Activity Report 2012

Project-Team TAO

Machine Learning and Optimisation

Table of contents

Project-Team TAO

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

Creation of the Project-Team: November 04, 2004 .

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, until Aug. 2012, 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 until Aug. 2012, HdR]

Engineers

Jean-Joseph Christophe
Hassen Doghmen [until September 2012]
Lovro Ilijasic [since September 2012]
Julien Nauroy

PhD Students

Ouassim Ait Elhara [Started October 2012]
Riad Akrour [SYMBRION, Started November 2010]
Ludovic Arnold [Started January 2010]
Rémi Bardenet [Defended December 2012]
Jérémie Bensadon [Started October 2012]
Zyed Bouzarkouna [CIFRE IFP, Defended April 2012]
Djalel Benbouzid [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]
Weijia Wang [Started October 2010]

Mouadh Yagoubi [CIFRE PSA, Defended July 2012]
Guohua Zhang [Co-Tutelle, Started September 2012]

Post-Doctoral Fellows

Yohei Akimoto
Verena Heidrich-Meisner [until August 2012]
Mostepha Redouane Khouadjia
Nadjib Lazaar [until September 2012]
Manuel Loth
David Meunier [until August 2012]
Mustafa Misir [since October 2012]
Anthony Mouraud [until November 2012]
Sébastien Rebecchi [until August 2012]

Administrative Assistant

Marie-Carol Lopes [Tao]

2. Overall Objectives

2.1. Presentation

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, feb. 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. investigating how specific prior knowledge and requirements can be accommodated in machine learning thanks to evolutionary computation (EC) and more generally Stochastic Optimization;

2. investigating how statistical machine learning can be used to interpret, study and enhance evolutionary computation;

3. facing diversified and real-world applications, requiring and suggesting new integrated ML/EC approaches.

2.3. Highlights of the Year

- Energy management is becoming one of the main focuses, and the most important applicative focus of TAO UCT-SIG. Underlying the various debates, ranging from climate change to nuclear power and integration of renewable energies in the grid (including transportation and storage), is a burning need for scenario simulation, evaluation and optimization. The scientific challenges concern the handling of continuous and discrete uncertainties (e.g. ecological impacts or emergence of future technologies) with a long term horizon. Our commitment is to provide principled studies of various investment scenarios in economical and ecological terms, including a rigorous handling of uncertainties. Specifically,

  – We actively worked to develop collaborations between European and Taiwanese experts of energy management (organization of a forum in Taiwan, http://top.twman.org/2012frtw, of meetings between French companies and Taiwanese academic visitors in Limoges and Paris http://www.lri.fr/ teytaud/france2012.html).

  – We developed a Ilab (collaboration between Inria Saclay-IDF / Artelys) on energy, involving our common participation to the European project Citines http://www.citines.com, aimed at optimal energy management at the scale of a city or an industrial area. We also successfully applied for an ADEME project named POST, aimed at the long term (2050) optimization of the power grid in Europe and North Africa and raising hard stochastic stock management issues. Another critical issue concerns the representation of strategies enabling to combine the good long term properties of direct policy search, and the efficiency of combinatorial optimization tools for structured problems.
    Additionally, a collaboration with Inria-Chile is under discussion. We are also working on creating a company in Taiwan, working with tools from the French industry.

– We also participated in several energy-related European meetings, including companies (section 8.5.1).

- Games remain a key and cool showcase to demonstrate the efficiency of our algorithms:
Our meta-learning approach in Monte-Carlo Tree Search (MCTS) was illustrated by playing 12 games against professional players in even conditions in 7x7; it won 7 games (6/6 win with the easy side and 1/6 win with the difficult side). We achieved the best performances so far on small board minesweeper, demonstrating the efficiency of MCTS on one-player stochastic games. In collaboration with Olivier Buffet (Loria), we scaled up previous implementations to large boards, demonstrating the efficiency of Monte-Carlo Tree Search as a tool for improving existing heuristics. For illustrating the pedagogical properties of simulation-based approaches, we developed tools for generating nice test cases in games and automatically checking the opponent level.
Besides, we realized experimental biological measurements (neuro-imagery, skin conductivity) on amateur and professional players, for further comparison and analysis.

- One of the main fundamental milestones on the TAO research agenda has been achieved by the OPT-SIG, bridging the gap between practice and theory in stochastic optimization through information-geometric optimization (IGO). IGO is devised as a canonical way to turn any smooth parametric family of probability distributions on an arbitrary, discrete or continuous search space X into a continuous-time black-box optimization method on X. Rooted on the Fisher metric, IGO shows invariance properties under various parameterizations of the distribution family [71], [19], [20]. IGO covers the state-of-art CMA-ES (invariant w.r.t. monotonous transformations of the objective function and linear transformations of the coordinate space) as a special case where the probability distribution is Gaussian.

This paper got the *excellent paper award* (international track) at TAAI conference (given to 3/55 papers).
BEST PAPER AWARD :
[36] **Strategic Choices: Small Budgets and Simple Regret in TAAI**. C.-W. CHOU, P.-C. CHOU, C.-S. LEE, D. LUPIEN SAINT-PIERRE, O. TEYTAUD, M.-H. WANG, L.-W. WU, S.-J. YEN.

3. Scientific Foundations

3.1. Introduction

This section describes TAO main research directions at the crossroad of Machine Learning and Evolutionary Computation. Since 2008, TAO has been structured in several special interest groups (SIGs) to enable the agile investigation of long-term or emerging theoretical or applicative issues. 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.

The first two SIGs consolidate the key TAO scientific pillars, while others evolve and adapt to new topics.

The first one, OPT-SIG, addresses stochastic continuous optimization, taking advantage of the fact that TAO is acknowledged the best French research group and one of the top international groups in evolutionary computation from a theoretical and algorithmic standpoint. A main priority on the OPT-SIG research agenda is to provide theoretical and algorithmic guarantees for the current world best continuous stochastic optimizer, CMA-ES, ranging from convergence analysis (Youhei Akimoto's and Verena Heidrich-Meister's post-docs) to a rigorous benchmarking methodology. Incidentally, this benchmark platform has been acknowledged since 2008 as "the" international continuous optimization benchmark, and its extension is at the core of the ANR project NumBBO (starting end 2012). Another priority is to address the current limitations of CMA-ES in terms of high-dimensional or expensive optimization (respectively Ouassim Ait El Hara's and Ilya Loshchilov's PhDs). Mouadh Yagoubi and Zyed Bouzarkouna (resp. CIFRE PSA and IFP-EN) PhD's have continued the EC tradition of industrial breakthrough applications, on which the EC fame solidly relies.

The second SIG, UCT-SIG, benefits from the MoGo expertise and its past and present world records in the domain of computer-Go, establishing the international visibility of TAO in sequential decision making. Since 2010, UCT-SIG resolutely moves to address the problems of energy management from a fundamental and applied perspective. On the one hand, energy management offers a host of challenging issues, ranging from long-horizon policy optimization to the combinatorial nature of the search space, from the modeling of prior knowledge to non-stationary environment to name a few. On the other hand, the energy management issue can hardly be tackled in a pure academic perspective: tight collaborations with industrial partners are needed to access the true operational constraints. Such international and national collaborations have been started by Olivier Teytaud during his one-year stay in Taiwan, and witnessed by the FP7 STREP Citines, the pending ADEME Post contract, and the Ilab with Artelys.

A third SIG, DIS-SIG, is devoted to the modeling and optimization of (large scale) distributed systems. DIS-SIG pursues and extends the goals of the former *Autonomic Computing* SIG, initiated by Cécile Germain and investigating the use of statistical machine learning for large scale computational architectures, from data acquisition (the Grid Observatory in the European Grid Initiative) to grid management and fault detection. More generally, how to model and manage network-based activities has been acknowledged a key topic *per se* in the last months, including: i) the modeling of multi-agent systems and the exploitation of simulation results in the SimTools RNSC network frame; ii) the management of the core communication topology for distributed SAT solving, in the *Microsoft-TAO project* framework. Further extensions are planned in the context of the TIMCO FUI project (started end 2012); the challenge is not only to port ML algorithms on massively distributed architectures, but to see how these architectures can inspire new ML criteria and methodologies.

A fourth SIG, CRI-SIG, focuses on the design of learning and optimization criteria. It elaborates on the lessons learned from the former *Complex Systems* SIG, showing that the key issue in challenging applications is to design the objective itself. Such targeted criteria are pervasive in the study and building of autonomous cognitive systems, ranging from intrinsic rewards in robotics to the notion of saliency in vision and image understanding. The desired criteria can also result from fundamental requirements, such as scale invariance in a statistical physics perspective, and guide the algorithmic design. Additionally, the criteria can also be domain-driven and reflect the expert priors concerning the structure of the sought solution (e.g. spatio-temporal consistency); the challenge is to formulate such criteria in a mixed convex/non differentiable objective function, amenable to tractable optimization.

The activity of the former *Crossing the chasm* SIG gradually decreased after the completion of the 2 PhD theses funded by the Microsoft/Inria joint lab (Adapt project) and devoted to hyper-parameter tuning. Indeed hyper-parameter tuning is still present in TAO, chiefly for continuous optimization (OPT-SIG, section 3.3) and AI planning (CRI-SIG, section 3.4).

3.2. Optimal Decision Making under Uncertainty

Participants: Olivier Teytaud [correspondent], Jean-Joseph Christophe, Adrien Couëtoux, Hassen Doghmen, Jérémie Decock, Nicolas Galichet, Manuel Loth, Marc Schoenauer, Michèle Sebag.

The UCT-SIG works on sequential optimization problems, where a decision has to be made at each time step along a finite time horizon, and the underlying problem involves uncertainties along an either adversarial or stochastic setting.

Application domains include energy management at various time scales and more generally planning, on the one hand, and games (Go, MineSweeper, NoGo [12]) on the other hand.

The main advances done this year include:

- In the domain of computer Go some new performances have been realized [16] and survey papers have been published in Communications of the ACM [10] and as a chapter [65].

- The extension of Upper Confidence Trees to continuous or large domains (states and/or actions) and to domains with high expertise or strong structure has been done [37], [31], [38].

- The extension of Upper Confidence Trees to multi-objective settings has been done, improving on scalarization-based multi-objective reinforcement learning [56].

- Anytime algorithms for discrete time control (note that the classical stochastic dynamic programming is by no means anytime) have been developed and integrated in the Metis software (section 5.1), based on the above cited results.

- Hybrid approaches combining upper confidence trees and e.g. direct policy search or domain specific approaches to yield robust performance w.r.t. long-term effects and take advantage of the combinatorial structure of the domain have been designed, specifically including problem-specific expertise in the playout phase in the domain of job-shop scheduling [53] or MineSweeper [54].

- Optimization algorithms for direct policy search have been designed [55].

- Within the European STREP Mash project, our favorite tools (in particular Monte-Carlo Tree Search) have been extended to difficult settings with no possibility to "undo" a decision [63]. The notion of "risk of exploration" has been investigated [59].

- An experimental analysis of bandit algorithms for small budget cases [36] got the excellent paper award at TAAI 2012.

- In collaboration with Christian Shulte (KTH, Stockholm), one of the main contributors to the well-known general-purpose CP solver *GECODE* (http://www.gecode.org/), and within the Microsoft-Inria joint lab- Adapt project, ideas from UCT have been integrated in GECODE and applied to the job-shop scheduling problem with good first results [52].

- The optimization of low-discrepancy sequences has been done, improving on the best results so far [7]; note that low-discrepancy sequences have been exploited in quite a few of our past works.

3.3. Continuous Optimization

Participants: Ouassim Ait ElHara, Yohei Akimoto, Anne Auger, Zyed Bouzarkouna, Alexandre Chotard, Nikolaus Hansen, Ilya Loshchilov, Verena Heidrich-Meisner, Yann Ollivier, Marc Schoenauer, Michèle Sebag, Olivier Teytaud, Mouadh Yagoubi.

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 methods for derivative-free continuous optimization. We work on different variants of the CMA-ES to improve it in various contexts as described below. We have proven the convergence of simplified variants of the CMA-ES algorithm using the theory of stochastic approximation providing the first proofs of convergence on composite of twice continuously differentiable functions and used Markov chain analysis for analyzing the step-size adaptation rule of the CMA-ES algorithm.

New algorithms based on surrogates and for constrained optimization. A new variant of CMA-ES to address constraint optimization has been designed [22]. In the Zyed Bouzarkouna's PhD, defended in April 2012, a CIFRE PhD in cooperation with IFP-EN (Institut Français du Pétrole - Energies Nouvelles), new variants of CMA-ES coupled with local-meta-models for expensive optimization have been proposed. They have been applied to well placement problem in oil industry [2]. In the context of his PhD thesis (to be defended in January 2013), Ilya Loshchilov has proposed different surrogate variants of CMA-ES based on ranking-SVM that preserve the invariance to monotonic transformation of the CMA-ES algorithm [51]. He has also explored new restart mechanisms for CMA-ES [47].

Benchmarking. We have continued our effort for improving standards in benchmarking and pursued the development of the COCO (COmparing Continuous Optimizers) platform. We have organized the ACM GECCO 2012 workshop on Black-Box-Optimization Benchmarking [2] and benchmarked different variants of the CMA-ES algorithms [27] [30] [29] [28] [48] [49] [50]. Our new starting ANR project NumBBO, centered on the COCO platform, aims at extending it for large-scale, expensive, constrained and multi-objective optimization.

---

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

Theoretical proofs of convergence. We have defined and analyzed a variant of CMA-ES being able to prove that its covariance matrix converges to the inverse Hessian on convex-quadratic functions [18]. We have analyzed the convergence of continuous time trajectories associated to step-size adaptive Evolution Strategies on monotonic C^2-composite Functions and proved the local convergence towards local minima [19]. This work has been the starting point for analyzing convergence of step-size adaptive ESs using the framework of stochastic approximation (work about to be submitted). In the context of his thesis, Alexandre Chotard has analyzed the step-size adaptation algorithm of CMA-ES on linear function using the theory of Markov Chains [35].

Multi-objective optimization. Mouadh Yagoubi has completed his PhD [3], a CIFRE cooperation with PSA (Peugeot-Citroen automotive industry). His work addressed the multi-disciplinary multi-objective optimization of a diesel engine, that led to propose some asynchronous parallelization of expensive objective functions (more than 2 days per evaluation for the complete 3D model) [57].

3.4. Designing criteria

Participants: Jamal Atif, Nicolas Bredèche, Cyril Furtlehner, Yoann Isaac, Victorin Martin, Jean-Marc Montanier, Hélène Paugam-Moisy, Marc Schoenauer, Michèle Sebag.

This new SIG, rooted on the claim that *What matters is the criterion*, aims at defining new learning or optimization objectives reflecting fundamental properties of the model, the problem or the expert prior knowledge.

A statistical physics perspective. In the context of the ANR project TRAVESTI (http://travesti.gforge.inria.fr) which ended this year, we have worked to address more specifically the inverse pairwise Markov random field (MRF) model. On one side we have formalized how the Ising model [64] can be used to perform travel time inference in this context (submitted to AISTATS). Extending the ordinary linear response theory in the vicinity of the so called "Bethe reference point" instead of the interaction-free reference point [67], we were able to provide explicit and tractable formulas to the Plefka's expansion and the related natural gradient at that point. These can be used into various possible algorithms for generating approximate solutions to the inverse Ising problem which we do investigate now. In parallel we have proposed also in [67], a method based on the "iterative proportional scaling" to learn both the factor graph and the couplings by selecting links one by one (submitted to AISTATS). In the Gaussian MRF case this can be implemented efficiently due to local transformations of the precision matrix after adding one link, which is unfortunately not possible in the Ising MRF. It is competitive with L_0 based approach in terms of precision and computational cost, while incidentally the L_1 based method potentially cheaper is not working well for this problem. The flexibility of the method offers in addition the possibility to combine it with spectral constraints like walk-summability with belief propagation or/and graph structure constrain to enforce compatibility with BP. With no additional computational cost we get a complete set of good trade-offs between likelihood and compatibility with BP. Concerning the analysis of the belief propagation we establish in [60] some sufficient condition for encoding a set of local marginals into a stable belief propagation fixed point. Following some work of last year concerning the modeling of congestion at the microscopic level we have finalized in [9] the analysis of a new family of queuing processes where the service rate is coupled stochastically to the number of clients leading a large deviation formulation of the fundamental diagram of traffic flow.

Multi-objective AI Planning. Within the ANR project DESCARWIN (http://descarwin.lri.fr), Mostepha-Redouane Kouadjia worked on the **multi-objective approach** to AI Planning using the Evolutionary Planner *Divide-and-Evolve*, that evolves a sequential decomposition of the problem at hand: each sub-problem is then solved in turn by some embedded classical planner [72]. Even though the embedded planner is single-objective, DaE can nevertheless handle multi-objective problems: Current work includes the implementation of the multi-objective version of DaE, the definition of some benchmark suite, and some first numerical experiments, comparing in particular the results of a full Pareto approach to those of the classical aggregation method. These works resulted in 3 conference

papers recently accepted, introducing a tunable benchmark test suite [45], demonstrating that the best quality measure for parameter tuning in this multi-objective framework is the hypervolume, even in the case of the aggregation approach [46], and comparing the evolutionary multi-objective approach with the aggregation method, the only method known to the AI Planning community [44]. The parameter-tuning algorithm designed for DAE, called Learn-and-Optimize, was published in the selection of papers from *Evolution Artificielle* conference [26]. Though originally designed for Evolutionary AI Planning, the method is applicable to domains where instances sharing similar characteristics w.r.t parameter tuning can be grouped in domains.

Image understanding. 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 issues about the "best" segmentation sequence to follow and/or how to avoid error propagation. Within the new sequential recognition framework proposed in [8], these issues are addressed as 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, specifically a saliency map, 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 relevance of the backtracking procedure to guarantee good and consistent results. In [70] we propose a method for simultaneously segmenting and recognizing objects in images, based on a structural representation of the scene and on a constraint propagation method. Within the ANR project LOGIMA, our goal is to address sequential object recognition as an abduction process [69]. Similar principles are at the core of Yoann Isaac's PhD (Digiteo Unsupervised Brain project), in collaboration with CEA LIST. The dictionary-learning approach used to decompose the EEG signal is required to comply with the structure of the data (e.g. spatio-temporal continuity; submitted).

Robotic value systems. Within the European SYMBRION IP, a key milestone toward autonomous cognitive agents has been to provide robots with internal or external rewards, yielding an interesting or competent behavior. Firstly, an objective-free setting referred to as open-ended evolution has been investigated [17], [5], where the criterion to be optimized is left implicit in the reproduction process. in the Secondly, preference-based reinforcement learning has been investigated in Riad Akrour's PhD, where the robot demonstrations are assessed by the expert and these assessments are used to learn a model of the expert's expectations. In [21], this work has been extended and combined with active learning to yield state-of-the art performances with few binary feedbacks from the expert. The hormone-based neural net controller first proposed by T. Schmickl et al. has been thoroughly analyzed and simplified in collaboration with Artificial Life Laboratory from Graz [34].

3.5. Distributed systems

Participants: Cécile Germain-Renaud [correspondent], Philippe Caillou, Dawei Feng, Nadjib Lazaar, Michèle Sebag.

The DIS-SIG explores the issues related to modeling and optimizing distributed systems, ranging from very large scale computational grids to multi-agent systems and distributed constraint solvers.

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 sociological) processes governing the grid behavior dispel the stationarity hypothesis. When the behavior can be modeled as a time series, an appealing class of models is a sequence of stationary processes separated by break points.The optimisation problem for structural break detection is difficult, because of high dimensionality and a complex objective function. Then, evolutionary algorithms are a method of choice. [15] revisits the optimisation strategy in the light of the general advances in evolutionary computation and the specific opportunities for a separable representation. The single-level optimisation problem is decoupled

into a bilevel optimisation. The upper level is the problem of finding the optimal number and location of the break points. The lower level optimises the autoregressive models given the number and locations of break points. At the upper level, our optimisation strategy exploits the state-of-the-art CMA-ES (Covariance Matrix Adaptation - Evolutionary Strategy) instead of the relatively straightforward Genetic Algorithm proposed in the classic AutoPARM fitting procedure for non-stationary time series. The associated representation addresses an important shortcoming of : the distance in the chromosomes space better maps to the distance in the model space. Furthermore, the representation becomes scalable. More precisely, it scales linearly with the length of the data set, independently of the cost of the objective function.

Fault management. Isolating users from the inevitable faults in large distributed systems is critical to Quality of Experience. Thus a significant part of the software infrastructure of large scale distributed systems collects information that will be exploited to discover if, where, and when the system is faulty. In the context of end-to-end probing as the class of monitoring techniques, minimizing the number of probes for a given discovery performance target is critical. While detection and diagnosis have the obvious advantage of providing an explanation of the failure, by exhibiting culprits, they strongly rely on a priori knowledge that is not available for massively distributed systems. Thus [39], [62] formulates the problem of probe selection for fault prediction based on end-to-end probing as a Collaborative Prediction (CP) problem, based on the reasonable assumption of an underlying factorial model. On an extensive experimental dataset from the EGI grid, the combination of the Maximum Margin Matrix Factorization approach to CP and Active Learning shows excellent performance, reducing the number of probes typically by 80% to 90%.

Multi-agent and games. The main research focus concerning multi-agent systems was on the observation and automatic description of multi-agent based simulations. Whereas usual parameter space exploration systems observe several experiments, the increasing complexity of simulations makes it harder to understand and describe what happens during a single experiment. It is simple to define global indicators to have an overview of the simulation or to follow individual agents, but the most interesting phenomena often occur at an intermediate level, where *groups of agents* are found. The group level is also suited to analyse and display the dynamics of the model. The online and agent-oriented analysis was achieved and its statistical soundness was assessed [6]. In collaboration with the CEA LIST laboratory, we developed a generic tool (SimAnalyer), which can be used online (with NetLogo) of offline (with Logs) to identify, describe, follow [33] and reproduce [58] clusters of agents in a simulation. To select the most interesting clusters and descriptive variables, new activity indicators reflecting the simulation dynamics have been designed [13].

Parallel SAT Solving. Recent Parallel SAT solvers use the so-called Conflict-Directed Clause Learning to exchange clauses between the different cores. However, when the number of cores increases, systematic clause sharing leads to communication saturation. Nadjib Lazaar's post-doc, funded by the Microsoft-Inria joint lab, investigated how the communication topology can be optimized online using a Multi-Armed Bandit setting [68], with an improvement of circa 10% (in number of problems solved) and over 50% (in computational time) over ManySAT 2.0, on the 2012 SAT and UNSAT problem suite.

4. Application Domains

4.1. Energy Management

Energy management, our prioritary application field, involves sequential decision making with:

- stochastic uncertainties (typically weather);
- both high scale combinatorial problems (as induced by nuclear power plants) and non-linear effects;
- high dimension (including hundreds of hydroelectric stocks);

- multiple time scales:
    - minutes (dispatching, ensuring the stability of the grid), essentially beyond the scope of our work, but introducing constraints for our time scales;
    - days (unit commitment, taking care of compromises between various power plants);
    - years, for evaluating marginal costs of long term stocks (typically hydroelectric stocks);
    - tenths of years, for investments.

Nice challenges also include:

- spatial distribution of problems; due to capacity limits we can not consider a power grid like Europe + North Africa as a single "production = demand" constraint; with extra connections we can equilibrate excess production by renewables for remote areas, but no in an unlimited manner.
- other uncertainties, which might be modelized by adversarial or stochastic frameworks (e.g. technological breakthroughs, decisions about ecological penalization).

We have several related projects (Citines, a European (FP7) project; in the near future we should start the Post project (Ademe); IOMCA, a ANR project). We have a collaboration with a company, Artelys, working on optimization in general, and in particular on energy management.

Technical challenges: Our work focuses on the combination of reinforcement learning tools, with their anytime behavior and asymptotic guarantees, with existing fast approximate algorithms; see 6.4. Our goal is to extend the state of the art by taking into account non-linearities which are often neglected in power systems due to the huge computational cost.

Related Activities:

- We are in the process of creating a Franco-Taiwanese company (maybe a taiwanese company using French software) for energy optimization in Taiwan.
- We have a joint team with Taiwan, namely the Indema associate team (see Section 8.4.1.1).
- We have a "Ilab" in progress with Artelys (see Section 5.1) for industrialization of our work.
- We organized various forums and meetings around Energy Management.

4.2. Air Traffic Control

Air Traffic Control has been an application field of Marc Schoenauer's work in the late 90s (PhD theses of F. Médioni in 98 and S. Oussedik in 2000). It was revived recently with Gaëtan Marceau-Caron's CIFRE PhD together with Thalès Air Systems (Areskin HAdjaz) and Thalès TRT (Pierre Savéant), around global optimization of the traffic in order to increase the capacity of the airspace without overloading the controllers. First results concern the on-line prediction of individual plane trajectories [42]: these results demonstrate that there is no hope to ever predict the global behavior of the traffic by considering planes individually. A new formulation of the problem is hence considered, where plane flows are modeled with Bayesian Networks [42]. On-going work is concerned with the simulation and optimization of air traffic based on this model.

5. Software

5.1. Metis

Participants: Olivier Teytaud [correspondent], Adrien Couëtoux, Jérémie Decock, Jean-Joseph Christophe.

Energy, Optimization, Planning

Many works in Energy Optimization, in particular in the case of high-scale sequential decision making, are based on one software per application, because optimizing the software eventually implies losing generality. Our goal is to develop with Artelys a platform, Metis, which can be used for several applications. In 2012 we interfaced existing codes in Artelys and codes developed in the Tao team; experiments have been performed and test cases have been designed. A main further work is the introduction of generic tools for stochastic dynamic programming into the platform, for comparison and hybridization with other tools from the UCT-SIG.

Our favorite challenge is the hybridization of "classical" tools (based on constraint satisfaction problems, or mixed integer linear programming or mixed integer quadratic programming), which are fast and accurate, with non-linear solvers which can take care of a sophisticated (non-linear) model.

5.2. MoGo

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

Go, Multi-armed bandit

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). Recent results include 7 wins out of 12 against professional players (in Brisbane, 2012). However, the work in the UCT-SIGhas now shifted to energy management.

5.3. CMA-ES: Covariance Matrix Adaptation Evolution Strategy

Participant: Nikolaus Hansen [correspondent].

Evolutionary Computation, stochastic optimization, real-parameter optimization

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

5.4. COmparing Continuous Optimizers

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

Evolutionary Computation, stochastic optimization, real-parameter optimization, benchmarking, derivative free optimization

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 is under continuous developement and has been used for the GECCO 2009, 2010 and 2012 workshops on "Black Box Optimization Benchmarking" (BBOB) (see Section 3.3), and serves as a basis for the test platform in the CSDL project.

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

5.5. MultiBoost

Participants: Balázs Kégl [correspondent], Djalel Benbouzid.

---

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

multi-class, multi-label classification

The MultiBoost package [4] provides a fast C++ implementation of multi-class/multi-label/multi-task boosting algorithms. It is based on ADABOOST.MH but it also implements popular cascade classifiers, ARC-GV, and FILTERBOOST. The package contains common multi-class base learners (stumps, trees, products, Haar filters). Further base learners and strong learners following the boosting paradigm can be easily implemented in a flexible framework.

Link: http://multiboost.org

5.6. Grid Observatory

Participants: Cécile Germain [correspondent], Julien Nauroy, Michèle Sebag.

Autonomic Computing, Green Computing 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 140 users are currently registered. The GO is supported by an Inria ADT (Action de Développement Technologique).

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. A tool has been made available to provide interactive visualization of the site activity (load, power usage and more) at the global and individual level. The GCO is supported by the CNRS PEPS program, and by University Paris-Sud through the MRM (Moyens de Recherche Mutualisés) program. The first barrier to improved energy efficiency is the the lack of overall data collection on the energy consumption of individual components of data centers. The Green Computing Observatory collects monitoring data on energy consumption of a large computing center, and publishes them through the Grid Observatory portal. These data include the detailed monitoring of the processors and motherboards, as well as global site information, such as overall consumption and external temperature, as global optimization is a promising way of research. A second barrier is making the collected data usable. The difficulty is to make the data readily consistent and complete, as well as understandable for further exploitation. For this purpose, the GCO opts for an ontological approach in order to rigorously define the semantics of the data (what is measured) and the context of their production (how are they acquired and/or calculated). The first achievements of these developments have been described in [40],and presented at the ICT-COST meeting and GreenDays@Lyon.

Link: http://grid-observatory.org

6. New Results

6.1. Realistic step sizes for optimization algorithms

Many theoretical results about objective improvement in the process of continuous optimization rely on the assumption that the steps of the algorithm are infinitesimally small, the only situation in which theoretical guarantees of improvement can be given. Y. Akimoto and Y. Ollivier have waived the necessity for such an assumption in a whole class of continuous optimization algorithms, thanks to the use of information geometry [20]. This takes theory closer to the practice of actual optimization algorithms.

6.2. Noisy Optimization Bounds with Constant Noise Variance

Many bounds in noisy evolutionary optimization are based on low variance assumptions (in particular, variance of noise converging to 0 close to the optima). Other bounds in the optimization literature consider difficult objective functions. We prove some new bounds, in the following setting[55]:

- without assuming that the variance is going to zero at the optimum;
- following some debates on the COCO mailing list (see 5.4), assuming that sampling far from the optimum (we had earlier results without this assumption; new results emphasize the contrast).

6.3. Extensions of Upper Confidence Trees

We developed extensions of Upper Confidence Trees to continuous or large domains (states and/or actions) and to domains with high expertise or strong structure[37], [31], [38] (incidentally realizing performances on MineSweeper); we recently submitted a proof of a variant of UCT with consistency proof in the continuous domains (both actions and random variables are allowed to be continuous). Another extension is to the difficult setting with no possibility to "undo" a decision or duplicate a state; see [63]. Yet another extension aims at multi-objective optimization [56].

6.4. Mixing myopic fast algorithms and asymptotically optimal algorithms

We made several works based on combining in sequential decision making:

- a fast algorithm providing quickly good heuristic results;
- an asymptotically optimal, too slow for real size problems.

Results are published in [31] and [38], outperforming the state of the art for MineSweeper in reasonable time; an application to energy has been done, and a new one is under work (see Section 4.1). We believe that this diea of combining fast approximate solutions and slow asymptotically optimal algorithms is a key for improving the state of the art in high dimensional combinatorial planning and that our results on MineSweeper and moderate size energy problem are a solid first step in this direction.

6.5. Adaptive Metropolis with Online Relabeling

In [23] we proposed a novel adaptive MCMC algorithm named AMOR (Adaptive Metropolis with Online Relabeling) for efficiently simulating from permutation-invariant targets occurring in, for example, Bayesian analysis of mixture models. An important feature of the algorithm is to tie the adaptation of the proposal distribution to the choice of a particular restriction of the target to a domain where label switching cannot occur. The algorithm relies on a stochastic approximation procedure for which we design a Lyapunov function that formally defines the criterion used for selecting the relabeling rule. This criterion reveals an interesting connection with the problem of optimal quantifier design in vector quantization which was only implicit in previous works on the label switching problem. In benchmark examples, the algorithm turns out to be fast-converging and efficient at selecting meaningful non-trivial relabeling rules to allow accurate parameter inference. In [24] the algorithm was applied to a synthetic mixture model inspired by the muonic water Cherenkov signal of the surface detectors in the Pierre Auger Experiment.

6.6. Reinforcement learning for frugal cascade learning

In [32] we propose an algorithm that builds sparse decision DAGs (directed acyclic graphs) from a list of base classifiers provided by an external learning method such as AdaBoost. The basic idea is to cast the DAG design task as a Markov decision process. Each instance can decide to use or to skip each base classifier, based on the current state of the classifier being built. The result is a sparse decision DAG where the base classifiers are selected in a data-dependent way. The method has a single hyperparameter with a clear semantics of controlling the accuracy/speed trade-off. The algorithm is competitive with state-of-the-art cascade detectors on three object-detection benchmarks, and it clearly outperforms them when there is a small number of base classifiers. Unlike cascades, it is also readily applicable for multi-class classification. Using the multi-class setup, we show on a benchmark web page ranking data set that we can significantly improve the decision speed without harming the performance of the ranker. Beside outperforming classical cascade designs on benchmark data sets, the algorithm also produces interesting deep structures where similar input data follows the same path in the DAG, and subpaths of increasing length represent features of increasing complexity.

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

- The ILAB *Metis* is a close partnership with the **SME ARTELYS** whose aim is to develop a generic optimisation platform for sequential decision making that could be used for different applications. See Section 5.1 for a detailed description.

7.2. Bilateral Grants with Industry

- **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.

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, Verena Heidrich-Meisner, Marc Schoenauer.
- **TIMCO**– 2012-2015 (432 kEur). FUI-System@tic (Région Ile de France grant).
  Participants: Cécile Germain, Marc Schoenauer, Lovro Ilijasic.

8.2. National Initiatives

- **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** – 2009-2013 (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: 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, Djalel Benbouzid, Nikolaus Hansen, Michèle Sebag, Cécile Germain
- **NUMBBO** 2012-2016 (290k for TAO). Analysis, Improvement and Evaluation of Numerical Black-box Optimizers, ANR project, Coordinator Anne Auger, Inria. Other partners: Dolphin, Inria Lille, Ecole des Mines de Saint-Etienne, TUD Dortmund
  Participants: Anne Auger, Nikolaus Hansen, Marc Schoenauer, Ouassim Ait ElHara
- **LOGIMA** 2012-2016 (136k for TAO). Logics, structural representations, mathematical morphology and uncertainty for semantic interpretation of images and videos, ANR project, Coordinator Céline Hudelot, MAS-ECP. Other partners: TAO , LTCI-Telecom ParisTech
  Local coordinator: Jamal Atif

8.2.1. Other

- **GCO** - 2011-2012) *Green Computing Observatory* PEPS-INS2I. Coordinator Cécile Germain. Participants: Cécile Germain, Julien Nauroy, Michèle Sebag.
- **GO** (2011-2013) *Observatoire de la Grille* Action de Développement technologique Inria. Coordinator Cécile Germain. Participants: Dawei Feng, Cécile Germain, Julien Nauroy, Michèle Sebag. Also funded by the France Grilles national initiative.

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: Almende, 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: January 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/

Abstract: The Mash project is about massive crowd-sourcing. It is based on several artificial applications. We however used the codes also for our favorite applications, because the original Mash applications have nearly no user, which make it hard to have massive crowd-sourcing; for our applications, we have a moderate number of users, but at least they are motivated. Our contributions are twofolds:

- Building solvers on top of existing expert solvers; this is quite related to our Metis platform (Section 5.1) and our work on Minesweeper and on the mixing of direct policy search and Monte-Carlo Tree Search;

- Adapting solvers for cases in which we can not "undo" on the problem, i.e. if we apply a decision, we can not come back to the previous time step; this makes planning much harder and slower. This is developped in [63].

8.3.1.3. CitInES

Title: City and Industry Energy Strategy

Type: COOPERATION (ICT)

Defi: Design of a decision support tool for sustainable, reliable and cost-effective e

Instrument: Specific Targeted Research Project (STREP)

Duration: October 2011 - March 2014

Coordinator: Artelys (France)

Others partners: Artelys (http://www.artelys.com), Inescp (http://www2.inescporto.pt/, Portugal), Ait (http://www.ait.ac.at/, Austria), Armines (Ecole des Mines, Paris http://www.ensmp.fr), Tupras (http://www.tupras.com.tr, Turkey), Ervet (http://www.ervet.it/, Italy), Schneider (http://www.schneider-electric.com), Cesena (Italy), Bologna (Italy)

See also: http://www.citines.com

Abstract: According to OECD, 67% of world energy is used by cities and 70% of CO_2 emissions come from cities. Therefore, optimizing urban energy investments is a key challenge for reducing polluting emissions and financial exposition to fuel price uncertainties. However, the definition of a sustainable, reliable and cost-effective energy strategy requires to simulate the whole energy chain (consumption, transport, distribution, storage, production) with different types of energy (electricity, gas, heat, wind, waste, etc.) and to assess the environmental and financial impacts of various long-term scenarios (fuel prices, consumption scenarios, etc.).

Local authorities facing this issue have today only partial answers to these questions (simulation of a given type of energy, of a part of the energy chain only or without any long-term risk assessment) and lack a global analysis.

The goal of the CitInES project is to design and develop decision-support software to help local authorities / industries to :

- Assess and compare energy strategies through detailed energy chain simulations

- Optimize local energy strategy to cost-effectively integrate green energy and reduce CO_2 emissions

- Define robust energy schemes to face fuel price uncertainties.

The CitInES project is financed by the European Commission, under 7th Framework Programme. It gathers:

- 4 high-level research centers (INESCP for electric system modelling, AIT for building and energy infrastructure planning, ARMINES for long-term energy strategies and Inria for optimization algorithms)
- 1 SME specialized in decision-support software in the energy field (Artelys, leader of the consortium)
- 2 well-known industrial groups (Schneider Electric for its expertise in electric systems; TUPRAS, Turkish refineries as end-user) and 1 national company (ERVET for its expertise in energy processes)
- 2 large cities (Cesena and Bologna as end-users).

8.3.1.4. EGI-Inspire

Title: Integrated Sustainable Pan-European Infrastructure for Researchers in Europe

Type: Research Infrastructures (CP-CSA)

Defi: European Grid Infrastructure Ecosystem

Instrument: Integrated Project (IP)

Duration: May 2010 - April 2014

Coordinator: EGI.eu (Netherlands)

Others partners: 50 institutions, coordinator for France: CNRS.

See also: http://www.egi.eu

Abstract: The EGI-InSPIRE project supports the transition from a project-based system to a sustainable pan-European e-Infrastructure, by supporting 'grids' of high-performance computing (HPC) and high-throughput computing (HTC) resources. EGI-InSPIRE supports the establishment of a sustainable model for a European Grid Infrastructure (EGI) that integrates resources contributed by national and domain-specific resource providers. Key to this process is a new organisation, EGI.eu, coordinator on behalf of the European resource provider community of the EGI-InSPIRE project. The EGI is a federation of independent national and domain specific resource providers, who support specific research communities and international collaborators both within Europe and worldwide.

8.3.2. Collaborations in European Programs, except FP7

Program: COST

Project acronym: Action IC0804

Project title: Energy Efficiency in Large Scale Distributed Systems

Duration: January 2009 - May 2013

Coordinator: IRIT

Other partners: see http://www.cost804.org

Abstract: The main objective of the Action is to foster original research initiatives addressing energy awareness/saving and to increase the overall impact of European research in the field of energy efficiency in distributed systems.

8.3.3. Collaborations with Major European Organizations

Partner 1: organisme 1, labo 1 (pays 1)

Sujet 1 (max. 2 lignes)

Partner 2: organisme 2, labo 2 (pays 2)

Sujet 2 (max. 2 lignes)

8.4. International Initiatives

8.4.1. Inria Associate Teams

8.4.1.1. INDEMA

Title: Intelligent Decision Making Mechanisms with Hidden Information, and Application to Electricity Generation

Inria principal investigator: Olivier Teytaud

International Partner (Institution - Laboratory - Researcher):

National University of Tainan (Taiwan) - Ontology Application and Software Engineering - Chang-Shing Lee

Duration: 2012 - 2014

See also: http://www.lri.fr/~teytaud/indema.html

The objectives of the project are three-folds:

- Objective 1: Designing consistent iterative realistic algorithms for partially observ- able 1-player or 2-player games. We mean:

  – consistent algorithms, in the sense that they are mathematically, provably, optimal asymptotically in the computation time.

  – iterative algorithms in the sense that when you give more time to the algorithm, it should be better; and with little time, it should do its best for replying something acceptable. This is also termed an anytime algorithm. Most algorithm which survive decades are iterative.

  – realistic algorithms; we mean that one can easily design a consistent iterative algorithm that will never work in practice in a real-world setting; so, additionally, we want an algorithm which looks reasonnable and we refer to the second objective for the assesment of this property

  We consider our work on MineSweeper[31] and combining Mcts and Dps [38] as realizations of this principle; we also apply this principle for real applications in the related Citines project.

- Objective 2: Impressive visible applications, e.g. applications in games or puzzles, because such games are very clear assessment tools. Possibilities include Minesweeper (on which we believe that much progress is still possible), Chinese Dark Chess, Kriegspiel, Phantom-Go, card games. Such nice results are critical for advertising and assessing our research. Since the beginning of the project, we had results on MineSweeper, Urban Rivals.

- Objective 3: Big industrial applications. Having both mathematics and visible real- izations in games and industrial applications might be considered as too much; yet, we have chosen to request the maximum possible funding and to include many people in the travelling; also, the persons in the project are all people working in related subjects, with various terminologies, and we already have concrete applications in mind, just far enough from our past activities for being new (we want to tackle in a principled manner partial observability which was somehow ignored in many past works) and close enough for strongly reducing the "warm up" time. In the fully observable case, we worked successfully for these three objectives and want to do the same in the partially observable case. More precisely, when working on real applications in the field of energy generation, we have seen that many problems are simplified so that they boil down to fully observable problems, but that this is a bad application; and our solvers must include some tricks for the partial observability. This is the main motivation for this project; we assume that mathematical analysis can be done on this (objective 1); that it will provide big results in games (objective 2) where many main programs are based on non-consistent algorithms. We believe

that requirements above (objective 1) and visible realizations will facilitate the migration to real-world application; also we point out that previous research projects involving us facilitated contacts with industry, in particular in the field of energy generation, which is a key point for this third objective. A roadmap for objective 3 is as follows:

– Check on simple versions of energy production problems whether the fully observable approximation is ok. We guess that in many cases it is not ok, and we want to clearly state to which extent (by how many percents) we loose in terms of loss function.

– Experiment our algorithms on real industrial problems. We will work both on Taiwan-centered and on Europe-Centered electricity generation problems in order to widen the scope of the analysis and so that both partners can be helpful in terms of applications in their own countries.

We have made papers related to energy management, including papers in very applied conferences. We are in the process of creating a company in Taiwan, hopefully during the 2nd semester of 2013. One student (Adrien Couëtoux) has spent 6 months there, another student has spent 5 months; Adrien just starts a second 6 months stay there.

8.4.2. Inria International Partners

8.4.2.1. Microsoft Research Cambridge

Within the *Microsoft-Inria Joint Lab*, the collaboration with Youssef Hamadi (Microsoft Research Cambridge), through the *Adapt* project, has been pursued, in spite of the departure of the 2 PhD students Alvaro Fialho and Alejandrao Arbelaez. Nadjib Lazaar and Manuel Loth have been hired as post-doc, and a new collaboration with Christian Shulte (KTH Stockholm) based on the use of Bandit algorithm within GECODE has recently given its first results [52] (see Section 3.2).

8.4.3. Participation In International Programs

- The UCT-SIGhas started a collaboration with Inria Chile around energy management; for the moment this is only preliminary discussions (a few face-to-face meetings in Paris, a visioconference with Inria Chile).

8.5. International Research Visitors

8.5.1. Visits of International Scientists

- Christian Shulte (Software and Computer Systems, School of Information and Communication Technology, KTH - Royal Institute of Technology in Stockholm, Sweden), Jan. 24-27, to initiate the generic implementation of Bandit algorithms in Gecode (see Section 3.2).

- Visits from a Taiwanese delegation, see the Franco-Taiwanese week. This included visits to Univ. Paris-Sud, to other universities (Paris-Nord, Limoges), to companies working around energy.

- One month visit from Cheng-Wei Chou, Taiwanese ph.D. student from National Dong-Hwa University.

- Francis Maes, Post-doc, Leuven University, Leuven, Nov. 20 to Dec. 21.

- One week visit from Muneki Yasuda associate professor in the department of Information science, Tohoku University.

   Xiangliang ZHANG (25-31 October 2012)

   Continued collaboration on large scale clustering.

   Institution: KAUST-King Abdullah University of Science and Technology (Saudi Arabia)

8.5.1.1. Internships

   Christopher DELGADO (from Apr 2012 until Sep 2012)

Subject: Designing lean classifiers for detectors and triggers

Institution: Massachusetts Institute of Technology (United States)

Gaurav MAHESHWARI (from Apr 2012 until Sep 2012)

Subject: Sampling-based statistical analysis in large-scale physics experiments

Institution: IIT HYDERABAD (India)

Mauro DI MASSO (from Mar 2012 until Sep 2012)

Subject: Evolutionary Adaptation and the Emergence of Speciation in a Population of Autonomous Robots

Institution: National University of Rosario (Argentina)

8.5.2. *Visits to International Teams*

- Olivier Teytaud has made a one-year visit (August 2011 to July 2012) in National University of Tainan, Taiwan, and to many other universities.

9. Dissemination

9.1. Scientific Animation

- Jamal Atif
  - Co-organiser of RTE workshop, a one-day RFIA associated workshop on Reasoning on Space and Time (http://www.lri.fr/~atif/RTE2012/).
  - Reviewer for Information Sciences, Fuzzy Sets and Systems, and many conferences in computer vision
- Anne Auger
  - Member of the ACM-SIGEVO Executive (Special Interest Group on Evolutionary Computation (was the International Society on Genetic and Evolutionary Algorithms before 2006);
  - THRaSH, *Theory of Randomized Search Heuristics workshop*, member of Steering Committee;
  - Editorial Board of *Evolutionary Computation*, MIT Press; Guest editor of Algorithmica special issue and ECJ special issue
  - Proceedings chair of the GECCO 2012 conference
  - Reviewer for IEEE Trans on Evolutionary Computation, ECJ, Journal of Global Optimization, JMLR, JOTA, MEME, PC member of PPSN, GECCO, EuroGP, FOGA
  - Co-organizer of GECCO 2012 Black-Box-Optimization-Benchmarking workshop.
- Nicolas Bredeche
  - Track chair ALIFE at GECCO 2012
  - PC member: Evostar 2012, Evorobot 2012, CEC 2012, SAB 2012
- Philippe Caillou
  - PC member at ECAI 2012
  - Coordinator of the SimTools Network (RNSC Network)
  - Reviewer for Simulation, Studia Informatica Universalis
- Cyril Furtlehner

- – Co-organiser of a one-day interdisciplinary workshop https://who.rocq.inria.fr/Jean-Marc. Lasgouttes/workshop/ on Information processing in complex systems with applications to traffic forecasting at Inria place d'Italie in Paris.
- Cécile Germain
  - – PC member for: IEEE/ACM Cluster, Cloud and Grid Computing (CCGRID) since 2009; Utility Cloud Computing (UCC) since 2011; IFIP International Conference on Network and Parallel Computing since 2005; SP Cloud workshop; EGEE/EGI user forum since 2008.
  - – Member (elected) of the University scientific council and board (Conseil Scientifique de l'Université et bureau).
- Nikolaus Hansen
  - – Editorial Board member of *Evolutionary Computation*, MIT Press;
  - – Co-organizer of the Dagstuhl Seminar 13271 Theory of Evolutionary Algorithms
  - – Tutorials at the LION and ACM-GECCO conferences
  - – Co-organizer of the ACM-GECCO 2012 Black-Box-Optimization-Benchmarking workshop
  - – PC member of most of the important conferences in the area of Evolutionary Computation
- Yann Ollivier
  - – Frequent refereeing for many of the main mathematical journals.
- Marc Schoenauer
  - – Elected 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.
  - – 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 (ACM-GECCO, PPSN, EvoStar, IEEE-CEC, SEAL, SAC, ...).
  - – Honorary Adjunct Professor, School of Computer Science, University of Adelaide, Australia (renewed until end 2015).
- Michèle Sebag
  - – 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.
  - – 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);
  - – PC member of the major conferences in ML and EC worldwide;

–   reviewer for ERC applications; FNRS (Belgium), NSWO (Netherlands).
–   member of CCSU: LRI, Paris-Sud; LPMA, Paris-Diderot.
- Olivier Teytaud
    –   Reviewer for various conferences and journals in optimization and machine learning.
    –   Organizer a forum in Taiwan (http://top.twman.org/2012frtw) for developing collaborations between European and Taiwanese experts of energy management.
    –   Organizer of an international meeting in France (combining academics and companies) as detailed in Section 8.5.1.

9.2. Invited talks

- Anne Auger: GECCO 2012 Tutorial, LION 2012 Tutorial, CMAP Seminar
- Cyril Furtlehner: 4th YSM-SPIP workshop in Sendai
- Yann Ollivier: Oxford (probability and statistics seminar), Orsay (probability and statistics seminar), IRCAM (information geometry seminar).
- Marc Schoenauer:
    –   18th CREST Open Workshop, *Managing and Optimising Multiplicity Computing*, UCL, Londond, UK, 22-23 March 2012;
    –   Complex Adaptive Systems Laboratory seminar, University College Dublin, Ireland, 25 May 2012;
    –   ECODAM, Doctoral Summer School on Evolutionary Computation and Data Mining, Faculty of Informatics, Iasi University, Romania, 18023 June 2012;
    –   Séminaire du département de génie de la production automatisée, Ecole Supérieure de Technologie, Montreal, Canada, 7 Dec. 2012.
- Michéle Sebag:
    –   Invited talk at the Spring Workshop on Mining and Learning, Bad Neuenahr, Germany, Apr. 18-20, 2012.
    –   Keynote speaker at the 36th Annual Conference of the German Classification Society, Hildesheim, August 1-3, 2012.
    –   Invited talk at the Turing's heritage session at ECAI 2012, August 27-31, 2012.
    –   Invited talk, Sixth "Starting Artificial Intelligence Research" Symposium, August 27, 2012.
    –   Invited tutorial, International Summer School on Resource-aware Machine Learning, Dortmund, Germany, Sept. 4-7, 2012.
    –   Invited tutorial, Constraint Programming 2012, Québec, Canada, Oct. 8-12, 2012.
- Olivier Teytaud:
    –   has been invited in several universities in Taiwan during his one-year stay in 2011-2012: Kaohsiung NUK; Hsinchu NDHU; Hsinchu NCTU; Hualien NDHU; Tainan NUTN.
    –   was invited at the Bielefeld "search methodologies" seminar 2012 (booklet).
    –   invited speaker at the "Unexpected Results" workshop at ECML-PKDD, Sept. 2012.

9.3. Teaching - Supervision - Juries

9.3.1. Teaching

Jamal Atif

Licence (ou équivalent) : approx. 192h (Computer science), L1, IUT Orsay, Univ. Paris-Sud, France.
Master (ou équivalent) : approx. 18h (Robotics and Autonomous agents), M2R, Univ. Paris-Sud, France.
"Directeur d'études" at Computer science department of IUT d'Orsay, Univ. Paris-Sud, France

Anne Auger

> Master : 12h (Optimization) M2R, Univ. Paris-Sud, France, 20 h (Stochastic Optimization) Ecole Centrale Paris

> Doctorat : 4h (Stochastic numerical optimization by Evolution Strategies) Evolution Artificielle Summer School, 2h (Covariance Matrix Adaptation Evolution Strategy - together with N. Hansen) ACM-GECCO Tutorial, 2h (CMA-ES tutorial) LION Conference invited tutorial

Nicolas Bredèche

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

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

Philippe Caillou

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

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

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

Michèle Sebag

> Licence (ou équivalent) : approx. 24h (Introduction to Machine Learning), L3 ENS-Cachan, France.

> Master (ou équivalent) : approx. 12h (Apprentissage Statistique et Optimisation, TC2), M2R, Univ. Paris-Sud, France.

> Master (ou équivalent) : approx. 12h (Apprentissage Statistique, Optimisation et Applications), M2R, Univ. Paris-Sud, France.

9.3.2. *Supervision*

PhD & HdR :

> Rémi Bardenet, *Méthodes d'échantillonnage pour l'Inférence et l'Optimisation en Physique des Particules*, Université Paris-Sud, Dec. 2012, B. Kegl [1].

> Zyed Bouzarkouna, *Optimisation de Puits Non Conventionnels : Type, Position et Trajectoire*, Université Paris-Sud, April 2012, A. Auger and M. Schoenauer [2]

> Mouadh Yagoubi, *Multi-objective parallel evolutionary algorithms : Application on Diesel Combustion*, Université Paris-Sud, July 2012, M. Schoenauer and Ludovic Thobois (ex-PSA) [3].

> L. Veysseire, *Courbure de Ricci grossière de processus markoviens*, Ecole normale supérieure de Lyon, Juillet 2012, Y. Ollivier.

> **PhDs in progress**

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

> Ouassim Ait Elhara, *Stochastic Large Scale Optimization*, Université Paris-Sud, October 2012, A. Auger and N. Hansen

> Ludovic Arnold, *Architectures Profondes pour la Vision Computationnelle*, Université Paris-Sud, Jan. 01., 2010, H. Paugam-Moisy and Ph. Tarroux (LIMSI) (defence in April 2013).

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

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

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

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

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

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

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

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

Ilya Loshchilov, *Rank-based Meta-models for Costly Optmization*, Université Paris-Sud, Oct. 01., 2009, M. Schoenauer and M. Sebag (defence in January 2013).

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

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

Jean-Marc Montanier, *Robotique Evolutionnaire pour l'Adaptation en Ligne d'un Essaim de Robot*, Université Paris-Sud, Oct. 01., 2009, N. Bredeche (defence in March 2013).

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

9.3.3. *Juries*

- Jamal Atif was a referee of the mid-term PhD committe of Yuan Yang (Telecom-ParisTech, April 2012)

- Philippe Caillou was in the PhD committee of An Vo Duc (LIP6 Paris, November 2012).

- Yann Ollivier was an external referee for the Habilitation of Benoît Kloeckner (Grenoble, December 2012).

- Marc Schoenauer was external reviewer for Eduardo Vellasques'PhD (Ecole de Technologie Supérieure, Université du Québec, Canada, Informatique), PhD committee member for Vincent Baudoui (Université de Toulouse, Mathématiques Appliquées), Gabriel Synaeve (Université de Grenoble, Informatique), Marco Montemuro (Université Pierre et Marie Curie, Mécanique), Giovanni Granato (Ecole Polytechnique, Mathématiques Appliquées), and HdR committee member for Daniel Delahaye (Université de Toulouse, Mathématiques Appliquées) and Frédéric Lardeux (Université d'Angers, Informatique). He was also external reviewer within the Tenure Committee for Prof. Dirk Arnold, Dalhousie University, Canada.

- Michèle Sebag was external reviewer for Michael Aupetit (HdR), Hai Le Son (PhD)

- Olivier Teytaud was in the PhD committee for Benoit Gandar (Clermont-Ferrand, November 2012).

9.4. Popularization

- Olivier Teytaud has written a popularization paper on Chess algorithms in Interstices.

- We made two demonstrations, in Taiwan, of parallel automatic player evaluation; the events were advertised in general audience newspapers (MoGo web page). Our Taiwanese partner has a motivation around PISA (international student assessment) and eTeaching.

- We also made in Brisbane a general audience demonstration of games against humans, with 7 wins out of 12 against professional players in fair games in 7x7 (MoGo web page).

- Yann Ollivier organizes a bi-monthly math seminar for undergrad students on Saturdays at Institut Henri Poincare (together with X. Caruso, I. Kortchemski, R. Mansuy and A. Taveneaux), with 100+ participants at each session.

- Yann Ollivier takes part in the organization of the European Union Contest for Young Scientists (science fair for high school students from 30+ countries organized by the European Commission).

- Yann Ollivier belongs to the scientific steering committee for the elaboration of a brochure *L'explosion des mathématiques* presenting a wide range of applications of mathematics, edited by the SMF and SMAI.

- Philippe Caillou was a teacher at the GAMA Winter School in Can Tho (Vietnam) in November 2012. The winter school objective was the diffusion of multi-agent based simulation methodology with the GAMA Platform to non-computer scientist in Can Tho Biology and Environment University.

10. Bibliography

Publications of the year

Doctoral Dissertations and Habilitation Theses

[1] R. BARDENET. *Towards adaptive learning and inference (applications to hyperparameter tuning and astroparticle physics)*, Université Paris Sud - Paris XI, December 2012.

[2] Z. BOUZARKOUNA. *Optimisation de placement des puits*, Université Paris Sud - Paris XI, April 2012, http://hal.inria.fr/tel-00690456.

[3] M. YAGOUBI. *Multi-objective parallel evolutionary algorithms : Application on Diesel Combustion*, Université Paris Sud - Paris XI, July 2012, http://tel.archives-ouvertes.fr/tel-00734108.

Articles in International Peer-Reviewed Journals

[4] D. BENBOUZID, R. BUSA-FEKETE, N. CASAGRANDE, F.-D. COLLIN, B. KÉGL. *Multiboost: a multi-purpose boosting package*, in "Journal of Machine Learning Research", 2012, vol. 13, p. 549-553, http://hal.inria.fr/in2p3-00698455.

[5] 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", 2012, vol. 18, n[o] 1, p. 101-129, http://hal.inria.fr/inria-00531450.

[6] P. CAILLOU, J. GIL-QUIJANO, X. ZHOU. *Automated observation of multi-agent based simulations: a statistical analysis approach*, in "Studia Informatica Universalis", 2013, page : to appear, http://hal.inria.fr/hal-00738384.

[7] F.-M. DE RAINVILLE, C. GAGNÉ, O. TEYTAUD, D. LAURENDEAU. *Evolutionary Optimization of Low-Discrepancy Sequences*, in "ACM Transactions on Modeling and Computer Simulation", 2012, vol. 22, n[o] 2, p. 9:1-9:25, http://hal.inria.fr/hal-00758158.

[8] 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", 2012, vol. 116, n[o] 1, p. 146-165, http://www.sciencedirect.com/science/article/pii/S1077314211001998.

[9] C. FURTLEHNER, J.-M. LASGOUTTES, M. SAMSONOV. *One-dimensional Particle Processes with Acceleration/Braking Asymmetry*, in "Journal of Statistical Physics", June 2012, vol. 147, n^o 6, p. 1113-1144 [*DOI :* 10.1007/S10955-012-0521-Y], http://hal.inria.fr/hal-00743369.

[10] S. GELLY, M. SCHOENAUER, M. SEBAG, O. TEYTAUD, L. KOCSIS, D. SILVER, C. SZEPESVARI. *The Grand Challenge of Computer Go: Monte Carlo Tree Search and Extensions*, in "Communications- ACM", 2012, vol. 55, n^o 3, p. 106-113, Invited paper, http://hal.inria.fr/hal-00695370.

[11] M. R. KHOUADJIA, B. SARASOLA, E. ALBA, L. JOURDAN, E.-G. TALBI. *A comparative study between dynamic adapted PSO and VNS for the vehicle routing problem with dynamic requests*, in "Applied Soft Computing", 2012, vol. 12, n^o 4, p. 1426-1439, http://hal.inria.fr/hal-00750715.

[12] C.-S. LEE, M.-H. WANG, Y.-J. CHEN, H. HAGRAS, M.-J. WU, O. TEYTAUD. *Genetic fuzzy markup language for game of NoGo*, in "Knowledge-Based Systems", 2012, vol. 34, p. 64-80, http://www.sciencedirect.com/science/article/pii/S0950705112000469.

[13] A. MUZY, F. VARENNE, B. P. ZEIGLER, J. CAUX, P. COQUILLARD, L. TOURAILLE, D. PRUNETTI, P. CAILLOU, O. MICHEL, D. HILL. *Refounding of Activity Concept ? Towards a Federative Paradigm for Modeling and Simulation*, in "Simulation: Transactions of the Society for Modeling and Simulation International", 2012, To appear [*DOI :* 10.1177/0037549712457852], http://hal.inria.fr/hal-00738218.

[14] Y. OLLIVIER, C. VILLANI. *A Brunn–Minkowski inequality on the discrete hypercube*, in "SIAM J. Discrete Math.", 2012, vol. 26, n^o 3, p. 983–996, http://dx.doi.org/10.1137/11085966X.

[15] T. ÉLTETO, N. HANSEN, C. GERMAIN-RENAUD, P. BONDON. *Scalable structural break detection*, in "Applied Soft Computing", 2012 [*DOI :* 10.1016/J.ASOC.2012.06.002], http://hal.inria.fr/hal-00711843.

Articles in Non Peer-Reviewed Journals

[16] C.-S. LEE, O. TEYTAUD, M.-H. WANG, S.-J. YEN. *Computational Intelligence Meets the Game of Go @ IEEE WCCI 2012*, in "Computational Intelligence Magazine", November 2012, Invited paper, http://hal.inria.fr/hal-00744982.

Invited Conferences

[17] N. BREDECHE, J.-M. MONTANIER. *Environment-driven Open-ended Evolution with a Population of Autonomous Robots*, in "Evolving Physical Systems Workshop, ALIFE XIII", East Lansing, États-Unis, 2012, p. 7-14, http://hal.inria.fr/hal-00731422.

International Conferences with Proceedings

[18] Y. AKIMOTO. *Analysis of a Natural Gradient Algorithm on Monotonic Convex-Quadratic-Composite Functions*, in "Proc. Genetic and Evolutionary Computation Conference (ACM-GECCO)", Philadelphia, États-Unis, April 2012, http://hal.inria.fr/hal-00688909.

[19] Y. AKIMOTO, A. AUGER, N. HANSEN. *Convergence of the Continuous Time Trajectories of Isotropic Evolution Strategies on Monotonic C2-composite Functions*, in "PPSN - 12th International Conference on Parallel Problem Solving from Nature - 2012", Taormina, Italie, Springer, June 2012, http://hal.inria.fr/hal-00710214.

[20] Y. AKIMOTO, Y. OLLIVIER. *Objective Improvement in Information-Geometric Optimization*, in "Foundations of Genetic Algorithms (FOGA 2013)", F. NEUMANN, K. DEJONG (editors), 2013, To appear.

[21] R. AKROUR, M. SCHOENAUER, M. SEBAG. *APRIL: Active Preference-learning based Reinforcement Learning*, in "ECML PKDD 2012", Bristol, Royaume-Uni, PETER FLACH ET AL. (editor), LNCS, Springer Verlag, September 2012, vol. 7524, p. 116-131, http://hal.inria.fr/hal-00722744.

[22] D. ARNOLD, N. HANSEN. *A (1+1)-CMA-ES for Constrained Optimisation*, in "Proc. Genetic and Evolutionary Computation Conference (ACM-GECCO)", Philadelphia, États-Unis, ACM, 2012, p. 297-304 [*DOI : 10.1145/2330163.2330207*], http://hal.inria.fr/hal-00696268.

[23] R. BARDENET, O. CAPPÉ, G. FORT, B. KÉGL. *Adaptive Metropolis with online relabeling*, in "Fifteenth International Conference on Artificial Intelligence and Statistics", La Palma, Espagne, N. LAWRENCE, M. GIROLAMI (editors), 2012, vol. 22, p. 91-99, http://hal.inria.fr/in2p3-00698479.

[24] R. BARDENET, B. KÉGL. *An adaptive Monte-Carlo Markov chain algorithm for inference from mixture signals*, in "14th International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT 2011)", Uxbridge, Royaume-Uni, L. TEODORESCU, D. BRITTON, N. GLOVER, G. HEINRICH, J. LAURET, A. NAUMANN, T. SPEER, P. TEIXEIRA-DIAS (editors), 2012, vol. 368, 012044 [*DOI : 10.1088/1742-6596/368/1/012044*], http://hal.inria.fr/in2p3-00714497.

[25] Z. BOUZARKOUNA, D. Y. DING, A. AUGER. *Well Placement Optimization under Uncertainty with CMA-ES Using the Neighborhood*, in "13th European Conference on the Mathematics of Oil Recovery ECMOR 2012", Biarritz, France, EAGE, September 2012, http://hal.inria.fr/hal-00727635.

[26] M. BRENDEL, M. SCHOENAUER. *Learn-and-Optimize: a Parameter Tuning Framework for Evolutionary AI Planning*, in "Artificial Evolution", Angers, France, JIN-KAO HAO ET AL. (editor), LNCS, Springer Verlag, September 2012, vol. 7401, p. 159-170, http://hal.inria.fr/inria-00632378.

[27] D. BROCKHOFF, A. AUGER, N. HANSEN. *Comparing Mirrored Mutations and Active Covariance Matrix Adaptation in the IPOP-CMA-ES on the Noiseless BBOB Testbed*, in "GECCO Companion", Philadelphia, PA, États-Unis, July 2012, p. 297-303, only loosely reviewed [*DOI : 10.1145/2330784.2330827*], http://hal.inria.fr/hal-00746120.

[28] D. BROCKHOFF, A. AUGER, N. HANSEN. *On the Effect of Mirroring in the IPOP Active CMA-ES on the Noiseless BBOB Testbed*, in "GECCO Companion", Philadelphia, PA, États-Unis, July 2012, p. 277-284, only loosely reviewed [*DOI : 10.1145/2330784.2330824*], http://hal.inria.fr/hal-00746123.

[29] D. BROCKHOFF, A. AUGER, N. HANSEN. *On the Impact of a Small Initial Population Size in the IPOP Active CMA-ES with Mirrored Mutations on the Noiseless BBOB Testbed*, in "GECCO Companion", Philadelphia, PA, États-Unis, July 2012, p. 285-290, only loosely reviewed [*DOI : 10.1145/2330784.2330825*], http://hal.inria.fr/hal-00746122.

[30] D. BROCKHOFF, A. AUGER, N. HANSEN. *On the Impact of Active Covariance Matrix Adaptation in the CMA-ES With Mirrored Mutations and Small Initial Population Size on the Noiseless BBOB Testbed*, in "GECCO Companion", Philadelphia, PA, États-Unis, July 2012, p. 291-296, only loosely reviewed [*DOI : 10.1145/2330784.2330826*], http://hal.inria.fr/hal-00746121.

[31] O. BUFFET, C.-S. LEE, W. LIN, O. TEYTAUD. *Optimistic Heuristics for MineSweeper*, in "International Computer Symposium", Hualien, Taïwan, Province De Chine, 2012, http://hal.inria.fr/hal-00750577.

[32] R. BUSA-FEKETE, D. BENBOUZID, B. KÉGL. *Fast classification using sparse decision DAGs*, in "29th International Conference on Machine Learning (ICML 2012)", Edinburgh, Royaume-Uni, June 2012, http://hal.inria.fr/in2p3-00711150.

[33] P. CAILLOU, J. GIL-QUIJANO. *SimAnalyzer: automated description of groups dynamics in agent-based simulations*, in "AAMAS 2012", Valencia, Espagne, V. CONITZER, M. WINIKOFF, L. PADGHAM, W. VAN DER HOEK (editors), International Foundation for Autonomous Agents and Multiagent Systems, June 2012, p. 1353-1354, http://hal.inria.fr/hal-00738182.

[34] L. CAZENILLE, N. BREDECHE, H. HAMANN, J. STRADNER. *Impact of Neuron Models and Network Structure on Evolving Modular Robot Neural Network Controllers*, in "Proc. Genetic and Evolutionary Computation Conference (ACM-GECCO)", Philadelphia, États-Unis, 2012, p. 89-96, http://hal.inria.fr/hal-00731411.

[35] A. CHOTARD, A. AUGER, N. HANSEN. *Cumulative Step-size Adaptation on Linear Functions: Technical Report*, in "International Conference on Parallel Problem Solving from Nature", Springer, June 2012, http://hal.inria.fr/hal-00704903.

[36] *Best Paper*
C.-W. CHOU, P.-C. CHOU, C.-S. LEE, D. LUPIEN SAINT-PIERRE, O. TEYTAUD, M.-H. WANG, L.-W. WU, S.-J. YEN. *Strategic Choices: Small Budgets and Simple Regret*, in "TAAI", Hualien, Taïwan, Province De Chine, 2012, http://hal.inria.fr/hal-00753145.

[37] A. COUËTOUX, H. DOGHMEN, O. TEYTAUD. *Improving the exploration in Upper Confidence Trees*, in "Learning and Intelligent OptimizatioN (LION'6)", Paris, France, Y. HAMADI, M. SCHOENAUER (editors), LNCS, Sringer Verlag, 2012, vol. 7219, p. 366-380, http://hal.inria.fr/hal-00745208.

[38] A. COUËTOUX, H. DOGHMEN, O. TEYTAUD. *Learning a move-generator for Monte-Carlo Tree Search*, in "International Computer Symposium", 2012, http://hal.inria.fr/hal-00759822.

[39] D. FENG, C. GERMAIN-RENAUD, T. GLATARD. *Distributed Monitoring with Collaborative Prediction*, in "12th IEEE International Symposium on Cluster, Cloud and Grid Computing (CCGrid'12)", Ottawa, Canada, IEEE, May 2012, e-pub ahead of print, http://hal.inria.fr/hal-00673148.

[40] C. GERMAIN-RENAUD, F. FURST, T. JACOB, M. JOUVIN, G. KASSEL, J. NAUROY, G. PHILIPPON. *The Green Computing Observatory: a data curation approach for green IT*, in "Selected papers from EGI Community Forum 2012", Berlin, Germany, Proceedings of Science, 2012, http://pos.sissa.it/archive/conferences/162/059/EGICF12-EMITC2_059.pdf.

[41] A. HADJAZ, G. MARCEAU, P. SAVÉANT, M. SCHOENAUER. *Increasing Air Traffic: What is the Problem?*, in "Proceedings of the SESAR Innovation Days (2012)", EUROCONTROL, November 2012, http://hal.inria.fr/hal-00766017.

[42] A. HADJAZ, G. MARCEAU, P. SAVÉANT, M. SCHOENAUER. *Online Learning for Ground Trajectory Prediction*, in "Proceedings of the SESAR Innovation Days (2012)", EUROCONTROL, November 2012, http://hal.inria.fr/hal-00766049.

[43] I. HEGEDS, R. BUSA-FEKETE, R. ORMÁNDI, M. JELASITY, B. KÉGL. *Peer-to-peer multi-class boosting*, in "18th International European Conference on Parallel and Distributed Computing (Euro-Par 2012)", Rhodes Island, Grèce, C. KAKLAMANIS, T. PAPATHEODOROU, P. G. SPIRAKIS (editors), 2012, vol. 7484, p. 389-400, ISBN: 978-3-642-32819-0 (Print) 978-3-642-32820-6 (Online), http://hal.inria.fr/in2p3-00726735.

[44] M. R. KHOUADJIA, M. SCHOENAUER, V. VIDAL, J. DRÉO, P. SAVÉANT. *Multi-Objective AI Planning: Comparing Aggregation and Pareto Approaches*, in "Proc. EvoCOP'13", M. MIDDENDORF, C. BLUM (editors), LNCS, Springer Verlag, 2013, To appear.

[45] M. R. KHOUADJIA, M. SCHOENAUER, V. VIDAL, J. DRÉO, P. SAVÉANT. *Multi-Objective AI Planning: Evaluating DAE-YAHSP on a Tunable Benchmark*, in "7th International Conference on Evolutionary Multi-Criterion Optimization", Sheffield, Royaume-Uni, R. C. PURSHOUSE, P. J. FLEMING, C. M. FONSECA (editors), LNCS, Springer Verlag, March 2013, To appear, http://hal.inria.fr/hal-00750560.

[46] M. R. KHOUADJIA, M. SCHOENAUER, V. VIDAL, J. DRÉO, P. SAVÉANT. *Quality Measures of Parameter Tuning for Aggregated Multi-Objective Temporal Planning*, in "Proc. LION'7", P. PARDALOS, G. NICOSIA (editors), LNCS, Springer Verlag, 2013, To appear.

[47] I. LOSHCHILOV, M. SCHOENAUER, M. SEBAG. *Alternative Restart Strategies for CMA-ES*, in "Parallel Problem Solving From Nature", Taormina, Italie, July 2012, http://hal.inria.fr/hal-00713415.

[48] I. LOSHCHILOV, M. SCHOENAUER, M. SEBAG. *Black-box optimization benchmarking of IPOP-saACM-ES and BIPOP-saACM-ES on the BBOB-2012 noiseless testbed*, in "GECCO Companion", Philadelphia, États-Unis, April 2012, only loosely reviewed, http://hal.inria.fr/hal-00690444.

[49] I. LOSHCHILOV, M. SCHOENAUER, M. SEBAG. *Black-box optimization benchmarking of IPOP-saACM-ES on the BBOB-2012 noisy testbed*, in "GECCO Companion", Philadelphia, États-Unis, April 2012, only loosely reviewed, http://hal.inria.fr/hal-00690543.

[50] I. LOSHCHILOV, M. SCHOENAUER, M. SEBAG. *Black-box Optimization Benchmarking of NIPOP-aCMA-ES and NBIPOP-aCMA-ES on the BBOB-2012 Noiseless Testbed*, in "GECCO Companion", Philadelphia, États-Unis, October 2012, only loosely reviewed, http://hal.inria.fr/hal-00737409.

[51] I. LOSHCHILOV, M. SCHOENAUER, M. SEBAG. *Self-Adaptive Surrogate-Assisted Covariance Matrix Adaptation Evolution Strategy*, in "Proc. Genetic and Evolutionary Computation Conference (ACM-GECCO)", Philadelphia, États-Unis, April 2012, http://hal.inria.fr/hal-00686570.

[52] M. LOTH, M. SEBAG, Y. HAMADI, M. SCHOENAUER, C. SCHULTE. *Hybridizing Constraint Programming and Monte-Carlo Tree Search: Application to the Job Shop problem*, in "Proc. LION'7", P. PARDALOS, G. NICOSIA (editors), LNCS, Springer Verlag, 2013, To appear.

[53] T. RUNARSSON, M. SCHOENAUER, M. SEBAG. *Pilot, Rollout and Monte Carlo Tree Search Methods for Job Shop Scheduling*, in "Learning and Intelligent OptimizatioN (LION'6)", Paris, France, Y. HAMADI, M. SCHOENAUER (editors), LNCS, Sringer Verlag, 2012, vol. 7219, p. 160-174, http://hal.inria.fr/hal-00736968.

[54] M. SEBAG, O. TEYTAUD. *Combining Myopic Optimization and Tree Search: Application to MineSweeper*, in "Learning and Intelligent OptimizatioN (LION'6)", Paris, France, Y. HAMADI, M. SCHOENAUER (editors), LNCS, Sringer Verlag, 2012, vol. 7219, p. 220-224, http://hal.inria.fr/hal-00712417.

[55] O. TEYTAUD, J. DECOCK. *Noisy Optimization Complexity*, in "FOGA - Foundations of Genetic Algorithms XII - 2013", Adelaide, Australie, Feb. 2013, http://hal.inria.fr/hal-00755663.

[56] W. WANG, M. SEBAG. *Multi-objective Monte-Carlo Tree Search*, in "Asian Conferenc on Machine Learning (ACML 2012)", Paris, France, S. C. HOI, W. BUNTINE (editors), JMLR Workshop and Conference Proceedings, 2012, vol. 25, p. 507-522.

[57] M. YAGOUBI, M. SCHOENAUER. *Asynchronous Master/Slave MOEAs and Heterogeneous Evaluation Costs*, in "Genetic and Evolutionary Computation Conference (GECCO 2012)", États-Unis, T. SOULE, J. H. MOORE (editors), ACM Press, 2012, p. 1007-1014, http://hal.inria.fr/hal-00689965.

National Conferences with Proceeding

[58] P. CAILLOU, J. GIL-QUIJANO. *Description automatique de dynamiques de groupes dans des simulations à base d'agents*, in "JFSMA 2012", France, October 2012, page : to appear, http://hal.inria.fr/hal-00738184.

[59] N. GALICHET, M. SEBAG. *Exploration prudente : une approche par méthode de Monte-Carlo arborescente contrainte*, in "RFIA 2012 (Reconnaissance des Formes et Intelligence Artificielle)", Lyon, France, January 2012, p. 978-2-9539515-2-3, Session "Posters", http://hal.inria.fr/hal-00656575.

[60] V. MARTIN, J.-M. LASGOUTTES, C. FURTLEHNER. *Local stability of Belief Propagation algorithm with multiple fixed points*, in "STAIRS'12 - Sixth "Starting Artificial Intelligence Research" Symposium", Montpellier, France, IOS Press, 2012, http://hal.inria.fr/hal-00719204.

Conferences without Proceedings

[61] R. BUSA-FEKETE, G. SZARVAS, T. ÉLTETO, B. KÉGL. *An apple-to-apple comparison of Learning-to-rank algorithms in terms of Normalized Discounted Cumulative Gain*, in "20th European Conference on Artificial Intelligence (ECAI 2012) : Preference Learning: Problems and Applications in AI Workshop", Montpellier, France, August 2012, http://hal.inria.fr/in2p3-00726760.

[62] D. FENG, C. GERMAIN-RENAUD, T. GLATARD. *Efficient fault monitoring with Collaborative Prediction*, in "Journées scientifiques mésocentres et France Grilles", Paris, France, France Grilles, October 2012, e-pub ahead of print, http://hal.inria.fr/hal-00758025.

[63] J.-B. HOOCK, J. BIBAI. *Solving a Goal-planning task in the MASH project*, in "The 2012 Conference on Technologies and Applications of Artificial Intelligence (TAAI 2012)", Tainan, Taïwan, Province De Chine, 2012, http://hal.inria.fr/hal-00738073.

Scientific Books (or Scientific Book chapters)

[64] C. FURTLEHNER. *An Ising Model for Road Traffic Inference*, in "From Hamiltonian Chaos to Complex Systems: a Nonlinear Physics Approach", X. LEONCINI, M. LEONETTI (editors), Springer, December 2012, http://hal.inria.fr/hal-00743351.

[65] C.-S. LEE, M.-H. WANG, O. TEYTAUD. *Fuzzy Ontologies for the Game of Go*, in "On Fuzziness. A Homage to Lotfi A. Zadeh", R. SEISING, E. TRILLAS, C. MORAGA, S. TERMINI (editors), Springer, 2012, p. 367-372, http://hal.inria.fr/hal-00758147.

Books or Proceedings Editing

[66] J.-K. HAO, P. LEGRAND, P. COLLET, N. MONMARCHÉ, E. LUTTON, M. SCHOENAUER (editors). *Artificial Evolution 2011*, LNCS, Springer Verlag, November 2012, vol. 7401, http://hal.inria.fr/hal-00643404.

Research Reports

[67] C. FURTLEHNER, Y. HAN, J.-M. LASGOUTTES, V. MARTIN. *Pairwise MRF Calibration by Perturbation of the Bethe Reference Point*, Inria, October 2012, n$^{\text{o}}$ RR-8059, 35, http://hal.inria.fr/hal-00743334.

[68] N. LAZAAR, Y. HAMADI, S. JABBOUR, M. SEBAG. *Cooperation control in Parallel SAT Solving: a Multi-armed Bandit Approach*, Inria, September 2012, n$^{\text{o}}$ RR-8070, 18, http://hal.inria.fr/hal-00733282.

Other Publications

[69] J. ATIF, C. HUDELOT, I. BLOCH. *Explanatory reasoning for image understanding using formal concept analysis and description logics*, 2013, Under review in IEEE Transactions in Man Systems and Cybernetics-Part A.

[70] O. NEMPONT, J. ATIF, I. BLOCH. *A constraint propagation approach to structural model based image segmentation and recognition*, 2013, Under review in Information Sciences.

References in notes

[71] L. ARNOLD, A. AUGER, N. HANSEN, Y. OLLIVIER. *Information-Geometric Optimization Algorithms: A Unifying Picture via Invariance Principles*, in "ArXiv e-prints", June 2011, http://adsabs.harvard.edu/abs/2011arXiv1106.3708A.

[72] J. BIBAI, P. SAVÉANT, M. SCHOENAUER, V. VIDAL. *An Evolutionary Metaheuristic Based on State Decomposition for Domain-Independent Satisficing Planning*, in "ICAPS 2010", R. BRAFMAN, H. GEFFNER, J. HOFFMANN, H. KAUTZ (editors), AAAI Press, May 2010, p. 15-25, http://hal.archives-ouvertes.fr/docs/00/45/62/92/PDF/icaps10.pdf.