on significantly more complex learning to rank algorithms than ours.

6.4. Continuous Optimization

Participants: Yohei Akimoto, Anne Auger, Zyed Bouzarkouna, Alexandre Chotard, Nikolaus Hansen, Ilya Loshchilov, Verena Heidrich-Meisner, Raymond Ros, Marc Schoenauer, Olivier Teytaud, Fabien Teytaud.

Our main expertise in continuous optimization is on stochastic search algorithms. We address theory, algorithm design and applications. The methods we investigate are adaptive techniques able to learn iteratively parameters of the distribution used to sample solutions. The Covariance Matrix Adaptation Evolution Strategy (CMA-ES) is nowadays one of the most powerful method for continuous optimization without derivatives. We work on different variants of the CMA-ES to improve it in various contexts as described below. In addition we have contributed to give an information geometry perspective to stochastic optimization unifying both continuous and discrete algorithms using a family of probability distribution parametrized by continuous parameters. The framework proposed in this context allow to retrieve many existing stochastic optimization algorithms when instantiated with different family of probability distributions including the CMA-ES when using gaussian distributions [85]. We have moreover clarified important design principles based on invariances [85], [11].

New algorithms based on derandomization and for mixed-integer optimization A new variant of CMA-ES to address problem with mixed-integer variables (vectors with both discrete and continuous variables) has been designed [81]. New algorithms using new selection schemes combined with derandomization have been designed and thoroughly theoretically and empirically investigated [16], [17]. A local search algorithm using an adaptive coordinate descent has been proposed [45]. We have as well investigated how to inject solutions in CMA-ES so as to improve performances if an oracle provide good solutions [82].

Distributed optimization We have proposed simple modifications of evolutionary algorithms so that they reach asymptotically the optimal $\log(\lambda)$ speed-up with λ processors and the linear speed-up $\Theta(\lambda)$ for a number of processors of order at most the dimensionality of the problem (for a pointwise solution); in particular bounding the selected population size in the Self-Adaptation algorithm and variants of this idea for fastening the decrease of the step-size in a relevant manner. All these works and more are part of Olivier Teytaud's HDR thesis [3], defended on April 22., and also build the first part of Fabien Teytaud's PhD [2], defended on December 8.. We wrote a chapter on lower bounds for distributed derivative free optimization in [79].

Benchmarking We have continued our effort for improving standards in benchmarking and pursued the development of the COCO - COmparing Continuous Optimizers platform [80]. We are organizing for the GECCO 2012 conference the Black-Box-Optimization Workshop[4].

Optimization with meta-models and surrogate We have investigated optimization using a coupling of CMA-ES and surrogates and applied it for the optimization of well placement [6]. We have proposed a new meta-model CMA-ES for the optimization of partially separable functions [19] and shown that it improves performances for solving the well placement problem [19].

Hyperparameter optimization In [18] we present hyper-parameter optimization results on tasks of training neural networks and deep belief networks (DBNs). We optimize hyper-parameters using random search and two new greedy sequential methods based on the expected improvement criterion. The sequential algorithms are applied to the most difficult DBN learning problems and find significantly better results than the best previously reported.

Multi-objective optimization We have investigated theoretically multi-objective algorithms based on the hypervolume [4] and proposed new selection operators based on tournament and multi-armed bandit framework [46].

Multimodal optimization We have shown in [58] a simple algorithm (a $(1 + 1)$-ES with quasi-random restart and murder operator), which, at least in the sequential case, performs as efficiently as much more tricky algorithms.

Mathematical bounds for noisy optimization The paper [27] shows upper and lower confidence bounds and/or experiment algorithms in the noisy optimization setting; in particular we compared an optimization algorithm based on bandits and an surrogate-model version; whereas the bandit approach is much faster if the noise decreases quickly to zero around the optimum, the surrogate-model version is faster if the noise does not decrease to zero.

---

[4]see http://coco.gforge.inria.fr/doku.php?id=bbob-2012

6.5. Optimal Decision Making

Participants: Olivier Teytaud, David Auger, Michèle Sebag, Cyril Furtlehner, Jean-Baptiste Hoock, Nataliya Sokolovska, Fabien Teytaud, Hassen Doghmen, Jean-Joseph Christophe, Jérémie Decock.

- Monte-Carlo Tree Search (MCTS) and Upper Confidence Trees (UCT) are main areas of the team. In particular, we ultra-weakly solved 7x7 Go by winning 20 games out of 20 against professional players in 7x7 Go, thanks to a Meta-Monte-Carlo-Tree Search[24]. The wins were with komi 9.5 as white, and 8.5 as black, suggesting that the ideal komi in 7x7 is 9.We also applied this algorithm to the recent "NoGo" framework, aimed at challenging MCTS for a game which looks like Go but with very different goals; our paper [26] was the first one applying MCTS to NoGo and now all strong programs use the MCTS approach for NoGo. We extended RAVE (Rapid Action Value Estimates) to the continuous settings [29]. In his PhD [2], Fabien Teytaud proposed several generic improvements of MCTS, including the use of (fast) decisive and anti-decisive moves for games, and applied it to the game of Havannah. An industrial application (to energy management) is proposed in [71]. A MCTS version for partially observable problems with bounded horizon was proposed in [86]. This version is proposed for the two-player case, but for simulations starting at the root; a version in the one-player case, starting from an arbitrary state (and therefore much more efficient for large horizon) is proposed in [30]. This work is extended by a belief state estimation by constraint satisfaction problems in [62]. Other developments and research around MCTS/UCT are described in the MoGo module.

- A related important algorithm is Nested Monte-Carlo; we got state of the art results for some traveling salesman variants with a very simple algorithm in [56].

- Fundamental analysis of partially observable games: we proved in [5] that partially observable games are undecidable (result also presented in the BIRS 2010 workshop and the Bielefeld seminar on Search Methodologies), even in the case of finite state spaces and deterministic transitions. This unexpected result is a priori a contradiction with known decidability results; this emphasizes the subtle difference between the classical decision problem (the existence of a strategy winning certainly, whatever may be the strategy of the opponent), which is used is most analysis, and the choice of the move with optimal winning probability. We pointed out that the relevant decision problem is, with no doubt, the latter; that the other decision problem has just been used because it is equivalent to choosing optimal play in the case of fully observable games; and, most importantly, that partially observable games are in fact undecidable, even in the finite deterministic case. On the other hand, on restricted settings, we have shown by some simple lemmas lower and upper bounds on the value of some partially observable games [63]. We extended Monte-Carlo Tree Search to the case of short-term partial information in [61]; this was successfully applied to the Urban Rival game, a widely played internet card game (now 17 millions of registered users) from a French company.

- Tuning of strategies: tuning strategies is a noisy optimization problem in which the convenient "variance of noise decreasing to zero around the optimum" usually does not hold. We have shown that in such a setting, the local bandit-style algorithms are slower than surrogate models; this is detailed in the continuous optimization part.

- We organized various computer-Go events, as due to the fame of our program MoGo we are often invited for such events; reports can be found in [95].

- We developed the "double progressive widening" trick, which is aimed at making consistent an algorithm from the finite case to the continuous stochastic case; we got good results in [60] on Q-Learning (with no mathematical proof) and on MCTS [28] (mathematical proof to be submitted soon).

- We have also worked on Nash equilibria of Matrix Games, where we proposed an algorithm for finding Nash equilibria faster when the Nash equilibrium is sparse [48], [47]; a mathematical proof is ready and will be submitted soon.

- Some works are in progress around applications of previous tools to active learning; active learning has also been investigated through conditional random fields in [59].

- Another related work, with motivations from autonomous robotics, combines the exploration of the search space through UCT, with an explicit model of the safe regions explored so far, called *Deja-Vu*. The Deja-Vu is used to constrain the exploration, mostly in the random phase, and is updated from the current explorations [67].

- The Ilab "Metis" just started; it's an Ilab between Tao, the Inria-Saclay team MaxPlus, and the SME Artelys http://www.artelys.com for a joint work on numerical libraries in Energy Management.

6.6. Large and Deep Networks

Participants: Ludovic Arnold, Sylvain Chevallier, Anthony Mouraud, Hélène Paugam-Moisy, Sébastien Rebecchi, Michèle Sebag.

Deep networks: RBM or AA The two main families of deep networks are implemented and studied by TAO: stacked RBM (Restricted Boltzmann Machines) and stacked AA (Auto-Associators).

Learning sparse features for deep networks Inspired by the theory of compressed sensing and beyond the common methods based on dictionary learning, we have proposed to learn sparsity and accuracy simultaneously by alternating two constraints on the weights of an Auto-Associator [55].

Spiking neuron networks for swarm robotics The model "SpikeAnts" [91] has been applied to a spatial robotic environment [23], in collaboration with Nicolas Bredeche (see Section 6.2, and has demonstrated even more its interest in the context of swarm robotics.

7. Contracts and Grants with Industry

7.1. Contracts with Industry

- **IFP** – 2008-2011 (24 kEur). *Optimisation de puits non-conventionnels: type, position et trajectoire*, side-contract to Zyed Bouzarkouna's CIFRE Ph.D.;
  Participants: Anne Auger, Zyed Bouzarkouna, Marc Schoenauer.

- **PSA** – 2009-2012 (45 kEur). , side-contract to Mouadh Yagoubi's CIFRE Ph.D.;
  Participants: Marc Schoenauer, Mouadh Yagoubi.

- **THALES** – 2011-2014 (40 kEur). , side-contract to Gaetan Marceau-Caron's CIFRE Ph.D.;
  Participants: Marc Schoenauer, Gaetan Marceau-Caron.

- **EXQIM** – 2011-2014 (40 kEur). , side-contract to Moez Hammami's CIFRE Ph.D.;
  Participants: Michèle Sebag, Moez Hammami.

- **ILAB METIS** – 2012- . Participants: O. Teytaud, M. Schoenauer, S. Gaubert (team Max-Plus), J.-J. Christophe (engineer), J. Decock (ph.D.).

7.2. Grants with Industry

- **ILAB METIS** – 2012- . Participants: O. Teytaud, M. Schoenauer, S. Gaubert (team Max-Plus), J.-J. Christophe (engineer), J. Decock (ph.D.).

8. Partnerships and Cooperations

8.1. Regional Initiatives

- **JASMIN**– 2010-2012 (205 kEur). DRIRE *programme FEDER*.
  Participants: CADLM, Intercim, TAO (Michèle Sebag).

- **CSDL**– 2009-2012 (290 kEur). FUI System@tic (Région Ile de France grant). *Complex System Design Lab*
  Participants: Anne Auger, Nikolaus Hansen, Ilya Loshchilov, Raymond Ros, Marc Schoenauer.

8.2. National Initiatives

- **INNOV NATION**– 2009-2011 (69 kEur). Fonds de Compétitivité des Entreprises. *Simulation multi-agent de diffusion et d'évolution d'idées dans un réseau social dynamique (application à un serious game).*
  Participants: Philippe Caillou, Samuel Thiriot.

- **OMD2** – 2009-2012 (131 kEur).*Optimisation Multi-Disciplinaire Distribuée*, ANR *programme COSINUS* Coordinator Maryan Sidorkiewicz, RENAULT Technocentre;
  Participants: Anne Auger, Yohei Akimoto, Nikolaus Hansen, Marc Schoenauer, Olivier Teytaud.

- **SyDiNMaLaS** – 2009-2012 (158 kEur).*Integrating Symbolic Discovery with Numerical Machine Learning for Autonomous Swarm Control*, ANR *programme BLANC* Coordinator Michèle Sebag, CNRS;
  Participants: David Meunier, Marc Schoenauer, Michèle Sebag.

- **TRAVESTI** – 2009-2012 (206 kEur).*Estimation du volume de Trafic par Interface Spatio-temporelle*, ANR *programme SYSCOMM 2008* Coordinator Cyril Furtlehner, INRIA;
  Participants: Anne Auger, Cyril Furtlehner, Victorin Martin, Maxim Samsonov.

- **ASAP** – 2009-2012 (178 kEur).*Apprentissage Statistique par une Architecture Profonde*, ANR *programme DEFIS 2009* Coordinator Alain Rakotomamonjy, LITIS, Université de Rouen, France;
  Participants: Sylvain Chevallier, Hélène Paugam-Moisy, Sébastien Rebecchi, Michèle Sebag.

- **IOMCA** 2010-2013 (264 kEur). Including Ontologies in Monte-Carlo Tree Search and Applications, ANR international project coordinated by O. Teytaud (Tao, INRIA).
  Participants: Adrien Couëtoux, O. Teytaud.

- **EXPLORA** 2010-2012 (289 kEur, to be shared with Inria Lille). EXPLOitation pour l'Allocation efficace de Ressources. Applications l'optimisation. ANR Project coordinated by R. Munos (INRIA Lille).
  Participants: David Auger, Olivier Teytaud.

- **DESCARWIN** 2010-2013 (201 kEur). Coordinateur P. Savéant, Thalès.
  Participants: Matthias Brendel, Mostepha-Redouane Khouadjia, Marc Schoenauer.

- **SIMINOLE** 2010-2014 (1180k, 250k for TAO). Large-scale simulation-based probabilistic inference, optimization, and discriminative learning with applications in experimental physics, ANR project, Coordinator B. Kégl (CNRS LAL).
  Participants: Balázs Kégl, Rémi Bardenet, Nikolaus Hansen, Michèle Sebag, Cécile Germain

8.3. European Initiatives

8.3.1. FP7 Projects

8.3.1.1. SYMBRION

Title: Symbiotic Evolutionary Robots Organisms

Type: COOPERATION (ICT)

Defi: Embedded systems design

Instrument: Integrated Project (IP)

Duration: February 2008 - January 2013

Coordinator: Universität Stuttgart (Germany)

Others partners: Vereniging voor christelijk hoger onderwijs, wetenschappelijk onderzoek en patiëntenzorg, Netherlands; Universität Graz, Austria; Universität Karlsruhe, Germany; Vlaams Interuniversitair Instituut Voor biotechnologie VZW, Blegium; University of the West of England, Bristol, United Kingdom; Eberhard Karls Universität Tübingen, Germany; University of York, United Kingdom; Université libre de Bruxelles, Belgium; INRIA, France.

See also: http://symbrion.eu

8.3.1.2. MASH

Title: Massive Sets of Heuristics For Machine Learning

Type: COOPERATION (ICT)

Defi: Cognitive Systems and Robotics

Instrument: Specific Targeted Research Project (STREP)

Duration: December 2010 - December 2012

Coordinator: IDIAP Research Institute (Switzerland)

Others partners: Centre National de la Recherche Scientifique, France; Weierstrass-Institut fur Angewandte Analysis Und Stochastik, Part of Furschungsverbund Berlin E.V, Germany; INRIA, France; Ceske Vysoke Uceni Technicke V Praze,Czech Republic.

See also: http://mash-project.eu/

8.3.1.3. CITINES

Title: Design of a decision support tool for sustainable, reliable and cost-effective energy strategies in cities and industrial complexes

Type: COOPERATION (ICT)

Defi: Smart Cities and Communities

Instrument: Specific Targeted Research Project (STREP)

Duration: October 2011 - March 2014

Coordinator: Artelys SA (France)

Others partners: Austrian Institute of Technology, Austria; INESC Porto, Portugal; ARMINES (CMA), France; SCHNEIDER ELECTRIC, France; City of Cesena, Italie; City of Bologna, Italy; Tupras - Turkish Petroleum Refineries Corporation, Turkey; ERVET, Italy; INRIA, France.

See also: Artelys Web site

8.4. Contracts managed by CNRS or Paris-Sud University

- **PASCAL2**, Network of Excellence, 2008-2013 (34 kE in 2008, 70 kE in 2009). Coordinator John Shawe-Taylor, University of Southampton. M. Sebag is manager of the Challenge Programme.

- **EGI** FP7 Infrastructure - 2010-2013 (48 kEur) Participants: Cécile Germain, Michèle Sebag, Davy Feng, Julien Nauroy

- **Grille Paris-Sud** MRM (Moyens de Recherche Mutualisés) 2010-2011 (23KE). Coordinator Balázs Kégl Participants: Cécile Germain, Michèle Sebag, Xiangliang Zhang, Julien Perez, Davy Feng,Julien Nauroy.

- **DIGIBRAIN**– 2007-2011(48 kEur). DIGITEO grant, coordinator Jean-Denis Muller CEA LIST, France
  Participants: Cédric Gouy-Pailler, Michèle Sebag.

- **Unsupervised-Brain**– 2011-2014(5 kEur). DIGITEO grant, coordinator Michèle Sebag LRI Université Paris Sud, France
  Participants: Yoann Isaac, Cédric Gouy-Pailler, Michèle Sebag.
- **MetaModel** – 2008-2011 (150 kEur). Advanced methodologies for modeling interdependent systems - applications in experimental physics, ANR "jeune chercheur" grant, coordinator Balázs Kégl
  Participants: Michèle Sebag, Cécile Germain, Robert Busa-Fekete

8.5. International Initiatives

8.5.1. INRIA International Partners

- NUTN (National University of Tainan, Taiwan). Collaboration of Olivier Teytaud around MoGo (see Invitation section below).
- University of Iceland. Prof. Thomas Philip Runarsson was invited for one month in TAO (October 2011) to work on bandit-based choice of heuristics in combinatorial optimization [57].

8.5.2. Visits of International Scientists

8.5.2.1. Invitations

- Olivier Teytaud (CR1) is invited researcher in NUTN (National University of Tainan, Taiwan) for one year.
- Adrien Couëtoux (ph.D. student) is in internship in NUTN (National University of Tainan, Taiwan) for 6 months.
- Jérémie Decock (ph.D. student) will go to NUTN (National University of Tainan, Taiwan) for 5 months.
- Jean-Baptiste Hoock (ph.D. student) has spent 12 days in Univ. Potsdam (October 2011) for the Mash project.

9. Dissemination

9.1. Animation of the scientific community

- Anne Auger
  - THRaSH, *Theory of Randomized Search Heuristics workshop*, member of Steering Committee;
  - Editorial Board of *Evolutionary Computation*, MIT Press;
  - Olivier Teytaud: Committee at EvoStar, ICML, Lion, Advances in Computer Games.

- Nicolas Bredeche
  - PC member at GECCO 2011, RIVF 2011, EuroGP 2011
  - Co-organiser of DevLeaNN workshop, a two-day workshop on Development and Learning in Artificial Neural Networks, Paris, France (http://devleann.iscpif.fr).
  - New Horizons in Evolutionary Robotics. co-editor, Springer [75].

- Philippe Caillou
  - PC member at EPIA 2011, V2CS 2011
  - Coordinator of the SimTools Network (RNSC Network)

- Nikolaus Hansen

–   Editorial Board of *Evolutionary Computation*, MIT Press;

- Marc Schoenauer
    –   Member of ACM-SIGEVO Executive since 2003 (Special Interest Group on Evolutionary Computation (was the International Society on Genetic and Evolutionary Algorithms before 2006); member of ACM-GECCO Business Committee (2012-2013).
    –   Parallel Problem Solving from Nature, Member of Steering Committee (since 1998);
    –   Co-chair with Youssef Hamadi (MSR Cambridge) of the LION'6 conference (Learning and Intelligent OptimizatioN) in Paris, January 2012.
    –   "Invited Speaker co-Chair" of IEEE-CEC (Congress on Evolutionary Computation) 2011, New-Orleans, USA.
    –   Editorial Board of *Evolutionary Computation*, MIT Press (Editor in Chief, 2002-2009); *Genetic Programming and Evolvable Machines*, Springer Verlag; *Applied Soft Computing*, Elsevier; *Natural Computing Series*, Springer Verlag.
    –   PC member of all important conferences in the area of Evolutionary Computation

- Michèle Sebag
    –   Member of the European Machine Learning and Knowledge Discovery from Databases Steering Committee since 2010; ECCAI Fellow since 2011;
    –   Workshop Chair of ECAI 2012 (Montpellier, August 2012);
    –   Co-organization with Einoshin Suzuki of the International Workshop on LEarning and data Mining for Robots (LEMIR) at IEEE-Int. Conf. on Data Mining, Vancouver dec. 2011;
    –   Area chair ICML11, Area chair ECML11;
    –   PC member of ILP11, GECCO11, reviewer on 3PhDs and 1 HdR; reviewer for ERC, CNRS, INRIA-Lille applications; member of LRI CCSU;
    –   Member of the CoNRS; Senior Advisory Board CHIST-ERA; member of the CSFRS (Conseil Supérieur de la Formation et Recherche Stratégique);
    –   Pattern Analysis, Statistical Learning and Computational Modelling NoE, Member of Steering Committee (PASCAL 2004-2008; PASCAL2, 2008-);
    –   Editorial Board of *Machine Learning Journal*, Springer Verlag; *Genetic Programming and Evolvable Machines*, Springer Verlag.

- Cécile Germain-Renaud
    –   PC member for: IEEE/ACM Cluster, Cloud and Grid Computing (CCGRID) since 2009; IFIP International Conference on Network and Parallel Computing since 2005; SP Cloud workshop; EGEE/EGI user forum since 2008.
    –   Member (elected) of the baord of the Faculty of Science (Conseil d'UFR) and University scientific board (Conseil Scientifique de l'Université).

9.2. Invited talks

- Marc Schoenauer: EVOLVE, Luxembourg (May).
- Michèle Sebag: KAUST, Saoudi Arabia (Feb.); U. Zurich, Switzerland (March); U. York, UK (Nov.); NeuroComp; ML@INRIA.

9.3. Teaching

Nicolas Bredeche

Licence: approx. 80h (Artificial Life), L2, Univ. Paris-Sud, France.

Master: approx. 120h (Evolutionary Computation, Artificial Intelligence), L2, Univ. Paris-Sud, France. Including 15h Evolutionary Robotics, M2R.

Philippe Caillou

Licence: approx. 192h (Computer science for managers), L1, IUT Sceaux, Univ. Paris-Sud, France.

Master: approx. 27h (Multi-Agents Systems), M2R, Univ. Paris-Sud, France.

Master: 3h (Multi-Agent Based Simulation), M2R, Univ. Paris-Dauphine, France.

Cécile Germain-Renaud

Licence: approx. 120h (Computer Architecture, head of Licence) L2, L3, Polytech, Univ. Paris-Sud, France.

Master: approx. 50h (Parallelism), M1 Computer Science, Univ. Paris-Sud, France.

Master: 3h (Multi-Agent Based Simulation), M2R, Univ. Paris-Dauphine, France.

Michèle Sebag

Master: 18h (Machine Learning), M2R, Univ. Paris-Sud, France.

PhD & HdR

HdR: O. Teytaud, *Artificial Intelligence and Optimization with Parallelism*, Université Paris-Sud, April 22. 2011 [3]

PhD: A. Arbelaez, *Learning During Search*, Université Paris-Sud, May 31., 2011, Y. Hamadi and M. Sebag [1]

PhD: F. Teytaud, *Introduction of Statistics in Optimization*, Université Paris-Sud, Dec. 8., 2011, M. Schoenauer and O. Teytaud [2]

PhDs in progress

R.Akrour, *Autonomous Robotics based on Information Theory*, Université Paris-Sud, Nov. 02., 2010, M. Sebag

L. Arnold, *Architectures Profondes pour la Vision Computationnelle*, Université Paris-Sud, Jan. 01., 2010, H. Paugam-Moisy and Ph. Tarroux (LIMSI)

R. Bardenet, *Méthodes d'Échantillonnage pour l'Inférence et l'Optimisation en Physique des Particules*, Université Paris-Sud, Nov. 13., 2009, B. Kegl

Z. Bouzarkouna, *Optimisation de Puits Non Conventionnels : Type, Position et Trajectoire*, Université Paris-Sud, Dec.01., 2008, M. Schoenauer

A. Chotard, *Enhancement and Analysis of Evolution Strategies*, Université Paris-Sud, Oct. 01., 2011, A. Auger and N. Hansen

A. Couëtoux *Monte-Carlo Tree Search and other Reinforcement Learning methods for Energy Management Applications*, Université Paris-Sud, Sept. 01.,2010, O. Teytaud

F. Dawei *Détection et diagnostic d'anomalies dans les systèmes globalisés à grande échelle*, Université Paris-Sud, Oct. 01., 2010, C. Germain

J. Decock *Comparison and Combination of Control and Reinforcement Learning methods for Energy Management Applications*, Université Paris-Sud, Oct. 03., 2011, O. Teytaud

N. Galichet *Integrity Preserving Policy Learning*, Université Paris-Sud, Oct. 01., 2011, M. Sebag

M. Hammami, *Traitement de Données Financières Haute Fréquence: Exploration de Méthodes de Construction Inductive à Grandes Echelle*, Université Paris Diderot, May 01., 2011, M. Sebag

J.-B. Hoock, *Goal Planning with Massive Sets of Heuristics*, Université Paris-Sud, Nov. 01.,2009, O Teytaud

Y. Isaac, *Apprentissage Génératif pour les Interfaces Cerveau-Machine*, Université Paris-Sud, Oct. 03., 2011, C. Gouy-Pallier (CEA) and M. Sebag

I. Loshchilov, *Rank-based Meta-models for Costly Optmization*, Université Paris-Sud, Oct. 01., 2009, M. Schoenauer and M. Sebag

G. Marceau-Caron, *Optimisation Globale du Trafic Aérien*, Université Paris-Sud, May 11., 2011, A. Hadjaz (Thalès Air Systems), P. Savéant (Thalès R&D) and M. Schoenauer

V. Martin, *Modélisation Probabiliste et Inférence par Propagation de Croyances : Application au Trafic Routier*, Université Paris-Sud, Dec. 01., 2009, A. de la Fortelle and JM Lasgouttes (INRIA Rocquencourt)

J.-M. Montanier, *Robotique Evolutionnaire pour l'Adaptation en Ligne d'un Essaim de Robot*, Université Paris-Sud, Oct. 01., 2009, N. Bredeche

W. Wang, *Théorie de l'Information pour l'Apprentissage statistique en Robotique Embarquée*, Université Paris-Sud, Oct.01., 2010, M. Sebag

M. Yagoubi, *Optimisation multi-objectif d'un bloc-moteur diésel*, Université Paris-Sud, Feb. 09., 2009, M. Schoenauer and L. Thobois (PSA)

10. Bibliography

Publications of the year

Doctoral Dissertations and Habilitation Theses

[1] A. ARBELAEZ. *Learning During Search*, Université Paris-Sud XI, Orsay - France, May 2011.

[2] F. TEYTAUD. *Introduction of Statistics in Optimization*, Universite Paris-Sud XI, Orsay - France, December 2011.

[3] O. TEYTAUD. *Artificial Intelligence and Optimization with parallelism*, Univeristé Paris-Sud, April 2011, Habilitation à Diriger des Recherches.

Articles in International Peer-Reviewed Journal

[4] A. AUGER, J. BADER, D. BROCKHOFF, E. ZITZLER. *Hypervolume-based Multiobjective Optimization: Theoretical Foundations and Practical Implications*, in "Theoretical Computer Science", December 2011 [*DOI :* 10.1016/J.TCS.2011.03.012], http://hal.inria.fr/inria-00638989/en.

[5] D. AUGER, O. TEYTAUD. *The Frontier of Decidability in Partially Observable Recursive Games*, in "International Journal on Fundations of Computer Science", 2012, Accepted.

[6] Z. BOUZARKOUNA, D. Y. DING, A. AUGER. *Well Placement Optimization with the Covariance Matrix Adaptation Evolution Strategy and Meta-Models*, in "Computational Geosciences", September 2011, p. 1-18, http://hal.inria.fr/hal-00628126/en.

[7] N. BREDECHE, J.-M. MONTANIER, W. LIU, A. WINFIELD. *Environment-driven Distributed Evolutionary Adaptation in a Population of Autonomous Robotic Agents*, in "Mathematical and Computer Modelling of Dynamical Systems", 2011, http://hal.inria.fr/inria-00531450/en.

[8] A. DEVERT, N. BREDECHE, M. SCHOENAUER. *Robustness and the Halting Problem for Multi-Cellular Artificial Ontogeny*, in "IEEE Transactions on Evolutionary Computation", 2011, http://hal.inria.fr/inria-00566879/en.

[9] T. ELTETO, C. GERMAIN-RENAUD, P. BONDON, M. SEBAG. *Towards Non-Stationary Grid Models*, in "Journal of Grid Computing", December 2011 [*DOI :* 10.1007/S10723-011-9194-Z], http://hal.inria.fr/inria-00616279/en.

[10] G. FOUQUIER, J. ATIF, I. BLOCH. *Sequential model-based segmentation and recognition of image structures driven by visual features and spatial relations*, in "Computer Vision and Image Understanding", 2011, accepted, toappear.

[11] N. HANSEN, R. ROS, N. MAUNY, M. SCHOENAUER, A. AUGER. *Impacts of Invariance in Search: When CMA-ES and PSO Face Ill-Conditioned and Non-Separable Problems*, in "Applied Soft Computing", 2011, vol. 11, p. 5755-5769 [*DOI :* 10.1016/J.ASOC.2011.03.001], http://hal.inria.fr/inria-00583669/en.

International Conferences with Proceedings

[12] R. AKROUR, M. SCHOENAUER, M. SEBAG. *Preference-Based Policy Learning*, in "Machine Learning and Knowledge Discovery in Databases", D. GUNOPULOS, T. HOFMANN, D. MALERBA, M. VAZIRGIANNIS (editors), LNCS, Springer Verlag, 2011, vol. 6911, p. 12-27, http://hal.inria.fr/inria-00625001/en.

[13] A. ARBELAEZ, Y. HAMADI. *Improving Parallel Local Search for SAT*, in "Learning and Intelligent Optimization, Fifth International Conference, LION 2011", COELLO COELLE, CARLOS A. (editor), LNCS 6683, Springer Verlag, 2011, p. 46-60.

[14] J. ATIF, C. HUDELOT, I. BLOCH. *Abduction in Description Logics using Formal Concept Analysis and Mathematical Morphology: application to image interpretation*, in "Concept Lattices and Applications (CLA2011)", Nancy, Paris, October 2011, p. 405-408.

[15] I. ATSONIOS, O. BEAUMONT, N. HANUSSE, Y. KIM. *On Power-Law Distributed Balls in Bins and its Applications to View Size Estimation*, in "ISAAC", Yokohama, Japan, December 2011, http://hal.inria.fr/inria-00618785/en.

[16] A. AUGER, D. BROCKHOFF, N. HANSEN. *Analyzing the Impact of Mirrored Sampling and Sequential Selection in Elitist Evolution Strategies*, in "Foundations of Genetic Algorithms (FOGA 2011)", Schwarzenberg, Austria, April 2011, p. 127-138 [*DOI :* 10.1145/1967654.1967666], http://hal.inria.fr/inria-00587507/en.

[17] A. AUGER, D. BROCKHOFF, N. HANSEN. *Mirrored Sampling in Evolution Strategies With Weighted Recombination*, in "Genetic and Evolutionary Computation Conference (GECCO 2011)", Dublin, Ireland, ACM, July 2011, p. 861-868 [*DOI :* 10.1145/2001576.2001694], http://hal.inria.fr/inria-00612522/en.

[18] J. BERGSTRA, R. BARDENET, Y. BENGIO, B. KÉGL. *Algorithms for Hyper-Parameter Optimization*, in "25th Annual Conference on Neural Information Processing Systems (NIPS 2011)", Granada, Spain, November 2011, http://hal.inria.fr/hal-00642998/en.

[19] *Best Paper*
Z. BOUZARKOUNA, A. AUGER, D. Y. DING. *Local-Meta-Model CMA-ES for Partially Separable Functions*, in "Genetic and Evolutionary Computation Conference (GECCO 2011)", Dublin, Ireland, July 2011, p. 869-876, http://hal.inria.fr/hal-00588977/en.

[20] Z. BOUZARKOUNA, D. Y. DING, A. AUGER. *Partially Separated Meta-models with Evolution Strategies for Well Placement Optimization*, in "73rd EAGE Conference & Exhibition incorporating SPE EUROPEC", Vienne, Austria, May 2011, http://hal.inria.fr/hal-00588983/en.

[21] R. BUSA-FEKETE, B. KÉGL, T. ELTETO, G. SZARVAS. *A Robust Ranking Methodology based on Diverse Calibration of AdaBoost*, in "European Conference on Machine Learning (ECML 2011)", Athens, Greece, November 2011, http://hal.inria.fr/hal-00643000/en.

[22] R. BUSA-FEKETE, B. KÉGL, T. ELTETO, G. SZARVAS. *Ranking by calibrated AdaBoost*, in "Yahoo! Learning to Rank Challenge", Haifa, Israel, June 2011, http://hal.inria.fr/hal-00643001/en.

[23] S. CHEVALLIER, N. BREDECHE, H. PAUGAM-MOISY, M. SEBAG. *Emergence of Temporal and Spatial Synchronous Behaviors in a Foraging Swarm*, in "ECAL 2011", Paris, France, T. LENAERTS, M. GIACOBINI, H. BERSINI, P. BOURGINE, M. DORIGO, R. DOURSAT (editors), LCNS, Springer, August 2011, p. 125-132.

[24] C.-W. CHOU, P.-C. CHOU, H. DOGHMEN, C.-S. LEE, T.-C. SU, F. TEYTAUD, O. TEYTAUD. *Towards a solution of 7x7 Go with Meta-MCTS*, in "Proc. Advances in Computer Games", 2011, Accepted.

[25] P.-C. CHOU, H. DOGHMEN, C.-S. LEE, F. TEYTAUD, O. TEYTAUD, H.-C. WANG, M.-H. WANG, S.-J. YEN, W.-L. WU. *Computational and Human Intelligence in Blind Go*, in "Computational Intelligence and Games", Seoul, Korea, Democratic People'S Republic Of, September 2011, http://hal.inria.fr/inria-00625849/en.

[26] C.-W. CHOU, O. TEYTAUD, S.-J. YEN. *Revisiting Monte-Carlo Tree Search on a Normal Form Game: NoGo*, in "EvoGames 2011", Turino, Italy, Lecture Notes in Computer Science, Springer-Verlag, April 2011, vol. 6624, p. 73-82 [*DOI :* 10.1007/978-3-642-20525-5], http://hal.inria.fr/inria-00593154/en.

[27] R. COULOM, P. ROLET, N. SOKOLOVSKA, O. TEYTAUD. *Handling Expensive Optimization with Large Noise*, in "Foundations of Genetic Algorithms", Austria, ACM, January 2011, p. 61-68, http://hal.inria.fr/hal-00517157/en.

[28] A. COUËTOUX, J.-B. HOOCK, N. SOKOLOVSKA, O. TEYTAUD, N. BONNARD. *Continuous Upper Confidence Trees*, in "LION'11: Proc. 5th International Conference on Learning and Intelligent OptimizatioN", Italy, LNCS 6683, Springer Verlag, January 2011, p. 433-445, http://hal.inria.fr/hal-00542673/en.

[29] A. COUËTOUX, M. MILONE, M. BRENDEL, H. DOGHMEN, M. SEBAG, O. TEYTAUD. *Continuous Rapid Action Value Estimates*, in "The 3rd Asian Conference on Machine Learning (ACML2011)", Taoyuan, Taiwan, Province Of China, C.-N. HSU, W. S. LEE (editors), Workshop and Conference Proceedings, JMLR, 2011, vol. 20, p. 19-31, http://hal.inria.fr/hal-00642459/en.

[30] A. COUËTOUX, M. MILONE, O. TEYTAUD. *Consistent Belief State Estimation, with Application to Mines*, in "Proc. TAAI 2011", 2011.

[31] A. COUËTOUX, O. TEYTAUD, H. DOGHMEN. *Improving exploration in Upper Confidence Trees*, in "Proc. LION'6", Y. HAMADI, M. SCHOENAUER (editors), LNCS, Springer Verlag, 2012.

[32] S. CUSSAT-BLANC, N. BREDECHE, H. LUGA, Y. DUTHEN, M. SCHOENAUER. *Artificial Gene Regulatory Network and Spatial Computation: A Case Study*, in "European Conference on Artificial Life", Paris, France, 2011, http://hal.inria.fr/inria-00601778/en.

[33] S. Cussat-Blanc, N. Bredeche, H. Luga, Y. Duthen, M. Schoenauer. *Artificial Gene Regulatory Networks and Spatial Computation: A Case Study*, in "ECAL", Paris, France, August 2011, http://hal.inria.fr/inria-00601816/en.

[34] Á. Fialho, Y. Hamadi, M. Schoenauer. *Optimizing Architectural and Structural Aspects of Buildings towards Higher Energy Efficiency*, in "GECCO 2011 Workshop on GreenIT Evolutionary Computation", Dublin, Ireland, July 2011, http://hal.inria.fr/inria-00591930/en.

[35] C. Furtlehner, Y. Han, J.-M. Lasgouttes, V. Martin, F. Moutarde. *Propagation of information on undirected dependency graphs for road traffic inference*, in "Chaos, Complexity and Transport, CCT'11", Marseille, France, 2011, http://hal.inria.fr/hal-00648681/en/.

[36] C. Furtlehner, J.-M. Lasgouttes, M. Samsonov. *The Fundamental Diagram on the Ring Geometry for Particle Processes with Acceleration/Braking Asymmetry*, in "Traffic and Granular Flow 2011", Moscou, Russie, Fédération De, 2011, http://hal.inria.fr/hal-00646988/en/.

[37] C. Germain-Renaud, A. Cady, P. Gauron, M. Jouvin, C. Loomis, J. Martyniak, J. Nauroy, G. Philippon, M. Sebag. *The Grid Observatory*, in "IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing", Newport Beach, United States, IEEE Computer Society Press, May 2011, http://hal.inria.fr/inria-00586502/en.

[38] C. Germain-Renaud, F. Fürst, M. Jouvin, G. Kassel, J. Nauroy, G. Philippon. *The Green Computing Observatory: a data curation approach for green IT*, in "International Conference on Cloud and Green Computing", Sydney, Australia, December 2011, http://hal.inria.fr/inria-00632423/en.

[39] V. Heidrich-Meisner, C. Igel. *Non-linearly increasing resampling in racing algorithms*, in "European Symposium on Artificial Neural Networks", Bruges, Belgium, M. Verleysen (editor), Evere, Belgium: d-side publications, April 2011, p. 465-470, http://hal.inria.fr/inria-00633006/en.

[40] B. Helmstetter, C.-S. Lee, F. Teytaud, O. Teytaud, M.-H. Wang, S.-J. Yen. *Random positions in Go*, in "Computational Intelligence and Games", Seoul, Korea, Democratic People'S Republic Of, September 2011, http://hal.inria.fr/inria-00625815/en.

[41] J.-B. Hoock, O. Teytaud. *Progress Rate in Noisy Genetic Programming for Choosing* λ, in "Artificial Evolution", Angers, France, J.-K. Hao, et al. (editors), 2011, http://hal.inria.fr/inria-00622150/en.

[42] Y. Kim, C. Germain-Renaud. *Characterizing E-Science File Access Behavior via Latent Dirichlet Allocation*, in "4th IEEE International Conference on Utility and Cloud Computing (UCC 2011)", Melbourne, Australia, IEEE, December 2011, http://hal.inria.fr/inria-00617914/en.

[43] A. Kuno, J.-M. Montanier, S. Takano, N. Bredeche, M. Schoenauer, M. Sebag, E. Suzuki. *On-board Evolutionary Algorithm and Off-line Rule Discovery for Column Formation in Swarm Robotics*, in "IEEE/ACM/WIC International Conference on Intelligent Agent Technology", Lyon, France, 2011, http://hal.inria.fr/inria-00601785/en.

[44] Z. Lewkovicz, S. Thiriot, P. Caillou. *How detailed should social networks be for labor market's models ?*, in "SNAMAS@AISB 2011", York, United Kingdom, April 2011, http://hal.inria.fr/inria-00579620/en.

[45] I. LOSHCHILOV, M. SCHOENAUER, M. SEBAG. *Adaptive Coordinate Descent*, in "Genetic and Evolutionary Computation Conference (GECCO 2011)", Dublin, Ireland, April 2011, http://hal.inria.fr/inria-00587534/en.

[46] I. LOSHCHILOV, M. SCHOENAUER, M. SEBAG. *Not all parents are equal for MO-CMA-ES*, in "Evolutionary Multi-Criterion Optimization 2011 (EMO 2011)", Ouro Preto, Brazil, February 2011, http://hal.inria.fr/inria-00565282/en.

[47] D. LUPIN SAINT-PIERRE, Q. LOUVEAUX, O. TEYTAUD. *Online Sparse Bandit for Card Games*, in "Proc. Advances in Computer Games", 2011, Accepted.

[48] D. LUPIN SAINT-PIERRE, Q. LOUVEAUX, O. TEYTAUD. *Online Sparse Bandits*, in "The 3rd Asian Conference on Machine Learning (ACML2011)", Taoyuan, Taiwan, Province Of China, 2011, http://hal.inria.fr/hal-00642461/en.

[49] B. MATTHIAS, M. SCHOENAUER. *Instance-Based Parameter Tuning and Learning for Evolutionary AI Planning*, in "Workshop on Planning and Learning at 21st ICAPS", Freiburg, Germany, June 2011, http://hal.inria.fr/inria-00632368/en.

[50] B. MATTHIAS, M. SCHOENAUER. *Instance-based parameter tuning for evolutionary AI planning*, in "Workshops Proc. of Genetic and Evolutionary Computation Conference", Dublin, Ireland, ACM Press, July 2011, p. 591-599, http://hal.inria.fr/inria-00632375/en.

[51] B. MATTHIAS, M. SCHOENAUER. *Learn-and-Optimize: a Parameter Tuning Framework for Evolutionary AI Planning*, in "Artificial Evolution", Angers, France, J.-K. HAO, ET AL. (editors), October 2011, http://hal.inria.fr/inria-00632378/en.

[52] J.-M. MONTANIER, N. BREDECHE. *Emergence of Altruism in Open-ended Evolution in a Population of Autonomous Agents*, in "GECCO", Dublin, Ireland, 2011, http://hal.inria.fr/inria-00601791/en.

[53] J.-M. MONTANIER, N. BREDECHE. *Surviving the Tragedy of Commons: Emergence of Altruism in a Population of Evolving Autonomous Agents*, in "European Conference on Artificial Life", Paris, France, 2011, http://hal.inria.fr/inria-00601776/en.

[54] J. PEREZ, B. KÉGL, C. GERMAIN-RENAUD. *Non-Markovian Reinforcement Learning for Reactive Grid scheduling*, in "Conférence Francophone d'Apprentissage", Chambéry, France, PRESSES UNIVERSITAIRES DES ANTILLES ET DE LA GUYANE (editor), Publibook, May 2011, http://hal.inria.fr/inria-00586504/en.

[55] S. REBECCHI, H. PAUGAM-MOISY, M. SEBAG. *Learning sparse features with an auto-associator*, in "DevLeaNN", 2011.

[56] A. RIMMEL, F. TEYTAUD, T. CAZENAVE. *Optimization of the Nested Monte-Carlo Algorithm on the Traveling Salesman Problem with Time Windows*, in "Evostar", Turin, Italy, February 2011, http://hal.inria.fr/inria-00563668/en.

[57] T. RUNARSSON, M. SCHOENAUER, M. SEBAG. *Pilot, Rollout and Monte Carlo Tree Search Methods for Combinatorial Optimization*, in "Proc. LION'6", Y. HAMADI, M. SCHOENAUER (editors), LNCS, Springer Verlag, 2012.

[58] M. Schoenauer, F. Teytaud, O. Teytaud. *A Rigorous Runtime Analysis for Quasi-Random Restarts and Decreasing Stepsize*, in "Artificial Evolution", Angers, France, J.-K. Hao, et al. (editors), October 2011, http://hal.inria.fr/inria-00625855/en.

[59] N. Sokolovska. *Aspects of Semi-Supervised and Active Learning in Conditional Random Fields*, in "ECML PKDD 2011", Greece, September 2011, p. 273-288, http://hal.inria.fr/hal-00624831/en.

[60] N. Sokolovska, O. Teytaud, M. Milone. *Q-Learning with Double Progressive Widening : Application to Robotics*, in "ICONIP 2011", China, September 2011, p. 103-112, http://hal.inria.fr/hal-00624832/en.

[61] O. Teytaud, S. Flory. *Upper Confidence Trees with Short Term Partial Information*, in "EvoGames 2011", Turino, Italy, Lecture Notes in Computer Science, Springer, 2011, vol. 6624, p. 153-162 [*DOI : 10.1007/978-3-642-20525-5*], http://hal.inria.fr/inria-00585475/en.

[62] O. Teytaud, M. Sebag. *Combining Myopic Optimization and Tree Search: Application to MineSweeper*, in "Proc. LION'6", Y. Hamadi, M. Schoenauer (editors), LNCS, Springer Verlag, 2012.

[63] F. Teytaud, O. Teytaud. *Lemmas on Partial Observation, with Application to Phantom Games*, in "Computational Intelligence and Games", Seoul, Korea, Democratic People'S Republic Of, September 2011, http://hal.inria.fr/inria-00625794/en.

[64] S. Thiriot, Z. Lewkovicz, P. Caillou, J.-D. Kant. *Referral hiring and labor markets: a computational study*, in "Artificial Economics 2011", The Hague, Netherlands, S. Osinga, G. J. Hofstede, T. Verwaart (editors), LNEMS, Springer-Verlag, September 2011, vol. 652, p. 15-25, http://hal.inria.fr/inria-00579625/en.

[65] M. Yagoubi, L. Thobois, M. Schoenauer. *Asynchronous Evolutionary Multi-Objective Algorithms with heterogeneous evaluation costs*, in "IEEE Congress on Evolutionary Computation, CEC 2011", New Orleans, LA, United States, June 2011, p. 21-28, http://hal.inria.fr/hal-00625318/en.

National Conferences with Proceeding

[66] J. Atif, C. Hudelot, I. Bloch. *Abduction dans les logiques de description : apport de l'analyse formelle de concepts et de la morphologie mathématique*, in "Représentation et Raisonnement sur le Temps et l'Espace (Atelier RTE 2011)", Chambery, France, May 2011, p. 405-408.

[67] N. Galichet, M. Sebag. *Exploration prudente: une approche par méthode de Monte-Carlo arborescente contrainte*, in "Proc. Reconnaissance des Formes et Intelligence Artificielle", January 2012, to appear.

[68] C. Germain-Renaud. *Modèles comportementaux de la grille : enjeux et exemples*, in "Premières journées scientifiques France-Grilles", Lyon, France, September 2011, http://hal.inria.fr/inria-00632438/en.

Conferences without Proceedings

[69] R. Bardenet, B. Kégl, G. Fort. *Relabelling MCMC Algorithms in Bayesian Mixture Learning*, in "Snowbird Learning Workshop", Fort Lauderdale, United States, April 2011, http://hal.inria.fr/in2p3-00590956/en.

[70] P. Caillou, E.-P. Gallié, V. Mérindol, T. Weil. *Caractérisation et typologie du contexte initial des pôles*, in "7e congrés de l'Académie de l'entrepreneuriat et de l'innovation", Paris, France, October 2011, http://hal.inria.fr/hal-00653037/en/.

[71] A. COUËTOUX, O. TEYTAUD, N. BONNARD, N. OMONT, O. RATIER. *Monte Carlo Tree Search appliqué à la gestion de stocks*, in "ROADEF 2011", France, March 2011, p. N241, p.I-149, http://hal.inria.fr/hal-00623668/en.

[72] E.-P. GALLIÉ, V. MÉRINDOL, P. CAILLOU, T. WEIL. *Les pôles de compétitivité français en fonction de leur contexte initial d'émergence : essai de caractérisation*, in "EvoReg 2011", Strasbourg, France, October 2011, http://hal.inria.fr/hal-00653036/en/.

[73] X. ZHOU, P. CAILLOU, J. GIL-QUIJANO. *Automated observation of complex systems simulations*, in "V2CS 2011", Paris, France, November 2011, http://hal.inria.fr/hal-00644639/en.

Scientific Books (or Scientific Book chapters)

[74] A. ARBELAEZ, Y. HAMADI, M. SEBAG. *Continuous Search in Constraint Programming*, in "Autonomous Search", Y. HAMADI, E. MONFROY, F. SAUBION (editors), Springer-Verlag, 2011.

[75] S. DONCIEUX, N. BREDECHE, J.-B. MOURET. *New Horizons in Evolutionary Robotics*, Springer, 2011, http://hal.inria.fr/inria-00566890/en.

[76] S. DONCIEUX, J.-B. MOURET, N. BREDECHE, V. PADOIS. *Evolutionary Robotics: Exploring New Horizons*, in "New Horizons in Evolutionary Robotics", Studies in Computational Intelligence, Springer, 2011, p. 3-25, http://hal.inria.fr/inria-00566896/en.

[77] J.-K. HAO, P. LEGRAND, P. COLLET, N. MONMARCHÉ, E. LUTTON, M. SCHOENAUER. *Artificial Evolution 2011*, LNCS, Springer Verlag, 2012, To appear, http://hal.inria.fr/hal-00643404/en.

[78] J.-M. MONTANIER, N. BREDECHE. *Embedded Evolutionary Robotics: The (1+1)-Restart-Online Adaptation Algorithm*, in "New Horizons in Evolutionary Robotics", Studies in Computational Intelligence, Springer, 2011, p. 155-169, http://hal.inria.fr/inria-00566898/en.

[79] O. TEYTAUD. *Lower Bounds for Evolution Strategies*, in "Theory of Randomized Search Heuristics", A. AUGER, B. DOERR (editors), Series on Theoretical Computer Science, World Scientific, May 2011, vol. 1, p. 327-354, http://hal.inria.fr/inria-00593179/en.

Research Reports

[80] N. HANSEN, S. FINCK, R. ROS. *COCO - COmparing Continuous Optimizers : The Documentation*, INRIA, May 2011, n° RT-0409, http://hal.inria.fr/inria-00597334/en.

[81] N. HANSEN. *A CMA-ES for Mixed-Integer Nonlinear Optimization*, INRIA, October 2011, n° RR-7751, http://hal.inria.fr/inria-00629689/en.

[82] N. HANSEN. *Injecting External Solutions Into CMA-ES*, INRIA, October 2011, n° RR-7748, http://hal.inria.fr/inria-00628254/en.

[83] V. MARTIN, J.-M. LASGOUTTES, C. FURTLEHNER. *The Role of Normalization in the Belief Propagation Algorithm*, INRIA, January 2011, n° RR-7514, http://hal.inria.fr/inria-00558444/en.

Other Publications

[84] R. AKROUR, M. SCHOENAUER, M. SEBAG. *Preference-based Reinforcement Learning*, 2011, NIPS Workshop on Choice Models and Preference Learning.

[85] L. ARNOLD, A. AUGER, N. HANSEN, Y. OLLIVIER. *Information-Geometric Optimization Algorithms: A Unifying Picture via Invariance Principles*, 2011, submitted, http://hal.inria.fr/hal-00601503/en.

[86] D. AUGER. *Multiple Tree for Partially Observable Monte-Carlo Tree Search*, 2011, submitted, http://hal.inria.fr/hal-00563480/en.

[87] E. DESCAMPS, C. FURTLEHNER, M. SCHOENAUER. , 2011, Work in progress.

[88] C. FURTLEHNER, J.-M. LASGOUTTES, M. SAMSONOV. *One-dimensional Particle Processes with Acceleration/Braking Asymmetry*, 2011, arXiv:1109.1761, submitted to J. Stat. Phys..

[89] V. MARTIN, C. FURTLEHNER, J.-M. LASGOUTTES. *Encoding Dependencies between real-valued observables with a binary latent MRF*, 2011, submitted to AISTAT.

References in notes

[90] C. CANDAN, J. DRÉO, P. SAVÉANT, V. VIDAL. *Parallel Divide-and-Evolve: Experiments with OpenMP on a Multicore Machine*, in "Proc. GECCO", N. KRANOGOR (editor), ACM Press, 2011, p. 1571-1579.

[91] S. CHEVALLIER, H. PAUGAM-MOISY, M. SEBAG. *SpikeAnts: a spiking neuron network modelling the emergence of organization in a complex system*, in "Advances in Neural Information Processing Systems 23", Vancouver, December 2010, p. 114-119.

[92] B. FREY, D. DUECK. *Clustering by passing messages between data points*, in "Science", 2007, vol. 315, p. 972–976.

[93] J.-B. HOOCK, C.-S. LEE, A. RIMMEL, F. TEYTAUD, O. TEYTAUD, M.-H. WANG. *Intelligent Agents for the Game of Go*, in "IEEE Computational Intelligence Magazine", November 2010, http://hal.inria.fr/inria-00544758/en/.

[94] J. O. KEPHART, D. M. CHESS. *The vision of autonomic computing*, in "Computer", 2003, vol. 36, p. 41-50.

[95] C.-S. LEE, M.-H. WANG, O. TEYTAUD, Y.-L. WANG. *The Game of Go @ IEEE WCCI 2010*, in "IEEE Computational Intelligence Magazine", 2010, vol. 5, n^o 4, p. 6-7 [*DOI :* 10.1109/MCI.2010.938371], http://hal.inria.fr/inria-00632302/en/.

[96] K. LI, Á. FIALHO, S. KWONG. *Multi-Objective Differential Evolution with Adaptive Control of Parameters and Operators*, in "LION'11: Proceedings of the 5th International Conference on Learning and Intelligent OptimizatioN", C. A. COELLO COELLO (editor), LNCS 6683, Springer Verlag, January 2011, p. 473-4887.

[97] G. MOORE. *Crossing the Chasm: Marketing and Selling High-Tech Products to Mainstream Customer*, Collins Business Essentials, 1991.

[98] I. RECHENBERG. *Evolutionstrategie: Optimierung Technisher Systeme nach Prinzipien des Biologischen Evolution*, Fromman-Hozlboog Verlag, 1973.

[99] I. RISH, M. BRODIE, S. MA, N. ODINTSOVA, A. BEYGELZIMER, G. GRABARNIK, K. HERNANDEZ. *daptive diagnosis in distributed dystems*, in "IEEE Transactions on Neural Networks (special issue on Adaptive Learning Systems in Communication Networks)", 2005, vol. 16, p. 1088-1109.