IN PARTNERSHIP WITH:
CNRS

Université Paris-Sud (Paris 11)

Activity Report 2014

Project-Team TAO

Machine Learning and Optimisation

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, Inria, Senior Researcher, HdR]
Anne Auger [Inria, Researcher]
Cyril Furtlehner [Inria, Researcher]
Nikolaus Hansen [Inria, Senior Researcher, HdR]
Balázs Kégl [CNRS, Senior Researcher at LAL, associate member, HdR]
Yann Ollivier [CNRS, Researcher, HdR]
Michèle Sebag [CNRS, Senior Researcher, HdR]
Olivier Teytaud [Inria, Researcher, HdR]

Faculty Members

Jamal Atif [Univ. Paris-Sud, Associate Professor, until Sep 2014, HdR]
Philippe Caillou [Univ. Paris-Sud, Associate Professor]
Florence d'Alché Buc [Univ. Évry, Professor, en demie-délégation until Aug 2014, HdR]
Aurélien Decelle [Univ. Paris-Sud, Associate Professor, from Sep 2014]
Blaise Hanczar [Univ. Paris-Descartes, Associate Professor, en délégation until Sep 2014]
Cécile Germain-Renaud [Univ. Paris-Sud, Professor, HdR]

Engineers

Emmanuel Benazera [Inria, from Apr 2014]
Jean-Joseph Christophe [Inria, until Aug 2014]
Lovro Ilijasic [Inria, granted by Conseil Régional d'Ile-de-France]
Julien Nauroy [Inria, until Sep 2014, granted by Conseil Régional d'Ile-de-France]
Karima Rafes [Inria granted by Conseil Régional d'Ile-de-France]

PhD Students

Ouassim Ait Elhara [Univ. Paris-Sud]
Riad Akrour [Inria, Post-Doc since Oct 2014]
Sandra Cecilia Astete Morales [Inria]
Asma Atamna [Inria]
Nacim Belkhir [Thales, CIFRE, from Apr 2014]
Jeremy Bensadon [Univ. Paris-Sud]
Vincent Berthier [Univ. Paris-Sud]
Antoine Bureau [Inria, from Oct 2014]
Marie-Liesse Cauwet [Inria]
Alexandre Chotard [Univ. Paris-Sud]
Jérémie Decock [Inria, until Nov 2014]
Dawei Feng [Univ. Paris-Sud, until Mar 2014]
Nicolas Galichet [Univ. Paris-Sud]
François Gonard [Inst. de Recherche Technologique SystemX, from Nov 2014]
Moez Hammami [Doctorant, until Apr 2014]
Yoann Isaac [CEA, until Oct 2014]
Jialin Liu [Inria]
Gaetan Marceau-Caron [Thalès (CIFRE), then PostDoc Inria since Oct 2014]
Basile Mayeur [Univ. Paris-Sud, from Sep 2014]

Sourava Prasad Mishra [Doctorant]
Thomas Schmitt [Inria, from Oct 2014]
Weijia Wang [Univ. Paris-Sud, until Jul 2014]
Guohua Zhang [Bourse Gvt Chinois]

Post-Doctoral Fellows

Marco Bressan [Inria, granted by Augure]
Jean-Baptiste Hoock [Univ. Paris-Sud, until Aug 2014]
David Lupien Saint Pierre [University des 3 Rivieres, Canada, from June 2014 to December 2014]

Visiting Scientists

Daria La Rocca [Univ. Roma 3, until Oct 2014]
Luigi Malago [University of Shinshu, Japan, Visiting collaborator, from Sep 2014]
Baptiste Roziere [ENS Lyon, until Aug 2014]

Administrative Assistant

Olga Mwana Mobulakani [Inria]

Others

Salim Abdelfettah [Inria, Internship, from Apr 2014 until Sep 2014]
Constance Deperrois [Inria, Internship, from Mar 2014 until Jul 2014]
Arthur Leroy [Inria, Internship, from Apr 2014 until Aug 2014]
Alexandre Quemy [INSA Rouen, Internship, until Oct 2014]
Chang-Min Wu [Inria, Internship, from Aug 2014 until Sep 2014]

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.

[1] MIT Technological Review, feb. 2001.

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.

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.

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 the 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 and constraint handling (respectively Ouassim Ait El Hara's, Ilya Loshchilov's PhDs and Asma Atamna's).

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 three stays (1 year, 6 months, 6 months) in Taiwan, and witnessed by the FP7 STREP Citines, the ADEME Post contract, and the METIS I-lab with SME Artelys.

The **E-Science** SIG (E-S-SIG) replaces and extends the former *Distributed systems* SIG, that was devoted to the modeling and optimization of (large scale) distributed systems, and itself was extending the goals of the original *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. Indeed, 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 are still being 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. But these activities have become more and more application-driven, from High Energy Physics for the highly distributed computation to the Social Sciences for the multi-agents approaches – hence the change of focus of this SIG. A major result of this theme is the creation of the Paris-Saclay Center for Data Science, co-chaired by Balázs Kégl, and the organization of the Higgs-ML challenge (http://higgsml.lal.in2p3.fr/), most popular challenge ever on the Kaggle platform.

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 6.3), AI planning (CRI-SIG, Section 6.5) and Air Traffic Control Optimization (Section 4.2). Recent work addressing algorithm selection using Collaborative

Filtering algorithms (CRI-SIG, Section 6.5) 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 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 had several related projects (Citines, a European (FP7) project; IOMCA, a ANR project), and we now work on the POST project, a ADEME BIA about investments in power systems. We have a collaboration with a company, Artelys, working on optimization in general, and in particular on energy management; this is a Inria ILAB.

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. We study various modelling errors, such as bias due to finite samples, linearization, and propose corrections.

Related Activities:

- 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. In particular, the Crystal tool is adopted by the European Community (http://www.artelys.com/news/120/90/Energy-The-European-Commission-Chooses-Artelys-Crystal)
- 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 since 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 (Areski 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. A new formulation of the problem, modeling the plane flows with Bayesian Networks, has been proposed in the Air Traffic Control community in 2013. In 2014, the corresponding stochastic multi-objective optimization problem has been tackled by Evolutionary Algorithms, leading to a general approach to uncertainty handling in Multi-Objective Evolutionary Algorithms [38], [59]. All details in Gaëtan's PhD [4].

5. New Software and Platforms

5.1. METIS

Participants: Olivier Teytaud [correspondent], Jérémie Decock, Jean-Joseph Christophe, Vincent Berthier, Marie-Liesse Cauwet, Jialin Liu, Sandra Cecilia Astete Morales.

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.

Metis is the Artelys/Tao contribution to Crystal, which is at the heart of the Post project, which is selected by the European Commission for a 4-years project for energy modeling http://www.artelys.com/news/120/90/Energy-The-European-Commission-Chooses-Artelys-Crystal.

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) in 7x7, and recently an optimization of the random seed which brings a significant improvement in Go and (unpublished) on the difficult case of phantom-Go. However, the work in the UCT-SIG has now shifted to energy management.

5.3. CMA-ES: Covariance Matrix Adaptation Evolution Strategy

Participants: Emmanuel Benazera, Nikolaus Hansen [correspondent].

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

The Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [65] is considered to be state-of-the-art in continuous domain evolutionary computation [64], 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, C++11, 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 6.3). 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

Participant: Balázs Kégl [correspondent].

Keywords: Multi-class, Multi-label Classification.

The MultiBoost package [63] 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 [27]. Its implementation is developed within the FUI project TIMCO.

Link: http://grid-observatory.org

5.7. Platforms

5.7.1. io.datascience

This Data as a Service (DaaS) platform is developed in the context of the Center for Data Science and the TIMCO project. Its overall goals is to exploit the advances in semantic web techniques for efficient sharing and usage of scientific data. A related specific software is the Tester for Triplestore (TFT) software suite, which benchmarks the compliance of sparql databases wrt the RDF standard and publishes the results through the SparqlScore service. TFT has been selected for the Semantic Web Challenge [42].

Links: https://io.datascience-paris-saclay.fr, https://github.com/BorderCloud/TFT, http://sparqlscore.com

6. New Results

6.1. Highlights of the Year

- The European commission has chosen Crystal-Supergrids (http://www.artelys.com/news/120/90/Energy-The-European-Commission-Chooses-Artelys-Crystal) for energy modeling and planning in Europe. Crystal-Supergrids is based on the Post project, an ADEME project between Artelys and Inria-TAO.
- The HiggsML challenge was the all-time most popular challenge organized by Kaggle. Cécile Germain-Renaud, Balázs Kégl and Marc Schoenauer were part of the organizing committee.
- Creation of the Center for Data Science, an interdisciplinary institute of the Université Paris-Saclay. Co-chaired by Balázs Kégl, with more than 250 permanent researchers in 35 laboratories, the CDS organizes continued cross-fertilization of machine learning and domain sciences.
- Best Paper Award at PPSN .

BEST PAPER AWARD :
[36] **Maximum Likelihood-based Online Adaptation of Hyper-parameters in CMA-ES in 13th International Conference on Parallel Problem Solving from Nature**. I. LOSHCHILOV, M. SCHOENAUER, M. SEBAG, N. HANSEN.

6.2. Optimal Decision Making under Uncertainty

Participants: Olivier Teytaud [correspondent], Jean-Joseph Christophe, Jérémie Decock, Nicolas Galichet, Marc Schoenauer, Michèle Sebag, Weijia Wang.

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.

After several years of success in the domain of GO, the most prominent application domain here is now energy management, at various time scales, and more generally planning. Furthermore, the work in this SIG has also lead to advances in continuous optimization at large, that somehow overlap with the work in the OPT-SIG (see 6.3).

The main advances done this year include:

Bandit-based Algorithms Active learning for the identification of biological dynamical systems has been tackled using Multi-Armed Bandit algorithms [35]. Weijia Wang's PhD [5] somehow summarizes the work done in TAO regarding Multi-objective Reinforcement Learning with MCTS algorithm. Differential Evolution was applied as an alternative to solve non-stationary Bandit problems [45].

Continuous optimization: parallelism, real-world, high-dimension and cutting-plane methods Our work in continuous optimization extends testbeds as follows: (i) including higher dimension (many testbeds in evolutionary algorithms consider dimension ≤ 40 or ≤ 100) (ii) taking into account computation time and not only the number of function evaluations (this makes a big difference in high dimension) (iii) including real world objective functions (iv) including parallelism, in particular, parallel convergence rates for differential evolution and particle swarm optimization [21]. We have a parallel version of cutting plane methods, which use more than black-box evaluations of the objective functions - we keep in mind that some of our black-box methods, on the other hand, also do not need convexity or the existence of a gradient.

Noisy optimization We have been working on noisy optimization in discrete and continuous domains. In the discrete case, we have shown the impact of heavy tails, and we have shown that resampling can solve some published open problems in an anytime manner. In the continuous case, we have shown [16] that a classical evolutionary principle (namely the step-size proportional to the distance to the optimum) implies that the optimal rates can not be reached - more precisely, we can have simple regret at best $O(1/\sqrt{number\ of\ fitness\ evaluations})$ in the simple case of an additive noise, whereas some published algorithms reached $O(1/number\ of\ fitness\ evaluations)$. One of the most directly applicable of our works is bias correction when the objective function $f(x)$ has the form $f(x) = \mathbb{E}_\omega f(x, \omega)$ and is approximated by $f(x) = \frac{1}{N}\sum_{i=1}^{N} f(x, \omega_i)$ for a given finite sample $\omega_1, \cdots, \omega_N$. We have also worked on portfolios of noisy optimizers [20], [34].

Discrete-time control with constrained action spaces. While Direct Policy Search is a reliable approach for discrete time control, it is not easily applicable in the case of a constrained high-dimensional action space. In the past, we have proposed DVS (Direct Value Search) for such cases [54]. The method is satisfactory, and we have additional mathematical results; in particular we prove positive results for non-Markovian, non-convex problems, and we prove a polynomial-time decision making and, simultaneously, exact asymptotic consistency for a non-linear transition [24]. Related work [60] also proposes to directly learn the value function, in a RL context, using some trajectories known to be bad.

Games. While still lightly contributing to the game of GO with our taiwanese partners [8], we obtained significant improvements in randomized artificial intelligence algorithms by decomposing the variance of the result into (i) the random seed (ii) the other random contributions such as the random seed of the opponent and/or the random part in the game. By optimizing our probability distribution on random seeds, we get significant improvements in e.g. phantom Go. This is basically a simple tool for learning opening books [44].

Adversarial bandits. High-dimensional adversarial bandits lead to two main drawbacks: (i) computation time (ii) highly mixed nature of the obtained solution. We developped methods which focus on sparse solution. Provably consistent, these methods are faster when the Nash equilibrium is sparse, and provides highly sparse solutions[17].

6.3. Continuous Optimization

Participants: Ouassim Ait Elhara, Asma Atamna, Anne Auger, Alexandre Chotard, Nikolaus Hansen, Yann Ollivier, Marc Schoenauer, Michèle Sebag, Olivier Teytaud, Luigi Malago, Emmanuel Benazera.

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 are well recognized in the field and were invited to write a book chapter on the design of continuous stochastic search[50].

Online adaptation of CMA-ES hyperparameters CMA-ES uses clever mechanisms to adapt the covariance matrix and the step-size, based on the evolution path. However, these mechanisms in turn use learning parameters, that were adjusted by trial-and-error in the seminal algorithm. However, thanks to the invariance properties of the algorithm, these values have been demonstrated to be very robust. An original mechanism has been proposed to adapt these hyper-parameters online, maximizing the likelihood of the selected sample at time to adapt the hyperparameters at time t-1. The corresponding paper published at PPSN received the Best Paper Award [36].

Linear Time and Space Complexity CMA-ES for Large-Scale Optimization We have been proposing a large-scale version of CMA-ES where the covariance matrix is restriced to a linear numbers of parameters. The update for the covariance matrix has been derived using the Information Geometric Optimization (IGO) framework and cumulation concepts borrowed from the original CMA

have been additionally included [14]. This work is part of a joint project between the TAO team and Shinshu university in Japan funded by the Japanese governement. In this context, Luigi Malago is visiting the team working on extending the proposed algorithm to a richer model.

Evaluation of Black-Box Optimizers We have been focusing on appraising the performance of step-size adaptation mechanisms for stochastic adaptive algorithms. We have shown that a too restrictive choice of test functions for the design of a method leads to misleading conclusions and proposed a thorough framework for evaluating step-size mechanism [29]. We have been pursuing our effort for *thorough and rigorous* benchmarking of black-box algorithms by organizing two more Black-Box-Optimization Benchmarking workshops that will take place at CEC 2015 and GECCO 2015. Those workshops are based on the platform COCO that we develop in the context of the ANR NumBBO project.

Theoretical Analysis of Stochastic Adaptive Algorithms We have analyzed the CSA-ES algorithm using resampling for **constrained optimization** on a linear function with a linear contraint. We have studied the behavior of the algorithm and proven success of failure of the algorithm depending on internal parameters of the algorithm [22]. We have extended a previous work on a linear function from the use of standard normal distribution to more general ones [23]. The published paper has been invited for an extension in an ECJ special issue. The extended paper had been submitted in december 2014. We have been providing a general methodology to prove the linear convergence of Comparison-based Step-size Adaptive Randomized Search on scaling-invariant functions by analyzing the stability of underlying Markov chains [57].

CMA-ES Library Besides our continuous work on implementations of CMA-ES (see e.g. github, PyPI), we have created a new library in C++11 (libcmaes). As part of the ANR SIMINOLE project, the library has been coupled with ROOT, the data analysis framework used at CERN, and generally in physics.

6.4. Applications to E-science

Participants: Cécile Germain-Renaud [correspondent], Marco Bressan, Philippe Caillou, Dawei Feng, Cyril Furtlehner, Blaise Hanczar, Karima Rafes, Balázs Kégl, Michèle Sebag.

The E-S-SIG explores the issues related to applications to E-Science, starting with modeling and optimizing very large scale computational grids, in particular in the context of Physics, to social sciences modelling with multi-agent systems.

The Higgs boson Machine Learning challenge The HiggsML challenge [2] has been set up to promote collaboration between high-energy physicists and computer scientists. The challenge, hosted by Kaggle, has drawn a remarkably large audience (with 1700+ teams it is one of the all-time most popular Kaggle challenges) and large coverage both in the social networks and in the media.

The goal of the challenge is to improve the procedure that classifies events produced by the decay of the Higgs boson versus events produced by other (background) processes, based on a training set of 250,000 examples. The challenge is a premier: it is the first time that a CERN experiment (ATLAS) made public such a large set of the official event and detector simulations. It also features a unique formal objective representing an approximation of the median significance (AMS) of a discovery (counting) test, which generates interesting algorithmic/theoretical questions beyond the usual challenges of finding and tuning the best classification algorithm [55].

A follow-up, the HEPML workshop was organized at NIPS14 [3], reporting on the results and the winning algorithms. The dataset and a software toolkit are available from the CERN Data Portal [4]

[2] https://www.kaggle.com/c/higgs-boson
[3] http://nips.cc/Conferences/2014/Program/event.php?ID=4292
[4] http://opendata.cern.ch

The Center for Data Science is a Lidex of the Université Paris-Saclay (UPSay), headed by Balazs Kégl and Arnak Dalalyan, gathering over 52 research teams and 34 labs with the goal of designing and applying automated methods to analyze massive and complex scientific datasets in order to extract useful information. Data science projects require expertise from a vast spectrum of disciplines (statistics, signal processing, machine learning, data mining, data visualization, high performance computing), besides the mastery of the scientific domain where the data originate from.

The goal of CDS is to establish an institutionalized agora in which scientists can find each other, exchange ideas, initiate and nurture interdisciplinary projects, and share their experience on past data science projects. To foster synergy between data analysts and data producers CDS organizes actions to provide initial resources for helping collaborations to get off the ground, to mitigate the non-negligible risk taken by researchers venturing into interdisciplinary data science projects, and to encourage the use of unconventional forms of information transmission and dissemination essential in this communication-intensive research area. The CDS fits perfectly in the recent surge of similar initiatives, both at the international and at the national level, and it has the potential to make the University Paris-Saclay one of the international fore-runners of data science [5].

Fault management As Lamport formulated decades ago, fault management in distributed systems exemplifies the unreachability of exact prior knowledge. Real-world large scale system add a supplementary complexity, which is non-stationarity.

- [12] models the system state and its ruptures (non-stationarity) through the flow of jobs as a stream (scalability), with a traceability goal (interpretability). These new streaming approaches involve self-calibration of the model based on scale invariance.
- D. Feng's PhD thesis [3] 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. [26] extends the matrix completion/compressd sensing setup to a sequential (tensor) context. We propose and evaluate a new algorithme, *Sequential Matrix Factorization* (SMF) that combines matrix completion with a self-calibrating exploration/exploitation balancing heuristic. Its active learning version (SMFA) exhibits superior performance over state-of-the-art methods.

Distributed system observation The work on distributed system automated analaysis and description[7] has been persued thru the continued development of the GAMA multi-agent framework https://code.google.com/p/gama-platform/wiki/GAMA. Philipps Caillou is associated to the new young researcher ANR ACTEUR, coordinated by Patrick Taillandier (IDEES, Rouen university), which will give an additional strucuture for further collaborations.

Identifying leaders in Social Networks The Modyrum contract with the SME Augure (funding Marco Bressan's Post-doc) aims at providing criteria to identify the trend leaders from blogs, tweets and other web-site posts. The same methods is being applied to fashion leaders in business as well as to opinion leaders in politics.

6.5. Designing criteria

Participants: Jamal Atif, Aurélien Decelle, Cyril Furtlehner, Yoann Isaac, Alexandre Quemy, Yann Ollivier, Marc Schoenauer, Michèle Sebag.

This 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 With motivating applications in large scale inference problems like traffic congestions we are pursuing our quest of practical solutions to inverse problems like in [39] where a method is proposed to invert a Gaussian Markov random field with topological and spectral constraints well suited to subsequent use of belief propagation as inference algorithm

[5] http://www.datascience-paris-saclay.fr/en

(see https://who.rocq.inria.fr/Jean-Marc.Lasgouttes/star-ips for the implementation). A more specific model for traffic inference has also been developped in [11]. A method adapted to the generalized belief propagation framework, aiming at adressing directly and systematically the loop corrections without loss of scalability is about to be completed.

Multi-objective ATC The new Bayesian approach of Air Traffic Control belongs to this SIG, but was described in the Section 4.2. Main publications are Gaétan Marceau's PhD [4] and the corresponding PPSN paper [38], [59].

Programming by Feedback Riad Akrour's PhD work on Preference Based Learning [1] culminated with the addition of a model for the user's competence in the interactive learning loop. In the resulting original paradigm, the user is sequentially proposed a series of behaviors and is only asked "Hot-or-cold" questions. The *Programming by Feedback* paradigm [15] will hopefully initiate a general way to allow non-digitally-proficient users to nevertheless control the behavior of software-based agents in their environment.

Multi-objective AI Planning This activity had almost stopped since the end of the DESCARWIN ANR project. However, a productive intership resulted in some new benchmarks in the ZenoTravel domain together with an exact solver ensuring the knowledge of the true Pareto front [41], [40].

Algorithm Selection Algorithm Selection can be viewed as a Collaborative Filtering problem, in which a problem "likes" an algorithm that is able to solve it. Initiated during Mustafa Misir's ERCIM postdoc in 2013, this idea has also been applied for Process Management [43], and is the basis of François Gonards's PhD funded by IRT SystemX in the context of aeronautics and car industry.

Outlier rejection in classification An original approach based on One-Class SVM has been proposed during Blaise Hanczar's on year *delegation* at TAO [28].

Learning sparse representations by auto-encoders Auto-encoders (AE) are a widely used tool for unsupervised learning, which consists of a neural network trained to reconstruct its own input via smaller-dimensional layers. The usual training criterion is the reconstruction error, however, the usual justification for AE is to learn a more compact data representation. In [62] we formalize this latter criterion using Minimum Description Length (MDL) and establish a comparison with the traditional reconstruction criterion. The MDL criterion has an interpretation as a denoising reconstruction and fully determines an optimal noise level, contrary to the literature on denoising AEs. More surprisingly, AE (aka Auto-associators) can also be used to learn sparse representations in the context of supervised learning [51].

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

- **Thalès Air Systems** (corr. Areski Hadjaz)– 2011-2014 (45 kEuros)
  Related to Gaétan Marceau-Caron's CIFRE PhD
  Participants: Marc Schoenauer
- **Thalès Research & Technology** (corr. Johann Dreo)– 2014-2017 (30 kEuros)
  Related to Nacim Belkhir's CIFRE PhD
  Participants: Marc Schoenauer
- **Modyrum (Modélisation Dynamique d'un Réseau Médiatique, related to Marco Bressan's postdoc** SME Augure – 2013-2015 (150 kEuros)
  Participants: Philippe Caillou, Cyril Furtlehner, Michèle Sebag
- **I-Lab METIS (A general framework for decision making with uncertainty plus energy-specific applications)** ARTELYS – 2011-2014 (40 kEuros)
  Related to Jérémie Decock's PhD
  Participants: Jérémie Decock, Jean-Joseph Christophe, Olivier Teytaud.

8. Partnerships and Cooperations

8.1. Regional Initiatives

- **TIMCO (Technology for In Memory Computing applications)** FUI Project – 2012-2015 (432 kEuros)
  Coordinator: Bull SA
  Participants: Cécile GERMAIN(WP Algorithm adaptation: the paradigm shif coordinator)
- **ROM - Réduction de modèles et optmisisation multiphysiques** – 2014 (73 kEuros).
  Coordinator: IRT System X
  Participants: Marc Schoneuaer, François Gonard (PhD)
- **ISN** – 2013-2016 (105 kEuros).
  Related to Thomas Schmitt's PhD - A Collaborative Filtering Approach to Matching Job Openings and Job Seekers
  Participants: Michèle Sebag, Thomas Schmitt
- **AutoML - An empirical approach to Machine Learning** – 2014-2017 (104 kEuros).
  Related to Sourava Mishra's PhD.
  Participants: Michèle Sebag, Balazs Kégl, Sourava Mishra

8.2. National Initiatives

- **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, Nikolaus Hansen, Emmanuel Benazera, 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
- **ACTEUR** – 2014-2018 (236kEuros). Cognitive agent development for urban simulations, ANR project, Coordinator P. Taillandier (IDEES, Univ Rouen).

  Participants: Philippe Caillou

8.2.1. Other

- **POST** – 2014-2018 (1,220 MEuros, including 500 kEuros for a 'private' cluster). Platform for the optimization and simulation of trans-continental grids
  ADEME (Agence de l'Environnement et de la Maîtrise de l'Energie)
  Coordinator: ARTELYS
  Participants: Olivier Teytaud, Marie-Liesse Cauwet, Jérémie Decock, Sandra Cecilia Astete Morales, David L. Saint-Pierre, J. Decock
- **E-LUCID** 2014-2017 (194 kEuros)
  Coordinator: Thales Communications & Security S.A.S
  Participants: Marc Schoenauer, Cyril Furtlehner

8.3. European Initiatives

8.3.1. FP7 & H2020 Projects

8.3.1.1. CitInES

Type: FP7

Defi: ICT for a low carbon economy

Instrument: Specific Targeted Research Project

Objectif: ICT systems for energy efficiency

Duration: October 2011 - March 2014

Coordinator: Artelys

Partners: Artelys (SME, France), Inria (Tao), AIT (Austria), Tupras (4 refineries, Turkey), Cesena (City, Italy), Ervet (Italy), Inesc-Porto (Portugal), Armines (France), Bologna (City, Italy)

Inria contact: Olivier Teytaud

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

8.3.1.2. EGI-Inspire

Type: FP7

Defi: e-Infrastructure

Instrument: CP-CSA

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

Duration: May 2010 - April 2014

Coordinator: EGI.eu foundation

Partner: Université Paris XI (France)

Inria contact: Cécile Germain-Renaud

Abstract: To support European science and innovation, a longer lasting operational model is now needed - both for coordinating the infrastructure itself and for delivering integrated services that cross national borders.The EGI-InSPIRE project will support the transition from a project-based system to a sustainable pan-European e-Infrastructure.

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 El ectricity Generation

International Partner (Institution - Laboratory - Researcher):

NUTN (TAIWAN)

Duration: 2012 - 2014

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

The objective of the project is three-fold:

- Objective 1: Designing consistent iterative realistic algorithms for partially observable 1-player or 2-player games.
  • Consistent algorithms (provably, asymptotically optimal in the computation time).
  • Iterative a.k.a. anytime algorithms, improving its results as the computational time allowed increases and requiring little time to yield a decent answer. Most algorithms which survive decades are iterative.
  • Realistic algorithms, i.e. suited to real-world settings.

- Objective 2: Impressive visible applications, e.g. applications in games or puzzles, such as Minesweeper (on which we believe that much progress is still possible), Chinese Dark Chess, Kriegspiel, Phantom-Go, or card games. Games and puzzles offer nice frameworks to assess and make our research highly visible.
- Objective 3: Big industrial applications. Having both mathematics and visible realizations in games and industrial applications might be considered as too ambitious. Yet, our strategy is to tackle e.g. the field of energy generation because: i) it is close from our past activities (thus reducing the warm-up time), yet with a new challenge, partial observability; ii) in real applications, many problems are simplified so that they boil down to fully observable problems, (e.g. through including tricks in the solvers); iii) our former achievements facilitate our contact with industry. Formally, 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; that these results will translate to real-world application.
  Our roadmap is:
  • Check on simple versions of energy production problems whether the fully observable approximation holds true. We guess that in many cases it does not; the next point is to assess the loss of performance incurred;
  • Experiment our algorithms on real industrial problems, considering both Taiwan-centered and Europe-Centered electricity generation problems in order to widen the scope of the analysis, enforcing the applicability of the approach.

8.4.2. Inria International Partners

8.4.2.1. Declared Inria International Partners

- Shinshu University (Professor Akimoto, Professor Tanaka, Professor Aguire). Partnership officialized via MOU signature between Inria and Shinshu University. Joint project funded by the Japanese governement.
- Dortmund University through the funded ANR project NumBBO.

8.5. International Research Visitors

8.5.1. Visits of International Scientists

- Holger Hoos, Professor, Dept of Computer Science, University of British Columbia, from Oct. 1. to Dec. 31., funded by Microsoft-Inria Joint Lab.
- Daria La Rocca, PhD student at University Roma 3, Italy, from Oct. 2013 until Oct. 2014
- Luigi Malago, Post-Doc at University Shinshu, Japan, since Sept. 2014 (see Section 8.4.2.1).

8.5.2. Visits to International Teams

8.5.2.1. Research stays abroad

- Olivier Teytaud, National University of Tainan and Dong Hwa University in Hualien, Taiwan (6 months).
- Jialin Liu, National University of Tainan and Dong Hwa University in Hualien, Taiwan (6 months).
- Marie-Liesse Cauwet, National University of Tainan and Dong Hwa University in Hualien, Taiwan (3 months).
- Constance Deperrois, National University of Tainan and Dong Hwa University in Hualien, Taiwan (1 month).
- Baptiste Roziere, National University of Tainan and Dong Hwa University in Hualien, Taiwan (2 months).

- Sandra Cecilia Astete Morales, National University of Tainan and Dong Hwa University in Hualien, Taiwan (1 month).
- Vincent Berthier, National University of Tainan and Dong Hwa University in Hualien, Taiwan (4 months).
- David L. Saint-Pierre, National University of Tainan and Dong Hwa University in Hualien, Taiwan (4 months).

9. Dissemination

9.1. Promoting Scientific Activities

9.1.1. Management positions in scientific organisations

9.1.1.1. 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). Anne Auger, member of Executive Board since 2011.
- 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, Michèle Sebag, member since 2010;
- ECCAI Fellow, Michèle Sebag, since 2011;
- LION (Learning and Intelligent OptimizatioN), Marc Schoenauer, member of the Steering Committee since 2012,
- Marc Schoenauer, Honorary Adjunct Professor, School of Computer Science, University of Adelaide, Australia (2009-2015).

9.1.1.2. national

- Michèle Sebag, member of the CoNRS; Senior Advisory Board CHIST-ERA; member of the CSFRS (Conseil Supérieur de la Formation et Recherche Stratégique);
- EA – Association Evolution Artificielle: Marc Schoenauer, founding president, now member of Advisory Committee. Anne Auger, member of Executive Commitee since 2008.

9.1.1.3. 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 et au Conseil Académique de l'université since 2013, directeur des études à l'IUT de Sceaux since 2009
- Cécile Germain-Renaud, elected memeber of the Scientific Council and of its board. University officer for scientific computing. Deputy head of the computer science departement, in charge of research.
- Michèle Sebag, membre élu du Conseil du Laboratoire de Recherche en Informatique et membre de la CCSU 27 since 2004.

9.1.1.4. Université Paris-Saclay

- Michèle Sebag, member of the Senate at Université Paris-Saclay; Member of the Executive committee and responsible for the DataSense axis, DigiCosme Labex;
- Marc Schoenauer, Member of the Executive committee and chair of the Research Committee, DigiCosme Labex;

9.1.1.5. 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* (VP for Research).

9.1.2. Scientific events organisation

9.1.2.1. general chair, scientific chair

- Anne Auger: ES/EP GECCO 2014 track chair

9.1.2.2. Member of the organizing committee

- HEPML workshop at NIPS 2014: Cecile Germain-Renaud and Balázs Kégl.
- Michèle Sebag, Senior Program Committee, ECAI 2014; Area Chair, IJCAI 2015.

9.1.2.3. Others

- Anne Auger: Organization of an invited mini-symposium on Function-Value-Free Optimization, SIAM Conference on Optimization, San Diego, USA.
- Anne Auger: Organization of invited session on Recent Advances on Continuous Randomized black-box optimization, PGMO-COPI'14 conference, Paris.
- Anne Auger, Nikolaus Hansen: Organization of workshops BBOB, IEEE Conference on Evolutionary Computation 2015, Sendai, Japan and GECCO 2015 Conference Madrid, Spain.
- Anne Auger: Women@GECCO workshop 2014, Vancouver Canada

9.1.3. Scientific events selection

9.1.3.1. Member of the conference Program Committee

All TAO members are reviewers for the main conferences in Machine Learning (IJCAI, ICML, ECAI, ECML-PKDD, ...) and Evolutionary Computation (ACM GECCO, PPSN, IEEE CEC, EvoStar, ...).

9.1.4. Journal

9.1.4.1. Member of the Editorial Board

- Anne Auger: Evolutionary Computation, MIT Press.
- Nikolaus Hansen: Evolutionary Computation, MIT Press.
- Marc Schoenauer: Evolutionary Computation, MIT Press, Advisory Board; Kluwer Genetic Programming and Evolvable Machines, Editorial Board; JMLR, Area Editor.
- Michèle Sebag: Machine Learning, Springer Verlag.

9.1.4.2. Reviewer

Similarly, all TAO members do occasional reviews for the main journals in ML and EC.

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

Licence: Jamal Atif, Computer Science, approx. 192h, L1, IUT Orsay, Univ. Paris-Sud.

Licence: Anne Auger, Stochastic Optimization, 20h, L1-3, Ecole Centrale Paris.

Licence: Philippe Caillou, Computer science for Managers, approx. 192h, L1, IUT Sceaux, Univ. Paris-Sud.

Licence: Aurélien Decelle, Computer Architecture, 30h, L3, Univ. Paris-Sud

Licence: Aurélien Decelle, Machine Learning and Artificial Life, 41h, L2, Univ. Paris-Sud

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

Licence: Michèle Sebag, Artificial Life, 10h, L2, Univ Paris-Sud.

Licence: Michèle Sebag, Introduction to Machine Learning, approx. 24h, L3, ENS-Cachan.

Master : Cécile Germain-Renaud, Parallelisme, approx. 50h, M1, U. Paris-Sud.

Master : Michèle Sebag, Machine Learning, 25h, M1, MPRI.

Master : Anne Auger, Optimisation, 12h, M2 Recherche, U. Paris-sud.

Master : Michèle Sebag, Machine Learning, 24h, M2 Recherche, U. Paris-sud.

Master : Michèle Sebag, Evolutionary Robotics, 15h, M2 Recherche, U. Paris-sud.

Master : Philippe Caillou, Multi-Agents Systems, 27h, M2 Recherche, U. Paris-sud.

Master : Philippe Caillou, Multi-Agent Based Simulation, 3h, M2 Recherche, U. Paris-Dauphine.

Master : Yann Ollivier, Deep learning, 4h, M2 Recherche, Telecom/Polytech.

9.2.2. Supervision

PhD: Dawei Feng, *Efficient End-to-End Monitoring for Fault Management in Distributed Systems*, 27/03/2014, Germain-Renaud [3].

PhD: Weijia Wang, *Multi-objective sequential decision making*, 11/07/2014, Michèle Sebag [5].

PhD: Gaetan Marceau Caron, *Optimization and Uncertainty Handling in Air Traffic Management*, 22/09/2014, Marc Schoenauer [4].

PhD: Jérémie Decock, *Hybridization of dynamic optimization methodologies*, 26/11/2014, Olivier Teytaud [2].

PhD: Riad Akrour, *Robust Preference-based Reinforcement Learning*, 30/09/2014, Marc Schoenauer and Michèle Sebag [1].

PhD in progress: Ouasssim Ait Elhara, *Stochastic Black-Box Optimization and Benchmarking in Large Dimension*, 1/10/2012, Anne Auger and Nikolaus Hansen

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, *Large-scale black-box optimization*, 1/10/2013, Olivier Teytaud

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: 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 Michèle Sebag

PhD in progress: Jialin Liu, *Portfolios of Noisy Optimization Algorithms*, 14/03/2013, Marc Schoenauer and Olivier Teytaud

PhD in progress: Pierre-Yves Massé, *Gradient methods in statistical learning*, 1/10/2014, Yann Ollivier

PhD in progress: Guohua Zhang, *Curiosity-Driven Navigation in Evolutionary Robotics*, 1/9/2011, Michèle Sebag

PhD in progress: Nacim Belkhir, *Ajustement automatique de paramètres pour les métaheuristiques*, 1/05/2014, Marc Schoenauer

Basile Mayeur, *Direct Value Learning*, 1/10/2014, Michèle Sebag

Antoine Bureau, *Programming by Feedback*, 1/10/2014, Michèle Sebag and Marc Schoenauer

Thomas Schmitt, *A Collaborative Filtering Approach to Matching Job Openings and Job Seekers*, 1/11/2014, Michèle Sebag

Sourava Mishra, *AutoML: an empirical approach to Machine Learning*, 1/10/2014, Balázs Kégl and Michèle Sebag

François Gonard, *Sélection d'algorithmes pour la conception de structures*, 1/11/2014, Marc Schoenauer

9.2.3. Tutorials and Summer Schools

- Anne Auger and Nikolaus Hansen: *Evolution Strategies and CMA-ES* at the ACM GECCO, July 2014.
- Nikolaus Hansen: *Introduction to Information Geometry: Stochastic Optimization* at the Information Geometry in Learning and Optimization summer school, September 2014, Copenhagen.

9.2.4. Juries

- Marc Schoenauer, Reviewer for Adrien Goeffon's HDR, Université d'Angers, Oct. 2014 ; Richard Alligier's PhD, Institut National Polytechnique de Toulouse, Oct. 2014 ; Salwa Belaqziz, Université de Marrakech, May 2014. Jury member for Steve Oudot's HDR, Université Paris-Sud, ; Jérémie Garcia's PhD, Université Paris-Sud, June 2014 (pdt of the jury) ; Mickael Buchet's PhD, Université Paris-Sud, Nov. 2014 (pdt of the jury).
- Michele Sebag, Reviewer for Jonathan Grizou, U. Bordeaux; Gwenael Bothorel, U. Toulouse. Jury member for Bilal Piot, U. Lorraine (pdt of the jury).
- Olivier Teytaud, Reviewer for Manel Tagorti's ph.D. (to be defended early 2015, Loria) and Sothea Hong (Irstea Montpellier).

9.3. Popularization

Yann Ollivier, co-organizes the European Union Contest for Young Scientists (science fair for high school students from 30+ countries organized by the European Commission).

10. Bibliography

Publications of the year

Doctoral Dissertations and Habilitation Theses

[1] R. AKROUR. *Robust Preference Learning-based Reinforcement Learning*, Université Paris-Sud, September 2014, https://hal.inria.fr/tel-01111276

[2] J. DECOCK. *Hybridization of dynamic optimization methodologies*, Université Paris Sud, November 2014, https://hal.inria.fr/tel-01103935

[3] D. FENG. *Efficient end-to-end monitoring for fault management in distributed systems*, Université Paris Sud - Paris XI, March 2014, https://tel.archives-ouvertes.fr/tel-01017083

[4] G. MARCEAU CARON. *Optimization and Uncertainty Handling in Air Traffic Management*, Paris-Sud XI, September 2014, https://hal.inria.fr/tel-01080370

[5] W. WANG. *Multi-objective sequential decision making*, Université Paris Sud - Paris XI, July 2014, https://tel.archives-ouvertes.fr/tel-01057079

Articles in International Peer-Reviewed Journals

[6] N. J.-B. BRUNEL, Q. CLAIRON, F. D'ALCHÉ-BUC. *Parametric Estimation of Ordinary Differential Equations with Orthogonality Conditions*, in "Journal of American Statistics Association", 2014, vol. 109, n° 505, pp. 173–185 [*DOI :* 10.1080/01621459.2013.841583], https://hal.archives-ouvertes.fr/hal-00867370

[7] 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° 6, https://hal.inria.fr/hal-00927587

[8] C.-S. LEE, M.-H. WANG, M.-J. WU, O. TEYTAUD, S.-J. YEN. *T2FS-based Adaptive Linguistic Assessment System for Semantic Analysis and Human Performance Evaluation on Game of Go*, in "IEEE Transactions on Fuzzy Systems", 2014, 22 p. , https://hal.inria.fr/hal-01059822

[9] M. SEBAG. *A tour of Machine Learning: an AI perspective*, in "AI Communications", 2014, vol. 27, n° 1, pp. 11-23 [*DOI :* 10.3233/AIC-130580], https://hal.inria.fr/hal-01109768

[10] Y. SHOGO, M. OHZEKI, A. DECELLE. *Detection of Cheating by Decimation Algorithm*, in "Journal of the Physical Society of Japan", January 2015, vol. 84, 024801 [*DOI :* 10.7566/JPSJ.84.024801], https://hal.archives-ouvertes.fr/hal-01105415

[11] K. SHUN, Y. MUNEKI, C. FURTLEHNER, K. TANAKA. *Traffic data reconstruction based on Markov random field modeling* , in "Inverse Problems", 2014, vol. 30, n° 2, 15 p. , https://hal.inria.fr/hal-01096947

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

International Conferences with Proceedings

[13] C. ADAM-BOURDARIOS, G. COWAN, C. GERMAIN, I. GUYON, B. KEGL, D. ROUSSEAU. *The ATLAS Higgs Boson Machine Learning Challenge*, in "International Conference on High Energy Physics(ICHEP) Conference", Valencia, Spain, July 2014, forthcoming, https://hal.inria.fr/hal-01111177

[14] Y. AKIMOTO, A. AUGER, N. HANSEN. *Comparison-Based Natural Gradient Optimization in High Dimension*, in "Genetic and Evolutionary Computation Conference GECCO'14", Vancouver, Canada, ACM, July 2014, https://hal.inria.fr/hal-00997835

[15] R. AKROUR, M. SCHOENAUER, M. SEBAG, J.-C. SOUPLET. *Programming by Feedback*, in "International Conference on Machine Learning", Pékin, China, June 2014, https://hal.inria.fr/hal-00980839

[16] S. ASTETE-MORALES, M.-L. CAUWET, O. TEYTAUD. *Evolution Strategies with Additive Noise: A Convergence Rate Lower Bound*, in "Foundations of Genetic Algorithms", Aberythswyth, United Kingdom, 2015, 9 p. , https://hal.inria.fr/hal-01077625

[17] D. AUGER, J. LIU, S. RUETTE, D. L. SAINT-PIERRE, O. TEYTAUD. *Sparse Binary Zero-Sum Games*, in "Asian Conference on Machine Learning", Ho-Chi-Minh-Ville, Vietnam, 2014, vol. 29, 16 p. , https://hal.inria.fr/hal-01077627

[18] A. BUREAU, M. SEBAG. *Bellmanian Bandit Network*, in "Autonomously Learning Robots, at NIPS", Montréal, Canada, Gerhard Neumann (TU-Darmstadt) and Joelle Pineau (McGill University) and Peter Auer (Uni Leoben) and Marc Toussaint (Uni Stuttgart), December 2014, https://hal.inria.fr/hal-01102970

[19] M.-L. CAUWET. *Noisy Optimization: Convergence with a Fixed Number of Resamplings*, in "EvoStar", Granada, Spain, April 2014, https://hal.inria.fr/hal-00976063

[20] M.-L. CAUWET, J. LIU, O. TEYTAUD. *Algorithm Portfolios for Noisy Optimization: Compare Solvers Early*, in "Learning and Intelligent Optimization Conference", Florida, United States, February 2014, https://hal.inria.fr/hal-00926638

[21] M.-L. CAUWET, O. TEYTAUD, S.-Y. CHIU, K.-M. LIN, S.-J. YEN, D. L. SAINT-PIERRE, F. TEYTAUD. *Parallel Evolutionary Algorithms Performing Pairwise Comparisons*, in "Foundations of Genetic Algorithms", Aberythswyth, United Kingdom, 2015, 15 p. , https://hal.inria.fr/hal-01077626

[22] A. CHOTARD, A. AUGER, N. HANSEN. *Markov Chain Analysis of Evolution Strategies on a Linear Constraint Optimization Problem*, in "IEEE Congress on Evolutionary Computation", Beijing, China, A. HUSSAIN, Z. ZENG, N. ZHANG (editors), http://www.ieee-wcci2014.org/committees.htm, July 2014, https://hal.inria.fr/hal-00977379

[23] A. CHOTARD, M. HOLEŇA. *A Generalized Markov-Chain Modelling Approach to (1,λ)-ES Linear Optimization*, in "13th International Conference on Parallel Problem Solving from Nature", Ljubljana, Slovenia, Lecture Notes in Computer Science, Springer, September 2014, vol. 8672, pp. 902 - 911 [*DOI :* 10.1007/978-3-319-10762-2_89], https://hal.inria.fr/hal-01091494

[24] J.-J. CHRISTOPHE, J. DECOCK, O. TEYTAUD. *Direct model predictive control*, in "European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)", Bruges, Belgium, April 2014, https://hal.inria.fr/hal-00958192

[25] Y. COADOU, B. KÉGL. *Boosted Decision Tree and application*, in "IN2P3 School of Statistics 2014", Autrans, France, May 2014, https://hal.inria.fr/in2p3-00997259

[26] D. FENG, C. GERMAIN-RENAUD, J. NAUROY. *Sequential fault monitoring*, in "Cloud and Autonomic Computing", London, United Kingdom, IEEE, September 2014, https://hal.inria.fr/hal-01064161

[27] C. GERMAIN, J. NAUROY, K. RAFES. *The Grid Observatory 3.0 - Towards reproducible research and open collaborations using semantic technologies* , in "EGI Community Forum 2014", Helsinki, Finland, May 2015, https://hal.inria.fr/hal-01104235

[28] B. HANCZAR, M. SEBAG. *Combination of One-Class Support Vector Machines for Classification with Reject Option*, in "Machine Learning and Knowledge Discovery in Databases - Part I", Nancy, France, T. CALDERS, F. ESPOSITO, E. HÜLLERMEIER, R. MEO (editors), Machine Learning and Knowledge Discovery in Databases - Part I, September 2014, vol. 8724, pp. 547 - 562 [*DOI :* 10.1007/978-3-662-44848-9_35], https://hal.inria.fr/hal-01109774

[29] N. HANSEN, A. ATAMNA, A. AUGER. *How to Assess Step-Size Adaptation Mechanisms in Randomised Search*, in "13th International Conference on Parallel Problem Solving from Nature", Ljubljana, Slovenia, T. BARTZ-BEIELSTEIN (editor), LNCS, Springer, September 2014, vol. 8672, pp. 60-69, https://hal.inria.fr/hal-00997294

[30] B. KÉGL. *Center for data science*, in "Paris-Saclay Center for Data Science Kick-off Meeting", Orsay, France, June 2014, https://hal.inria.fr/in2p3-01020019

[31] B. KÉGL. *Real time multivariate classifiers*, in "Trigger, Online and Offline Computing Workshop", Geneve, Switzerland, September 2014, http://hal.in2p3.fr/in2p3-01070943

[32] B. KÉGL, D. ROUSSEAU, C. GERMAIN, I. GUYON, G. COWAN. *Introduction to the HEPML Workshop and the HiggsML challenge*, in "HEPML workshop at NIPS14 - Neural Information Processing Systems Conference", Montreal, Canada, December 2014, http://hal.in2p3.fr/in2p3-01100982

[33] J. LIU, D. L. SAINT-PIERRE, O. TEYTAUD. *A mathematically derived number of resamplings for noisy optimization*, in "Companion - Genetic and Evolutionary Computation Conference (GECCO 2014)", Vancouver, Canada, ACM, July 2014, pp. 61-62, https://hal.inria.fr/hal-00979442

[34] J. LIU, O. TEYTAUD. *Meta online learning: experiments on a unit commitment problem*, in "European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning", Bruges, Belgium, April 2014, https://hal.inria.fr/hal-00973397

[35] A. LLAMOSI, A. MEZINE, F. D'ALCHÉ-BUC, V. LETORT, M. SEBAG. *Experimental Design in Dynamical System Identification: A Bandit-Based Active Learning Approach*, in "Machine Learning and Knowledge Discovery in Databases - Part II", Nancy, France, Springer Verlag, September 2014, vol. 8725, pp. 306 - 321 [*DOI :* 10.1007/978-3-662-44851-9_20], https://hal.inria.fr/hal-01109775

[36] *Best Paper*
I. LOSHCHILOV, M. SCHOENAUER, M. SEBAG, N. HANSEN. *Maximum Likelihood-based Online Adaptation of Hyper-parameters in CMA-ES*, in "13th International Conference on Parallel Problem Solving from Nature", Ljubljana, Slovenia, September 2014, https://hal.inria.fr/hal-01003504.

[37] L. MALAGÒ, G. PISTONE. *Information Geometry of Gaussian Distributions in View of Stochastic Optimization*, in "Foundations of Genetic Algorithms XIII", Aberystwyth, United Kingdom, Jun He, Thomas Jansen, Gabriela Ochoa and Christine Zarges, January 2015, https://hal.inria.fr/hal-01108986

[38] G. MARCEAU, M. SCHOENAUER. *Racing Multi-Objective Selection Probabilities*, in "13th International Conference on Parallel Problem Solving from Nature", Ljubljana, Slovenia, September 2014, 1 p. , https://hal.archives-ouvertes.fr/hal-01009907

[39] V. MARTIN, C. FURTLEHNER, Y. HAN, J.-M. LASGOUTTES. *GMRF Estimation under Topological and Spectral Constraints*, in "7th European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases", Nancy, France, T. CALDERS, F. ESPOSITO, E. HÜLLERMEIER, R. MEO (editors), Lecture Notes in Computer Science, Springer Berlin Heidelberg, September 2014, vol. 8725, pp. 370-385 [*DOI :* 10.1007/978-3-662-44851-9_24], https://hal.archives-ouvertes.fr/hal-01065607

[40] A. QUEMY, M. SCHOENAUER. *True Pareto Fronts for Multi-Objective AI Planning Instances*, in "European Conference on Combinatorial Optimization - EvoCOP", Copenhague, Denmark, F. CHICANO, G. OCHOA (editors), LNCS, Springer Verlag, April 2015, forthcoming, https://hal.archives-ouvertes.fr/hal-01109777

[41] A. QUEMY, M. SCHOENAUER, V. VIDAL, J. DRÉO, P. SAVÉANT. *Solving Large MultiZenoTravel Benchmarks with Divide-and-Evolve*, in "Learning and Intelligent OptimizatioN - LION 9", Lille, France, C. DHAENENS, L. JOURDAN, M.-E. MARMION (editors), LNCS, Springer Verlag, January 2015, 7 p. , https://hal.archives-ouvertes.fr/hal-01109776

[42] K. RAFES, J. NAUROY, C. GERMAIN. *TFT, Tests For Triplestores* , in "Semantic Web Challenge, part of the International Semantic Web Conference", Riva Del Garda, Italy, October 2014, https://hal.inria.fr/hal-01104252

[43] J. RIBEIRO, J. CARMONA, M. MISIR, M. SEBAG. *A Recommender System for Process Discovery*, in "Business Process Management", Eindhoven, Netherlands, S. SADIQ, P. SOFFER, H. VÖLZER (editors), Proc. 12th Business Process Management, Springer Verlag, September 2014, vol. 8659, pp. 67 - 83 [*DOI :* 10.1007/978-3-319-10172-9_5], https://hal.inria.fr/hal-01109766

[44] D. L. SAINT-PIERRE, O. TEYTAUD. *Nash and the Bandit Approach for Adversarial Portfolios*, in "CIG 2014 - Computational Intelligence in Games", Dortmund, Germany, Computational Intelligence in Games, IEEE, August 2014, 7 p. [*DOI :* 10.1109/CIG.2014.6932897], https://hal.inria.fr/hal-01077628

[45] D. L. ST-PIERRE, J. LIU. *Differential Evolution Algorithm Applied to Non-Stationary Bandit Problem*, in "2014 IEEE Congress on Evolutionary Computation (IEEE CEC 2014)", Beijing, China, July 2014, https://hal.inria.fr/hal-00979456

[46] G. ZHANG, M. SEBAG. *Coupling Evolution and Information Theory for Autonomous Robotic Exploration*, in "13th International Conference on Parallel Problem Solving from Nature", Ljubliana, Slovenia, T. BARTZ-BEIELSTEIN, J. BRANKE, B. FILIPIC, J. SMITH (editors), Lecture Notes in Computer Science, Springer Verlag, September 2014, vol. 8672, pp. 852 - 861 [*DOI :* 10.1007/978-3-319-10762-2_84], https://hal.inria.fr/hal-01109770

National Conferences with Proceedings

[47] B. KÉGL. *Introduction to MVA approach*, in "IN2P3 School of Statistics 2014", Autrans, France, May 2014, https://hal.inria.fr/in2p3-00997253

Conferences without Proceedings

[48] B. KÉGL. *La science des données pour les données de la science*, in "9th Journées Informatique de l'IN2P3-IRFU", Le Grau du Roi, France, October 2014, http://hal.in2p3.fr/in2p3-01076155

Scientific Books (or Scientific Book chapters)

[49] *CMA-ES: A Function Value Free Second Order Optimization Method*, 2014, https://hal.inria.fr/hal-01110313

[50] 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, 2014, pp. 145-180, https://hal.inria.fr/hal-00808450

[51] S. REBECCHI, H. PAUGAM-MOISY, M. SEBAG. *Learning Sparse Features with an Auto-Associator*, in "Growing Adaptive Machines", T. KOWALIW, N. BREDECHE, R. DOURSAT (editors), Studies in Computational Intelligence, Springer Verlag, 2014, vol. 557, pp. 139 - 158 [*DOI :* 10.1007/978-3-642-55337-0_4], https://hal.inria.fr/hal-01109773

Books or Proceedings Editing

[52] Y. OLLIVIER, H. PAJOT, C. VILLANI (editors). *Optimal Transportation* , London Mathematical Society Lecture Note Series, Cambridge University PressGrenoble, France, 2014, vol. 413, https://hal.archives-ouvertes.fr/hal-01104763

Research Reports

[53] A. CHOTARD, M. HOLENA. *A Generalized Markov-Chain Modelling Approach to* $(1, \lambda)$*-ES Linear Optimization: Technical Report*, June 2014, https://hal.inria.fr/hal-01003015

Scientific Popularization

[54] J. DECOCK, J.-J. CHRISTOPHE, O. TEYTAUD. *Optimization of Energy Policies Using Direct Value Search*, May 2014, 9èmes Journées Francophones de Planification, Décision et Apprentissage (JFPDA'14), https://hal.inria.fr/hal-00997562

Other Publications

[55] C. ADAM-BOURDARIOS, G. COWAN, C. GERMAIN, I. GUYON, B. KÉGL, D. ROUSSEAU. *Learning to discover: the Higgs boson machine learning challenge*, May 2014 [*DOI :* 10.7483/OPENDATA.ATLAS.MQ5J.GHXA], https://hal.inria.fr/hal-01104487

[56] O. AIT ELHARA, A. AUGER, N. HANSEN. *Large-Scale Optimization of Low Effective and Low Epsilon-Effective Dimension Problems*, February 2015, https://hal.inria.fr/hal-01112850

[57] A. AUGER, N. HANSEN. *Linear Convergence of Comparison-based Step-size Adaptive Randomized Search via Stability of Markov Chains*, May 2014, https://hal.inria.fr/hal-00877160

[58] N. LIM, F. D'ALCHÉ-BUC, C. AULIAC, G. MICHAILIDIS. *Operator-valued Kernel-based Vector Autoregressive Models for Network Inference*, March 2014, https://hal.archives-ouvertes.fr/hal-00872342

[59] G. MARCEAU, M. SCHOENAUER. *Racing Multi-Objective Selection Probabilities*, June 2014, Extended preprint of PPSN 2014 paper, https://hal.inria.fr/hal-01002854

[60] B. MAYEUR, R. AKROUR, M. SEBAG. *Direct Value Learning: a Rank-Invariant Approach to Reinforcement Learning*, October 2014, https://hal.archives-ouvertes.fr/hal-01090982

[61] U.-M. O'REILLY, A. ESPARCIA-ALCAZAR, A. AUGER, C. DOERR, A. EKART, G. OCHOA. *Women@GECCO 2014*, 2014, 2 p. , Summary of the Women@GECCO meeting [*DOI :* 10.1145/2598394.2611386], http://hal.upmc.fr/hal-01086538

[62] Y. OLLIVIER. *Auto-encoders: reconstruction versus compression*, 2014, https://hal.archives-ouvertes.fr/hal-01104268

References in notes

[63] 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

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

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