IN PARTNERSHIP WITH:
CNRS
Université Paris-Sud (Paris 11)

Activity Report 2013

Project-Team TAO

Machine Learning and Optimization

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

RESEARCH CENTER
Saclay - Île-de-France

THEME
Optimization, machine learning and statistical methods

Table of contents

Project-Team TAO

Keywords: Machine Learning, Statistical Learning, Stochastic Algorithms, Optimization

Creation of the Project-Team: 2004 November 04.

1. Members

Research Scientists
- Marc Schoenauer [Team leader, Senior Researcher, Inria, HdR]
- Anne Auger [Junior Researcher, Inria]
- Cyril Furtlehner [Junior Researcher, Inria]
- Nikolaus Hansen [Junior Researcher, Inria, HdR]
- Yann Ollivier [Junior Researcher, CNRS, HdR]
- Michèle Sebag [Team co-leader, Senior Researcher, CNRS, HdR]
- Olivier Teytaud [Junior Researcher, Inria, HdR]

Faculty Members
- Jamal Atif [Associate Professor, Univ Paris XI, HdR]
- Philippe Caillou [Associate Professor, Univ Paris XI]
- Florence d'Alché Buc [Professor, Univ Évry, en délégation until Aug 2013, HdR]
- Cécile Germain-Renaud [Professor, Univ Paris XI, HdR]

External Collaborators
- Balázs Kégl [Senior Researcher, CNRS, HdR]
- Hélène Paugam-Moisy [Professor, Université Lyon 2, HdR]

Engineers
- Jean-Joseph Christophe [Inria, I-lab METIS]
- Lovro Ilijasic [Inria, Conseil Régional d'Ile-de-France, from Sep 2013]
- Julien Nauroy [Inria, Conseil Régional d'Ile-de-France, from Sep 2013]

PhD Students
- Ouassim Ait Elhara [Univ Paris XI]
- Riad Akrour [Inria, FP7 IP SYMBRION, until Oct 2013 ; ATER Univ Paris XI, since Nov 2013]
- Ludovic Arnold [Inria, ANR ASAP, until Apr 2013]
- Asma Atamna [Inria, ANR NumBBO, from November 2013]
- Jeremy Bensadon [Univ Paris XI]
- Vincent Berthier [Univ Paris XI, from Sep 2013]
- Marie-Liesse Cauwet [Inria, Chile until August 2013, ADEME POST, from Sep 2013]
- Sandra Cecilia Astete Morales [Inria, ADEME POST,from Sep 2013]
- Alexandre Chotard [Univ Paris XI]
- Adrien Couëtoux [Inria, Artelys SA, until Aug 2013, Defended Sep 2013]
- Jérémie Decock [Inria, Inria]
- Dawei Feng [Univ Paris XI]
- Nicolas Galichet [Univ Paris XI]
- Yoann Isaac [UCEA, Unsupervised Brain project]
- Jialin Liu [Inria, ADEME POST, from March 2013]
- Ilya Loshchilov [Inria, FUI CSDL then FRM, until January 2013]
- Gaetan Marceau-Caron [CIFRE, Thales]
- Jean-Marc Montanier [Univ Paris XI, until Feb 2013]
- Weija Wang [Univ Paris XI]
- Yifan Yang [Telecom Paristech, from Oct 2013]
- Guohua Zhang [Chinese Scholarship]

Post-Doctoral Fellows

Yohei Akimoto [Inria, ANR SIMINOLE, until Mar 2013]
Marco Bressan [Inria, Modyrum (Augure), from Feb 2013]
Mostepha Redouane Khouadjia [Inria, ANR DESCARWIN, until Mar 2013]
Manuel Loth [Microsoft-Inria until October 2013]
David Lupien Saint Pierre [Inria, ADEME POST, from Sep 2013]
Victorin Martin [Inria, ANR SIMINOLE, until Aug 2013]
Mustafa Misir [Inria, ERCIM, until Oct 2013]

Visiting Scientists

Daria La Rocca [PhD Ministère de la Défense Italien, from Nov 2013]
Martin Pilat [PhD Charles Univ. Prague, from Apr 2013 until Jun 2013]

Administrative Assistants

Marie-Carol Lopes [Inria, team assistant until Mar 2013]
Olga Mwana Mobulakani [Inria, team assistant from Jul 2013]

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, predict, understand, and 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 is based on 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 are 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

2.3.1. Extensions of Multi-Armed Bandits and Monte-Carlo Tree Search

Risk Avoidance Exploration might exert a toll on the agent/system safety in real-world contexts (e.g., controlling a power system or a robot). Risk adverse criteria have been pioneered in MAB, together with multi-objective reinforcement learning – see [12] and [19].

Continuous Options The Rapid Action Value Estimate (RAVE) has been extended to continuous settings [27].

2.3.2. Information Theory and Natural Gradient

Information-geometric Optimization: convergence results. Theoretical guarantees have been obtained for continuous optimization algorithms in the framework of *information geometry* (IGO). Previous improvement guarantees for gradient descent-based methods were valid only for infinitesimally small step sizes. Information geometry and using the *natural gradient* provide improvement guarantees for finite step sizes as is the case in practice [22]. Along the same lines, geodesics in statistical manifolds have been used for estimation of distribution optimization algorithms.

Neural Network Training is a hard optimization problem, sensitive to the problem representation and the optimization trajectory. Within a Riemannian geometry framework, the use of intrinsic Riemannian gradient has been shown to support an affine transformation-invariant optimization approach, with significant robustness improvements at the same cost as the state of the art [66]. This Riemannian approach has been applied to recurrent neural nets, with very satisfactory results on difficult symbolic sequences with non-local dependencies [65]. In the related field of stacked restricted Boltzman machines, we have shown that the layer-wise approach supporting the celebrated deep learning approach yields *globally optimal* results provided the inference (as opposed to generative) model is rich enough, with quantitative estimates [60]. This result is the first of its kind on layerwise deep learning.

3. Research Program

3.1. The Four Pillars of TAO

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 and 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 **Stochastic Continuous Optimization** SIG (OPT-SIG) takes 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 state-of-the-art continuous stochastic optimizer, CMA-ES, ranging from convergence analysis (Youhei Akimoto's post-docs) to a rigorous benchmarking methodology. Incidentally, this benchmark platform COCO has been acknowledged since 2009 as “the“ international continuous optimization benchmark, and its extension is at the core of the ANR project NumBBO (started 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).

The **Optimal Decision Making under Uncertainty** SIG (UCT-SIG) benefits from the MoGo expertise (see Section 5.2 and the team previous activity reports) 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 ADEME Post contract, and the METIS I-lab with SME Artelys.

The **Distributed systems** 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-Renaud 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*, including the modeling of multi-agent systems and the exploitation of simulation results in the SimTools RNSC network frame. Further extensions have been developped 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.

The **Designing Criteria** 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 often 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. As a matter of fact, though not a major research topic any more, hyper-parameter tuning has become pervasive in TAO, chiefly for continuous optimization (OPT-SIG, Section 6.1), AI planning (CRI-SIG, Section 6.4) and Air Traffic Control Optimization (Section 4.2). Recent work addressing algorithm selection using Collaborative Filtering algorithms (CRI-SIG, Section 6.4) can (and will) indeed be applied to hyper-parameter tuning for optimization algorithms.

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 cannot consider a power grid like Europe + North Africa as a single “production = demand” constraint; with extra connections we can balance excess production by renewables for remote areas, although to a limited extent.
- Other uncertainties, which might be modeled 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), and POST, a ADEME project about investments in power systems. We have a collaboration with the SME Artelys, that works on optimization in general, and on energy management in particular.

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 6.2. 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 an I-lab in progress with Artelys (see Section 5.1) in order to ensure the transfer of our work.
- We have 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 1998 and S. Oussedik in 2000). It was revived recently with Gaëtan Marceau-Caron's CIFRE PhD together with Thalès Air Systems (Areski Hadjaz) and Thalès TRT (Pierre Savéant), tackling the global optimization of the traffic in order to increase the capacity of the airspace without overloading the controllers. A new formulation of the problem, modeling the plane flows with Bayesian Networks, has been proposed to the Air Traffic Control community [48], [50]. The goal of the optimization is to minimize the cumulated delays of all flights, while maintaining a reasonnable level of congestion in all sectors. These objectives are computed using Monte-Carlo simulations of the Bayesian network, and Evolutionary Algorithms are used to address the resulting stochastic multi-objective optimization problem [49].

5. Software and Platforms

5.1. METIS

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

Keywords: 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) models.

5.2. MoGo

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

Keywords:

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) and outpeforming professional players in 7x7. However, the work in the UCT-SIG has now shifted to energy management.

5.3. CMA-ES: Covariance Matrix Adaptation Evolution Strategy

Participant: Nikolaus Hansen [correspondent].

Keywords: Evolutionary Computation, Stochastic Optimization, Real-parameter Optimization.

The Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [72] is considered to be state-of-the-art in continuous domain evolutionary computation [69], and in stochastic optimization at large. It has been shown to be highly competitive on different problem classes even with deterministic continuous algorithms using numerically computed gradients (see the results published on COCO platform). 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.

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

5.4. COmparing Continuous Optimizers

Participants: Nikolaus Hansen [correspondent], Anne Auger, Marc Schoenauer, Ouassim Ait Elhara, Asma Atamna.

Keywords: 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, C, Java, and Python. The post-processing code is provided in Python. The code is under continuous developement and has been used for the GECCO 2009, 2010, 2012, and 2013 workshops on "Black Box Optimization Benchmarking" (BBOB) (see Section 6.1). It is now undergoing major changes thanks to the ANR project NumBBO that will add constraint handling and multi-objective benchmarks to the existing platform.

Link: http://coco.gforge.inria.fr/ and http://numbbo.gforge.inria.fr/

5.5. MultiBoost

Participants: Balázs Kégl [correspondent], Djalel Benbouzid.

Keywords: Multi-class, Multi-label Classification.

The MultiBoost package [68] 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-Renaud [correspondent], Julien Nauroy, Michèle Sebag.

Keywords: Autonomic Computing, Green Computing.

The Grid Observatory (GO) 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 since 2008. These data are stored on the grid, and made accessible through a web portal without the need of grid credentials. The GO is fully integrated with the evolution of EGI monitoring. More than 250 users are currently registered. The acquisition has been extended to the University cloud StratusLab hosted by the VirtualData center.

The Green Computing Observatory (GCO) monitors the VirtualData center; it collects data on energy consumption and publishes the data through the Grid Observatory. These data include the detailed monitoring of the processors and motherboards, as well as global site information. The first results on energy saving opportunities have been presented at the Green Days@Luxembourg meeting.

In order to make the GO data readily consistent and complete, as well as understandable for further exploitation, an original approach has been designed, based on a flexible data schema built in collaboration with the users. Its implementation is developed within the FUI project TIMCO. The GO has been supported by an Inria ADT (Action de Développement Technologique) up to September 2013, and by University Paris Sud through the MRM (Moyens de Recherche Mutualisés) program. Stabilization through the VirtualData initiative is currently explored.

Link: http://grid-observatory.org

6. New Results

6.1. Continuous Optimization

Participants: Ouassim Ait Elhara, Yohei Akimoto, Asma Atamna, Anne Auger, Alexandre Chotard, Nikolaus Hansen, Ilya Loshchilov, Yann Ollivier, Marc Schoenauer, Michèle Sebag, Olivier Teytaud.

Our expertise in continuous optimization is focused on stochastic search algorithms. We address theory, algorithm design, and applications. The methods we investigate are adaptive techniques that are able to learn iteratively the parameters of the distribution used to sample (new) 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 previously proven the convergence of simplified variants of the CMA-ES algorithm using the theory of stochastic approximation, and have provided the first proofs of convergence on composite of twice continuously differentiable functions. More recently, we used Markov chain analysis for analyzing the step-size adaptation rules of evolution strategies related to the CMA-ES algorithm.

Surrogate models for CMA-ES. In the context of his PhD thesis defended in January 2013 [4], 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. As a follow-up, he has proposed an original over-exploitation mechanism in case of accurate surrogate [44]. Several of these models have entered the BBOB-2013 workshop [43]. Further research direction using the KL divergence between successive distributions as a trigger for a new learning phase has been proposed [45].

Step-size adaptive methods. We have proposed a new step-size adaptation mechanism that can loosely be interpreted as a new variant of the 1/5 success rule for comma (non-elitist) strategies and which is applicable with a large population size [21]. The rule uses the success of the median fitness of the current population compared to a (different) fitness percentile from the previous population.

Principles of Stochastic Optimization. Based on the framework of *information geometry* (IGO), theoretical guarantees have been obtained for continuous optimization algorithms: using the *natural gradient* provides improvement guarantees even when using finite step sizes [22]. We have considered the principles of designing effective stochastic optimization algorithms from the bottom-up and the top-down perspective [56]. The top-down perspective takes the information-geometrical view-point and largely confirms the bottom-up construction.

Benchmarking. We have continued our effort for improving standards in benchmarking and pursued the development of the COCO (COmparing Continuous Optimizers) platform (see Section 5.4). We have organized the ACM GECCO 2013 workshop on Black-Box-Optimization Benchmarking [2] and benchmarked different surrogate-based variants of the CMA-ES algorithm [26], [44], [43]. Our new starting ANR project NumBBO, centered on the COCO platform, aims at extending it for large-scale, expensive, constrained and multi-objective optimization.

Theoretical proofs of convergence. We have established the connection between convergence of comparison based step-size adaptive randomized search and the stability analysis of some underlying Markov chains. This connection heavily exploits invariance properties of the algorithm. In a first paper we establish this connection for scaling-invariant functions and prove sufficient conditions for linear convergence expressed in terms of stability conditions [63]. We have proven, using this defined methodology, the linear convergence of a famous algorithm introduced independently by several researchers and known as the (1+1)-ES with one-fifth success rule [62]. In [32], we have proven the linear convergence of a modified evolutionary algorithm without assuming quasi-convexity.

[2] see http://coco.gforge.inria.fr/doku.php?id=bbob-2013

6.2. Optimal Decision Making under Uncertainty

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

The most prominent application domain is now energy management, at various time scales, and more generally, planning in uncertain environments. The main advances done this year include:

- A work on metagaming/investment [12], where a macroscopic decision has to be made (e.g., investment decisions, which plants should be built) prior to operational decisions (e.g., unit commitment policy, i.e., the operational management of the system). This is a key part of our activity for 2014.
- Bandit problems with risk [36]. Bandit problems are quite related to metagaming problems (they correspond to the unstructured case).
- A theoretical work on the consistency of Monte Carlo Tree Search / Upper Confidence Tree in continuous domains [27]. A non-trivial extension was necessary for proving such a consistency.
- Noisy optimization is a key part of our work [61], as it is crucial for direct policy search or more generally for dynamic optimization:
  - We have proven lower bounds under "locality assumptions" [33], which are usually informally assumed by some practicionners for justifying the use of evolutionary algorithms.
  - In cases with strong noise (variance not decreasing to zero around the optimum) we proved log-log convergence for simple rules for choosing the number of resamplings [23].
- Several submissions are joint works with Ailab, National Dong Hwa University, Hualien, Taiwan. The drafts can be found at http://www.lri.fr/~teytaud/indema.html.
- 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 for the choice of the variable values during the exploration of the constraint tree. The most critical issue lied in the definition of a meaningful reward for a given node (variable = value) that could cope with the multiple restarts of the search: the deeper the failure, the larger the reward (and hence this work also pertains to the CRI-SIG(Section 6.4). Initial results have been obtained with job-shop scheduling problems [47] and more extensive results have been obtained on 3 benchmarks of the CP community [46].

6.3. Distributed Systems

Participants: Cécile Germain-Renaud [correspondent], Philippe Caillou, Dawei Feng, Cyril Furtlehner, Victorin Martin, 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 large scale traffic management.

Fault management. As Lamport formulated decades ago, fault management in distributed systems exemplifies the unreachability of exact prior knowledge. Real-world large scale systems additionally face the non-stationarity issue.

[20] models the system state and its ruptures (non-stationarity) through the flow of jobs as a stream (scalability), with a traceability goal (interpretability), and addresses a key difficulty in Data Streaming, which is timely detection of a change in the generative process underlying the data stream *drift*. A statistical model based on spatial distance and time frequency is proposed, together with adaptive thresholding. Theoretical and experimental validation show the robustness of the method.

D. Feng's PhD 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 [13]. Monitoring large scale distributed systems differ from CP's usual applications (personalized recommendation), in two major ways. On the brighter side, while users cannot be queried for specific recommendations, probes can be launched at will. On the downside, firstly the distribution of the probe results is highly skewed, faults being a small fraction of the total population, and secondly, some of the faults are transient. Amongst the numerous approaches addressing CP, Minimum Margin Matrix Factorization (MMMF) is easily amenable to active learning, which addresses fault sparsity both at spatial (skewed distributions) and temporal (transients) level. From extensive experiments on real-world data, we have shown that modelling probe-based fault prediction as a CP task and addressing this task through MMMF is an extremely efficient strategy for fault prediction. Comparative analysis and experiments motivate the critical advantage of active learning. It offers a scalable alternative to direct AUC optimization. Similarly, comparison with bias-aware methods (Mixed Membership Matrix Factorization) indicates that the capacity of actively selecting the most informative probes provides the most efficient method to capture the time variability of the system.

Multi-agent and games. Resuming earlier work, our goal is to provide an automated abstract description of simulation results. Data mining methods are used to identify groups in complex simulations [11]. Using activity indicators to identify the most interesting agent groups [17], the groups and their evolution are described through one or several simulations [10]. To facilitate the dissemination of the algorithms, we have participated in the development of a generic multiagent platform (GAMA), in collaboration with IRD, University of Rouen (IDEES), and University of Toulouse (IRIT) [34], [35].

A statistical physics perspective. With motivating applications in large scale traffic congestion inference problems, we have

- Settled a method for encoding real data with pairwise dependencies into an Ising copula [58] suitable to infer real-valued data (travel time) from the computation of its corresponding latent binary state (congested/not congested) probabilities.
- In parallel we have investigated the inverse Ising problem, proposing among other methods a loop analysis based on a duality transformation, leading to a dual belief propagation algorithm runing on the dual graph formed by the network of independent cycles. This aims at finding an MRF to represent pairwise correlated data, close to a dependence tree, able to take into account most important loops [14], [15].

6.4. Designing Criteria

Participants: Jamal Atif, Yoann Isaac, Mostepha Redouane Khouadjia, Hélène Paugam-Moisy, Marc Schoenauer, Michèle Sebag.

This recently created 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.

Image understanding. We continued our effort on the development of model-based image understanding approaches (e.g. [71]). In [18], we have proposed a method for simultaneously segmenting and recognizing objects in images, based on a structural representation of the scene and on a constraint propagation method. Theoretical guaranties have been provided along with a quantitative assessment on healthy and pathological brain structures in magnetic resonance images. Within the ANR project LOGIMA (collaboration with ECP, Telecom PariTech and TU Dresden), our goal is to introduce a new lattice-based representation and reasoning framework suited for dealing with spatial objects in the presence of uncertainty. This framework associates under the aegis of general lattice theory ingredients from mathematical morphology, description logics and formal concept analysis. A first development of this framework can be found in [7] where it has been exploited for the definition of abductive reasoning services and applied to high-level image understanding. Several theoretical issues have been raised in the development of this new framework. Some of them were tackled in [25], [24], [30]. In [25], we have shown how mathematical morphology operators can ben defined on general concept lattices. Explicit join-commuting and meet-commuting operators are defined either from particular valuations on the corresponding lattice or from the decomposition of their elements. In [24], we extended our mathematical morphology based adductive reasoning to multivalued logics, hence allowing us to deal with several uncertainty and imprecision phenomena. In [30], metrics between bipolar information - where the information is represented by a positive/preference part and a negative/constraint part - have been introduced based on particular dilations.

Structured learning. With motivations in bio-informatics and brain computer interfaces, the goal is to take into account priors about the spatio-temporal structure of the underlying phenomenon in order to propose a generative model of the data.

In the context of Yoann Isaac's PhD (Digiteo Unsupervised Brain project), in collaboration with CEA LIST, the goal is to design a representation endowed with appropriate invariance properties. Specifically, within the framework of sparse dictionary coding, we have introduced new priors allowing us to capture both spatial and temporal regularity of multivariate brainwave signals [54]. The learning/optimization criterion, while being multivariate, contains several non-differentiable terms, raising new optimization issues; the proposed approach extends the classical split Bregman iterations algorithm to the multi-dimensional case with several non-differentiable terms [37].

In the context of regulatory gene networks, the challenge is to combine probabilistic inference (does a gene regulate another one) with relational learning (the set of genes is organized in a network). Ensemble learning approaches have been used to cope with the imbalanced nature of the data, e.g., bagging Markov logic networks or boosting operator-valued kernel-based regressors [55], [64]. Another issue, regarding the indeterminacy of the models due to the data sparsity, is addressed through prior-guided regularization beyond model sparsity such as orthogonality [8] or stability [16].

In the domain of medical imaging, the exploitation of computational tomography data to model tumor physiology is hindered by the huge noise level; the multi-task setting is leveraged to provide a better robustness to noise [51].

Robotic value systems. Within the European SYMBRION IP, investigations on the preference-based reinforcement learning were continued 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 [67], this work had been extended and combined with active learning to yield state-of-the art performances with few binary feedbacks from the expert. The work has first concentrated this year on handling the noise due to expert's mistakes [53], and bridging the reality gap when porting the algorithms on real robots (e-pucks and one Nao robot) – these results will be published in Riad Akrour's PhD dissertation, to be defended in Spring 2014.

Algorithm Selection as Collaborative Filtering. The crucial issue when addressing algorithm selection problems is to be able to come up with features that can describe the problems: with representative features, algorithm selection amounts to supervised learning. However, except for some rare domaines (e.g., SAT, [73]), no satisfying set of features exists. However, algorithm selection can also be viewed as a recommendation problem, and tackled by collaborative filtering: users more or less like movies, and similarly, instances like algorithms as much as these algorithms are efficient in solving it. Applying collaborative filtering leads to designing a latent feature space in which the representation of the problems is highly adapted to the algorithm selection problem. A critical issue in collaborative filtering is the 'cold start' problem, that is making recommendations for a brand new user/problem instance. This issue has been handled by a surrogate model of the latent factors, mapping the initial features onto the latent ones. The *Algorithm Recommender System* has been successfully applied to 3 different domains: Satisfiability, Constraint Programming, and Continuous Black-Box Optimization (data from the COCO platform, see Section 5.4) [59].

Social Networks with insider information. The analysis of social networks based on the contents and structure of information exchanges most often pertains to descriptive learning, e.g., explaining the growth of the communication graph or investigating the sensitivity of existing algorithms to hyper-parameters [31]. In the Modyrum context (coll. SME Augure), a supervised learning perspective is investigated, taking advantage of the fact that experts already know part of the sought results in some specific domains of interest. Based on e.g., Twitter and blogs data, the goal is to define generic features and supervised learning algorithms, enabling to characterize the targets of interest depending on the public relation focus.

Multi-objective AI Planning. Within the ANR project DESCARWIN (http://descarwin.lri.fr), Mostepha-Redouane Kouadjia continued his work on the **multi-objective approach** to AI Planning using the Evolutionary Planner *Divide-and-Evolve* (DaE), that evolves a sequential decomposition of the problem at hand: each sub-problem is then solved in turn by some embedded classical planner [70]. 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 [39], 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 [41], and comparing the evolutionary multi-objective approach with the aggregation method, the only method known to the AI Planning community [38]. A sum-up of these recent results have been published at IJCAI [40].

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

- Thalès Air Systems (corr. Areski Hadjaz), related to Gaétan Marceau-Caron's CIFRE PhD, May 2011 - May 2014, 15kEuros per year.
- Modyrum (*Modélisation Dynamique d'un Réseau Médiatique*, related to Marco Bressan's postdoc), SME Augure, started Feb. 2013, 150kEuros.

  Participants: Marco Bressan, Cyril Furtlehner, Michèle Sebag.
- I-Lab METIS (*A general framework for decision making with uncertainty plus energy-specific applications*, related to Jérémie Decock's PhD, ARTELYS-Inria, Sept.2011 - Aug.2014, 40kEuros.

  Participants: Jérémie Decock, Jean-Joseph Christophe, Olivier Teytaud.

8. Partnerships and Cooperations

8.1. Regional Initiatives

- **TIMCO** – 2012-2015 (432kEuros for TAO). Technology for In Memory Computing. A NUMA architecture with very large memory. FUI project, Coodinator Bull SA, WP *Algorithm adaptation: the paradigm shif* coordinator Cécile Germain.

8.2. National Initiatives

- **ASAP** – 2009-2013 (178 kEuros). *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 kEuros). Including Ontologies in Monte-Carlo Tree Search and Applications, ANR international project coordinated by O. Teytaud (TAO, Inria).
  Participants: Adrien Couëtoux, O. Teytaud.
- **DESCARWIN** – 2010-2013 (201 kEuros). Coordinateur P. Savéant, Thalès R&T.
  Participants: Mostepha-Redouane Khouadjia, Marc Schoenauer.
- **SIMINOLE** – 2010-2014 (1180kEuros, 250kEuros 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-Renaud
- **NUMBBO** – 2012-2016 (290kEuros for TAO). Analysis, Improvement and Evaluation of Numerical Blackbox Optimizers, ANR project, Coordinator Anne Auger, Inria. Other partners: Dolphin, Inria Lille, Ecole des Mines de Saint-Etienne, TU Dortmund
  Participants: Anne Auger, Nikolaus Hansen, Marc Schoenauer, Ouassim Ait ElHara
- **LOGIMA** – 2012-2016 (136kEuros 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

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

8.3. European Initiatives

8.3.1. FP7 Projects

- **SYMBRION**

  Type: COOPERATION (Integrated Project)

  Program: Embedded systems design

  Instrument: Integrated Project

  Objective: FET proactive: Pervasive adaptation

  Duration: February 2008 - July 2013

  Coordinator: Sergey Kornienko and Paul Levi, Stuttgart University (Germany).

  Partners: Universität Stuttgart (USTUTT), Universität Graz (IZG), Vrije Universiteit (VU), Universität Karlsruhe (UNIKARL), Flanders Institute for Biotechnology (VIB), University of the West of England, Bristol (UWE), Eberhard Karls Universität Tübingen (UT), University of York (UY), Université Libre de Bruxelles (CENOLI), and Inria-TAO.

  Inria contact: M. Schoenauer

  Abstract: SYMBRION, an FP7 IP (Integrated Project), involving 10 partners from Robotics (Electronics and Mechanics), Evolutionary Biology, and Computer Science (working on bio-inspired complex systems). Integrating hardware and software design, Symbrion IP aims at designing autonomous swarm robots. The software will involve both time-scales of evolutionary learning and on-line learning, in direct connection with TAO research themes.
- **CitInES**

  Type: COOPERATION (STREP)

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

Instrument: Specific Targeted Research Project

Objective: ICT systems for energy efficiency

Duration: October 2011 - March 2014

Coordinator: Artelys (SME)

Other Partners: AIT (Austria), INESC Porto (Portugal), ARMINES (France), Schneider Electric SAS (France), Comune di Cesena (Italy), Comune di Bologna (Italy), TUPRAS (Turkey), ERVET (Italy)

Inria contact: Olivier Teytaud

Abstract: The overall objective of CitInES is to design and demonstrate a multi-scale multi-energy decision-making tool to optimise the energy efficiency of cities or large industrial complexes by enabling them to define sustainable, reliable and cost-effective long-term energy strategies. Demonstrations will take place in two cities in Italy, Cesena and Bologna, and in one oil refinery in Turkey, Tupras. Innovative energy system modelling and optimization algorithms will be designed to allow end-users to optimize their energy strategy through detailed simulations of local energy production, storage, transport, distribution and consumption, including demand side management and coordination functionalities enabled by smart grid technologies. All energy vectors (electricity, gas, heat...), usages (heating, air conditioning, lighting, transportation...) and sectors (residential, industrial, tertiary, urban infrastructure) will be considered to draw a holistic map of the city/industry energy behaviour. Energy strategy analyses will encompass advanced long-term risk analysis. As economic and technical situations are constantly evolving, a relevant energy strategy should be robust to different prospective scenarios. Hence, a diversified energy portfolio will allow city and industry authorities to react more efficiently to fuel price stresses and to decrease their exposition to a given energy solution. The expected impacts on end-users are threefold : 1) to assess the economic and environmental impacts of urban planning scenarios in terms of energy; 2) to optimise their local energy strategy to cost-effectively reduce CO2 emissions, including usage of local renewable energies, electric mobility integration, multi-energy coordination, smart grid integration and demand-side management; and 3) to assess financial and environmental long-term risks and propose robust energy schemes to face fuel and CO2 price uncertainties. The developed software will also be used as a communication tool for end-users to facilitate consultations between actors and to promote local authority decisions towards citizens. CitInES methodology will be demonstrated by optimizing long-term energy strategies for the two partner cities and for the partner oil refinery. The proposed strategies will be assessed and compared to initial end-user strategies to measure energy and CO2 emission savings.

- **EGI**

Program: Collaborative Project and Coordination and Support Action (CP-CSA)

Project acronym: EGI-Inspire

Project title: European Grid Infrastructures

Duration: May 2010 - April 2014

Coordinator: Steven Newhouse EGI.eu

Other Partners: 40 in Europe and 8 more worldwide (details on http://www.egi.eu)

Inria contact: Cécile Germain

Abstract: Collaborative effort involving more than 50 institutions in over 40 countries. Its mission is to establish a sustainable European Grid Infrastructure (EGI). EGI-InSPIRE is ideally placed to join together the new Distributed Computing Infrastructures (DCIs) such as clouds, supercomputing networks and desktop grids, for the benefit of user communities within the European Research Area.

- **Network of Excellence PASCAL**

  Type: COOPERATION (FP7)

  Program: Pattern Analysis, Statistical Modelling and Computational Learning

  Instrument:

  Objective: PASCAL is a Network of Excellence funded by the European Union. It has established a distributed institute that brings together researchers and students across Europe, and is now reaching out to countries all over the world.

  Duration: March 2008 - July 2013

  Coordinator: John Shawe-Taylor, (Scientific coordinator), University College London, UK and Steve Gunn (Operational), University of Southampton, UK

  Other Partners:

  Inria contact: Michèle Sebag

  Abstract: PASCAL is developing the expertise and scientific results that will help create new technologies such as intelligent interfaces and adaptive cognitive systems. To achieve this, it supports and encourages collaboration between experts in Machine Learning, Statistics and Optimization. It also promotes the use of Machine Learning in many relevant application domains such as Machine vision, Speech, Haptics, Brain-Computer Interface, User-modeling for computer human interaction, Multimodal integration, Natural Language Processing, Information Retrieval, Textual Information Access.

- **MASH**

  Program: Investigation of the design of complex learning systems to increase the performance of artificial intelligence

  Project acronym: MASH

  Project title: Massive Sets of Heuristics

  Duration: October 2010 - June 2013

  Coordinator: Idiap Research Institute (Martigny, Switzerland)

  Other Partners: Heudiasyc laboratory (CNRS and UTC, Compiègne, France), University of Potsdam (Germany), Center for Machine Perception of the Czech Technical University, Pragua.

  Inria contact: Olivier Teytaud

  Abstract: The goal of the MASH project is to create new tools for the collaborative development of large families of feature extractors. It aims at starting a new generation of learning software with great prior model complexity. The project is structured around this web platform. It comprises collaborative tools, such as a wiki-based documentation and a forum, and an experiment center to run and analyze experiments continuously. The applications targeted by the project are classical vision problems, and goal-planning in a 3D video game and with a real robotic arm. The scientific issues to be tackled along the course of the project are numerous, from standard Machine Learning questions such as learning and prediction with very large feature spaces and tight computational constraints, to original problems related to clustering in a functional space.

8.3.2. Collaborations in European Programs, except FP7

Program: COST

Project acronym:IC0804

Project title: Energy efficiency in large scale distributed systems

Duration: January 2009 - May 2013

Coordinator:Jean-Marc Pierson IRIT

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

Abstract: The COST Action IC0804 proposes realistic energy-efficient alternate solutions to share IT distributed resources. While much effort is nowadays put into hardware specific solutions to lower energy consumptions, a complementary approach is necessary at the distributed system level, i.e., middleware, network and applications. The Action characterizes the energy consumption and energy efficiencies of these components.

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: National University of Tainan (Taiwan)

Duration: 2012 - 2014

See also: http://www.lri.fr/~teytaud/taiwanday2.html

The objectives of the project are threefolds:

- Objective 1: Designing consistent iterative realistic algorithms for partially observable 1-player or 2-player games. We mean:
  d • 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.
- 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.
- Objective 3: Big industrial applications. Having both mathematics and visible realizations 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 ?eld of energy generation, we have seen that many problems are simpli?ed 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 ?eld 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.

Some continuously updated and more detailed descriptions of several works in progress can be found at http://www.lri.fr/~teytaud/indema.html.

8.4.2. Inria International Partners

- On-going collaboration with Christian Schulte (KTH, Stockholm), one of the main developers of the GECODE Constraint Programming platform (see Section 6.2).
- Shinshu University, Faculty of Engineering, project *Global Research on the Framework of Evolutionary Solution Search to Accelerate Innovation*, from the "Strategic Young Researcher Overseas Visits Program for Accelerating Brain Circulation" program, in which TAO and DOLPHIN (Inria Lille) are partner labs and will host Japanese students in the forthcoming 4 years.

8.4.3. Inria International Labs

Olivier Teytaud, 10 days in Inria Chile: meetings with several companies and institutes. They were followed by video-conferences with Endesa and email discussions between our partner Artelys and Cedec-Sing.

8.5. International Research Visitors

8.5.1. Visits of International Scientists

- Visit of a Taiwanese delegation, about power systems and E-learning mainly; more information at http://www.lri.fr/~teytaud/france2013.html. Contact: O. Teytaud.
- Visit of a delegation of Shinshu University, Faculty of Engineering, inuding Dr. Miura (University Trustee and VP), Dr. Tanaka, coordinator of the *Global Research on the Framework of Evolutionary Solution Search to Accelerate Innovation* project (see above).

8.5.2. Visits to International Teams

- **Olivier Teytaud**, invited research fellow, National Dong Hwa University, 4.5 months, 2013.

9. Dissemination

9.1. Scientific Animation

9.1.1. Management positions in scientific organisations

International

- THRaSH, *Theory of Randomized Search Heuristics workshop*: Anne Auger, member of Steering Committee
- ACM SIGEVO (Special Interest Group on Evolutionary Computation), Marc Schoenauer, member of Executive Board since 2000, officer (Secretary) since 2012; member of Business Committe (2011-2013).

- Parallel Problem Solving from Nature: Marc Schoenauer, Member of Steering Committee, (since 1998).
- PASCAL NoE (Pattern Analysis, Statistical modelling, Computational Learning), Michèle Sebag, member of the Steering Committee (PASCAL 2004-2008; PASCAL2, 2009-2013).
- European Machine Learning and Knowledge Discovery from Databases Steering Committee, MichèSebag, member since 2010;
- ECCAI Fellow, Michè Sebag, since 2011;
- Marc Schoenauer, Honorary Adjunct Professor, School of Computer Science, University of Adelaide, Australia (2009-2015).

National

- Michè Sebag, member of the CoNRS; Senior Advisory Board CHIST-ERA; member of the CSFRS (Conseil Supérieur de la Formation et Recherche Stratégique); member of the Senate at Université Paris-Saclay; responsible for the DataSense axis in the DigiCosme Labex;
- EA – Association Evolution Artificielle: Marc Schoenauer, founding president, now member of Advisory Committee. Anne Auger, member of Executive Commitee since 2008.
- SimTools Network (RNSC Network). Philippe Caillou, coordinator since 2011.

Université Paris-Sud

- Jamal Atif, "Directeur d'études" at Computer Science department of IUT d'Orsay ; membre de la CCSU 27 (membre du Bureau) since 2012; membre élu au conseil d'Institut, IUT d'Orsay ; membre du Bureau du département Informatique de l'IUT d'Orsay since 2011.
- Anne Auger, membre du Conseil du Laboratoire de Recherche en Informatique since 2012;
- Philippe Caillou, membre élu du Conseil Scientifique de l'université since 2013, directeur des études à l'IUT de Sceaux since 2009
- Cécile Germain, membre élu du Conseil Scientifique de l'université since 2012, membre du Bureau, chargée de mission à l'Informatique Scientifique.
- Michèle Sebag, membre élu du Conseil du Laboratoire de Recherche en Informatique et membre de la CCSU 27 since 2004.
- Olivier Teytaud, représentant des B pour le comité d'évaluation du LRI.

Inria Saclay

- Anne Auger, membre de la Commission de Suivi Doctoral ; représentante du centre de Saclay à la Commission des Jeunes Chercheurs.
- Marc Schoenauer, Délégué Scientifique (aka VP Research) since 2010.
- Olivier Teytaud, TAO representative at CUMI since 2008.

9.1.2. Organisation of Conferences and Scientific Events

- *BBOB Black-Box Optimization Benchmarking* workshop at the ACM GECCO Genetic and Evolutionary Computation Conference, 2013, Anne Auger and Nikolaus Hansen, co-organizers.
- Dagstuhl Seminar 13271 *Theory of Evolutionary Algorithms*, 2013: Nikolaus Hansen, co-organizer.
- ACM-GECCO, ES-EP track, Anne Auger co-chair.
- Franco-Taiwanese 2013 meeting on AI, E-learning and power systems: https://www.lri.fr/~teytaud/france2013.html

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

- L1 (IUT Orsay, Univ. Paris-Sud), Computer Science, Jamal Atif, approx. 192h since 2010.
- L1-3 (Ecole Centrale Paris), Stochastic Optimization, Anne Auger, 20h since 2011.
- L1 (IUT Sceaux, Univ. Paris-Sud), Computer science for Managers, Philippe Caillou, approx. 192h since 2011.
- L2, L3 (Polytech, Univ. Paris-Sud), Computer Architecture, Cécile Germain-Renaud, head of Licence, approx. 120h since 2009.
- L2 Univ Paris-Sud, Vie Artificielle, Michèle Sebag, 10h in 2013.
- L3 (ENS-Cachan), Introduction to Machine Learning, Michèle Sebag, approx. 24h since 2011.
- M1 Computer Science (U. Paris-Sud), Parallelisme, Cécile Germain-Renaud, approx. 50h since 2009.
- M1 Computer Science MPRI, Machine Learning, Michèle Sebag, 25h since 2013.
- Master 2 Recherche (U. Paris-sud), Optimisation, Anne Auger, 12h since 2011.
- Master 2 Recherche (U. Paris-Sud), Machine Learning, Michèle Sebag, 24h since 2011.
- Master 2 Recherche (U. Paris-Sud), Evolutionary Robotics, MichèSebag, 15h since 2013.
- Master 2 Recherche (U. Paris-Sud), Multi-Agents Systems, Philippe Caillou, 27h since 2011.
- Master 2 Recherche Paris-Dauphine, Multi-Agent Based Simulation, Philippe Caillou, 3h since 2011.

9.2.2. Supervision

- HdR: Jamal Atif, *Quelques contributions à l'interprétation d'images, à l'apprentissage statistique et à la cartographie cérébrale*, Université Paris-Sud, 31/10/2013
- PhD: Ludovic Arnold, *Learning Deep Representations: toward a new understanding of the deep learning paradigm*, 25/06/2013, Hélène Pagam-Moisy et Philippe Tarroux [1]
- PhD: Jean-Baptiste Hoock, *Contributions to Simulation-based High-dimensional Sequential Decision Making*, 10/4/2013, Olivier Teytaud [3]
- PhD: Ilya Loshchilov, *Surrogate-Assisted Evolutionary Algorithms*, 6/1/2013 [4]
- PhD: Jean-Marc Montanier, *Environment-driven Distributed Evolutionary Adaptation for Collective Robotic Systems*, 1/3/2013, Nicolas Bredèche [6]
- PhD: Victorin Martin, *Modélisation Probabiliste et Inférence par l'Algorithme Belief Propagation*, 23/5/2013, Cyril Furtlehner and Jean-Marc Lasgouttes [5]
- PhD: Adrien Couétoux, *Monte Carlo Tree Search (MCTS) for Continuous and Stochastic Sequential Decision Making Problems*, 30/9/2013, Olivier Teytaud [2]
- PhD in progress: Ouasssim Ait Elhara, *Stochastic Black-Box Optimization and Benchmarking in Large Dimension*, 1/10/2012, Anne Auger and Nikolaus Hansen

- PhD in progress: Riad Akrour, *Preference Based Reinforcement Learning*, 1/11/2010, Marc Schoenauer and Michèle Sebag
- Sandra Cecilia Astete Morales, *Random Processes for Optimization with Risk*, 1/9/2013, Olivier Teytaud
- PhD in progress: Asma Atamna, *Analysis, Improvement and Benchmarking of Constraint Handling for Stochastic Blackbox Continuous Optimization*, 1/11/2013, Anne Auger and Nikolaus Hansen
- PhD in progress: Jérémy Bensadon, *Information theory for learning and optimization*, 1/10/2012, Yann Ollivier
- PhD in progress: Vincent Berthier, *Unsupervised Learning and Brain Wave Data*, 1/10/2013, Jamal Atif and Michèle Sebag
- PhD in progress: Marie-Liesse Cauwet, *Noisy Optimization for Artificial Intelligence*, 1/9/2013, Olivier Teytaud
- PhD in progress: Alexandre Chotard, *Enhancement and Analysis of Evolution Strategies*, 01.10.2011, Anne Auger and Nikolaus Hansen
- PhD in progress: Jérémie Decock, *Large Scale Constrained Direct Policy Search and Applications to Power Systems*, 3/10/2011, Olivier Teytaud
- PhD in progress: Dawei Feng, *Détection et diagnostic d'anomalies dans les systèmes répartis à grande échelle*, 1/10/2010, Germain-Renaud
- PhD in progress: Nicolas Galichet, *Risk-Aware Reinforcement Learning*, 1/10/2011, Michl`e Sebag
- PhD in progress: Yoann Isaac, *Une approche non-supervisée pour la passage à l'échelle des interfaces cerveau-machine*, 1/10/2011, Jamal Atif and Michl`e Sebag
- PhD in progress: Jialin Liu, *Portfolios of Noisy Optimization Algorithms*, 14/03/2013, Marc Schoenauer and Olivier Teytaud
- PhD in progress: Gaetan Marceau Caron, *Global Multi-objective Optimization in Air Traffic Control*, 14/12/2010, Marc Schoenauer
- PhD in progress: Weija Wang, *Multi-Objective Reinforcement Learning*, 1/10/2010, Michèle Sebag
- PhD in progress: Yifan Yang, *Jamal*, 1/9/2013, Jamal Atif
- PhD in progress: Guohua Zhang, *Curiosity-Driven Navigation in Evolutionary Robotics*, 1/9/2011, Michèle Sebag

9.2.3. Juries

- Marc Schoenauer, reviewer for Pawan Kumar Mudigonda's HDR, Ecole Normale Supérieure de Cachan, Nov. 2013
- Marc Schoenauer, external reviewer for Martin Pilat's PhD, Charles University, Prague, Czech Republic, July 2013 ; external reviewer for Tianjun Liao's PhD, Ecole Polytechnique de l'Université Libre de Bruxelles, Belgian, June 2013; jury member for Una Blenic's PhD at Université d'Angers, France, April 2013; external jury member for José Manuel Garcia-Nieto's PhD, University of Malaga, Spain, January 2013.
- Michèle Sebag, reviewer for Yuan Yang, Telecom ParisTech, June 2013; Emilie Morvant, Université Aix Marseille, June 2013; Floarea Serban, University of Zurich, Switzerland, 2013; Darko Cerepnalkoski, Josef Stefan Institute, Slovenia, Sept. 2013.
- Marc Schoenauer, reviewer for Irish Science Foundation, Project *Multi-core Attribute Grammatical Evolution* at University Limerick, July 2013
- Michèle Sebag, jury de recrutement CR Inria Lille, April 2013; jury de recrutement TelecomParisTech, June 2013; jury de promotion TelecomParisTech, Oct. 2013.
- Jamal Atif, Michèle Sebag, Olivier Teytaud, members of juries for MDC and PR hiring at Université Paris-Sud.

- Michèle Sebag, jury AAP Digiteo, March 2013.

9.3. Popularization

- Michèle Sebag is interviewed at the *Palais de la Découverte* about Artificial Intelligence, and the video, by "Société de production Stand Alone Media", is visible on YouTube at http://www.youtube.com/watch?v=uEW32KikKJ8; talk in "Questions de science et enjeux citoyens" (QSEC), opération culturelle de la région Île-de-France (Ulm, May 2013);
- Yann Ollivier co-organized a bi-monthly math seminar for undergrad students on Saturdays at Institut Henri Poincare, 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 was part of the scientific steering committee for the booklet *L'explosion des mathématiques* presenting a wide range of applications of mathematics, edited by the SMF and SMAI (planned diffusion: 10,000–20,000 copies).

10. Bibliography

Publications of the year

Doctoral Dissertations and Habilitation Theses

[1] L. ARNOLD. , *Learning Deep Representations : Toward a better new understanding of the deep learning paradigm*, Université Paris Sud - Paris XI, June 2013, http://hal.inria.fr/tel-00842447

[2] A. COUETOUX. , *Monte Carlo Tree Search pour les problèmes de décision séquentielle en milieu continus et stochastiques*, Université Paris Sud - Paris XI, September 2013, http://hal.inria.fr/tel-00927252

[3] J.-B. HOOCK. , *Contributions to Simulation-based High-dimensional Sequential Decision Making*, Université Paris Sud - Paris XI, April 2013, http://hal.inria.fr/tel-00912338

[4] I. LOSHCHILOV. , *Surrogate-Assisted Evolutionary Algorithms*, Université Paris Sud - Paris XI and Institut national de recherche en informatique et en automatique - Inria, January 2013, http://hal.inria.fr/tel-00823882

[5] V. MARTIN. , *Modélisation probabiliste et inférence par l'algorithme Belief Propagation*, Ecole Nationale Supérieure des Mines de Paris, May 2013, http://hal.inria.fr/tel-00867693

[6] J.-M. MONTANIER. , *Environment-driven Distributed Evolutionary Adaptation for Collective Robotic Systems*, Université Paris Sud - Paris XI, March 2013, http://hal.inria.fr/tel-00811496

Articles in International Peer-Reviewed Journals

[7] J. ATIF, C. HUDELOT, I. BLOCH. *Explanatory reasoning for image understanding using formal concept analysis and description logics*, in "IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans: Systems and Humans", October 2013, pp. 1–19 [*DOI :* 10.1109/TSMC.2013.2280440], http://hal.inria.fr/hal-00862563

[8] N. BRUNEL, Q. CLAIRON, F. D'ALCHÉ-BUC. *Parametric Estimation of Ordinary Differential Equations with Orthogonality Conditions*, in "Journal of American Statistics Association", 2013, (to appear), http://hal.inria.fr/hal-00867370

[9] R. BUSA-FEKETE, B. KÉGL, T. ELTETO, G. SZARVAS. *Tune and mix: learning to rank using ensembles of calibrated multi-class classifiers*, in "Machine Learning", 2013, vol. 93, pp. 261-292, ANR-2010-COSI-002 [*DOI :* 10.1007/S10994-013-5360-9], http://hal.inria.fr/in2p3-00869803

[10] P. CAILLOU, J. GIL-QUIJANO. *Description automatique de dynamiques de groupes dans des simulations à base d'agents*, in "Revue d'Intelligence Artificielle", January 2014, vol. 27, n^{o} 6, http://hal.inria.fr/hal-00927587

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

[12] C.-W. CHOU, P.-C. CHOU, J.-J. CHRISTOPHE, A. COUETOUX, P. DE FREMINVILLE, N. GALICHET, C.-S. LEE, J. LIU, D. L. ST-PIERRE, M. SEBAG, O. TEYTAUD, M.-H. WANG, L.-W. WU, S.-J. YEN. *Strategic Choices in Optimization*, in "Journal of Information Sciences and Engineering", 2013, http://hal.inria.fr/hal-00863577

[13] D. FENG, C. GERMAIN-RENAUD, T. GLATARD. *Efficient Distributed Monitoring with Active Collaborative Prediction*, in "Future Generation Computer Systems", 2013, vol. 29, n^{o} 8, pp. 2272-2283 [*DOI :* 10.1016/J.FUTURE.2013.06.001], http://hal.inria.fr/hal-00784038

[14] C. FURTLEHNER. *Approximate Inverse Ising models close to a Bethe Reference Point*, in "Journal of Statistical Mechanics: Theory and Experiment", September 2013, P09020 p. [*DOI :* 10.1088/1742-5468/2013/09/P09020], http://hal.inria.fr/hal-00865085

[15] C. FURTLEHNER. *Pairwise MRF Models Selection for Traffic Inference*, in "Interdisciplinary Information Sciences", August 2013, vol. 19, n^{o} 1, pp. 17-22 [*DOI :* 10.4036/IIS.2013.17], http://hal.inria.fr/hal-00865089

[16] G. MICHAILIDIS, F. D'ALCHÉ-BUC. *Autoregressive models for gene regulatory network inference: Sparsity, stability and causality issues*, in "Mathematical Biosciences", December 2013, vol. 246, n^{o} 2, pp. 326–334 [*DOI :* 10.1016/J.MBS.2013.10.003], http://hal.inria.fr/hal-00909809

[17] 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", February 2013, vol. 89, n^{o} 2, pp. 156-177 [*DOI :* 10.1177/0037549712457852], http://hal.inria.fr/hal-00738218

[18] O. NEMPONT, J. ATIF, I. BLOCH. *A constraint propagation approach to structural model based image segmentation and recognition*, in "Information Sciences", October 2013, vol. 246, pp. 1-27, http://hal.inria.fr/hal-00862455

[19] W. WANG, M. SEBAG. *Hypervolume indicator and dominance reward based multi-objective Monte-Carlo Tree Search*, in "Machine Learning", May 2013, vol. 92, n^{o} 2-3, pp. 403-429 [*DOI :* 10.1007/S10994-013-5369-0], http://hal.inria.fr/hal-00852048

[20] X. ZHANG, C. FURTLEHNER, C. GERMAIN-RENAUD, M. SEBAG. *Data Stream Clustering with Affinity Propagation*, in "IEEE Transactions on Knowledge and Data Engineering", 2014, http://hal.inria.fr/hal-00862941

International Conferences with Proceedings

[21] O. AIT ELHARA, A. AUGER, N. HANSEN. *A Median Success Rule for Non-Elitist Evolution Strategies : Study of Feasibility*, in "Genetic and Evolutionary Computation Conference", Amsterdam, Netherlands, C. BLUM, E. ALBA (editors), ACM Press, March 2013, pp. 415-422, http://hal.inria.fr/hal-00801414

[22] Y. AKIMOTO, Y. OLLIVIER. *Objective Improvement in Information-Geometric Optimization*, in "Foundations of Genetic Algorithms XII", Adelaide, Australia, January 2013, http://hal.inria.fr/hal-00752489

[23] S. ASTETE-MORALES, J. LIU, O. TEYTAUD. *Noisy optimization convergence rates*, in "Proceeding of the fifteenth annual conference companion on Genetic and evolutionary computation conference companion", Amsterdam, Netherlands, ACM, 2013, pp. 223–224 [*DOI :* 10.1145/2464576.2464687], http://hal.inria.fr/hal-00863584

[24] J. ATIF, I. BLOCH, F. DISTEL, C. HUDELOT. *A fuzzy extension of explanatory relations based on mathematical morphology*, in "EUSFLAT 2013", Milano, Italy, September 2013, pp. 244–351, http://hal.inria.fr/hal-00862605

[25] J. ATIF, I. BLOCH, F. DISTEL, C. HUDELOT. *Mathematical morphology operators over concept lattices*, in "ICFCA", Dresden, Germany, Springer, May 2013, vol. LNAI 7880, pp. 28–43, http://hal.inria.fr/hal-00862621

[26] A. AUGER, D. BROCKHOFF, N. HANSEN. *Benchmarking the Local Metamodel CMA-ES on the Noiseless BBOB'2013 Test Bed*, in "GECCO (Companion), workshop on Black-Box Optimization Benchmarking (BBOB'2013)", Amsterdam, Netherlands, July 2013, http://hal.inria.fr/hal-00825840

[27] D. AUGER, A. COUETOUX, O. TEYTAUD. *Continuous Upper Confidence Trees with Polynomial Exploration - Consistency*, in "ECML/PKKD 2013", Prague, Czech Republic, H. BLOCKEEL, K. KERSTING, S. NIJSSEN, F. ŽELEZNÝ (editors), LNCS, Springer Verlag, September 2013, vol. 8188, pp. 194-209, http://hal.inria.fr/hal-00835352

[28] R. BARDENET. *Monte Carlo methods*, in "IN2P3 School of Statistics (SOS2012)", Autrans, France, T. DELEMONTEX, A. LUCOTTE (editors), 2013, vol. 55, ISBN:978-2-7598-1032-1 [*DOI :* 10.1051/EPJCONF/20135502002], http://hal.inria.fr/in2p3-00846142

[29] R. BARDENET, M. BRENDEL, B. KÉGL, M. SEBAG. *Collaborative hyperparameter tuning*, in "30th International Conference on Machine Learning (ICML 2013)", Atlanta, United States, S. DASGUPTA, D. MCALLESTER (editors), 2013, vol. 28, pp. 199-207, http://hal.inria.fr/in2p3-00907381

[30] I. BLOCH, J. ATIF. *Distance to bipolar information from morphological dilation*, in "European Society for Fuzzy Logic and Technology (EUSFLAT)", Milano, Italy, September 2013, pp. 266–273, http://hal.inria.fr/hal-00862603

[31] M. BRESSAN, E. PESERICO, L. PRETTO. *The power of local information in PageRank*, in "International conference on World Wide Web (WWW)", Rio de Janeiro, Brazil, 2013, http://hal.inria.fr/hal-00862816

[32] J. DECOCK, O. TEYTAUD. *Linear Convergence of Evolution Strategies with Derandomized Sampling Beyond Quasi-Convex Functions*, in "EA - 11th Biennal International Conference on Artificial Evolution - 2013", Bordeaux, France, Lecture Notes in Computer Science, Springer, August 2013, http://hal.inria.fr/hal-00907671

[33] J. DECOCK, O. TEYTAUD. *Noisy Optimization Complexity Under Locality Assumption*, in "FOGA - Foundations of Genetic Algorithms XII - 2013", Adelaide, Australia, February 2013, http://hal.inria.fr/hal-00755663

[34] A. DROGOUL, E. AMOUROUX, P. CAILLOU, B. GAUDOU, A. GRIGNARD, N. MARILLEAU, P. TAILLANDIER, M. VAVASEUR, D.-A. VO, J.-D. ZUCKER. *GAMA: A Spatially Explicit, Multi-level, Agent-Based Modeling and Simulation Platform*, in "Practical Applications of Agents and Multi-Agent Systems", Spain, Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2013, vol. 7879, pp. 271-274, http://hal.inria.fr/hal-00834494

[35] A. DROGOUL, E. AMOUROUX, P. CAILLOU, B. GAUDOU, A. GRIGNARD, N. MARILLEAU, P. TAILLANDIER, M. VAVASEUR, D.-A. VO, J.-D. ZUCKER. *GAMA: multi-level and complex environment for agent-based models and simulations (demonstration)*, in "international conference on Autonomous agents and multi-agent systems", United States, 2013, pp. 1361-1362, http://hal.inria.fr/hal-00834498

[36] N. GALICHET, M. SEBAG, O. TEYTAUD. *Exploration vs Exploitation vs Safety: Risk-averse Multi-Armed Bandits*, in "Asian Conference on Machine Learning 2013", Canberra, Australia, C. S. ONG, T. B. HO (editors), Journal of Machine Learning Research : Workshop and Conference Proceedings, November 2013, vol. 29, pp. 245-260, http://hal.inria.fr/hal-00924062

[37] Y. ISAAC, Q. BARTHÉLEMY, C. GOUY-PAILLER, J. ATIF, M. SEBAG. *Multi-dimensional sparse structured signal approximation using split bregman iterations*, in "38th International Conference on Acoustics, Speech, and Signal Processing (ICASSP)", Vancouver, Canada, IEEE, May 2013, pp. 3826-3830, http://hal.inria.fr/hal-00862645

[38] M. R. KHOUADJIA, M. SCHOENAUER, V. VIDAL, J. DRÉO, P. SAVÉANT. *Multi-Objective AI Planning: Comparing Aggregation and Pareto Approaches*, in "EvoCOP – 13th European Conference on Evolutionary Computation in Combinatorial Optimisation", Vienna, Austria, M. MIDDENDORF, C. BLUM (editors), LNCS, Springer Verlag, March 2013, vol. 7832, pp. 202-213, http://hal.inria.fr/hal-00820634

[39] M. R. KHOUADJIA, M. SCHOENAUER, V. VIDAL, J. DRÉO, P. SAVÉANT. *Multi-Objective AI Planning: Evaluating DAE-YAHSP on a Tunable Benchmark*, in "EMO'13 - 7th International Conference on Evolutionary Multi-Criterion Optimization", Sheffield, United Kingdom, R. C. PURSHOUSE, P. J. FLEMING, C. M. FONSECA (editors), LNCS, Springer Verlag, March 2013, vol. 7811, pp. 36-50, http://hal.inria.fr/hal-00750560

[40] M. R. KHOUADJIA, M. SCHOENAUER, V. VIDAL, J. DRÉO, P. SAVÉANT. *Pareto-Based Multiobjective AI Planning*, in "IJCAI 2013", Beijing, China, F. ROSSI (editor), IJCAI/AAAI, August 2013, http://hal.inria.fr/hal-00835003

[41] M. R. KHOUADJIA, M. SCHOENAUER, V. VIDAL, J. DRÉO, P. SAVÉANT. *Quality Measures of Parameter Tuning for Aggregated Multi-Objective Temporal Planning*, in "LION7 - Learning and Intelligent OptimizatioN Conference", Catania, Italy, P. PARDALOS, G. NICOSIA (editors), LNCS, Springer Verlag, March 2013, To appear, http://hal.inria.fr/hal-00820617

[42] B. KÉGL. *Introduction to multivariate discrimination*, in "IN2P3 School of Statistics (SOS2012)", Autrans, France, T. DELEMONTEX, A. LUCOTTE (editors), 2013, vol. 55, ISBN:978-2-7598-1032-1 [*DOI :* 10.1051/EPJCONF/20135502001], http://hal.inria.fr/in2p3-00846125

[43] I. LOSHCHILOV, M. SCHOENAUER, M. SEBAG. *BI-population CMA-ES Algorithms with Surrogate Models and Line Searches*, in "Workshop Proceedings of the (GECCO) Genetic and Evolutionary Computation Conference", Amsterdam, Netherlands, ACM, April 2013, 8 p. , http://hal.inria.fr/hal-00818596

[44] I. LOSHCHILOV, M. SCHOENAUER, M. SEBAG. *Intensive Surrogate Model Exploitation in Self-adaptive Surrogate-assisted CMA-ES (saACM-ES)*, in "Genetic and Evolutionary Computation Conference (GECCO 2013)", Amsterdam, Netherlands, C. BLUM, E. ALBA (editors), April 2013, pp. 439-446, http://hal.inria.fr/hal-00818595

[45] I. LOSHCHILOV, M. SCHOENAUER, M. SEBAG. *KL-based Control of the Learning Schedule for Surrogate Black-Box Optimization*, in "Conférence sur l'Apprentissage Automatique", Lille, France, August 2013, http://hal.inria.fr/hal-00851189

[46] M. LOTH, M. SEBAG, Y. HAMADI, M. SCHOENAUER. *Bandit-based Search for Constraint Programming*, in "International Conference on Principles and Practice of Constraint Programming", Uppsala, Sweden, C. SCHULTE (editor), LNCS, Springer Verlag, September 2013, vol. 8124, pp. 464-480, http://hal.inria.fr/hal-00863451

[47] 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 "Learning And Intelligent Optimization Conference", Catania, Italy, January 2013, http://hal.inria.fr/hal-00863453

[48] G. MARCEAU, P. SAVÉANT, M. SCHOENAUER. *Computational Methods for Probabilistic Inference of Sector Congestion in Air Traffic Management*, in "Interdisciplinary Science for Innovative Air Traffic Management", Toulouse, France, July 2013, http://hal.inria.fr/hal-00862243

[49] G. MARCEAU, P. SAVÉANT, M. SCHOENAUER. *Multiobjective Tactical Planning under Uncertainty for Air Traffic Flow and Capacity Management*, in "IEEE Congress on Evolutionary Computation", Cancun, Mexico, C. A. C. COELLO, Y. JIN (editors), IEEE Press, June 2013, pp. 1548-1555, http://hal.inria.fr/hal-00862223

[50] G. MARCEAU, P. SAVÉANT, M. SCHOENAUER. *Strategic Planning in Air Traffic Control as a Multi-objective Stochastic Optimization Problem*, in "ATM Seminar 2013", Chicago, United States, June 2013, http://hal.inria.fr/hal-00862186

[51] B. ROMAIN, V. LETORT, O. LUCIDARME, L. ROUET, F. D'ALCHÉ-BUC. *A multi-task learning approach for compartmental model parameter estimation in DCE-CT sequences*, in "16th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2013)", Nagoya, Japan, Lecture Notes in Computer Science, September 2013, vol. 8150, pp. 271–278 [*DOI :* 10.1007/978-3-642-40763-5_34], http://hal.inria.fr/hal-00832184

[52] B. SZORENYI, R. BUSA-FEKETE, I. HEGEDÜS, R. ORMANDI, M. JELASITY, B. KÉGL. *Gossip-based distributed stochastic bandit algorithms*, in "30th International Conference on Machine Learning (ICML 2013)", Atlanta, United States, S. DASGUPTA, D. MCALLESTER (editors), 2013, vol. 28, pp. 19-27, http://hal.inria.fr/in2p3-00907406

Conferences without Proceedings

[53] R. AKROUR, M. SCHOENAUER, M. SEBAG. *Interactive Robot Education*, in "ECML/PKDD Workshop on Reinforcement Learning with Generalized Feedback: Beyond Numeric Rewards", Berlin, Germany, J. FUERNKRANZ, E. HÜLLERMEIER (editors), September 2013, http://hal.inria.fr/hal-00931347

[54] Y. ISAAC, Q. BARTHÉLEMY, J. ATIF, C. GOUY-PAILLER, M. SEBAG. *Régularisations spatiales pour la décomposition de signaux EEG sur un dictionnaire temps-fréquence*, in "Colloque Gretsi XXIV", France, September 2013, http://hal.inria.fr/hal-00862707

[55] N. LIM, Y. SENBABAOGLU, G. MICHAILIDIS, F. D'ALCHÉ-BUC. *Boosting an operator-valued kernel model and application to network inference*, in "Workshop on Machine Learning for System Identification", Atlanta, United States, June 2013, http://hal.inria.fr/hal-00844490

Scientific Books (or Scientific Book chapters)

[56] N. HANSEN, A. AUGER. *Principled Design of Continuous Stochastic Search: From Theory to Practice*, in "Theory and Principled Methods for the Design of Metaheuristics", Y. BORENSTEIN, A. MORAGLIO (editors), Natural Computing Series, Springer, 2013, http://hal.inria.fr/hal-00808450

[57] Y. OLLIVIER. *A visual introduction to Riemannian curvatures and some discrete generalizations*, in "Analysis and Geometry of Metric Measure Spaces: Lecture Notes of the 50th Séminaire de Mathématiques Supérieures (SMS), Montréal, 2011", G. DAFNI, R. MCCANN, A. STANCU (editors), AMS, 2013, pp. 197-219, http://hal.inria.fr/hal-00858008

Research Reports

[58] V. MARTIN, J.-M. LASGOUTTES, C. FURTLEHNER. , *Using Latent Binary Variables for Online Reconstruction of Large Scale Systems*, Inria, December 2013, n^{o} RR-8435, 34 p. , http://hal.inria.fr/hal-00922106

[59] M. MISIR, M. SEBAG. , *Algorithm Selection as a Collaborative Filtering Problem*, December 2013, 43 p. , http://hal.inria.fr/hal-00922840

Other Publications

[60] L. ARNOLD, Y. OLLIVIER. , *Layer-wise learning of deep generative models*, 2013, http://hal.inria.fr/hal-00794302

[61] S. ASTETE-MORALES, M.-L. CAUWET, A. COUETOUX, J. DECOCK, J. LIU, O. TEYTAUD. *Noisy Optimization*, in "Dagstuhl seminar 13271", Dagstuhl, Germany, 2013, Dagstuhl seminar 13271, http://hal.inria.fr/hal-00844305

[62] A. AUGER, N. HANSEN. , *Linear Convergence on Positively Homogeneous Functions of a Comparison Based Step-Size Adaptive Randomized Search: the (1+1) ES with Generalized One-fifth Success Rule*, October 2013, http://hal.inria.fr/hal-00877161

[63] A. AUGER, N. HANSEN. , *On Proving Linear Convergence of Comparison-based Step-size Adaptive Randomized Search on Scaling-Invariant Functions via Stability of Markov Chains*, November 2013, http://hal.inria.fr/hal-00877160

[64] N. LIM, F. D'ALCHÉ-BUC, C. AULIAC, G. MICHAILIDIS. , *Operator-valued Kernel-based Vector Autoregressive Models for Network Inference*, October 2013, http://hal.inria.fr/hal-00872342

[65] Y. OLLIVIER. , *Efficient Riemannian training of recurrent neural networks for learning symbolic data sequences*, 2013, Preliminary version. 2nd version: recurrent tensor-square differential metric added, more thorough experiments, title changed, http://hal.inria.fr/hal-00857980

[66] Y. OLLIVIER. , *Riemannian metrics for neural networks*, 2013, (3rd version: tensor square derivative metric discussed), http://hal.inria.fr/hal-00857982

References in notes

[67] R. AKROUR, M. SCHOENAUER, M. SEBAG. *APRIL: Active Preference-learning based Reinforcement Learning*, in "ECML PKDD 2012", PETER FLACH ET AL (editor), LNCS, Springer Verlag, September 2012, vol. 7524, pp. 116-131

[68] 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, pp. 549-553, http://hal.inria.fr/in2p3-00698455

[69] H.-G. BEYER. *Evolution Strategies*, in "Scholarpedia", 2007, vol. 2, n^{o} 8, 1965 p. [*DOI :* 10.4249/SCHOLARPEDIA.1965], http://www.scholarpedia.org/article/Evolution_strategies

[70] 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, pp. 15-25, http://hal.archives-ouvertes.fr/docs/00/45/62/92/PDF/icaps10.pdf

[71] 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", January 2012, vol. 116, n^{o} 1, pp. 146–165, Impact factor: 1,340

[72] N. HANSEN, A. OSTERMEIER. *Completely Derandomized Self-Adaptation in Evolution Strategies*, in "Evolutionary Computation", 2001, vol. 9, n^{o} 2, pp. 159-195

[73] F. HUTTER, Y. HAMADI, H. H. HOOS, K. LEYTON-BROWN. *Performance Prediction and Automated Tuning of Randomized and Parametric Algorithms*, in "Principles and Practice of Constraint Programming (CP'06)", 2006, pp. 213–228