Activity Report 2019

Project-Team TAU

Tackling the under-specified

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

Table of contents

Project-Team TAU

Creation of the Project-Team: 2019 July 01

Keywords:

Computer Science and Digital Science:
        A3.3.3. - Big data analysis
        A3.4. - Machine learning and statistics
        A3.5.2. - Recommendation systems
        A8.2. - Optimization
        A8.6. - Information theory
        A9.2. - Machine learning
        A9.3. - Signal analysis

Other Research Topics and Application Domains:
        B1.1.4. - Genetics and genomics
        B4. - Energy
        B7.2.1. - Smart vehicles
        B9.1.2. - Serious games
        B9.5.3. - Physics
        B9.5.5. - Mechanics
        B9.5.6. - Data science
        B9.6.10. - Digital humanities

1. Team, Visitors, External Collaborators

Research Scientists
    Marc Schoenauer [Team leader, Inria, Senior Researcher, HDR]
    Michele Alessandro Bucci [Inria, Starting Research Position, from Oct 2019]
    Guillaume Charpiat [Inria, Researcher]
    Cyril Furtlehner [Inria, Researcher]
    Flora Jay [CNRS, Researcher]
    Michèle Sebag [CNRS, Senior Researcher, HDR]
    Paola Tubaro [CNRS, Researcher, HDR]

Faculty Members
    Cécile Germain [Univ Paris-Sud, Emeritus, HDR]
    Isabelle Guyon [Univ Paris-Sud and Inria, Professor and *Chaire d'Excellence UPSaclay-Inria*, HDR]
    Nicolas Spyratos [Univ Paris-Sud, Emeritus, HDR]
    Philippe Caillou [Univ Paris-Sud, Associate Professor]
    Aurélien Decelle [Univ Paris-Saclay, Associate Professor]
    Francois Landes [Univ Paris-Saclay, Associate Professor]

Post-Doctoral Fellows
    Jean Cury [Univ Paris-Sud]
    Ksenia Gasnikova [Inria, from Feb 2019]
    Saumya Jetley [Inria, from Oct 2019]

PhD Students

Eleonore Bartenlian [Ministère de l'Enseignement Supérieur et de la Recherche]
Victor Berger [Inria]
Leonard Blier [Facebook]
Benjamin Donnot [RTE, until Jan 2019]
Mandar Chandorkar [CWI, until Oct 2019]
Balthazar Donon [RTE, from oct 2018]
Guillaume Doquet [Inria, until Dec 2019]
Victor Estrade [Univ Paris-Sud, until Sep 2019]
Loris Felardos Saint Jean [Inria]
Giancarlo Fissore [Univ Paris-Sud]
Julien Girard [CEA]
Diviyan Kalainathan [Univ Paris-Sud]
Armand Lacombe [Inria, from Oct 2019]
Wenzhuo Liu [IRT System X, from Nov 2019]
Zhengying Liu [Ecole Polytechnique]
Nizam Makdoud [Thales]
Marc Nabhan [Renault]
Adrian Pol [CERN]
Herilalaina Rakotoarison [Inria]
Théophile Sanchez [Univ Paris-Sud]
Lisheng Sun [Univ Paris-Sud]
Corentin Tallec [Ecole polytechnique, until Sep 2019]
Marion Ullmo [Univ Paris-Sud]
Pierre Wolinski [Ecole Normale Supérieure Paris]

Technical staff
Victor Alfonso Naya [Inria, Engineer, from Dec 2019]
Berna Bakir Batu [Inria, Engineer]
Laurent Basara [Inria, Engineer]
Etienne Brame [Inria, Engineer, until Jan 2019]
Raphael Jaiswal [Inria, Engineer, until Nov 2019]
Adrien Pavao [Univ Paris-Sud, Engineer, from Nov 2019]
Jonas Renault [CNRS, Engineer]

Interns and Apprentices
Marvin Timotheus Baumann [Inria, from Jun 2019 until Jul 2019]
Moncef Belguechi [Inria, from Apr 2019 until Jul 2019]
Ritik Dutta [Inria, from May 2019 until Aug 2019]
Nathan Grinsztajn [Inria, from Apr 2019 until Sep 2019]
Remi Perrier [Univ Paris-Sud, from May 2019 until Aug 2019]
Nidina Sapkota [from Jun 2019 until Jul 2019]
Luca Veyrin Forrer [Inria, from Apr 2019 until Aug 2019]
Zhen Xu [Inria, from Apr 2019 until Sep 2019]

Administrative Assistant
Laurence Fontana [Inria]

Visiting Scientist
Burak Yelmen [from Feb 2019 until Apr 2019]

External Collaborators
Yann Ollivier [Facebook AI Research, Paris, HDR]
Olivier Goudet [Université d'Angers, HDR]
Véronique Ventos [Nukkai, Paris, HDR]
Wafa Medjitna [from Jun 2019 until Nov 2019]

2. Overall Objectives

2.1. Presentation

Since its creation in 2003, TAO activities had constantly but slowly evolved, as old problems were being solved, and new applications arose, bringing new fundamental issues to tackle. But recent abrupt progresses in Machine Learning (and in particular in Deep Learning) have greatly accelerated these changes also within the team. It so happened that this change of slope also coincided with some more practical changes in TAO ecosystem: following Inria 12-years rule, the team definitely ended in December 2016. The new team TAU (for **TA**ckling the **U**nderspecified) has been proposed, and formally created in July 2019. At the same time important staff changes took place, that also justify even sharper changes in the team focus. During the year 2018, the second year of this new era for the (remaining) members of the team, our research topics have now stabilized around a final version of the TAU project.

Following the dramatic changes in TAU staff during the years 2016-2017 (see the 2017 activity report of the team for the details), the research around continuous optimization has definitely faded out in TAU (while the research axis on hyperparameter tuning has focused on Machine Learning algorithms), the Energy application domain has slightly changed direction under Isabelle Guyon's supervision (Section 4.2), after the completion of the work started by Olivier Teytaud, and a few new directions have emerged, around the robustness of ML systems (Section 3.1.2). The other research topics have been continued, as described below.

2.2. Context and overall goal of the project

Building upon the expertise in machine learning (ML) and optimization of the TAO team, the TAU project tackles some **under-specified challenges behind the New Artificial Intelligence wave**. The simultaneous advent of massive data and massive computational power, blurring the boundaries between data, structure, knowledge and common sense, seemingly makes it possible to fulfill all promises of the good old AI, now or soon.

This makes NewAI under-specified in three respects. A first dimension regards the relationships between AIs and human beings. The necessary conditions for AI sysytems to be accepted by mankind and/or contribute to the common good are yet to be formally defined; it is hard to believe that a general and computable definition of "ethical behavior" can be set once for all. Some of these necessary conditions (explainable and causal modeling; unbiased data and models; model certification) can nevertheless be cast as ambitious and realistic goals for public research.

A second dimension regards the relationships between AI, data and knowledge. Closed world AIs can manage and acquire sufficient data to reach human-level performances from scratch [81]. In open worlds however, prior knowledge is used in various ways to overcome the lack of direct interactions with the world, e.g. through i) exploiting domain-dependent data invariances in intension or in extension (ranging from convolution to domain augmentation); ii) taking advantage of the low-rank structure (generative learning) or known properties (equivariant learning) of the observed data; iii) leveraging diverse domains and datasets, assumedly related to each other (domain adaptation; multi-task learning). A general and open question is how available prior knowledge can be best leveraged by an AI algorithm, all the more so as domains with small to medium-size data are considered.

A third dimension regards the intrinsic limitations of AI in terms of information theory. Long established theories, e.g. rooted in Occam's razor, currently hardly account for the practical leaps of deep learning, where the solution dimension outnumbers the input dimension. Beyond trials-and-errors, a long-term goal is to characterize the learning landscape w.r.t. order parameters to be defined, and *a priori* estimate the regions of problem instances where it is likely/possible/unlikely to learn accurate models.

The above under-specified AI issues define three core research pillars (Section 3), examining three interdependent aspects of AI:

I. The first pillar aims to answer the question of what 'good AI' means, and how to build it. More specifically, our goal is to advance the state of the art concerning robust learning (against adversarial attacks), causal modeling (aimed to support explanations and prescriptions), and unbiased models in the sense of prescribed neutrality constraints (including the assessment and repair of the data).

II. The second pillar tackles the "innate *vs* acquired" question: how to best combine available human knowledge, and agnostic machine learning. TAU will examine this question focusing on domains with spatial and temporal multi-scale structure, as pervasive in natural sciences (where domain knowledge is expressed using PDEs, or through powerful compact representations as in signal processing), taking advantage of the pluri-disciplinary expertise and scientific collaborations of the TAU members.

III. The third pillar aims to understand the learning landscape. In the short term, it tackles the so-called Auto-☆ issue of automatically selecting and configuring an algorithm portfolio for a problem instance. This issue governs the knowledge transfer from research labs to industry [95], [94], all the more so as massive computational resources are at stake. In the medium term, our goal is to integrate hyper-parameters and model structure in the learning criteria, using information theory and/or bilevel programming [96]. In the long-term, our goal is to establish a phase diagram of the learning landscape, through i) determining order parameters; ii) relating the different regions defined along these order parameters, to the quality of the optimal solution, and the probability of finding a good approximation thereof. These goals are aligned with the unique scientific expertise of TAU in statistical physics and in information theory, and benefit from our decade-long expertise in Auto-☆.

The above research pillars will take inspiration and be validated with three application domains (Section 4):

1.Energy management encompasses a variety of scientific problems related to research pillars I. (fair learning, privacy-compliant modelling, safety-related guarantees) and II. (spatio-temporal multi-scale modelling, distributional learning). It is also a strategic application for the planet, where TAU benefits from the TAO expertise and the long established relationships with Artelys (ILab Metis) and RTE.

2.Computational Social Sciences offer questions and methodological lessons about how to address these questions in a common decency spirit, along research pillar I. On-going studies at TAU include the learning and randomized assessment of prescriptive models for Human Resources (hiring and vocational studies; quality of life at work and economic performance) and nutrition habits (in relation with social networks and health), where i) learned models must be unbiased although data are undoubtedly biased; ii) prior knowledge must be accounted for and the interpretation of the learned models is mandatory; iii) causal modelling is key as models are deployed for *prescription* and self-fulfilling prophecies must be avoided at all costs.

3.Optimal data-driven design considers several physical/simulated phenomenons, ranging from high-energy physics to space weather, from population biology to medical imaging, from signal processing to certification of autonomous vehicle controllers, with: i) medium-size data; ii) extensive prior knowledge, notably concerning the symmetries and properties of the sought models; iii) computationally expensive simulators. All three characteristics are relevant to pillars II and III.

3. Research Program

3.1. Toward Good AI

As discussed by [141], the topic of ethical AI was non-existent until 2010, was laughed at in 2016, and became a hot topic in 2017 as the AI disruptivity with respect to the fabric of life (travel, education, entertainment, social networks, politics, to name a few) became unavoidable [138], together with its expected impacts on the nature and amount of jobs. As of now, it seems that the risk of a new AI Winter might arise from legal[1] and societal[2] issues. While privacy is now recognized as a civil right in Europe, it is feared that the GAFAM, BATX

---

[1] For instance, the (fictitious) plea challenge proposed to law students in Oct. 2018 considered a chain reaction pileup occurred among autonomous and humanly operated vehicles on a highway.
[2] For instance related to information bubbles and nudge [100], [155].

and others can already capture a sufficient fraction of human preferences and their dynamics to achieve their commercial and other goals, and build a Brave New Big Brother (BNBB, a system that is openly beneficial to many, covertly nudging, and possibly dictatorial).

The ambition of TAU is to mitigate the BNBB risk along several intricated dimensions, and build i) causal and explainable models; ii) fair data and models; iii) provably robust models.

3.1.1. *Causal modeling and biases*

Participants: Isabelle Guyon, Michèle Sebag, Philippe Caillou, Paola Tubaro
PhD: Diviyan Kalainathan
Collaboration: Olivier Goudet (Université d'Angers), David Lopez-Paz (Facebook)

The extraction of causal models, a long goal of AI [139], [117], [140], became a strategic issue as the usage of learned models gradually shifted from *prediction* to *prescription* in the last years. This evolution, following Auguste Comte's vision of science (*Savoir pour prévoir, afin de pouvoir*) indeed reflects the exuberant optimism about AI: Knowledge enables Prediction; Prediction enables Control. However, although predictive models can be based on correlations, prescriptions can only be based on causal models[3].

Among the research applications concerned with causal modeling, predictive modeling or collaborative filtering at TAU are all projects described in section 4.1 (see also Section 3.4), studying the relationships between: i) the educational background of persons and the job openings (FUI project JobAgile and DataIA project Vadore); ii) the quality of life at work and the economic performance indicators of the enterprises (ISN Lidex project Amiqap) [119] ; iii) the nutritional items bought by households (at the level of granularity of the barcode) and their health status, as approximated from their body-mass-index (IRS UPSaclay Nutriperso); iv) the actual offer of restaurants and their scores on online rating systems. In these projects, a wealth of data is available (though hardly sufficient for applications ii), iii and iv))) and there is little doubt that these data reflect the imbalances and biases of the world as is, ranging from gender to racial to economical prejudices. Preventing the learned models from perpetuating such biases is essential to deliver an AI endowed with common decency.

In some cases, the bias is known; for instance, the cohorts in the Nutriperso study are more well-off than the average French population, and the Kantar database includes explicit weights to address this bias through importance sampling. In other cases, the bias is only guessed; for instance, the companies for which Secafi data are available hardly correspond to a uniform sample as these data have been gathered upon the request of the company trade union.

3.1.2. *Robustness of Learned Models*

Participants: Guillaume Charpiat, Marc Schoenauer, Michèle Sebag
PhDs: Julien Girard, Marc Nabhan, Nizham Makhoud
Collaboration: Zakarian Chihani (CEA); Hiba Hage, and Yves Tourbier (Renault); Jérôme Kodjabachian (Thalès THERESIS)

Due to their outstanding performances, deep neural networks and more generally machine learning-based decision making systems, referred to as MLs in the following, have been raising hopes in the recent years to achieve breakthroughs in critical systems, ranging from autonomous vehicles to defense. The main pitfall for such applications lies in the lack of guarantees for MLs robustness.

---

[3]One can predict that it rains based on the presence of umbrellas in the street; but one cannot induce rainfall by going out with an umbrella. Likewise, the presence of books/tablets at home and the good scores of children at school are correlated; but offering books/tablets to all children might fail to improve their scores *per se*, if both good scores and books are explained by a so-called confounder variable, like the presence of adults versed in books/tablets at home.

Specifically, MLs are used when the mainstream software design process does not apply, that is, when no formal specification of the target software behavior is available and/or when the system is embedded in an open unpredictable world. The extensive body of knowledge developed to deliver guarantees about mainstream software − ranging from formal verification, model checking and abstract interpretation to testing, simulation and monitoring − thus does not directly apply either. Another weakness of MLs regards their dependency to the amount and quality of the training data, as their performances are sensitive to slight perturbations of the data distribution. Such perturbations can occur naturally due to domain or concept drift (e.g. due to a change in light intensity or a scratch on a camera lens); they can also result from intentional malicious attacks, a.k.a adversarial examples [156].

These downsides, currently preventing the dissemination of MLs in safety-critical systems (SCS), call for a considerable amount of research, in order to understand when and to which extent an MLs can be certified to provide the desired level of guarantees.

Julien Girard's PhD (CEA scholarship), started in Oct. 2018, co-supervised by Guillaume Charpiat and Zakaria Chihani (CEA), is devoted to the extension of abstract interpretation to deep neural nets, and the formal characterization of the transition kernel from input to output space achieved by a DNN (robustness by design, coupled with formally assessing the coverage of the training set). This approach is tightly related to the inspection and opening of black-box models, aimed to characterize the patterns in the input instances responsible for a decision – another step toward explainability.

On the other hand, experimental validation of MLs, akin statistical testing, also faces three limitations: i) real-world examples are notoriously insufficient to ensure a good coverage in general; ii) for this reason, simulated examples are extensively used; but their use raises the *reality gap* issue [128] of the distance between real and simulated worlds; iii) independently, the real-world is naturally subject to domain shift (e.g. due to the technical improvement and/or aging of sensors). Our collaborations with Renault tackle such issues in the context of the autonomous vehicle (see Section 7.1.3).

3.2. Hybridizing numerical modeling and learning systems

Participants: Alessandro Bucci, Guillaume Charpiat, Cécile Germain, Isabelle Guyon, Marc Schoenauer, Michèle Sebag
PhD: Théophile Sanchez, Loris Felardos, Wenzhuo Liu

In sciences and engineering, human knowledge is commonly expressed in closed form, through equations or mechanistic models characterizing how a natural or social phenomenon, or a physical device, will behave/evolve depending on its environment and external stimuli, under some assumptions and up to some approximations. The field of numerical engineering, and the simulators based on such mechanistic models, are at the core of most approaches to understand and analyze the world, from solid mechanics to computational fluid dynamics, from chemistry to molecular biology, from astronomy to population dynamics, from epidemiology and information propagation in social networks to economy and finance.

Most generally, numerical engineering supports the simulation, and when appropriate the optimization and control[4] of the phenomenons under study, although several sources of discrepancy might adversely affect the results, ranging from the underlying assumptions and simplifying hypotheses in the models, to systematic experiment errors to statistical measurement errors (not to mention numerical issues). This knowledge and know-how are materialized in millions of lines of code, capitalizing the expertise of academic and industrial labs. These softwares have been steadily extended over decades, modeling new and more fine-grained effects through layered extensions, making them increasingly harder to maintain, extend and master. Another difficulty is that complex systems most often resort to hybrid (pluridisciplinary) models, as they involve many components interacting along several time and space scales, hampering their numerical simulation.

---

[4]Note that the causal nature of mechanistic models is established from prior knowledge and experimentations.

At the other extreme, machine learning offers the opportunity to model phenomenons from scratch, using any available data gathered through experiments or simulations. Recent successes of machine learning in computer vision, natural language processing and games to name a few, have demonstrated the power of such agnostic approaches and their efficiency in terms of prediction [123], inverse problem solving [170], and sequential decision making [162], [81], despite their lack of any "semantic" understanding of the universe. Even before these successes, Anderson's claim was that *the data deluge [might make] the scientific method obsolete* [70], as if a reasonable option might be to throw away the existing equational or software bodies of knowledge, and let Machine Learning rediscover all models from scratch. Such a claim is hampered among others by the fact that not all domains offer a wealth of data, as any academic involved in an industrial collaboration around data has discovered.

Another approach will be considered in TAU, investigating how existing mechanistic models and related simulators can be partnered with ML algorithms: i) to achieve the same goals with the same methods with a gain of accuracy or time; ii) to achieve new goals; iii) to achieve the same goals with new methods.

Toward more robust numerical engineering: In domains where satisfying mechanistic models and simulators are available, ML can contribute to improve their accuracy or usability. A first direction is to refine or extend the models and simulators to better fit the empirical evidence. The goal is to finely account for the different biases and uncertainties attached to the available knowledge and data, distinguishing the different types of *known unknowns*. Such *known unknowns* include the model hyper-parameters (coefficients), the systematic errors due to e.g., experiment imperfections, and the statistical errors due to e.g., measurement errors. A second approach is based on learning a surrogate model for the phenomenon under study that incorporate domain knowledge from the mechanistic model (or its simulation). See Section 7.5 for case studies.

A related direction, typically when considering black-box simulators, aims to learn a model of the error, or equivalently, a post-processor of the software. The discrepancy between simulated and empirical results, referred to as *reality gap* [128], can be tackled in terms of domain adaptation [74], [99]. Specifically, the source domain here corresponds to the simulated phenomenon, offering a wealth of inexpensive data, and the target domain corresponds to the actual phenomenon, with rare and expensive data; the goal is to devise accurate target models using the source data and models.

Extending numerical engineering: ML, using both experimental and numerical data, can also be used to tackle new goals, that are beyond the current state-of-the-art of standard approaches. Inverse problems are such goals, identifying the parameters or the initial conditions of phenomenons for which the model is not differentiable, or amenable to the adjoint state method.

A slightly different kind of inverse problem is that of recovering the ground truth when only noisy data is available. This problem can be formulated as a search for the simplest model explaining the data. The question then becomes to formulate and efficiently exploit such a simplicity criterion.

Another goal can be to model the distribution of given quantiles for some system: The challenge is to exploit available data to train a generative model, aimed at sampling the target quantiles.

Examples tackled in TAU are detailed in Section 7.5. Note that the "Cracking the Glass Problem", described in Section 7.2.3 is yet another instance of a similar problem.

Data-driven numerical engineering : Finally, ML can also be used to sidestep numerical engineering limitations in terms of scalability, or to build a simulator emulating the resolution of the (unknown) mechanistic model from data, or to revisit the formal background.

When the mechanistic model is known and sufficiently accurate, it can be used to train a deep network on an arbitrary set of (space,time) samples, resulting in a meshless numerical approximation of the model [151], supporting by construction *differentiable programming* [125].

When no mechanistic model is sufficiently efficient, the model must be identified from the data only. Genetic programming has been used to identify systems of ODEs [149], through the identification of invariant quantities from data, as well as for the direct identification of control commands of nonlinear complex systems, including some chaotic systems [88]. Another recent approach uses two deep neural networks, one for the state of the system, the other for the equation itself [142]. The critical issues for both approaches include the scalability, and the explainability of the resulting models. Such line of research will benefit from TAU unique mixed expertise in Genetic Programming and Deep Learning.

Finally, in the realm of signal processing (SP), the question is whether and how deep networks can be used to revisit mainstream feature extraction based on Fourier decomposition, wavelet and scattering transforms [76]. E. Bartenlian's PhD (started Oct. 2018), co-supervised by M. Sebag and F. Pascal (Centrale-Supélec), focusing on musical audio-to-score translation [150], inspects the effects of supervised training, taking advantage from the fact that convolution masks can be initialized and analyzed in terms of frequency.

3.3. Learning to learn

According to Ali Rahimi's test of times award speech at NIPS 17, the current ML algorithms *have become a form of alchemy*. Competitive testing and empirical breakthroughs gradually become mandatory for a contribution to be acknowledged; an increasing part of the community adopts trials and errors as main scientific methodology, and theory is lagging behind practice. This style of progress is typical of technological and engineering revolutions for some; others ask for consolidated and well-understood theoretical advances, saving the time wasted in trying to build upon hardly reproducible results.

Basically, while practical achievements have often passed the expectations, there exist caveats along three dimensions. Firstly, excellent performances do not imply that the model has captured what was to learn, as shown by the phenomenon of adversarial examples. Following Ian Goodfellow, some well-performing models might be compared to *Clever Hans*, the horse that was able to solve mathematical exercizes using non verbal cues from its teacher [116]; it is the purpose of Pillar I. to alleviate the *Clever Hans* trap (section 3.1).

Secondly, some major advances, e.g. related to the celebrated adversarial learning [105], [99], establish proofs of concept more than a sound methodology, where the reproducibility is limited due to i) the computational power required for training (often beyond reach of academic labs); ii) the numerical instabilities (witnessed as random seeds happen to be found in the codes); iii) the insufficiently documented experimental settings. What works, why and when is still a matter of speculation, although better understanding the limitations of the current state of the art is acknowledged to be a priority. After Ali Rahimi again, *simple experiments, simple theorems are the building blocks that help us understand more complicated systems*. Along this line, [135] propose toy examples to demonstrate and understand the defaults of convergence of gradient descent adversarial learning.

Thirdly, and most importantly, the reported achievements rely on carefully tuned learning architectures and hyper-parameters. The sensitivity of the results to the selection and calibration of algorithms has been identified since the end 80s as a key ML bottleneck, and the field of automatic algorithm selection and calibration, referred to as AutoML or Auto-☆ in the following, is at the ML forefront.

TAU aims to contribute to the ML evolution toward a more mature stage along three dimensions. In the short term, the research done in Auto-☆ will be pursued (section 3.3.1). In the medium term, an information theoretic perspective will be adopted to capture the data structure and to calibrate the learning algorithm *depending on the nature and amount of the available data* (section 3.3.2). In the longer term, our goal is to leverage the methodologies forged in statistical physics to understand and control the trajectories of complex learning systems (section 3.3.3).

3.3.1. *Auto-**

Participants: Isabelle Guyon, Marc Schoenauer, Michèle Sebag
PhD: Guillaume Doquet, Zhengying Liu, Herilalaina Rakotoarison, Lisheng Sun
Collaboration: Olivier Bousquet, André Elisseeff (Google Zurich)

The so-called Auto-☆ task, concerned with selecting a (quasi) optimal algorithm and its hyper-parameters depending on the problem instance at hand, remained a key issue in ML for the last three decades [75], as well as in optimization at large [115], including combinatorial optimization and constraint satisfaction [122], [104] and continuous optimization [71]. This issue, tackled by several European projects along the decades, governs the knowledge transfer to industry, due to the shortage of data scientists. It becomes even more crucial as models are more complex and their training requires more computational resources. This has motivated several international challenges devoted to Auto-ML [113] (see also Section 3.4), including the AutoDL challenge series [129] launched in 2019[5] (see also Section 7.6).

Several approaches have been used to tackle Auto-☆ in the literature, and TAU has been particularly active in several of them. Meta-learning aims to build a surrogate performance model, estimating the performance of an algorithm configuration on *any* problem instance characterized from its meta-feature values [146], [104], [72], [71], [103]. Collaborative filtering, considering that a problem instance "likes better" an algorithm configuration yielding a better performance, learns to recommend good algorithms to problem instances [153], [137]. Bayesian optimization proceeds by alternatively building a surrogate model of algorithm performances on *the* problem instance at hand, and tackling it [95]. This last approach currently is the prominent one; as shown in [137], the meta-features developed for AutoML are hardly relevant, hampering both meta-learning and collaborative filtering. The design of better features is another long-term research direction, in which TAU has recently been [32], ans still is very active. more recent approach used in TAU [40] extends the Bayesian Optimization approach with a Multi-Armed Bandit algorithm to generate the full Machine Learning pipeline, competing with the famed AutoSKLearn [95] (see Section 7.2.1). These results are presented in Section 7.2.1.

3.3.2. *Information theory: adjusting model complexity and data fitting*

Participants: Guillaume Charpiat, Marc Schoenauer, Michèle Sebag
PhD: Corentin Tallec, Pierre Wolinski, Léonard Blier
Collaboration: Yann Ollivier (Facebook)

In the 60s, Kolmogorov and Solomonoff provided a well-grounded theory for building (probabilistic) models best explaining the available data [147], [108], that is, the shortest programs able to generate these data. Such programs can then be used to generate further data or to answer specific questions (interpreted as missing values in the data). Deep learning, from this viewpoint, efficiently explores a space of computation graphs, described from its hyperparameters (network structure) and parameters (weights). Network training amounts to optimizing these parameters, namely, navigating the space of computational graphs to find a network, as simple as possible, that explain the past observations well.

This vision is at the core of variational auto-encoders [121], directly optimizing a bound on the Kolmogorov complexity of the dataset. More generally variational methods provide quantitative criteria to identify super-fluous elements (edges, units) in a neural network, that can potentially be used for structural optimization of the network (Leonard Blier's PhD, started Oct. 2018).

The same principles apply to unsupervised learning, aimed to find the maximum amount of structure hidden in the data, quantified using this information-theoretic criterion.

The known invariances in the data can be exploited to guide the model design (e.g. as translation invariance leads to convolutional structures, or LSTM is shown to enforce the invariance to time affine transformations of the data sequence [157]). Scattering transforms exploit similar principles [76]. A general theory of how to detect *unknown* invariances in the data, however, is currently lacking.

The view of information theory and Kolmogorov complexity suggests that key program operations (composition, recursivity, use of predefined routines) should intervene when searching for a good computation graph. One possible framework for exploring the space of computation graphs with such operations is that of Genetic Programming. It is interesting to see that evolutionary computation appeared in the last two years among the best candidates to explore the space of deep learning structures [145], [126]. Other approaches might proceed by combining simple models into more powerful ones, e.g. using "Context Tree Weighting" [166] or switch

---

[5]https://autodl.chalearn.org/neurips2019

distributions [90]. Another option is to formulate neural architecture design as a reinforcement learning problem [73]; the value of the building blocks (predefined routines) might be defined using e.g., Monte-Carlo Tree Search. A key difficulty is the computational cost of retraining neural nets from scratch upon modifying their architecture; an option might be to use neutral initializations to support warm-restart.

3.3.3. *Analyzing and Learning Complex Systems*

Participants: Cyril Furtlehner, Aurélien Decelle, François Landes, Michèle Sebag
PhD: Giancarlo Fissore
Collaboration: Enrico Camporeale (CWI); Jacopo Rocchi (LPTMS Paris Sud), the Simons team: Rahul Chako (post-doc), Andrea Liu (UPenn), David Reichman (Columbia), Giulio Biroli (ENS), Olivier Dauchot (ESPCI), Hufei Han (Symantec).

Methods and criteria from statistical physics have been widely used in ML. In early days, the capacity of Hopfield networks (associative memories defined by the attractors of an energy function) was investigated by using the replica formalism [69]. Restricted Boltzmann machines likewise define a generative model built upon an energy function trained from the data. Along the same lines, Variational Auto-Encoders can be interpreted as systems relating the free energy of the distribution, the information about the data and the entropy (the degree of ignorance about the micro-states of the system) [165]. A key promise of the statistical physics perspective and the Bayesian view of deep learning is to harness the tremendous growth of the model size (billions of weights in recent machine translation netwowrks), and make them sustainable through e.g. posterior drop-out [136], weight quantization and probabilistic binary networks [131]. Such "informational cooling" of a trained deep network can reduce its size by several orders of magnitude while preserving its performance.

Statistical physics is among the key expertises of TAU, originally only represented by Cyril Furtlehner, later strenghtened by Aurélien Decelle's and François Landes' arrivals in 2014 and 2018. On-going studies are conducted along several directions.

Generative models are most often expressed in terms of a Gibbs distributions $P[S] = \exp(-E[S])$, where energy E involves a sum of building blocks, modelling the interactions among variables. This formalization makes it natural to use mean-field methods of statistical physics and associated inference algorithms to both train and exploit such models. The difficulty is to find a good trade-off between the richness of the structure and the efficiency of mean-field approaches. One direction of research pursued in TAU, [97] in the context of traffic forecasting, is to account for the presence of cycles in the interaction graph, to adapt inference algorithms to such graphs with cycles, while constraining graphs to remain compatible with mean-field inference.

Another direction, explored in TAO/TAU in the recent years, is based on the definition and exploitation of self-consistency properties, enforcing principled divide-and-conquer resolutions. In the particular case of the message-passing Affinity Propagation algorithm for instance [168], self-consistency imposes the invariance of the solution when handled at different scales, thus enabling to characterize the critical value of the penalty and other hyper-parameters in closed form (in the case of simple data distributions) or empirically otherwise [98].

A more recent research direction examines the quantity of information in a (deep) neural net along the random matrix theory framework [78]. It is addressed in Giancarlo Fissore's PhD, and is detailed in Section 7.2.3.

Finally, we note the recent surge in using ML to address fundamental physics problems: from turbulence to high-energy physics and soft matter as well (with amorphous materials at its core) [19]. TAU's dual expertise in Deep Networks and in statistical physics places it in an ideal position to significantly contribute to this domain and shape the methods that will be used by the physics community in the future. François Landes' recent arrival in the team makes TAU a unique place for such interdisciplinary research, thanks to his collaborators from the Simons Collaboration *Cracking the Glass Problem* (gathering 13 statistical physics teams at the international level). This project is detailed in Section 7.2.3.

Independently, François Landes is actively collaborating with statistical physicists (Alberto Rosso, LPTMS, Univ. Paris-Saclay) and physcists at the frontier with geophysics (Eugenio Lippiello, Second Univ. of Naples) [20]. A possible CNRS grant (80Prime) may finance a shared PhD, at the frontier between seismicity and ML (Alberto Rosso, Marc Schoenauer and François Landes).

3.4. Organisation of Challenges

Participants: Cécile Germain, Isabelle Guyon, Marc Schoenauer, Michèle Sebag

Challenges have been an important drive for Machine Learning research for many years, and TAO members have played important roles in the organization of many such challenges: Michèle Sebag was head of the challenge programme in the Pascal European Network of Excellence (2005-2013); Isabelle Guyon, as mentioned, was the PI of many challenges ranging from causation challenges [109], to AutoML [110]. The Higgs challenge [67], most attended ever Kaggle challenge, was jointly organized by TAO (C. Germain), LAL-IN2P3 (D. Rousseau and B. Kegl) and I. Guyon (not yet at TAO), in collaboration with CERN and Imperial College.

TAU was also particularly implicated with the ChaLearn Looking At People (LAP) challenge series in Computer Vision, in collaboration with the University of Barcelona [92] including the Job Candidate Screening Coopetition [91]; the Real Versus Fake Expressed Emotion Challenge (ICCV 2017) [163]; the Large-scale Continuous Gesture Recognition Challenge (ICCV 2017) [163]; the Large-scale Isolated Gesture Recognition Challenge (ICCV 2017) [163].

Other challenges have been organized in 2019, or are planned for the near future, detailed in Section 7.6. In particular, many of them now run on the Codalab platform, managed by Isabelle Guyon and maintained at LRI.

4. Application Domains

4.1. Computational Social Sciences

Participants: Philippe Caillou, Isabelle Guyon, Michèle Sebag, Paola Tubaro
Collaboration: Jean-Pierre Nadal (EHESS); Marco Cuturi, Bruno Crépon (ENSAE); Thierry Weil (Mines); Jean-Luc Bazet (RITM)

Computational Social Sciences (CSS) studies social and economic phenomena, ranging from technological innovation to politics, from media to social networks, from human resources to education, from inequalities to health. It combines perspectives from different scientific disciplines, building upon the tradition of computer simulation and modeling of complex social systems [102] on the one hand, and data science on the other hand, fueled by the capacity to collect and analyze massive amounts of digital data.

The emerging field of CSS raises formidable challenges along three dimensions. Firstly, the definition of the research questions, the formulation of hypotheses and the validation of the results require a tight pluridisciplinary interaction and dialogue between researchers from different backgrounds. Secondly, the development of CSS is a touchstone for ethical AI. On the one hand, CSS gains ground in major, data-rich private companies; on the other hand, public researchers around the world are engaging in an effort to use it for the benefit of society as a whole [124]. The key technical difficulties related to data and model biases, and to self-fulfilling prophecies have been discussed in section 3.1. Thirdly, CSS does not only regard scientists: it is essential that the civil society participate in the science of society [152].

TAO was involved in CSS for the last five years, and its activities have been strengthened thanks to P. Tubaro's and I. Guyon's expertises respectively in sociology and economics, and in causal modeling. Details are given in Section 7.3.

4.2. Energy Management

Participants: Isabelle Guyon, Marc Schoenauer, Michèle Sebag
PhD: Victor Berger, Benjamin Donnot, Balthazar Donon, Herilalaina Rakotoarison
Collaboration: Antoine Marot, Patrick Panciatici (RTE), Vincent Renault (Artelys)

Energy Management has been an application domain of choice for TAO since the end 2000s, with main partners SME Artelys (METIS Ilab Inria; ADEME project POST; on-going ADEME project NEXT), RTE (See.4C European challenge; two CIFRE PhDs), and, since Oct. 2019, IFPEN. The goals concern i) optimal planning over several spatio-temporal scales, from investments on continental Europe/North Africa grid at the decade scale (POST), to daily planning of local or regional power networks (NEXT); ii) monitoring and control of the French grid enforcing the prevention of power breaks (RTE); iii) improvement of house-made numerical methods using data-intense learning in all aspects of IFPEN activities (as described in Section 3.2).

Optimal planning over long periods of time amounts to optimal sequential decision under high uncertainties, ranging from stochastic uncertainties (weather, market prices, demand prediction) handled based on massive data, to non-stochastic uncertainties (e.g., political decisions about the nuclear policy) handled through defining and selecting a tractable number of scenarios. Note that non-anticipativity constraints forbid the use of dynamic programming-related methods; this led to propose the *Direct Value Search* method [77] at the end of the POST project.

The daily maintainance of power grids requires the building of approximate predictive models on the top of any given network topology. Deep Networks are natural candidates for such modelling, considering the size of the French grid ($\sim$ 10000 nodes), but the representation of the topology is a challenge when, e.g. the RTE goal is to quickly ensure the "n-1" security constraint (the network should remain safe even if any of the 10000 nodes fails). Existing simulators are too slow to be used in real time, and the size of actual grids makes it intractable to train surrogate models for all possible (n-1) topologies (see Section 7.4 for more details).

Furthermore, predictive models of local grids are based on the estimated consumption of end-customers: Linky meters only provide coarse grain information due to privacy issues, and very few samples of fine-grained consumption are available (from volunteer customers). A first task is to transfer knowledge from small data to the whole domain of application. A second task is to directly predict the peaks of consumption based on the user cluster profiles and their representativity (see Section 7.4.2).

4.3. Data-driven Numerical Modeling

Participants: Alessandro Bucci, Guillaume Charpiat, Cécile Germain, Isabelle Guyon, Flora Jay, Marc Schoenauer, Michèle Sebag
PhD and Post-doc: Victor Estrade, Loris Felardos, Adrian Pol, Théophile Sanchez, Wenzhuo Liu
Collaboration: D. Rousseau (LAL), M. Pierini (CERN)

As said (section 3.2), in domains where both first principle-based models and equations, and empirical or simulated data are available, their combined usage can support more accurate modelling and prediction, and when appropriate, optimization, control and design. This section describes such applications, with the goal of improving the time-to-design chain through fast interactions between the simulation, optimization, control and design stages. The expected advances regard: i) the quality of the models or simulators (through data assimilation, e.g. coupling first principles and data, or repairing/extending closed-form models); ii) the exploitation of data derived from different distributions and/or related phenomenons; and, most interestingly, iii) the task of optimal design and the assessment of the resulting designs.

The proposed approaches are based on generative and adversarial modelling [121], [106], extending both the generator and the discriminator modules to take advantage of the domain knowledge.

A first challenge regards the design of the model space, and the architecture used to enforce the known domain properties (symmetries, invariance operators, temporal structures). When appropriate, data from different distributions (e.g. simulated vs real-world data) will be reconciled, for instance taking inspiration from real-valued non-volume preserving transformations [85] in order to preserve the natural interpretation.

Another challenge regards the validation of the models and solutions of the optimal design problems. The more flexible the models, the more intensive the validation must be, as reminded by Leon Bottou. Along this way, generative models will be used to support the design of "what if" scenarios, to enhance anomaly detection and monitoring via refined likelihood criteria.

In the application case of dynamical systems such as fluid mechanics, the goal of incorporating machine learning into classical simulators is to speed up the simulations. Many possible tracks are possible for this; for instance one can search to provide better initialization heuristics to solvers (which make sure that physical constraints are satisfied, and which are responsible of most of the computational complexity of simulations) at each time step; one can also aim at predicting directly the state at $t + 100$, for instance, or at learning a representation space where the dynamics are linear (Koopman - von Neumann). The topic is very active in the deep learning community. To guarantee the quality of the predictions, concepts such as Liapunov coefficients (which express the speed at which simulated trajectories diverge from the true ones) can provide a suitable theoretical framework.

5. Highlights of the Year

5.1. Highlights of the Year

5.1.1. *Awards*

- Best Paper Award in Machine Learning at ECML-PKDD 2019 in Wurzburg to Guillaume Doquet and Michèle Sebag for their paper *Agnostic feature selection* . See Guillaume's PhD [13] for more details.
- Nacim Belkhir, Winner ACM-GECCO 2019 **BBComp single-objective**, **BBComp two-objective** and **three-objective** tracks. The winning program is a slighly modified version of the one Nacim wrote during his PhD in TAU in 2017 [71], co-supervised by Marc Schoenauer, Johann Dréo and Pierre Savéant (Thalès TRT).

5.1.2. *Visibility*

- Marc Schoenauer, expert seconding Guillaume Klossa, special advisor to European Vice-President Ansip, for the report *Toward European Media Sovereignty* giving strategic advice on the opportunities and challenges related to the use artificial intelligence, with a focus on the media sector.

BEST PAPER AWARD:

[32]

G. F. DOQUET, M. SEBAG. *Agnostic feature selection*, in "ECML PKDD 2019 - European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases", Würzburg, Germany, September 2019, https://hal.archives-ouvertes.fr/hal-02436824

6. New Software and Platforms

6.1. io.datascience

Input Output Data Science

KEYWORDS: Open data - Semantic Web - FAIR (Findable, Accessible, Interoperable, and Reusable)

FUNCTIONAL DESCRIPTION: io.datascience (Input Output Data Science) is the instance of the Linked Wiki platform developed specifically in Paris-Saclay University as part of its Center for Data Science.

The goal of io.datascience: to facilitate the sharing and use of scientific data. The technological concept of io.datascience: the exploitation of semantic web advances, and in particular wiki technologies.

(Findable, Accessible, Interoperable, and Reusable) (Wilkinson, M., and The FAIR Guiding Principles for Scientific Data Management and Stewardship, Nature Scientific Data 2016)

io.datascience is both a data sharing platform and a framework for further development. It realizes a practical implementation of FAIR (Findable, Accessible, Interoperable, and Reusable - Wilkinson, M., Nature Scientific Data 2016) principles through a user-centric approach.

- Partners: Border Cloud - Paris Saclay Center for Data Science - Université Paris-Sud

- Contact: Cécile Germain

- Publications: Data acquisition for analytical platforms: Automating scientific workflows and building an open database platform for chemical anlysis metadata - A platform for scientific data sharing - TFT, Tests For Triplestores : Certifying the interoperability of RDF database systems using a continuous delivery workflow - Une autocomplétion générique de SPARQL dans un contexte multi-services - Certifying the interoperability of RDF database systems - Transforming Wikipedia into an Ontology-based Information Retrieval Search Engine for Local Experts using a Third-Party Taxonomy - The Grid Observatory 3.0 - Towards reproducible research and open collaborations using semantic technologies

- URL: https://io.datascience-paris-saclay.fr/

6.2. Codalab

KEYWORDS: Benchmarking - Competition

FUNCTIONAL DESCRIPTION: Challenges in machine learning and data science are competitions running over several weeks or months to resolve problems using provided datasets or simulated environments. Challenges can be thought of as crowdsourcing, benchmarking, and communication tools. They have been used for decades to test and compare competing solutions in machine learning in a fair and controlled way, to eliminate "inventor-evaluator" bias, and to stimulate the scientific community while promoting reproducible science. See our slide presentation.

As of June 2019 Codalab exceeded 40,000 users, 1000 competitions (300 public), and had over 300 submissions per day. Some of the areas in which Codalab is used include Computer vision and medical image analysis, natural language processing, time series prediction, causality, and automatic machine learning. Codalab was selected by the Région Ile de France to organize its challenges in the next three years.

TAU is going to continue expanding Codalab to accommodate new needs. One of our current focus is to support use of challenges for teaching (i.e. include a grading system as part of Codalab) and support for hooking up data simulation engines in the backend of Codalab to enable Reinforcement Learning challenges and simulate interactions of machines with an environment. For the fith year, we are using Codalab for student projects. M2 AIC students create mini data science challenges in teams of 6 students. L2 math and informatics students then solve them as part of their mini projects. We are collaborating with RPI (New York, USA) and Université de Grenoble to use this platform as part of a curriculum of medical students. We created a special application called ChaGrade to grade homework using challenges. Our PhD. students are involved in co-organizing challenges to expose the research community at large with the topic of their PhD. This helps them formalizing a task with rigor and allows them to disseminate their research.

- Contact: Isabelle Guyon

- URL: http://competitions.codalab.org

6.3. Cartolabe

KEYWORD: Information visualization

FUNCTIONAL DESCRIPTION: The goal of Cartolabe is to build a visual map representing the scientific activity of an institution/university/domain from published articles and reports. Using the HAL Database, Cartolabe provides the user with a map of the thematics, authors and articles . ML techniques are used for dimensionality reduction, cluster and topics identification, visualisation techniques are used for a scalable 2D representation of the results.

NEWS OF THE YEAR: This year, Cartolabe was applied to the Grand Debat dataset (3M individual propositions from french Citizen, see https://cartolabe.fr/map/debat). The results were used to test both the scaling capabilities of Cartolabe and its flexibility to non-scientific and non-english corpuses. We also Added submap capabilities to display the result of a year/lab/word filtering as an online generated heatmap with only the filtered points to facilitate the exploration.

- Participants: Philippe Caillou, Jean-Daniel Fekete, Jonas Renault and Anne-Catherine Letournel
- Partners: LRI - Laboratoire de Recherche en Informatique - CNRS
- Contact: Philippe Caillou
- URL: http://www.cartolabe.fr/

7. New Results

7.1. Toward Good AI

7.1.1. Causal Modeling

Participants: Philippe Caillou, Isabelle Guyon, Michèle Sebag;
PhDs: Diviyan Kalainathan
Collaboration: David Lopez-Paz (Facebook).

The search for **causal models** relies on quite a few hardly testable assumptions, e.g. causal sufficiency [160]; it is a data hungry task as it has the identification of independent and conditionally independent pairs of variables at its core. A new approach investigated through the Cause-Effects Pairs (CEP) Challenge [112] formulates causality search as a supervised learning problem, considering the joint distributions of pairs of variables (e.g. (Age, Salary)) labelled with the proper causation relationship between both variables (e.g. Age "causes" Salary) and learning algorithms apt to learn from distributions have been proposed [114]. An edited book has been published [48], that somewhat summarizes the whole history of Cause-EFfect Paris research. Several chapters of this book have co-authors in TAU: *Evaluation methods of cause-effect pairs* [49], *Learning Bivariate Functional Causal Models* [47], *Discriminant Learning Machines* [51], and *Results of the Cause-Effect Pair Challenge* [50].

In D. Kalainathan's PhD [14] and O. Goudet's postdoc, the search for causal models has been tackled in the framework of generative networks [107], trained to minimize the Maximum Mean Discrepancy loss; the resulting Causal Generative Neural Network improves on the state of the art on the CEP Challenge. CGNN favorably compares with the state of the art w.r.t. usual performance indicators (AUPR, SID) on main causal benchmarks, though with a large computational cost.

An attempt to scale up causal discovery, we proposed the Structural Agnostic Model approach [14] [120]. Working directly on the observational data, this global approach implements a variant of the popular adversarial game [99] between a discriminator, attempting to distinguish actual samples from fake ones, obtained by generating each variable, given real values from all others. A sparsity L_1 penalty forces all generators to consider only a small subset of their input variables, yielding a sparse causal graph. SAM obtains state-of-the-art performances on causal benchmarks, and scales up to a few hundred variables.

An innovative usage of causal models is for educational training in sensitive domains, such as medicine, along the following line. Given a causal generative model, artificial data can be generated using a marginal distribution of causes; such data will enable students to test their diagnosis inference (with no misleading spurious correlations in principle), while forbidding to reverse-engineer the artificial data and guess the original data. Some motivating applications for causal modeling are described in section 4.1.

7.1.2. Explainability

Participants: Isabelle Guyon, François Landes, Marc Schoenauer, Michèle Sebag
PhD: Marc Nabhan

Causal modeling is one particular method to tackle explainability, and TAU has been involved in other initiatives toward explainable AI systems. Following the LAP (Looking At People) challenges, Isabelle Guyon and co-organizers have edited a book [143] that presents a snapshot of explainable and interpretable models in the context of computer vision and machine learning. Along the same line, they propose an introduction and a complete survey of the state-of-the-art of the explainability and interpretability mechanisms in the context of first impressions analysis [57].

The team is also involved in the proposal for the IPL HyAIAI (Hybrid Approaches for Interpretable AI), coordinated by the LACODAM team (Rennes) dedicated to the design of hybrid approaches that combine state of the art numeric models (e.g., deep neural networks) with explainable symbolic models, in order to be able to integrate high level (domain) constraints in ML models, to give model designers information on ill-performing parts of the model, to provide understandable explanations on its results. Kickoff took place in September 2019, and we are still looking for good post-doc candidates.

Note also that the on-going work on the identification of the border of the failure zone in the parameter space of the autonomous vehicle simulator [37] (Section 7.1.3) also pertains to explainability.

Finally, a completely original approach to DNN explainability might arise from the study of structural glasses (7.2.3), with a parallel to Graph Neural Networks (GNNs), that could become an excellent non-trivial example for developing explainability protocols.

7.1.3. Robustness of AI Systems

Participants: Guillaume Charpiat, Marc Schoenauer, Michèle Sebag
PhDs: Julien Girard, Marc Nabhan, Nizham Makhoud
Collaboration: Zakarian Chihani (CEA); Hiba Hage and Yves Tourbier (Renault); Johanne Cohen (LRI-GALAC) and Christophe Labreuche (Thalès)

As said (Section 3.1.2, TAU is considering two directions of research related to the certification of MLs. The first direction, related to formal approaches, is the topic of Julien Girard's PhD (see also Section 3.1.2). On the opposite, the second axis aims to increase the robustness of systems that can only be experimentally validated. Two paths are investigated in the team: assessing the coverage of the datasets (more particularly here, used to train an autonomous vehicle controller), topic of Marc Nabhan's CIFRE with Renault; and detecting flaws in the system by reinforcement learning, as done by Nizam Makdoud's CIFRE PhD with Thalès THERESIS.
Formal validation of Neural Networks
The topic of provable deep neural network robustness has raised considerable interest in recent years. Most research in the literature has focused on adversarial robustness, which studies the robustness of perceptive models in the neighbourhood of particular samples. However, other works have proved global properties of smaller neural networks. Yet, formally verifying perception remains uncharted. This is due notably to the lack of relevant properties to verify, as the distribution of possible inputs cannot be formally specified. With Julien Girard-Satabin's PhD thesis, we propose to take advantage of the simulators often used either to train machine learning models or to check them with statistical tests, a growing trend in industry. Our formulation [34] allows us to formally express and verify safety properties on perception units, covering all cases that could ever be generated by the simulator, to the difference of statistical tests which cover only seen examples. Along with this theoretical formulation, we provide a tool to translate deep learning models into standard logical formulae. As a proof of concept, we train a toy example mimicking an autonomous car perceptive unit, and we formally verify that it will never fail to capture the relevant information in the provided inputs.
Experimental validation of Autonomous Vehicle Command Statistical guarantees (e.g., less than 10^{-8} failure per hour of operation) are obtained by empirical tests, involving millions of kilometers of driving in all possible road, weather and traffic conditions as well as intensive simulations, the only way to full control of the driving conditions. The validation process thus involves 3 steps: i) making sure that all parts of the space of possible scenarios are covered by experiments/tests with sufficiently fine grain; ii) identify failures zones in the space of scenarios; iii) fix the controller flaws that resulted in these failures.

TAU is collaborating with Renault on step ii) within Marc Nabhan's CIFRE PhD (defense expected in Sept. 2020). The current target scenario is the insertion of a car on a motorway, the "drosophila" of autonomous car scenarios and the goal is the identification of the conditions of failures of the autonomous car controller. Only simulations are considered here, with one scenario being defined as a parameter setting of the in-house simulator SCANeR. The goal is the detection of as many failures as possible, running as few simulations as possible, and the identification of the borders of the failure zone using an as simple as possible description, thus allowing engineers to understand the reasons for the flaws. A first paper was published [37] proposing several approaches for the identification of failures. On-going work is concerned with a precise yet simple definition of the border of the failure zone.

Reinforcement Learning from Advice In the context of his CIFRE PhD with Thalès, Nizam Makdoud tests (in simulation) physical security systems using reinforcement learning to learn the best sequence of action that will break through the system. This lead him to propose an original approach called *LEarning from Advice* (LEA) that uses knowledge from several policies learned on different tasks. Whereas Learning by imitation uses the actions of the known policy, the proposed method uses the different Q-functions of the known policies. The main advantage of this strategy is its robustness to poor advice, as the policy then reverts to standard DDPG [127]. The results (submitted) demonstrate that LEA is able to learn faster than DDPG if given good-enough policies, and only slightly slower when given lousy advices.

Learning Multi-Criteria Decision Aids (Hierarchical Choquet models) In collaboration with Johanne Cohen (LRI-GALAC) and Christophe Labreuche (Thalès), the representation and data-driven elicitation of hierarchical Choquet models has been tackled. A specific neural architecture, enforcing by design the model constraints (monotonicity, additivity), and supporting the end-to-end training of the Multi-Criteria Decision aid, has been proposed in Roman Bresson's PhD. Under mild assumptions, an identifiability result (existence and unicity of the sought model in the neural space) is obtained. The approach is empirically validated and successfully compared to the state of the art.

7.2. Learning to Learn

7.2.1. *Auto-**

Participants: Guillaume Charpiat, Isabelle Guyon, Marc Schoenauer, Michèle Sebag
PhDs: Léonard Blier, Guillaume Doquet, Zhengying Liu, Herilalaina Rakotoarison, Lisheng Sun, Pierre Wolinski
Collaboration: Vincent Renault (SME Artelys); Yann Ollivier (Facebook)

Auto-☆ studies at TAU investigate several research directions.

As mentioned in Section 3.3, a popular approach for algorithm selection is collaborative filtering. In Lisheng Sun's PhD [15], active learning was used on top of the CofiRank algorithm for matrix factorization [164], improving the results and the time to solution of the recommendation algorithm. Furthermore, most real-world domains evolve with time, and an important issue in real-world applications is that of life-long learning, as static models can rapidly become obsolete. Another contribution in Lisheng's PhD is an extension of AutoSklearn that detects concept drifts and corrects the current model accordingly.

An original approach to Auto-☆, explored in Herilalaina Rakotoarison's PhD, extends and adapts Monte-Carlo Tree Search to explore the structured space of pre-processing + learning algorithm configurations, and gradually determine the best pipeline [40]; the resulting MOSAIC algorithm performs on par with AutoSklearn, the winner of Auto-☆ international competitions in the last few years.

Auto-☆ would be much easier if appropriate and affordable meta-features (describing datasets) were available. Taking inspiration from equivariant learning [144] and learning from distributions [130], on-going work aims to learn such meta-features, based on the OpenML archive [161].

A key building block in Auto-☆ lies in the data preparation and specifically variable selection. Guillaume Doquet's PhD [13] addresses the problem of agnostic feature selection, independently of any target variable. The point d'orgue of his work is Agnos [32] (Best Paper Award at ECML 2019), that combines an AutoEncoder with structural regularizations to sidestep the combinatorial optimization problem at the core of feature selection. The extensive experimental validation of AgnoS on the scikit-feature benchmark suite demonstrates its ability compared to the state of the art, both in terms of supervised learning and data compression.

Several works have focused on the adjustment of specific hyper-parameters for neural nets. Pierre Wolinski's PhD (to be defended in January 2020, publication submitted) studies three such hyper-parameters: i) network width (number of neurons in each layer); ii) regularizer importance in the objective function to minimize (factor balancing data term and regularizer); and iii) learning rate. Regarding the network width, it is adjusted during training thanks to a criterion quantifying each neuron's importance, naturally leading to a sparsification effect (as for L1 norm minimization). This study is actually extendable to not only layers' widths but also layers' connectivity (e.g., in modern networks where each layer may be connected to any other layer with 'skip' connections). Regarding the regularizer weight, it is formulated as a probabilistic prior from a Bayesian perspective, which leads to a particular value that the regularizer weight should have in order the network to satisfy some property. Regarding the learning rate, Pierre Wolinski and Leonard Blier [27] proposed to attach fixed learning rates to each neuron (picked randomly) and calibrate this learning rate distribution in such a way that neurons are sequentially active, learning in an optimally agile manner during a first learning phase, and being stable in later phases. This remove the need to tune the learning rate.

A last direction of investigation concerns the design of challenges, that contribute to the collective advance of research in the Auto-☆ direction. The team has been very active in the series of AutoML challenges [154], and steadily contributes to the organization of new challenges (Section 7.6).

7.2.2. *Deep Learning: Practical and Theoretical Insights*

Participants: Guillaume Charpiat, Marc Schoenauer, Michèle Sebag
PhDs: Léonard Blier, Corentin Tallec
Collaboration: Yann Ollivier (Facebook AI Research, Paris), the Altschuler and Wu lab. (UCSF, USA), Y. Tarabalka (Inria Titane)

Although a comprehensive mathematical theory of deep learning is yet to come, theoretical insights from information theory or from dynamical systems can deliver principled improvements to deep learning and/or explain the empirical successes of some architectures compared to others.

In his PhD [16], Corentin Tallec presents several contributions along these lines:

- In [158], it is shown that the LSTM structure can be understood from axiomatic principles, enforcing the *robustness of the learned model to temporal deformation (warpings) in the data*. The complex LSTM architecture, introduced in the 90's, has become the currently dominant architecture for modeling temporal sequences (such as text) in deep learning. It is shown that this complex LSTM architecture necessarily arises if one wants the model to be able to handle time warpings in the data (such as arbitrary accelerations or decelerations in the signal), and their complex equations can be derived axiomatically.

- In [132] (oral presentation at ICML 2018) the issue of mode dropping in adversarial generative models is tackled using information theory. The adversary (discriminator) task is set to predict the proportion of true and fake images in a set of images, via an information theory criterion, thus working at the level of the *overall distribution* of images. The discriminator is thus mande more able to detect statistical imbalances between the modes created by the generator, thereby reducing the mode dopping phenomenon. The proposed architecture, inspired from equivariant approaches, is provably able to detect all permutation-invariant statistics in a set of images.

- In [159], the problem of recurrent network training is tackled via the theory of dynamical systems by proposing a simple fully online solution avoiding the "time rewind" step, based on real-time, noisy but unbiased approximations of model gradients, which can be implemented easily in a black-box fashion on top of any recurrent model, and which is well-justified mathematically. The price to pay is an increase of variance.

- In [42], we identify sensitivity to time discretization of Deep RL in near continuous-time environments as a critical factor. Empirically, we find that Q-learning-based approaches collapse with small time steps. Formally, we prove that Q-learning does not exist in continuous time. We detail a principled way to build an off-policy RL algorithm that yields similar performances over a wide range of time discretizations, and confirm this robustness empirically.

Several other directions have been investigated: In [41], we introduce a multi-domain adversarial learning algorithm in the semi-supervised setting. We extend the single source H-divergence theory for domain adaptation to the case of multiple domains, and obtain bounds on the average- and worst-domain risk in multi-domain learning. This leads to a new loss to accommodate semi-supervised multi-domain learning and domain adaptation. We obtain state-of-the-art results on two standard image benchmarks, and propose as a new benchmark a novel bioimage dataset, CELL, in the domain of automated microscopy data, where cultured cells are imaged after being exposed to known and unknown chemical perturbations, and in which each dataset displays significant experimental bias.

Another direction regards the topology induced by a trained neural net, and how similar are two samples in the NN perspective. The definition proposed in [29] relies on varying the NN parameters and examining whether the impacts of this variation on both samples are aligned. The mathematical properties of this similarity measure are investigated and the similarity is shown to define a kernel on the input space. This kernel can be used to tractably estimate the sample density, and it leads to new directions for the statistical learning analysis of NN, e.g. in terms of additional loss (requiring that similar examples have a similar latent representation in the above sense) or in terms of resistance to noise. Specifically, a multimodal image registration task is presented where almost perfect accuracy is reached, despite a high label noise (see Section 7.5.2). Such an impressive self-denoising phenomenon can be explained and quantified as a noise averaging effect over the labels of similar examples.

7.2.3. Analyzing and Learning Complex Systems

Participants: Cyril Furtlehner, Aurélien Decelle, François Landes
PhDs: Giancarlo Fissore
Collaboration: Jacopo Rocchi (LPTMS Paris Sud); the Simons team: Rahul Chako (post-doc), Andrea Liu (UPenn), David Reichman (Columbia), Giulio Biroli (ENS), Olivier Dauchot (ESPCI).; Clément Vignax (EPFL); Yufei Han (Symantec).

The information content of a trained restricted Boltzmann machine (RBM) can be analyzed by comparing the singular values/vectors of its weight matrix, referred to as modes, to that of a random RBM (typically following a Marchenko-Pastur distribution) [83]. The analysis of a single learning trajectory is replaced by analyzing the distribution of a well chosen ensemble of models. In G. Fissore's PhD, the learning trajectory of an RBM is shown to start with a linear phase recovering the dominant modes of the data, followed by a non-linear regime where the interaction among the modes is characterized [84]. Although simplifying assumptions are required for a mean-field analysis in closed form of the above distribution, it nevertheless delivers some simple heuristics to speed up the learning convergence and to simplify the models.

This analysis will be extended along two directions: handling missing data [58]; and considering exactly solvable RBM (non-linear RBM for which the contrastive divergence can be computed in closed forms, e.g. using a spherical model) [54]. W.r.t. missing data, state of the art results have been obtained on semi-supervised tasks in the context of Internet-of-Things security, considering a high rate of missing inputs and labels). On the theoretical side exact generic RBM learning trajectories have been characterized, showing intriguing connections based on Bose-Einstein condensation mechanism associated to information storing. Our collaboration with J. Rocchi(LPTMS, Univ. Paris Sud) aims to characterize the landscape of RBMs learned from different initial conditions, and to relate this landscape to the number of parameters (hidden nodes) of the system.

An emerging research topic, concerns the interpretation of deep learning by means of Gaussian processes and associated neural tangent limit kernel in the thermodynamical limit obtained by letting layer's width go to infinity [118]. Various things are planned to be investigated on the basis of this theoretical tool in particular how this translate to RBM or DBM setting, and whether a double dip behaviour is to be expected as well for generative models.

As mentioned earlier, the use of ML to address fundamental physics problems is quickly growing. This leads to some methodological mistakes from newcomers, that have been investigated by Rémi Perrier (2 month internship). One example is the domain of glasses (how the structure of glasses is related to their dynamics), which is one of the major problems in modern theoretical physics. The idea is to let ML models automatically find the hidden structures (features) that control the flowing or non-flowing state of matter, discriminating liquid from solid states. These models can then help identifying "computational order parameters", that would advance the understanding of physical phenomena [19], on the one hand, and support the development of more complex models, on the other hand. More generally, attacking the problem of amorphous condensed matter by novel Graph Neural Networks (GNN) architectures is a very promising lead, regardless of the precise quantity one may want to predict. Currently GNNs are engineered to deal with molecular systems and/or crystals, but not to deal with amorphous matter. This second axis is currently being attacked in collaboration with Clément Vignac (PhD Student at EPFL), using GNNs. Furthermore, this problem is new to the ML community and it provides an original non-trivial example for engineering, testing and benchmarking explainability protocols.

7.3. Computational Social Sciences

Computational Social Sciences (CSS) is making significant progress in the study of social and economic phenomena thank to the combination of social science theories and new insight from data science. While the simultaneous advent of massive data and unprecedented computational power has opened exciting new avenues, it has also raised new questions and challenges.

Several studies are being conducted in TAU, about labor (labor markets, platform "micro-work", quality of life and economic performance), about nutrition (health, food, and socio-demographic issues), around Cartolabe, a platform for scientific information system and visual querying and around GAMA, a multi-agent based simulation platform.

7.3.1. Labor Studies

Participants: Philippe Caillou, Isabelle Guyon, Michèle Sebag, Paola Tubaro
PhDs: Diviyan Kalainathan, Guillaume Bied, Armand Lacombe
Post-Docs: Saumya Jetley
Engineers: Raphael Jaiswal, Victor Alfonso Naya
Collaboration: Jean-Pierre Nadal (EHESS); Marco Cuturi, Bruno Crépon (ENSAE); Antonio Casilli, Ulrich Laitenberger (Telecom Paris); Odile Chagny (IRES); Alessandro Delfanti (University of Toronto)

A first area of activity of TAU in Computational Social Sciences is the study of labor, from the functioning of the job market, to the rise of new, atypical forms of work in the networked society of internet platforms, and the quality of life at work.

Job markets Two projects deal with the domain of job markets and machine learning. The DATAIA project Vadore, in collaboration with ENSAE and Pôle Emploi, has two goals. First, to improve the recommendation of jobs for applicants (and the recommendation of applicants to job offers). The main originalities in this project are: i) to use both machine learning and optimal transport to improve the recommendation by learning a matching function for past hiring, and then to apply optimal transport-like bias to tackle market congestion (e.g. to avoid assigning many applicants to a same job offer); ii) to use randomized test on micro-markets (AB testing) in collaboration with Pôle Emploi to test the global impact of the algorithms.

The JobAgile project, BPI-PIA contract, coll. EHESS, Dataiku and Qapa, deals with low salary interim job recommendations. A main difference with the Vadore project relies on the high reactivity of the Qapa and Dataiku startups: i) to actually implement AB-testing; ii) to explore related functionalities, typically the recommendation of formations; iii) to propose a visual querying of the job market, using the Cartolabe framework (below).

The platform economy and digital labor

Another topic concerns the digital economy and the transformations of labor that accompany it. One part of the platform economy carries promises of social, not only techno-economic, innovation. If enthusiasms for a new "sharing economy" or "collaborative economy" have progressively faded away, values of decentralization, autonomy, and flatter coordination are still commonly associated to platforms. A conference paper by P. Tubaro studies how events constitute places where actors of the platform economy negotiate values and collectively drive forms of social change [43].

The platform economy and its effects on labor are also linked to the current developments of AI [22]. In collaboration with A.A. Casilli (Telecom ParisTech), P. Tubaro has received funding from the Union Force Ouvrière, from France Stratégie (a Prime Minister's service), and from MSH Paris-Saclay, to map "micro-work" in France (DiPLab project). The term micro-work refers to small, data-related tasks that are performed online against low remunerations, such as tagging objects in images, transcribing bits of text, and recording utterances aloud. Specialized platforms such as Amazon Mechanical Turk, Clickworker and Microworkers recruit online providers to execute these tasks for their clients, mostly for data-intensive production processes. In addition to poor working conditions and low pay, micro-work raises issues in terms of privacy and data protection, insofar as outside providers are entrusted with data that may include personal information [23].

The results of the DiPLab study were published in a report that attracted significant media attention [52], and presented as part of a large event on micro-work at the headquarters of France Stratégie in June 2019. Two articles relying on DiPlab results, one of which was first published as a working paper [62], [61], are now under review.

A joint franco-Canadian grant obtained by P. Tubaro and A. Delfanti (University of Toronto) enabled the creation of an "International Network on Digital Labor", aiming to bring together scholars interested in various forms of digital platform labor. Inauguration of the network involved the organization of two workshops, one in Paris (June 2019) and the other in Toronto (October 2019).

Further research on how digital platforms transform labor practices and affect the very definition of professions is being undertaken by P. Tubaro as part of a two-year (2018-2020) grant from DARES (French Ministry of Labor), in collaboration with O. Chagny of IRES, a union-funded think-tank), and A.A. Casilli (Telecom ParisTech).

A newly-obtained grant from ANR (with A.A. Casilli and U. Laitenberger, Telecom ParisTech) will enable P. Tubaro to further explore the global production networks that link AI developers and producers to data-related work across national boundaries, following outsourcing chains that extend from France to French-speaking African countries, and from Spain and the USA to parts of Latin America. This new project, entitled "The HUman Supply cHain behind smart technologies" (HUSH), will start in January 2020.

7.3.2. Health, food, and socio-demographic relationships

Participants: Philippe Caillou, Michèle Sebag, Paola Tubaro
Post-doc: Ksenia Gasnikova
Collaboration: Louis-Georges Soler, Olivier Allais (INRA)

Another area of activity concerns the relationships between eating practices, socio-demographic features and health.

The Nutriperso project (IRS Univ. Paris-Saclay, coll. INRA, CEA, CNRS, INSERM, Telecom ParisTech and Univ. Paris-Sud) aims to: i) determine the impact of food items on health (e.g., related to T2 diabetes); ii) identify alternative food items, admissible in terms of taste and budget, and better in terms of health; iii) emit personalized food recommendations> One project motivation is the fact that general recommendations (e.g., *Eat 5 fruit and vegetable per day*) are hardly effective on populations at risk. Based on the Kantar database, reporting the food habits of 20,000 households over 20 years, our challenge is to analyze the food purchases at an unprecedented fine-grained scale (at the barcode level), and to investigate the relationship between diets, socio-demographic features, and body mass index (BMI). The challenge also regards the direction of causality; while some diets are strongly correlated to high BMI, the question is to determine whether, e.g., sugar-free sodas are a cause of obesity, or a consequence thereof, or both a cause and a consequence. A main difficulty is the lack of control populations to assess a diet impact. Such a control population could be approximated in the case of the organic diet, showing a statistically significant impact of this diet on the BMI distribution. The question of finding confounders (e.g. based on wealth or education) or "backdoor" variables in under study.

7.3.3. *Scientific Information System and Visual Querying*

Participants: Philippe Caillou, Michèle Sebag
Engineers: Anne-Catherine Letournel, Jonas Renault
Collaboration: Jean-Daniel Fekete (AVIZ, Inria Saclay)

A third area of activity concerns the 2D visualisation and querying of a corpus of documents. Its initial motivation was related to scientific organisms, institutes or Universities, using their scientific production (set of articles, authors, title, abstract) as corpus. The Cartolabe project started as an Inria ADT (coll. TAO and AVIZ, 2015-2017). It received a grant from CNRS (coll. TAU, AVIZ and HCC-LRI, 2018-2019). Further extensions, as an open-source platform, are under submission at the time of writing.

The originality of the approach is to rely on the content of the documents (as opposed to, e.g. the graph of co-authoring and citations). This specificity allowed to extend Cartolabe to various corpora, such as Wikipedia, Bibliotheque Nationale de France, or the Software Heritage. Cartolabe was also applied in 2019 to the *Grand Debat* dataset: to support the interactive exploration of the 3 million propositions; and to check the consistency of the official results of the *Grand Debat* with the data.

Among its intended functionalities are: the visual assessment of a domain and its structuration (who is expert in a scientific domain, how related are the domains); the coverage of an institute expertise relatively to the general expertise; the evolution of domains along time (identification of rising topics). A round of interviews with beta-user scientists is under way since end 2019. Cartolabe usage raises questions at the crossroad of human-centered computing, data visualization and machine learning: i) how to deal with stressed items (the 2D projection of the item similarities poorly reflects their similarities in the high dimensional document space; ii) how to customize the similarity and exploit the users' feedback about relevant neighborhoods.

7.3.4. *Multi-Agent based simulation framework for social science*

Participants: Philippe Caillou
Collaboration: Patrick Taillandier (INRA), Alexis Drogoul and Nicolas Marilleau (IRD), Arnaud Grignard (MediaLab, MIT), Benoit Gaudou (Université Toulouse 1)

Since 2008, P. Caillou contributes to the development of the GAMA platform, a multi-agent based simulation framework. Its evolution is driven by the research projects using it, which makes it very well suited for social sciences studies and simulations.

The focus of the development team in 2019 was on the stability of the platform and on the documentation to provide a stable and well documented framework to the users.

7.4. Energy Management

7.4.1. *Power Grids Daily Management*

Participants: Isabelle Guyon, Marc Schoenauer

PhDs: Benjamin Donnot, Balthazar Donon
Collaboration: Antoine Marot, Patrick Panciatici (RTE), Olivier Teytaud (Facebook)

Benjamin Donnot's CIFRE PhD with RTE [12] dealt with Power Grid safety: The goal is to assess in real time the so-called "(n-1) safety" (see Section 4.2) of possible recovery actions modifying the topology of the grid after some problem occurred somewhere on the grid. However, the HADES simulator, that allows to compute the power flows in the whole network, is far too slow to simulate in real time all n-1 possible failures of a tentative topology. A simplified simulator is also available, but its accuracy is too poor to give good results. Deep surrogate models can be trained off-line for a given topology, based on the results of the slow simulator, with high-enough accuracy, but training as many models as possible failures (i.e., n-1) obviously doesn't scale up: the topology of the grid must be an input of the learned model, allowing to instantly compute the power flows at least for grid configurations close to the usual running state of the grid. A standard approach is the one-hot encoding of the topology, where n additional boolean inputs are added to the neural network, encoding the presence or absence of each line. Nevertheless, this approach poorly generalize to topologies outside the distribution of the ones used for training.

An original "guided dropout" approach was first proposed [87], in which the topology directly acts on the connections of the deep network: a missing line suppresses some connections. Whereas the standard dropout method disconnect random connections for every batch, in order to improve the generalization capacity of the network, the "guided dropout" method removes some connections based on the actual topology of the network. This approach is experimentally validated against the one-hot encoding on small subsets of the French grid (up to 308 lines). Interestingly, and rather surprisingly, even though only examples with a single disconnected line are used in the training set, the learned model is able of some additive generalization, and predictions are also accurate enough in the case 2 lines are disconnected. The guided dropout approach was later robustified [86] by learning to rapidly rank higher order contingencies including all pairs of disconnected lines, in order to prioritize the cases where the slow simulator is run: Another neural network is trained to rank all (n-1) and (n-2) contingencies in decreasing order of presumed severity.

The guided dropout approach has been further extended and generalized with the LEAP (Latent Encoding of Atypical Perturbation) architecture [30], [17], by crossing-out connections between the encoder and the decoder parts of the ResNet architecture. LEAP then performs transfer learning over spaces of distributions of topology perturbations, allowing to better handle more complex actions on the tolology, going beyond (n-1) and (n-2) perturbations by also including node-split, a current action in the real world. The LEAP approach was theoretically studied in the case of additive perturbations, and experimentally validated on an actual sub-grid of the French grid with 46 consumption nodes, 122 production nodes, 387lines and 192 substations.

LEAP is also the firs part of Balthazar Donon's on-going PhD, that currently develops using a completely different approach to approximate the power flows on a grid, i.e. that of Graph Neural Networks (GNNs). From a Power Grid perspective, GNNs can be viewed as including the topology in the very structure of the neural network, and learning some generic transfer function amongst nodes that will perform well on any topology. First results [31] use a loss based on a large dataset of actual power flows computed using the slow HADES simulator. The results indeed generalize to very different topologies than the ones used for training, in particular very different sizes of power grids. On-going work [56] removes the need to run HADES thanks to a loss that directly aims to minimize Kirshoff's law on all lines.

7.4.2. Local Grids Optimization, and the Modeling of Worst-case Scenarios

Participants: Isabelle Guyon, Marc Schoenauer, Michèle Sebag
PhDs: Victor Berger, Herilalaina Rakotoarison; **Post-doc**: Berna Batu
Collaboration: Vincent Renaut (Artelys)

One of the goals of the ADEME Next project, in collaboration with SME Artelys (see also Section 4.2), is the sizing and capacity design of regional power grids. Though smaller than the national grid, regional and urban grids nevertheless raise scaling issues, in particular because many more fine-grained information must be taken into account for their design and predictive growth.

Regarding the design of such grids, and provided accurate predictions of consumption are available (see below), off-the-shelf graph optimization algorithms can be used. Berna Batu is gathering different approaches. Herilalaina Rakotoarison's PhD tackles the automatic tuning of their parameters (see Section 7.2.1); while the Mosaic algorithm is validated on standard AutoML benchmarks [40], its application to Artelys' home optimizer at large Knitro is on-going, and compared to the state-of-the-art in parameter tuning (confidential deliverable).

In order to get accurate consumption predictions, V. Berger's PhD tackles the identification of the peak of energy consumption, defined as the level of consumption that is reached during at least a given duration with a given probability, depending on consumers (profiles and contracts) and weather conditions. The peak identification problem is currently tackled using Monte-Carlo simulations based on consumer profile- and weather-dependent individual models, at a high computational cost. The challenge is to exploit individual models to train a generative model, aimed to sampling the collective consumption distribution in the quantiles with highest peak consumption. The concept of *Compositional Variational Auto-Encoder* was proposed: it is amenable to multi-ensemblist operations (addition or subtraction of elements in the composition), enabled by the invariance and generality of the whole framework w.r.t. respectively, the order and number of the elements. It has been first tested on synthetic problems [26].

7.5. Data-driven Numerical Modelling

7.5.1. *High Energy Physics*

Participants: Cécile Germain, Isabelle Guyon
PhD: Victor Estrade, Adrian Pol
Collaboration: D. Rousseau (LAL), M. Pierini (CERN)

The role and limits of simulation in discovery is the subject of V. Estrade's PhD, specifically uncertainty quantification and calibration, that is how to handle the systematic errors, arising from the differences ("known unknowns") between simulation and reality, coming from uncertainty in the so-called nuisance parameters. In the specific context of HEP analysis, where relatively numerous labelled data are available, the problem is at the crosspoint of domain adaptation and representation learning. We have investigated how to directly enforce the invariance w.r.t. the nuisance in the sought embedding through the learning criterion (tangent back-propagation) or an adversarial approach (pivotal representation). The results [93] contrast the superior performance of incorporating a priori knowledge on a well separated classes problem (MNIST data) with a real case setting in HEP, in relation with the Higgs Boson Machine Learning challenge [68] and the TRrackML challenge [46]. More indirect approaches based on either incorporating variance reduction for the parameter of interest or constraining the representation in a variational auto-encoder farmework are currently considered.

Anomaly detection (AD) is the subject of A. Pol's PhD. Reliable data quality monitoring is a key asset in delivering collision data suitable for physics analysis in any modern large-scale high energy physics experiment. [21] focuses on supervised and semi-supervised methods addressing the identification of anomalies in the data collected by the CMS muon detectors. The combination of DNN classifiers capable of detecting the known anomalous behaviors, and convolutional autoencoders addressing unforeseen failure modes has shown unprecedented efficiency. The result has been included in the production suite of the CMS experiment at CERN. Recent work has focused on improving AD for the trigger system, which is the first stage of event selection process in most experiments at the LHC at CERN. The hierarchical structure of the trigger process called for exploiting the advances in modeling complex structured representations that perform probabilistic inference effectively, and specifically variational autoencoders. Previous works argued that training VAE models only with inliers is insufficient and the framework should be significantly modified in order to discriminate the anomalous instances. In this work, we exploit the deep conditional variational autoencoder (CVAE) and we define an original loss function together with a metric that targets hierarchically structured data AD [39], [64]. This results in an effective, yet easily trainable and maintainable model.

The highly visible TrackML challenge [46] is described in section 7.6.

7.5.2. *Remote Sensing Imagery*

Participants: Guillaume Charpiat
Collaboration: Yuliya Tarabalka, Armand Zampieri, Nicolas Girard, Pierre Alliez (Titane team, Inria Sophia-Antipolis)

The analysis of satellite or aerial images has been a long-time ongoing topic of research, but the remote sensing community moved only very recently to a principled vision of the tasks in a machine learning perspective, with sufficiently large benchmarks for validation. The main topics are the segmentation of (possibly multispectral) remote sensing images into objects of interests, such as buildings, roads, forests, etc., and the detection of changes between two images of the same place taken at different moments. The main differences with classical computer vision is that images are large (covering whole countries, typically cut into 5000×5000 pixels tiles), containing many small, potentially similar objects (and not one big object per image), that every pixel needs to be annotated (w.r.t. assigning a single label to a full image), and that the ground truth is often not reliable (spatially mis-registered, missing new constructions).

These last years, deep learning techniques took over classical approaches in most labs, adapting neural network architectures to the specifics of the tasks. This is due notably to the creation of several large scale benchmarks (including one by us [133] and, soon after, larger ones by GAFAM).

This year, we continued the work started in [167] about the registration of remote sensing images (RGB pictures) with cadastral maps (made of polygons indicating buildings and roads). We extended it in [33] to the case of real datasets, i.e. to noisy data. Indeed, in remote sensing, datasets are often large but of poor ground truth annotation quality. It turns out that, when training on datasets with noisy labels, one can still obtain accuracy scores far better than the noise variance in the training set, due to averaging effects over the labels of similar examples. To properly explain this, a theoretical study was conducted (cf. Section 7.2.2). Given any already trained neural network and its noisy training set, without knowing the real ground truth, we were then able to quantify this noise averaging effect [29].

We also tackled the problem of pansharpening, i.e. the one of producing a high-resolution color image, given a low-resolution color image and a high-resolution greyscale one [35], with deep convolutional neural networks as well.

7.5.3. *Space Weather Forecasting*

Participants: Cyril Furtlehner, Michèle Sebag
PhD: Mandar Chandorkar
Collaboration: Enrico Camporeale (CWI)

Space Weather is broadly defined as the study of the relationships between the variable conditions on the Sun and the space environment surrounding Earth. Aside from its scientific interest from the point of view of fundamental space physics phenomena, Space Weather plays an increasingly important role on our technology-dependent society. In particular, it focuses on events that can affect the performance and reliability of space-borne and ground-based technological systems, such as satellite and electric networks that can be damaged by an enhanced flux of energetic particles interacting with electronic circuits.[6]

Since 2016, in the context of the Inria-CWI partnership, a collaboration between TAU and the Multiscale Dynamics Group of CWI aims to long-term Space Weather forecasting. The goal is to take advantage of the data produced everyday by satellites surveying the sun and the magnetosphere, and more particularly to relate solar images and the quantities (e.g., electron flux, proton flux, solar wind speed) measured on the L1 libration point between the Earth and the Sun (about 1,500,000 km and 1 hour time forward of Earth). A challenge is to formulate such goals in terms of supervised learning problem, while the "labels" associated to solar images are recorded at L1 (thus with a varying and unknown time lag). In essence, while typical ML models aim to answer the question *What*, our goal here is to answer both questions *What* and *When*. This project has been articulated around Mandar Chandorkar's Phd thesis[11] which has been defended this year in Eindhoven. One

---

[6]After a recent survey conducted by the insurance company Lloyd's, an extreme Space Weather event could produce up to $2.6 trillion in financial damage.

of the main result that has been obtained concerns the prediction of solar wind impacting earth magnetosphere from solar images. In this context we encountered an interesting sub-problem related to the non deterministic travel time of a solar eruption to earth's magnetosphere. We have formalized it as the joint regression task of predicting the magnitude of signals as well as the time delay with respect to their driving phenomena. We have provided in[28] an approach to this problem combining deep learning and an original Bayesian forward attention mechanism. A theoretical analysis based on linear stability has been proposed to put this algorithm on firm ground. From the practical point of view, encouraging tests have been performed both on synthetic data and real data with results slightly better than those present in the specialized literature on a small dataset. Various extension of the method, of the experimental tests and of the theoretical analysis are planned.

7.5.4. *Genomic Data and Population Genetics*

Participants: Guillaume Charpiat, Flora Jay, Aurélien Decelle, Cyril Furtlehner
PhD: Théophile Sanchez – **PostDoc**: Jean Cury
Collaboration: Bioinfo Team (LRI), Estonian Biocentre (Institute of Genomics, Tartu, Estonia), Pasteur Institute (Paris), TIMC-IMAG (Grenoble)

Thanks to the constant improvement of DNA sequencing technology, large quantities of genetic data should greatly enhance our knowledge about evolution and in particular the past history of a population. This history can be reconstructed over the past thousands of years, by inference from present-day individuals: by comparing their DNA, identifying shared genetic mutations or motifs, their frequency, and their correlations at different genomic scales. Still, the best way to extract information from large genomic data remains an open problem; currently, it mostly relies on drastic dimensionality reduction, considering a few well-studied population genetics features.

We developed an approach that extracts features from genomic data using deep neural networks and combines them with a Bayesian framework to approximate the posterior distribution of demographic parameters. The key difficulty is to build flexible problem-dependent architectures, supporting transfer learning and in particular handling data with variable size. We designed new generic architectures, that take into account DNA specificities for the joint analysis of a group of individuals, including its variable data size aspects and compared their performances to state-of-the-art approaches [148]. In the short-term these architectures can be used for demographic inference or selection inference in bacterial populations (ongoing work with a postdoctoral researcher, J Cury, and the Pasteur Institute); the longer-term goal is to integrate them in various systems handling genetic data or other biological sequence data.

In collaboration with the Institute of Genomics of Tartu (Estonia; B Yelmen, 3-month visitor at LRI), we leveraged two types of generative neural networks (Generative Adversarial Networks and Restricted Boltzmann Machines) to learn the high dimensional distributions of real genomic datasets and create artificial genomes [66]. These artificial genomes retain important characteristics of the real genomes (genetic allele frequencies and linkage, hidden population structure, ...) without copying them and have the potential to be valuable assets in future genetic studies by providing anonymous substitutes for private databases (such as the ones hold by companies or public institutes like the Institute of Genomics of Tartu). Yet, ensuring anonymity is a challenging point and we measured the privacy loss by using and extending the Adversarial Accuracy score developed by the team for synthetic medical data [44].

In collaboration with TIMC-IMAG, we proposed a new factor analysis approach that process genetic data of multiple individuals from present-day and ancient populations to visualize population structure and estimate admixture coefficients (that is, the probability that an individual belongs to different groups given the genetic data). This method corrects the traditionally-used PCA by accounting for time heterogeneity and enables a more accurate dimension reduction of paleogenomic data [59].

7.5.5. *Sampling molecular conformations*

Participants: Guillaume Charpiat
PhD: Loris Felardos
Collaboration: Jérôme Hénin (IBPC), Bruno Raffin (InriAlpes)

Numerical simulations on massively parallel architectures, routinely used to study the dynamics of biomolecules at the atomic scale, produce large amounts of data representing the time trajectories of molecular configurations, with the goal of exploring and sampling all possible configuration basins of given molecules. The configuration space is high-dimensional (10,000+), hindering the use of standard data analytics approaches. The use of advanced data analytics to identify intrinsic configuration patterns could be transformative for the field.

The high-dimensional data produced by molecular simulations live on low-dimensional manifolds; the extraction of these manifolds will enable to drive detailed large-scale simulations further in the configuration space. This year, we studied how to bypass simulations by directly predicting, given a molecule formula, its possible configurations. This is done using Graph Neural Networks [89] in a generative way, producing 3D configurations. The goal is to sample all possible configurations, and with the right probability.

7.5.6. *Storm trajectory prediction*

Participants: Guillaume Charpiat
Collaboration: Sophie Giffard-Roisin (IRD), Claire Monteleoni (Boulder University), Balazs Kegl (LAL)

Cyclones, hurricanes or typhoons all designate a rare and complex event characterized by strong winds surrounding a low pressure area. Their trajectory and intensity forecast, crucial for the protection of persons and goods, depends on many factors at different scales and altitudes. Additionally storms have been more numerous since the 1990s, leading to both more representative and more consistent error statistics.

Currently, track and intensity forecasts are provided by numerous guidance models. Dynamical models solve the physical equations governing motions in the atmosphere. While they can provide precise results, they are computationally demanding. Statistical models are based on historical relationships between storm behavior and other parameters [82]. Current national forecasts are typically driven by consensus methods able to combine different dynamical models.

Statistical models perform poorly compared to dynamical models, although they rely on steadily increasing data resources. ML methods have scarcely been considered, despite their successes in related forecasting problems [169]. A main difficulty is to exploit spatio-temporal patterns. Another difficulty is to select and merge data coming from heterogeneous sensors. For instance, temperature and pressure are real values on a 3D spatial grid, while sea surface temperature or land indication rely on a 2D grid, wind is a 2D vector field, while many indicators such as geographical location (ocean, hemisphere...) are just real values (not fields), and displacement history is a 1D vector (time). An underlying question regards the *innate* vs *acquired* issue, and how to best combine physical models with trained models. The continuation of the work started last year [101] shows that with deep learning one can outperform the state-of-the-art in many cases [18].

7.6. Challenges

Participants: Cécile Germain, Isabelle Guyon, Adrien Pavao, Anne-Catherine Letournel, Michèle Sebag
PhD: Zhengying Liu, Lisheng Sun, Balthazar Donon
Collaborations: D. Rousseau (LAL), André Elisseeff (Google Zurich), Jean-Roch Vilmant (CERN), Antoine Marot and Benjamin Donnot (RTE), Kristin Bennett (RPI), Magali Richard (Université de Grenoble).

The TAUgroup uses challenges (scientific competitions) as a means of stimulating research in machine learning and engage a diverse community of engineers, researchers, and students to learn and contribute advancing the state-of-the-art. The TAUgroup is community lead of the open-source **Codalab** platform, hosted by Université Paris-Saclay. The project had grown in 2019 and includes now an engineer dedicated full time to administering the platform and developing challenges (Adrien Pavao), financed by a new project just starting with the Région Ile-de-France. This project will also receive the support of the Chaire Nationale d'Intelligence Artificielle of Isabelle Guyon for the next four years.

Following the highly successful ChaLearn **AutoML** Challenges (NIPS 2015 – ICML 2016 [111] – PKDD 2018 [113]), a series of challenges on the theme of **AutoDL** [129] was run in 2019 (see http://autodl.chalearn. org, addressing the problem of tuning the hyperparameters of Deep Neural Networks, including the topology of the network itself. Co-sponsored by Google Zurich, it required participants to upload their code on the Codalab platform. The series included two challenges in computer vision called **AutoCV** and **AutoCV2**, to promote automatic machine learning for image and video processing, in collaboration with University of Barcelona [45]. It also included challenges in speech processing (**AutoSpeech**), text processing (**AutoNLP**), weakly supervised learning (**AutoWeakly**) and times series (**AutoSeries**), co-organized with 4Paradigm. It culminated with launching the **AutoDL** challenge combining multiple modalities (presently on-going). The winners of each challenge open-sourced their code. GPU cloud resources were donated by Google. AutoDL was an official NeurIPS 2020 competition.

Part of the High Energy Physics activities of the team, **TrackML** [79], [80] first phase was run and co-sponsored by Kaggle, until September 2018. The second phase has been run on Codalab until March 2019, requiring code submission; algorithms were then ranked by combining accuracy and speed. The best submissions largely outperform the existing solutions. The challenge has been presented at NeurIPS [46], and at a CERN workshop[7]. I. Guyon and C. Germain are in the organizing committee, and M. Schoenauer is member of the Advisory Committee. The TAU team, in collaboration with CERN, has taken a leading role in stimulating both the ML and HEP communities to address the combinatorial complexity explosion created by the next generation of particle detectors.

A new challenge series in Reinforcement Learning was started with the company RTE France, one the theme "Learning to run a power network" [134] (**L2RPN**, http://l2rpn.chalearn.org). The goal is to test the potential of Reinforcement Learning to solve a real world problem of great practical importance: controlling electricity transportation in smart grids while keeping people and equipment safe. The first edition was run in Spring 2019 and was part of the official selection of the IJCNN 2019 conference. It ran on the Codalab platform coupled with the open source PyPower simulator of power grids interfaced with the Opengym RL framework, developed by OpenAI. In this gamified environment, the participants had to create a proper controller of a small grid of 14 nodes. Not all of them used RL, but some combinations of RL and human expertise proved to be competitive. In 2020, we will launch a new edition of the challenge with a more powerful simulator rendering the grid more realistic and capable of simulating a 118-node grid within our computational constraints. This competition was already accepted as part of the official program of IJCNN 2020.

The **HADACA** project (EIT Health) aims to run a series of challenges to promote and encourage innovations in data analysis and personalized medicine. Université de Grenoble organized a challenge on matrix factorization (https://www.medinfo-lyon.org/en/matrixen) using Codalab. The challenge gathered transdisciplinary instructors (researchers and professors), students, and health professionals (clinicians). The HADACA project contributed to create a large dataset to assess tumor heterogeneity in cancer research as well as developing innovative pedagogical methods to sensitize students to big data analysis in health. One of the products of HADACA is the ChaGrade platform (https://chagrade.lri.fr/), a tool allowing instructors to easily use challenges in the classroom, grading them as homework, and monitoring submissions and progress. HADACA will be pursued in 2020 by a sequel project also funded by EIT Health, called COMETH. The objective of COMETH will be to create an environment to conduct systematic benchmarks, based on Codalab. As a synergistic activity, TAU is also engaged in a collaboration with the Rensselaer Polytechnic Institute (RPI, New-York, USA) to use challenges in the classroom, as part of their health-informatics curriculum.

It is important to introduce **challenges in ML teaching**. This has been done (and is on-going) in I. Guyon's Licence and Master courses [38] : some assignments to Master students are to design small challenges, which are then given to Licence students in labs, and both types of students seem to love it. Codalab has also been used to implement reinforcement learning homework in the form of challenges by Victor Berger and Heri Rakotoarison for the class of Michèle Sebag. Along similar line, F. Landes proposed a challenge in the context of S. Mallat's course, at Collège de France. Finally, in collaboration with aiforgood.org, and Heri Rakotoarison has put in place a hackathon for the conference Data Science Africa (https://codalab.lri.fr/competitions/522)

---

[7]https://indico.cern.ch/event/813759/

In terms of dissemination, four books were published in 2019 in the Springer series on challenges in machine learning, see http://www.chalearn.org/books.html.

8. Bilateral Contracts and Grants with Industry

8.1. Bilateral Contracts with Industry

TAU will continue TAO policy about technology transfer, accepting any informal meeting following industrial requests for discussion (and we are happy to be too much solicited), and deciding about the follow-up based upon the originality, feasibility and possible impacts of the foreseen research directions, provided they fit our general canvas. This lead to the following 5 on-going CIFRE PhDs, with the corresponding side-contracts with the industrial supervisor, plus 3 other bilateral contracts. In particular, we now have a first "Affiliate" partner, the SME DMH, and hope to further develop in the future this form of transfer. Note that it can also sometimes lead to collaborative projects, as listed in the following sections.

- **DMH** 2019 (1 an, 45kEuros) related to consulting activities with DMH (Digital for Mental Health)[8].
  Coordinator: Aurélien Decelle and Simon Moulieras (DMH)
  Participants: Michèle Sebag

- **CIFRE Renault** 2017-2020 (45 kEuros), related to Marc Nabhan's CIFRE PhD *Sûreté de fonctionnement d'un véhicule autonome - évaluation des fausses détections au travers d'un profil de mission réduit*
  Coordinator: Marc Schoenauer and Hiba Hage (Renault)
  Participants: Marc Nabhan (PhD), Yves Tourbier (Renault)

- **BOBCAT** The new B-tO-B work intermediaries: comparing business models in the "CollaborATive" digital economy, 2018-2020 (100k euros), funded by DARES (French Ministry of Labor).
  Coordinator : Odile Chagny (IRES)
  Participants: Paola Tubaro and Antonio A. Casilli (Telecom Paris)

- **INDL-KW** International Network on Digital Labor - Kickoff Workshops, 2019 (10k euros), funded by CNRS and the University of Toronto.
  Coordinator: Paola Tubaro and Alessandro Delfanti (UoT)
  Participants: Antonio A. Casilli (Telecom Paris)

- **CIFRE Thalès** 2018-2021 (45 kEuros), with Thales Teresis, related to Nizam Makdoud's CIFRE PhD
  Coordinator: Marc Schoenauer and Jérôme Kodjabatchian
  Participants: Nizam Makdoud

- **CIFRE RTE** 2018-2021 (72 kEuros), with Réseau Transport d'Electricité, related to Balthazar Donon's CIFRE PhD
  Coordinator: Isabelle Guyon and Antoine Marot (RTE)
  Participants: Balthazar Donon, Marc Schoenauer

- **CIFRE FAIR** 2018-2021 (45 kEuros), with Facebook AI Research, related to Leonard Blier's CIFRE PhD
  Coordinator: Marc Schoenauer and Yann Olliver (Facebook)
  Participants: Guillaume Charpiat, Michèle Sebag, Léonard Blier

- **IFPEN** (Institut Français du Pétrole Energies Nouvelles) 2019-2023 (300 kEuros), to hire an Inria Starting Research Position (PhD + 4-6 years) to work in all topics mentioned in Section 3.2 relevant to IFPEN activity (see also Section 4.2). Started October 2019.
  Coordinator: Marc Schoenauer
  Participants: Alessandro Bucci, Guillaume Charpiat

---

[8]This "Affiliate" contract has been inspired by the affiliate program of Technion

9. Partnerships and Cooperations

9.1. National Initiatives

9.1.1. ANR

- **EPITOME** 2017-2020 (225kEuros), *Efficient rePresentatIon TO structure large-scale satellite iMagEs* (Section 7.5.2).
  Coordinator: Yuliya Tarabalka (Titane team, Inria Sophia-Antipolis)
  Participant: Guillaume Charpiat
- **HUSH** 2020-2023 (348kEuros), *The HUman Supply cHain behind smart technologies*.
  Coordinator: Antonio A. Casilli (Telecom Paris)
  Participant: Paola Tubaro

9.1.2. Others

- **Nutriperso** 2017-2020, 122 kEuros. Personalized recommendations toward healthier eating practices (Section 7.3.2).
  U. Paris-Saclay IRS (*Initiative de Recherche Stratégique*)
  Partners: INRA (coordinator), INSERM, Agro Paristech, Mines Telecom
  Participants: Philippe Caillou, Flora Jay, Michèle Sebag, Paola Tubaro
- **IRS CDS** 2017-2020, 75 kEuros. Personalized recommendations toward healthier eating practices
  U. Paris-Saclay IRS (*Initiative de Recherche Stratégique*)
  Partners: INRA (coordinator), INSERM, Agro Paristech, Mines Telecom
  Participants: Philippe Caillou, Flora Jay, Michèle Sebag, Paola Tubaro
- **PIA Adamme** 2015-2019 (258 kEuros) Machine Learning on a mass-memory architecture.
  Coordinator: Bruno Farcy (Bull SAS)
  Participants: Marc Schoenauer, Guillaume Charpiat, Cécile Germain-Renaud
- **NEXT** 2017-2021 (675 kEuros). Simulation, calibration, and optimization of regional or urban power grids (Section 4.2).
  ADEME (Agence de l'Environnement et de la Maîtrise de l'Energie)
  Coordinator: ARTELYS
  Participants Isabelle Guyon, Marc Schoenauer, Michèle Sebag, Victor Berger (PhD), Herilalaina Rakotoarison (PhD), Berna Bakir Batu (Post-doc)
- **DATAIA Vadore** 2018-2020 (105 kEuros) VAlorizations of Data to imprOve matching in the laboR markEt, with CREST (ENSAE) and Pôle Emploi (Section 7.3.1).
  Coordinator: Michèle Sebag
  Participants: Philippe Caillou, Isabelle Guyon
- **PIA JobAgile** 2018-2021 (379 kEuros) *Evidence-based Recommandation pour l'Emploi et la Formation* (Section 7.3.1).
  Coordinator: Michèle Sebag and Stéphanie Delestre (Qapa)
  Participants: Philippe Caillou, Isabelle Guyon
- **HADACA** 2018-2019 (50 kEuros), within EIT Health, for the organization of challenges toward personalized medicine (Section 7.6).
  Coordinator: Magali Richard (Inria Grenoble)
  Participants: Isabelle Guyon
- **IPL HPC-BigData** 2018-2022 (100 kEuros) High Performance Computing and Big Data (Section 7.5.5)
  Coordinator: Bruno Raffin (Inria Grenoble)
  Participants: Guillaume Charpiat, Loris Felardos (PhD)

- **ScGlass** 2016-2020 (10 M$), "Cracking the Glass problem" international collaboration on cracking the glass problem, funded by the Simons Fundation (NY, NYC, USA).
  Coordinator: 13 PIs around the world (see https://scglass.uchicago.edu/)
  Participants: (alumni, actively collaborating with members) François Landes

9.2. European Initiatives

9.2.1. Collaborations with Major European Organizations

CERN: collaboration with two major CERN experiments (ATLAS and CMS) on the role of machine learning at all stages of the scientific discovery process. C. Germain supervises a CERN-funded PhD.

9.3. International Initiatives

9.3.1. Inria International Labs

IIL CWI-Inria

Associate Team involved in the International Lab:

9.3.1.1. MDG-TAO

Title: Data-driven simulations for Space Weather predictions

International Partner (Institution - Laboratory - Researcher):

CWI (Netherlands) - Multiscale Dynamics Group - Enrico Camporeale

Start year: 2017

See also: http://pages.saclay.inria.fr/cyril.furtlehner/html/mdg-tao.html

We propose an innovative approach to Space Weather modeling: the synergetic use of state-of-the-art simulations with Machine Learning and Data Assimilation techniques, in order to adjust for errors due to non-modeled physical processes, and parameter uncertainties. We envision a truly multidisciplinary collaboration between experts in Computational Science and Data assimilation techniques on one side (CWI), and experts in Machine Learning and Data Mining on the other (Inria). Our research objective is to realistically tackle long-term Space Weather forecasting, which would represent a giant leap in the field. This proposal is extremely timely, since the huge amount of (freely available) space missions data has not yet been systematically exploited in the current computational methods for Space Weather. Thus, we believe that this work will result in cutting-edge results and will open further research topics in space Weather and Computational Plasma Physics.

10. Dissemination

10.1. Promoting Scientific Activities

10.1.1. Scientific Events: Organisation

10.1.1.1. General Chair, Scientific Chair

Isabelle Guyon - Competition co-chair, ECMLPKDD 2019

Michele Sebag - Area Chair NIPS 2017-2019, ICLR 2020

Flora Jay - PASADENA workshop co-chair, Paris, 2019

10.1.1.2. Member of Organizing Committees

Guillaume Charpiat - Organizing & scientific committee of the ForMaL summer school, at ENS Cachan, June 2019, and of WAISE (Second International Workshop on Artificial Intelligence Safety Engineering) held at SafeComp in September 2019.

Flora Jay - Organizing & scientific committee of Research Program "Ecosystem dynamics : stakes, data and models" at Institut Pascal, Paris-Saclay, 2019.

Isabelle Guyon - Advisory committee BayLearn 2019; Co-organizer NeurIPS 2019 workshop on Challenges in Machine Learning; Co-organizer NeurIPS 2019 NewInML workshop.

Marc Schoenauer - Steering Committee, Parallel Problem Solving from Nature (PPSN); Steering Committee, Learning and Intelligent OptimizatioN (LION).

Michele Sebag - President of Steering Committee, Eur. Conf. on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD).

10.1.2. Scientific Events: Selection

10.1.2.1. Reviewer

All TAU members are reviewers of the main conferences in their respective fields of expertise.

10.1.3. Journal

10.1.3.1. Member of the Editorial Boards

Isabelle Guyon - Action editor, *Journal of Machine Learning Research* (JMLR); series editor, *Springer series Challenges in Machine Learning* (CiML).

Marc Schoenauer - Advisory Board, *Evolutionary Computation Journal*, MIT Press, and *Genetic Programming and Evolutionary Machines*, Springer Verlag; Action editor, *Journal of Machine Learning Research* (JMLR).

Michèle Sebag - Editorial Board, Machine Learning, Springer Verlag.

Paola Tubaro - Associate Editorial Board *Sociology*, Sage; Co-editor, *Revue Française de Sociologie*, Presses de Sciences Po.

10.1.3.2. Reviewer - Reviewing Activities

All members of the team reviewed numerous articles for the most prestigious journals in their respective fields of expertise.

10.1.4. Invited Talks

Guillaume Charpiat - *Deep Learning for Satellite Imagery* in the Imagine team (ENPC), at Champs-sur-Marne, March 2019; and at LRDE lab (EPITA), at Kremlin-Bicêtre, April 2019; *Deep Learning for Storm Trajectory Prediction and Remote Sensing* at the seminar AI for Climate, at Jussieu (Paris), December 2019.

Cyril Furtlehner - *A machine learning approach to solar wind speed forecasting from solar images*, Machine Learning in Heliophysics conference, Amsterdam September 2019.

Flora Jay - *Machine Learning and Deep Learning for Population Genetics*, Statistics/Machine Learning at Paris-Saclay, Bures-sur-Yvette, Jan 2019; *Inferring Demography? A Deep Learning Approach for Population Genetic Data*, ALPHY 7-8 Feb. 2019; *Génomes présents, histoires d'antan: Apprentissage profond pour la génétique des populations*, Inria Seminar Unithé ou Café with G Charpiat, Feb 2019; *When AI & Big Data meet life sciences - Advances in research and ethical questions*, Round Table at YLRS2019, Paris, Jun 2019; *Inferring past history from genetic data using ABC and Deep Learning approaches*, Seminar at Lille University, Jan 2019; *Creating Artificial Human Genomes Using Generative Models*, Seminars at LCQB and Evolmol IBENS, Paris, Nov and Dec 2019.

Julien Girard - *Formal validation for machine learning*, DataIA day on Safety & AI, Palaiseau, Sept. 11th.

Isabelle Guyon - Invited talk at 2019 Sackler Colloquium of the Science of Deep Learning, and Keynote IJCNN 2019: "Neural network solvers for power transmission problems" (https://youtu.be/YBcRzIAFDYU).

Michèle Sebag - *Meta-Learning*, Kickoff of the Kompetenz Center in Machine Learning, Rhine Ruhr Region, Jan. 2019; *Structural Agnostic Models*, Oberwolfach Symposium on Causality, May 2019; *Artificial Intelligence & Causal Modeling*, CREST Big Data Applications Symposium, Tokyo, Sept. 24, 2019; *Some news and questions about AI and Machine Learning*, Arensberg Symposium, Leuven, Nov. 2019.

Marc Schoenauer - *Intelligence Artificielle Mythes et réalités*, audition par la Fédération Française du Bâtiment, Feb. 13, 2019; *Intelligence Artificielle : du congrès de Dartmouth au rapport Villani*, Séminaire IFPEN, Apr. 4. 2019; *Intelligence Artificielle : de Dartmouth à l'apprentissage profond et au rapport Villani*, Rencontres Franciliennes de Mécanique, May 28, 2019; *A gentle introduction to AI and (Deep) Learning*, Mexican-French Workshop, Aug. 27. 2019, Mexico; *Intelligence artificielle : certification, transparence, et impact sur la société*, Data Science Day, Mines Paris Tech, Sept. 18, 2019; *When Big Data and Machine Learning meet Partial Differential Equations*, CREST Big Data Applications Symposium, Tokyo, Sept. 25, 2019;

Paola Tubaro - *Sélectionné.e par une IA ? Algorithmes, inégalités, et les "humains dans la boucle"*, Centre D'Alembert, Orsay, 18/04/2019; *Dans la fabrique des algorithmes : plateformes, micro-travail et dynamiques d'externalisation*, INSEE, 14/05/2019; *The human labour that makes AI possible: An empirical study of micro-work in France*, Alan Turing Institute, London, 22/01/2019; *Que font les big data aux sciences sociales ? Retour sur une 'crise' annoncée*, EHESS, 21/02/2019.

10.1.5. Leadership within the Scientific Community

Isabelle Guyon - President and co-founder of ChaLearn, a non-for-profit organization dedicated to the organization of challenge.

Marc Schoenauer - Chair of ACM-SIGEVO (Special Interest Group on Evolutionary Computation), 2015-2019, now in Advisory Board; Founding President (since 2015) of SPECIES (Society for the Promotion of Evolutionary Computation In Europe and Surroundings), that organizes the yearly series of conferences *EvoStar*.

Michèle Sebag - Elected Chair of Steering Committee, ECML-PKDD; Board member, Institut de Convergence DataIA.

Paola Tubaro - Convenor of the Social Network Analysis Group of British Sociological Association; co-founder of European Network on Digital Labor.

10.1.6. Scientific Expertise

Guillaume Charpiat - Member of the Inria Saclay Commission Scientifique, and as such, jury committee for grants for PhD theses / postdocs / professor delegations; expertise for DigiCosme grants, ANRT CIFRE PhD grants; for GPU platforms: Jean Zay (GENCI) and Lab-IA (Saclay plateau). Discussion panel of the workshop day IA & Océan-Atmosphère-Climat (IMT Atlantique Rennes), February 6th. Panel of the Machine Learning workshop at CRiP ITES, Deauville, April 5th.

Cécile Germain - Evaluator for the H2020-2016-CNECT program; member of the DFG review panel within Germany's excellence strategy selection process.

Isabelle Guyon - Member of the NeurIPS foundation board.

Marc Schoenauer - Comité Scientifique IA, SCube (Scientipôle Savoirs & Société), Orsay; Scientific Committee, TrackML (see Section 7.6); Comité de sélection, Chaire ABEONA-ENS "Biais et Equité en IA"; Conseil Scientifique, IFPEN; Scientific Advisory Board, BCAM, Bilbao, Spain; Scientific Advisory Board, Tara Oceans, Paris.

Michèle Sebag - Hiring juries : LRI; Centrale-Supélec; ENS-Paris; UCA-Nice; U. Freiburg, Germany. Selection juries: Awards NIPS 2019. Propositions NSERC, Canada; Propositions Dpt STIC; Prix de thèse AFIA. Expert Committee from Finland's Minister of Economic Affairs, AI Strategy, February 2019.

10.1.7. Research Administration

Isabelle Guyon - Representative of UPSud in the DataIA *Institut de Convergence* Program Committee, University of Paris-Saclay. Responsible of Master AIC (becoming Paris-Saclay master in Artificial Intelligence).

Marc Schoenauer - Deputy Scientific Director of Inria (in French, Directeur Scientifique Adjoint, DSA), in charge of AI.

Michele Sebag - Deputy director of LRI, CNRS UMR 8623; elected member of the Research Council of Univ. Paris-Saclay; member of the STIC department council of Univ. Paris-Saclay; member of the Scientific Council of Labex AMIES, Applications des Mathématiques ds l'Industrie, l'Entreprise et la Société.

Paola Tubaro - Representative of CNRS in the DataIA *Institut de Convergence* Program Committee, University of Paris-Saclay; member of the Board, Maison des Sciences de l'Homme Paris-Saclay; member of CLIP, Institut Pascal, University of Paris-Saclay.

10.2. Teaching - Supervision - Juries

10.2.1. Teaching

Licence : Philippe Caillou, Computer Science for students in Accounting and Management, 192h, L1, IUT Sceaux, Univ. Paris Sud.

Licence : Aurélien Decelle, Computer Architecture, 42h, L2, Univ. Paris-Sud.

Licence : Aurélien Decelle, Introduction to Machine Learning, 57h, L2, Univ. Paris-Sud.

Licence : François Landes, Mathematics for Computer Scientists, 51h, L2, Univ. Paris-Sud.

Licence : François Landes, Intro to Machine Learning, 48h, L2, Univ. Paris-Sud.

Licence and Polytech : Cécile Germain, Computer Architecture

Licence : Isabelle Guyon: Introduction to Data Science, L1, Univ. Paris-Sud.

Licence : Isabelle Guyon, Project: Resolution of mini-challenges (created by M2 students), L2, Univ. Paris-Sud.

Master : François Landes, Machine Learning, 34h, M1 Polytech, U. Paris-sud.

Master : Aurélien Decelle, Machine Learning, 26h, M1, Univ. Paris-Sud.

Master : Aurélien Decelle, Probability and statistics, 26h, M1, Univ. Paris-Sud.

Master : Aurélien Decelle, Probability, statistics and information theory, M1, Univ. Paris-Sud.

Master : Guillaume Charpiat, Deep Learning in Practice, 24h, M2 Recherche, Centrale-Supelec + MVA.

Master : Isabelle Guyon, Project: Creation of mini-challenges, M2, Univ. Paris-Sud.

Master : Michèle Sebag, Machine Learning, 12h; Deep Learning, 9h; Reinforcement Learning, 12h; M2 Recherche, U. Paris-Sud.

Master : François Landes, Machine Learning, 22h, M2 Recherche, U. Paris-Sud.

Master : Paola Tubaro, Start -up project for engineering students, 24h, Telecom ParisTech.

Master : Paola Tubaro, Sociology of social networks, 24h, M2, EHESS/ENS.

Master : Paola Tubaro, Social and economic network science, 24h, M2, ENSAE.

Doctorate: Paola Tubaro, Research Methods, 12h, University of Insubria, Italy.

Summer school: Guillaume Charpiat, Machine Learning & Deep Learning Tutorial, 4h30, ForMaL, ENS Cachan, June 4th

10.2.2. Supervision

HdR - Paola Tubaro, *Decoding the platform society: Organizations, markets and networks in the digital economy*, 11/12/2019, Sciences Po Paris.

PhD - Benjamin DONNOT, *Deep learning methods for predicting flows in power grids : novel architectures and algorithms.*, 13/02/2019, Isabelle Guyon and Antoine Marot (RTE)

PhD - Corentin TALLEC, *Reinforcement Learning and Recurrent Neural Networks: Dynamical approaches*, 7/10/2019, Université Paris-Saclay, Yann Ollivier

PhD - Mandar CHANDORKAR, *Machine Learning in Space Weather*, 14/11/2019, University of Eindhoven, Enrico Camporeale, Cyril Furtlehner, and Michèle Sebag

PhD - Guillaume DOQUET, *Agnostic Feature Selection*, 29/11/2019, Université Paris-Saclay, Michèle Sebag

PhD - Diviyan KALAINATHAN, *Generative Neural Networks to Infer Causal Mechanisms: Algorithms and Applications*, 17/12/2019, Université Paris-Saclay, Michèle Sebag and Isabelle Guyon

PhD - Lisheng SUN, *Meta-Learning as a Markov Decision Process*, 19/12/2019, Université Paris-Saclay, Isabelle Guyon and Michèle Sebag

PhD in progress - Eléonore BARTENLIAN, *Deep Learning pour le traitement du signal*, 1/10/2018, Michèle Sebag and Frédéric Pascal (Centrale-Supélec)

PhD in progress - Victor BERGER, *Variational Anytime Simulator*, 1/10/2017, Michèle Sebag

PhD in progress - Guillaume BIED, *Valorisation des Données pour la Recherche d'Emploi*, 1/10/2019, Bruno Crepon (CREST-ENSAE) and Philippe Caillou

PhD in progress - Leonard BLIER, *Vers une architecture stable pour les systèmes d'apprentissage par renforcement*, 1/09/2018, Yann Ollivier (Facebook AI Research, Paris) and Marc Schoenauer

PhD in progress - Tony BONNAIRE, *Reconstruction de la toile cosmique*, from 1/10/2018, Nabila Aghanim (Institut d'Astrophysique Spatiale) and Aurélien Decelle

PhD in progress - Balthazar DONON, *Apprentissage par renforcement pour une conduite stratégique du système électrique*, 1/10/2018, Isabelle Guyon and Antoine Marot (RTE)

PhD in progress - Victor ESTRADE *Robust domain-adversarial learning, with applications to High Energy Physics,* 01/10/2016, Cécile Germain and Isabelle Guyon.

PhD in progress - Loris FELARDOS, *Neural networks for molecular dynamics simulations*, 1/10/2018, Guillaume Charpiat, Jérôme Hénin (IBPC) and Bruno Raffin (InriAlpes)

PhD in progress - Giancarlo FISSORE, *Statistical physics analysis of generative models*, 1/10/2017, Aurélien Decelle and Cyril Furtlehner

PhD in progress - Julien GIRARD, *Vérification et validation des techniques d'apprentissage automatique*, 1/10/2018, Zakarian Chihani (CEA) and Guillaume Charpiat

PhD in progress - Nicolas GIRARD, *Satellite image vectorization using neural networks*, 1/10/2017, Yuliya Tarabalka & Pierre Alliez (Inria Sophia-Antipolis) and Guillaume Charpiat

PhD in progress - Armand LACOMBE, *Recommandation de Formations: Application de l'apprentissage causal dans le domaine des ressources humaines*, 1/10/2019, Michele Sebag and Philippe Caillou

PhD in progress - Zhengying LIU, *Automation du design des reseaux de neurones profonds*, 1/10/2017, Isabelle Guyon

PhD in progress - Nizam MAKDOUD, *Motivations intrinsèques en apprentissage par renforcement. Application à la recherche de failles de sécurité*, 1/02/2018, Marc Schoenauer and Jérôme Kodjabachian (Thalès ThereSIS, Palaiseau).

PhD in progress - Marc NABHAN, *Sûreté de fonctionnement d'un véhicule autonome - évaluation des fausses détections au travers d'un profil de mission réduit*, 1/10/2017, Marc Schoenauer and Hiba Hage (Renault)

PhD in progress - Adrian POL *Machine Learning Anomaly Detection, with application to CMS Data Quality Monitoring,* 01/10/2016, Cécile Germain.

PhD in progress - Herilalaina RAKOTOARISON, *Automatic Algorithm Configuration for Power Grid Optimization*, 1/10/2017, Marc Schoenauer and Michèle Sebag

PhD in progress - Théophile SANCHEZ, *Reconstructing the past: deep learning for population genetics*, 1/10/2017, Guillaume Charpiat and Flora Jay

PhD in progress - Pierre WOLINSKI, *Learning the Architecture of Neural Networks*, from 1/9/2016, Yann Ollivier (Facebook AI Research, Paris) and Guillaume Charpiat

PhD in progress - Wenzhuo LIU, Machine Learning for Numerical Simulation of PDEs, from 1/11/2019, Mouadh Yagoubi (IRT SystemX) and Marc Schoenauer

PhD in progress - Marion ULLMO, *Reconstruction de la toile cosmique*, from 1/10/2018, Nabila Aghanim (Institut d'Astrophysique Spatiale) and Aurélien Decelle

10.2.3. Juries

Marc Schoenauer: Reviewer for Dennis Wilson, IRIT, Université Toulouse; PhD jury of Patricio Cerda Reyes, U. Paris-Saclay, 28/11/2019.

François Landes: PhD rapporteur of Martina Teruzzi (Condensed matter and Machine Learning PhD, at SISSA, Trieste, Italy).

Paola Tubaro: PhD jury of Sophie Balech, U. Paris Nanterre, 09/07/2019; PhD jury of Linda Rua, U. Paris Dauphine, 13/12/2019.

Guillaume Charpiat: half-way PhD committee ("à mi-parcours") of Rodrigo Daudt (ONERA), of Julie Rivet (EPITA) and of Nissim Zerbib (Institut de la Vision)

Isabelle Guyon: PhD jury Justine Falque, U. Paris-Saclay (29/11/2019).

Michele Sebag: Reviewer for A. Chemla (PhD U. Roma & IRCAM); Pierre Fournier (PhD ISIR).

10.3. Popularization

10.3.1. Articles and contents

- Michèle Sebag - *Il faut dissiper le malentendu sur "les prétentions infondées" de l'intelligence artificielle*, tribune du journal *Le Monde* du 12 août 2019.

- Paola Tubaro - BBC World service, *The 'microworkers' making your digital life possible*, 02/08/2019; Radio France Inter, *Travailleurs du clic, les soutiers du clavier*, 23/03/2019; Science & Vie, *La force de l'IA repose sur des travailleurs bien r'eels*, 24/10/2019; AFP/L'Express, *Les travailleurs du clic, petites mains invisibles de l'économie numérique*, 21/06/2019; CNRS Le Journal, *Ces microtravailleurs de l'ombre*, 24/05/2019; Le Monde, *Jobs du clic : la France compte plus de 250 000 micro-travailleurs*, 15/02/2019; L'Humanité Dimanche, *Travailleuse du clic, la double invisibilité*, 02/05/2019.

10.3.2. Interventions

- Guillaume Charpiat - *Génomes présents, histoires d'antan: Apprentissage profond pour la génétique des populations*, Inria Seminar Unithé ou Café with Flora Jay, Feb 2019.

- Marc Schoenauer - *Intelligence Artificielle Mythes et réalités*, Classe virtuelle de la Délégation Académique du Numérique Educatif, Feb. 2, 2019; *Intelligence Artificielle Mythes et réalités*, Médiathèque George Sand, Palaiseau, Feb. 9, 2019.

- Michèle Sebag - Journées Scientifiques Inria, Lyon 2019; Centre de Recherches Interdisciplinaires, Workshop on Artificial Intelligence and Women Empowerment (Oct. 2019); CRI, *Open AI: From big data to smart data* (Nov. 2019).

- Paola Tubaro - *Le Micro-Travail en France : derrière l'automatisation, de nouvelles précarités au travail ?* CFE-CGC Orange, Paris, 01/07/2019; *Le Micro-Travail en France*, Fondation Gabriel Péri, Paris, 26/06/2019; *Big data et société : vie privée, travail, inégalités*, Banque de France, 06/06/2019.

11. Bibliography

Major publications by the team in recent years

[1] C. ADAM-BOURDARIOS, G. COWAN, C. GERMAIN-RENAUD, I. GUYON, B. KÉGL, D. ROUSSEAU. *The Higgs Machine Learning Challenge*, in "Journal of Physics: Conference Series", December 2015, vol. 664, n⁰ 7 [*DOI :* 10.1088/1742-6596/664/7/072015], https://hal.inria.fr/hal-01745998

[2] L. DA COSTA, Á. FIALHO, M. SCHOENAUER, M. SEBAG. *Adaptive Operator Selection with Dynamic Multi-Armed Bandits*, in "Proc. Genetic and Evolutionary Computation Conference (GECCO)", ACM, 2008, pp. 913-920, ACM-SIGEVO 10-years Impact Award [*DOI :* 10.1145/1389095.1389272], https://hal.inria.fr/inria-00278542

[3] C. FURTLEHNER, A. DECELLE. *Cycle-based Cluster Variational Method for Direct and Inverse Inference*, in "Journal of Statistical Physics", August 2016, vol. 164, n⁰ 3, pp. 531–574, https://hal.inria.fr/hal-01214155

[4] S. GELLY, M. SCHOENAUER, M. SEBAG, O. TEYTAUD, L. KOCSIS, D. SILVER, C. SZEPESVARI. *The Grand Challenge of Computer Go: Monte Carlo Tree Search and Extensions*, in "Communications- ACM", 2012, vol. 55, n⁰ 3, pp. 106-113, https://hal.inria.fr/hal-00695370

[5] O. GOUDET, D. KALAINATHAN, P. CAILLOU, D. LOPEZ-PAZ, I. GUYON, M. SEBAG. *Learning Functional Causal Models with Generative Neural Networks*, in "Explainable and Interpretable Models in Computer Vision and Machine Learning", Springer Series on Challenges in Machine Learning, Springer International Publishing, 2018, https://arxiv.org/abs/1709.05321 [*DOI :* 10.1007/978-3-319-98131-4], https://hal.archives-ouvertes.fr/hal-01649153

[6] T. LUCAS, C. TALLEC, J. VERBEEK, Y. OLLIVIER. *Mixed batches and symmetric discriminators for GAN training*, in "ICML - 35th International Conference on Machine Learning", Stockholm, Sweden, July 2018, https://hal.inria.fr/hal-01791126

[7] E. MAGGIORI, Y. TARABALKA, G. CHARPIAT, P. ALLIEZ. *Convolutional Neural Networks for Large-Scale Remote Sensing Image Classification*, in "IEEE Transactions on Geoscience and Remote Sensing", 2017, vol. 55, n⁰ 2, pp. 645-657, https://hal.inria.fr/hal-01369906

[8] M. MISIR, M. SEBAG. *Alors: An algorithm recommender system*, in "Artificial Intelligence", 2017, vol. 244, pp. 291-314, Published on-line Dec. 2016, https://hal.inria.fr/hal-01419874

[9] Y. OLLIVIER, L. ARNOLD, A. AUGER, N. HANSEN. *Information-Geometric Optimization Algorithms: A Unifying Picture via Invariance Principles*, in "Journal of Machine Learning Research", 2017, vol. 18, n⁰ 18, pp. 1-65, https://hal.inria.fr/hal-01515898

[10] X. ZHANG, C. FURTLEHNER, C. GERMAIN-RENAUD, M. SEBAG. *Data Stream Clustering with Affinity Propagation*, in "IEEE Transactions on Knowledge and Data Engineering", 2014, vol. 26, n⁰ 7, 1 p. , https://hal.inria.fr/hal-00862941

Publications of the year

Doctoral Dissertations and Habilitation Theses

[11] M. CHANDORKAR. *Machine Learning in Space Weather*, Université of Eindhoven, November 2019, https://hal.inria.fr/tel-02430788

[12] B. DONNOT. *Deep learning methods for predicting flows in power grids : novel architectures and algorithms*, Université Paris-Saclay, February 2019, https://tel.archives-ouvertes.fr/tel-02045873

[13] G. F. DOQUET. *Agnostic Feature Selection*, Université Paris-Saclay/Université Paris-Sud, November 2019, https://hal.archives-ouvertes.fr/tel-02436845

[14] D. KALAINATHAN. *Generative Neural networks to infer Causal Mechanisms: Algorithms and applications*, Université Paris Sud (Paris 11) - Université Paris Saclay, December 2019, https://hal.inria.fr/tel-02435986

[15] L. SUN-HOSOYA. *Meta-Learning as a Markov Decision Process*, Université Paris-Saclay, December 2019, https://hal.archives-ouvertes.fr/tel-02422144

[16] C. TALLEC. *Recurrent Neural Networks and Reinforcement Learning: Dynamic Approaches*, Université Paris-Saclay, October 2019, https://hal.inria.fr/tel-02434367

Articles in International Peer-Reviewed Journals

[17] B. DONON, B. DONNOT, I. GUYON, Z. LIU, A. MAROT, P. PANCIATICI, M. SCHOENAUER. *LEAP Nets for System Identification and Application to Power Systems*, in "Neurocomputing", 2019, To appear, forthcoming, https://hal.inria.fr/hal-02422708

[18] S. GIFFARD-ROISIN, M. YANG, G. CHARPIAT, C. K. BONFANTI, B. KÉGL, C. MONTELEONI. *Tropical Cyclone Track Forecasting using Fused Deep Learning from Aligned Reanalysis Data*, in "Frontiers in Big Data", January 2020, https://arxiv.org/abs/1910.10566 , forthcoming, https://hal.archives-ouvertes.fr/hal-02329437

[19] F. P. LANDES, G. BIROLI, O. DAUCHOT, A. LIU, D. REICHMAN. *Attractive versus truncated repulsive supercooled liquids: The dynamics is encoded in the pair correlation function*, in "Physical Review E ", January 2020, vol. 101, n^o 1, https://arxiv.org/abs/1906.01103 [*DOI :* 10.1103/PHYSREVE.101.010602], https://hal.inria.fr/hal-02439295

[20] E. LIPPIELLO, G. PETRILLO, F. P. LANDES, A. ROSSO. *Fault Heterogeneity and the Connection between Aftershocks and Afterslip*, in "Bulletin of the Seismological Society of America", April 2019, vol. 109, n^o 3, pp. 1156-1163, https://arxiv.org/abs/1812.05862 [*DOI :* 10.1785/0120180244], https://hal.inria.fr/hal-02156407

[21] A. A. POL, G. CERMINARA, C. GERMAIN, M. PIERINI, A. SETH. *Detector Monitoring with Artificial Neural Networks at the CMS Experiment at the CERN Large Hadron Collider*, in "Computing and Software for Big Science", January 2019, vol. 3, n^o 1 [*DOI :* 10.1007/S41781-018-0020-1], https://hal.inria.fr/hal-02427998

[22] P. TUBARO, A. A. CASILLI. *Micro-work, artificial intelligence and the automotive industry*, in "Journal of Industrial and Business Economics = Economia e politica industriale", June 2019, pp. 1-13 [*DOI :* 10.1007/s40812-019-00121-1], https://hal.archives-ouvertes.fr/hal-02148979

[23] P. TUBARO. *La vie privée, un bien commun ?*, in "Regards croisés sur l'économie", 2019, vol. 23, n^o 2, pp. 129-137 [*DOI :* 10.3917/RCE.023.0129], https://hal.archives-ouvertes.fr/hal-02196333

[24] P. TUBARO. *Les Learning Analytics vus par la sociologie*, in "Distances et Médiations des Savoirs", December 2019, n^o 28, https://hal.archives-ouvertes.fr/hal-02418562

[25] P. TUBARO. *Whose results are these anyway? Reciprocity and the ethics of "giving back" after social network research*, in "Social Networks", November 2019 [*DOI :* 10.1016/J.SOCNET.2019.10.003], https://hal.archives-ouvertes.fr/hal-02360709

International Conferences with Proceedings

[26] V. BERGER, M. SEBAG. *From abstract items to latent spaces to observed data and back: Compositional Variational Auto-Encoder*, in "ECML PKDD 2019 - European Conference on Machine learning and knowledge discovery in databases", Würzburg, Germany, September 2019, https://arxiv.org/abs/2001.07910 , https://hal.inria.fr/hal-02431955

[27] L. BLIER, P. WOLINSKI, Y. OLLIVIER. *Learning with Random Learning Rates*, in "ECML PKDD 2019 - European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases", Würzburg, Germany, September 2019, https://arxiv.org/abs/1810.01322 - 16 pages, 4 figures, code available on GitHub, https://hal.archives-ouvertes.fr/hal-01888352

[28] M. CHANDORKAR, C. FURTLEHNER, B. PODUVAL, E. CAMPOREALE, M. SEBAG. *Dynamic Time Lag Regression: Predicting What and When*, in "ICLR 2020 - 8th International Conference on Learning Representations", Addis Abeba, Ethiopia, April 2020, https://hal.inria.fr/hal-02422148

[29] G. CHARPIAT, N. GIRARD, L. FELARDOS, Y. TARABALKA. *Input Similarity from the Neural Network Perspective*, in "NeurIPS 2019 - 33th Annual Conference on Neural Information Processing Systems", Vancouver, Canada, December 2019, https://hal.inria.fr/hal-02394647

[30] B. DONNOT, B. DONON, I. GUYON, Z. LIU, A. MAROT, P. PANCIATICI, M. SCHOENAUER. *LEAP nets for power grid perturbations*, in "ESANN 2019 - 27th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning", Bruges, Belgium, April 2019, https://arxiv.org/abs/1908.08314 , https://hal.archives-ouvertes.fr/hal-02268886

[31] B. DONON, B. DONNOT, I. GUYON, A. MAROT. *Graph Neural Solver for Power Systems*, in "IJCNN 2019 - International Joint Conference on Neural Networks", Budapest, Hungary, July 2019, https://hal.archives-ouvertes.fr/hal-02175989

[32] *Best Paper*
G. F. DOQUET, M. SEBAG. *Agnostic feature selection*, in "ECML PKDD 2019 - European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases", Würzburg, Germany, September 2019, https://hal.archives-ouvertes.fr/hal-02436824.

[33] N. GIRARD, G. CHARPIAT, Y. TARABALKA. *Noisy Supervision for Correcting Misaligned Cadaster Maps Without Perfect Ground Truth Data*, in "IGARSS 2019 - IEEE International Geoscience and Remote Sensing Symposium", Yokohama, Japan, July 2019, https://arxiv.org/abs/1903.06529 , https://hal.inria.fr/hal-02065211

[34] J. GIRARD-SATABIN, G. CHARPIAT, Z. CHIHANI, M. SCHOENAUER. *CAMUS: A Framework to Build Formal Specifications for Deep Perception Systems Using Simulators*, in "ECAI 2020 - 24th European Conference on Artificial Intelligence", Santiago de Compostela, Spain, June 2020, https://hal.inria.fr/hal-02440520

[35] A. KHALEL, O. TASAR, G. CHARPIAT, Y. TARABALKA. *Multi-Task Deep Learning for Satellite Image Pansharpening and Segmentation*, in "IGARSS 2019 - IEEE International Geoscience and Remote Sensing Symposium", Yokohama, Japan, July 2019, https://hal.inria.fr/hal-02276549

[36] A. MAROT, A. ROSIN, L. CROCHEPIERRE, B. DONNOT, P. PINSON, L. BOUDJELOUD-ASSALA. *Interpreting atypical conditions in systems with deep conditional Autoencoders: the case of electrical consumption*, in "ECML PKDD 2019 - European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases", Wurzburg, Germany, September 2019, https://hal.archives-ouvertes.fr/hal-02266963

[37] M. NABHAN, M. SCHOENAUER, Y. TOURBIER, H. HAGE. *Optimizing coverage of simulated driving scenarios for the autonomous vehicle*, in "IEEE ICCVE 2019 - International Conference on Connected Vehicles and Expo", Graz, Austria, November 2019, https://hal.inria.fr/hal-02433530

[38] A. PAVAO, D. KALAINATHAN, L. SUN-HOSOYA, K. P. BENNETT, I. GUYON. *Design and Analysis of Experiments: A Challenge Approach in Teaching*, in "NeurIPS 2019 - 33th Annual Conference on Neural Information Processing Systems", Vancouver, Canada, December 2019, https://hal.inria.fr/hal-02415639

[39] A. A. POL, V. BERGER, G. CERMINARA, C. GERMAIN, M. PIERINI. *Anomaly Detection With Conditional Variational Autoencoders*, in "ICMLA 2019 - 18th IEEE International Conference on Machine Learning and Applications", Boca Raton, United States, 18th International Conference on Machine Learning Applications, December 2019, https://hal.inria.fr/hal-02396279

[40] H. RAKOTOARISON, M. SCHOENAUER, M. SEBAG. *Automated Machine Learning with Monte-Carlo Tree Search*, in "IJCAI-19 - 28th International Joint Conference on Artificial Intelligence", Macau, China, International Joint Conferences on Artificial Intelligence Organization, August 2019, pp. 3296-3303 [*DOI :* 10.24963/IJCAI.2019/457], https://hal.inria.fr/hal-02300884

[41] A. SCHOENAUER SEBAG, L. HEINRICH, M. SCHOENAUER, M. SEBAG, L. WU, S. ALTSCHULER. *Multi-Domain Adversarial Learning*, in "ICLR 2019 - Seventh annual International Conference on Learning Representations", New Orleans, United States, Tara Sainath, May 2019, https://hal.inria.fr/hal-01968180

[42] C. TALLEC, L. BLIER, Y. OLLIVIER. *Making Deep Q-learning methods robust to time discretization*, in "ICML 2019 - Thirty-sixth International Conference on Machine Learning", Long Beach, United States, PMLR, 2019, n^o 97, pp. 6096-6104, https://arxiv.org/abs/1901.09732 , https://hal.inria.fr/hal-02435523

[43] P. TUBARO. *Emergent relational structures at a "sharing economy" festival*, in "COMPLEX NETWORKS 2018 - The 7th International Conference on Complex Networks and Their Applications", Cambridge, United Kingdom, Springer, 2019, vol. 2, pp. 559-571, https://hal.archives-ouvertes.fr/hal-01947911

[44] A. YALE, S. DASH, R. DUTTA, I. GUYON, A. PAVAO, K. P. BENNETT. *Privacy Preserving Synthetic Health Data*, in "ESANN 2019 - European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning", Bruges, Belgium, April 2019, https://hal.inria.fr/hal-02160496

Conferences without Proceedings

[45] Z. LIU, I. GUYON, J. C. S. JACQUES JUNIOR, M. MADADI, S. ESCALERA, A. PAVAO, H. J. ESCALANTE, W.-W. TU, Z. XU, S. TREGUER. *AutoCV Challenge Design and Baseline Results*, in "CAp 2019 - Conférence sur l'Apprentissage Automatique", Toulouse, France, July 2019, https://hal.archives-ouvertes.fr/hal-02265053

Scientific Books (or Scientific Book chapters)

[46] S. AMROUCHE, L. BASARA, P. CALAFIURA, V. ESTRADE, S. FARRELL, D. FERREIRA, L. FINNIE, N. FINNIE, C. GERMAIN, V. V. GLIGOROV, T. GOLLING, S. GORBUNOV, H. GRAY, I. GUYON, M. HUSHCHYN, V. INNOCENTE, M. KIEHN, E. MOYSE, J.-F. PUGET, Y. REINA, D. ROUSSEAU, A. SALZBURGER, A. USTYUZHANIN, J.-R. VLIMANT, J. S. WIND, T. XYLOURIS, Y. YILMAZ. *The Tracking Machine Learning Challenge: Accuracy Phase*, in "The NeurIPS '18 Competition", The Springer Series on Challenges in Machine Learning, Springer, November 2019, vol. 8, pp. 231-264 [*DOI :* 10.1007/978-3-030-29135-8_9], https://hal.inria.fr/hal-02427989

[47] O. GOUDET, D. KALAINATHAN, M. SEBAG, I. GUYON. *Learning Bivariate Functional Causal Models*, in "Cause Effect Pairs in Machine Learning", I. GUYON, A. STATNIKOV, B. B. BATU (editors), The Springer Series on Challenges in Machine Learning, Springer Verlag, October 2019, pp. 101-153 [*DOI :* 10.1007/978-3-030-21810-2_3], https://hal.inria.fr/hal-02433201

[48] I. GUYON, A. STATNIKOV, B. B. BATU. *Cause Effect Pairs in Machine Learning*, The Springer Series on Challenges in Machine Learning, Springer Verlag, 2019 [*DOI :* 10.1007/978-3-030-21810-2], https://hal.inria.fr/hal-02433195

[49] I. GUYON, O. GOUDET, D. KALAINATHAN. *Evaluation Methods of Cause-Effect Pairs*, in "Cause Effect Pairs in Machine Learning", I. GUYON, A. STATNIKOV, B. B. BATU (editors), The Springer Series on Challenges in Machine Learning, Springer Verlag, October 2019, pp. 27-99 [*DOI :* 10.1007/978-3-030-21810-2_2], https://hal.inria.fr/hal-02433198

[50] I. GUYON, A. STATNIKOV. *Results of the Cause-Effect Pair Challenge*, in "Cause Effect Pairs in Machine Learning", I. GUYON, A. STATNIKOV, B. B. BATU (editors), The Springer Series on Challenges in Machine Learning. Springer, Springer Verlag, October 2019, pp. 237-256 [*DOI :* 10.1007/978-3-030-21810-2_7], https://hal.inria.fr/hal-02433204

[51] D. KALAINATHAN, O. GOUDET, M. SEBAG, I. GUYON. *Discriminant Learning Machines*, in "Cause Effect Pairs in Machine Learning", I. GUYON, A. STATNIKOV, B. B. BATU (editors), The Springer Series on Challenges in Machine Learning, Springer Verlag, October 2019, pp. 155-189 [*DOI :* 10.1007/978-3-030-21810-2_4], https://hal.inria.fr/hal-02433203

Research Reports

[52] A. A. CASILLI, P. TUBARO, C. LE LUDEC, M. COVILLE, M. BESENVAL, T. MOUHTARE, E. WAHAL. *Micro-work in France. Behind Automation, New Forms of Precarious Labour?*, Projet de recherche DiPLab, May 2019, https://hal.archives-ouvertes.fr/hal-02139528

Other Publications

[53] T. BONNAIRE, N. AGHANIM, A. DECELLE, M. DOUSPIS. *T-ReX: a graph-based filament detection method*, December 2019, https://arxiv.org/abs/1912.00732 - Submitted to A&A, 16 pages, 16 figures, https://hal.archives-ouvertes.fr/hal-02420696

[54] A. DECELLE, C. FURTLEHNER. *Gaussian-Spherical Restricted Boltzmann Machines*, December 2019, https://arxiv.org/abs/1910.14544 - 27 pages, 8 figures, https://hal.archives-ouvertes.fr/hal-02420813

[55] A. DECELLE, S. HWANG, J. ROCCHI, D. TANTARI. *Inverse problems for structured datasets using parallel TAP equations and RBM*, December 2019, https://arxiv.org/abs/1906.11988 - working paper or preprint, https://hal.archives-ouvertes.fr/hal-02420755

[56] B. DONON, R. CLÉMENT, B. DONNOT, A. MAROT, I. GUYON, M. SCHOENAUER. *Neural Networks for Power Flow : Graph Neural Solver*, November 2019, working paper or preprint, https://hal.archives-ouvertes.fr/hal-02372741

[57] H. J. ESCALANTE, H. KAYA, A. A. SALAH, S. ESCALERA, Y. GUCLUTURK, U. GÜÇLÜ, X. BARÓ, I. GUYON, J. C. S. JACQUES JUNIOR, M. MADADI, S. AYACHE, E. VIEGAS, F. GURPINAR, A. S. WICAKSANA, C. C. S. LIEM, M. A. J. VAN GERVEN, R. VAN LIER. *Explaining First Impressions: Modeling, Recognizing, and Explaining Apparent Personality from Videos*, January 2019, https://arxiv.org/abs/1802.00745 - Accepted at Transaction on Affective Computing, https://hal.inria.fr/hal-01991652

[58] G. FISSORE, A. DECELLE, C. FURTLEHNER, Y. HAN. *Robust Multi-Output Learning with Highly Incomplete Data via Restricted Boltzmann Machines*, December 2019, https://arxiv.org/abs/1912.09382 - working paper or preprint, https://hal.archives-ouvertes.fr/hal-02420824

[59] O. FRANÇOIS, S. LIÉGEOIS, B. DEMAILLE, F. JAY. *Inference of population genetic structure from temporal samples of DNA*, December 2019, biorxiv [*DOI :* 10.1101/801324], https://hal.archives-ouvertes.fr/hal-02413974

[60] J. GIRARD-SATABIN, G. CHARPIAT, Z. CHIHANI, M. SCHOENAUER. *CAMUS: A Framework to Build Formal Specifications for Deep Perception Systems Using Simulators*, November 2019, https://arxiv.org/abs/1911.10735 - working paper or preprint, https://hal.inria.fr/hal-02374956

[61] C. LE LUDEC, P. TUBARO, A. A. CASILLI. *Combien de personnes micro-travaillent en France ? Estimer l'ampleur d'une nouvelle forme de travail*, February 2019, Il s'agit de la version française du document arxiv 1901.03889 (HAL 02012731). Publié dans les documents de travail i3, https://hal.archives-ouvertes.fr/hal-02021525

[62] C. LE LUDEC, P. TUBARO, A. A. CASILLI. *How many people microwork in France? Estimating the size of a new labor force*, January 2019, https://arxiv.org/abs/1901.03889 - 9 pages, 1 figure, 7 tables. The article is part of the DiPLab ("Digital Platform Labor") research project. For more information: http://diplab.eu/Il s'agit de la version anglaise du document hal-02021525, https://hal.archives-ouvertes.fr/hal-02012731

[63] Z. LIU, Z. XU, S. ESCALERA, I. GUYON, J. C. S. JACQUES JUNIOR, M. MADADI, A. PAVAO, S. TREGUER, W.-W. TU. *Towards Automated Computer Vision: Analysis of the AutoCV Challenges 2019*, November 2019, working paper or preprint, https://hal.archives-ouvertes.fr/hal-02386805

[64] A. A. POL, V. BERGER, G. CERMINARA, C. GERMAIN, M. PIERINI. *Trigger Rate Anomaly Detection with Conditional Variational Autoencoders at the CMS Experiment*, December 2019, Machine Learning and the Physical Sciences Workshop at the 33rd Conference on Neural Information Processing Systems (NeurIPS), Poster, https://hal.inria.fr/hal-02428005

[65] A. A. POL, G. CERMINARA, C. GERMAIN, M. PIERINI, A. SETH. *Detector monitoring with artificial neural networks at the CMS experiment at the CERN Large Hadron Collider*, January 2019, https://arxiv.org/abs/1808.00911 - working paper or preprint, https://hal.inria.fr/hal-01976256

[66] B. YELMEN, A. DECELLE, L. ONGARO, D. MARNETTO, F. MONTINARO, C. FURTLEHNER, L. PAGANI, F. JAY. *Creating Artificial Human Genomes Using Generative Models*, December 2019, biorxiv [*DOI : 10.1101/769091*], https://hal.archives-ouvertes.fr/hal-02413942

References in notes

[67] C. ADAM-BOURDARIOS, G. COWAN, C. GERMAIN, I. GUYON, B. KÉGL, D. ROUSSEAU. *How Machine Learning won the Higgs Boson Challenge*, in "Proc. European Symposium on ANN, CI and ML", 2016

[68] C. ADAM-BOURDARIOS, G. COWAN, C. GERMAIN-RENAUD, I. GUYON, B. KÉGL, D. ROUSSEAU. *The Higgs Machine Learning Challenge*, in "Journal of Physics: Conference Series", 2015, vol. 664, n^o 7 [*DOI : 10.1088/1742-6596/664/7/072015*], https://hal.inria.fr/hal-01745998

[69] D. J. AMIT, H. GUTFREUND, H. SOMPOLINSKY. *Statistical Mechanics of Neural Networks Near Saturation*, in "Annals of Physics", 1987, vol. 173, pp. 30-67

[70] C. ANDERSON. *The End of Theory: The Data Deluge Makes the Scientific Method Obsolete*, 2008, Wired Magazine, https://www.wired.com/2008/06/pb-theory/

[71] N. BELKHIR. *Per Instance Algorithm Configuration for Continuous Black Box Optimization*, Université Paris-Saclay, November 2017

[72] N. BELKHIR, J. DRÉO, P. SAVÉANT, M. SCHOENAUER. *Per instance algorithm configuration of CMA-ES with limited budget*, in "Proc. ACM-GECCO", 2017, pp. 681-688

[73] I. BELLO, B. ZOPH, V. VASUDEVAN, Q. V. LE. *Neural Optimizer Search with Reinforcement Learning*, in "34th ICML", 2017, pp. 459–468

[74] S. BEN-DAVID, J. BLITZER, K. CRAMMER, A. KULESZA, F. PEREIRA, J. W. VAUGHAN. *A theory of learning from different domains*, in "Machine Learning", 2018, vol. 79, n^o 1, pp. 151–175 [*DOI : 10.1007/S10994-009-5152-4*]

[75] J. BERGSTRA, R. BARDENET, Y. BENGIO, B. KÉGL. *Algorithms for Hyper-Parameter Optimization*, in "NIPS 25", 2011, pp. 2546–2554

[76] J. BRUNA, S. MALLAT. *Invariant Scattering Convolution Networks*, in "IEEE Trans. Pattern Anal. Mach. Intell.", 2013, vol. 35, n^o 8, pp. 1872–1886

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

[78] R. COUILLET, M. DEBBAH. *Random matrix methods for wireless communications*, Cambridge University Press, 2011

[79] DAVID ROUSSEAU ET AL.. *The TrackML challenge*, in "NIPS 2018 - 32nd Annual Conference on Neural Information Processing Systems", Montreal, Canada, December 2018, pp. 1-23, https://hal.inria.fr/hal-01745714

[80] DAVID ROUSSEAU ET AL.. *The TrackML Particle Tracking Challenge*, July 2018, https://hal.inria.fr/hal-01680537

[81] DAVID SILVER ET AL.. *Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm*, in "CoRR", 2017, vol. abs/1712.01815, http://arxiv.org/abs/1712.01815

[82] M. DEMARIA, M. MAINELLI, L. K. SHAY, J. A. KNAFF, J. KAPLAN. *Further improvements to the statistical hurricane intensity prediction scheme (SHIPS)*, in "Weather and Forecasting", 2005, vol. 20, n$^{\text{o}}$ 4, pp. 531–543

[83] A. DECELLE, G. FISSORE, C. FURTLEHNER. *Spectral dynamics of learning in restricted Boltzmann machines*, in "EPL (Europhysics Letters)", 2017, vol. 119, n$^{\text{o}}$ 6, 60001 p.

[84] A. DECELLE, G. FISSORE, C. FURTLEHNER. *Thermodynamics of Restricted Boltzmann Machines and Related Learning Dynamics*, in "J. Stat. Phys.", 2018, vol. 172, pp. 1576-1608

[85] L. DINH, J. SOHL-DICKSTEIN, S. BENGIO. *Density estimation using Real NVP*, in "Int. Conf. on Learning Representations (ICLR)", 2017

[86] B. DONNOT, I. GUYON, M. SCHOENAUER, A. MAROT, P. PANCIATICI. *Anticipating contingengies in power grids using fast neural net screening*, in "IEEE International Joint Conference on Neural Networks", Rio de Janeiro, Brazil, Proc. 2018 International Joint Conference on Neural Networks (IJCNN), IEEE, July 2018, https://arxiv.org/abs/1805.02608 , https://hal.archives-ouvertes.fr/hal-01783669

[87] B. DONNOT, I. GUYON, M. SCHOENAUER, A. MAROT, P. PANCIATICI. *Fast Power system security analysis with Guided Dropout*, in "26th European Symposium on Artificial Neural Networks", Bruges, Belgium, Electronic Proceedings ESANN 2018, April 2018, https://arxiv.org/abs/1801.09870 , https://hal.archives-ouvertes.fr/hal-01695793

[88] T. DURIEZ, S. L. BRUNTON, B. R. NOACK. *Machine Learning Control – Taming Nonlinear Dynamics and Turbulence*, Springer International Publishing, 2017

[89] D. K. DUVENAUD, D. MACLAURIN, J. AGUILERA-IPARRAGUIRRE, R. GÓMEZ-BOMBARELLI, T. HIRZEL, A. ASPURU-GUZIK, R. P. ADAMS. *Convolutional Networks on Graphs for Learning Molecular Fingerprints*, in "NIPS", 2015, pp. 2224–2232

[90] T. V. ERVEN, P. GRUNWALD, S. DE ROOIJ. *Catching up faster by switching sooner: a predictive approach to adaptive estimation with an application to the AIC-BIC dilemma: Catching up Faster*, in "J. Royal Statistical Society: B", 2012, vol. 74, n^o 3, pp. 361–417 [*DOI : 10.1111/J.1467-9868.2011.01025.X*]

[91] H. J. ESCALANTE, I. GUYON, S. ESCALERA, J. JACQUES, M. MADADI, X. BARÓ, S. AYACHE, E. VIEGAS, Y. GÜÇLÜTÜRK, U. GÜÇLÜ, M. A. J. VAN GERVEN, R. VAN LIER. *Design of an Explainable Machine Learning Challenge for Video Interviews*, in "IJCNN 2017 - 30th International Joint Conference on Neural Networks", Anchorage, AK, United States, Neural Networks (IJCNN), 2017 International Joint Conference on, IEEE, 2017, pp. 1-8, https://hal.inria.fr/hal-01668386

[92] S. ESCALERA, X. BARÓ, H. J. ESCALANTE, I. GUYON. *ChaLearn looking at people: A review of events and resources*, in "2017 International Joint Conference on Neural Networks (IJCNN)", 2017, pp. 1594-1601

[93] V. ESTRADE, C. GERMAIN, I. GUYON, D. ROUSSEAU. *Systematics aware learning: a case study in High Energy Physics*, in "ESANN 2018 - 26th European Symposium on Artificial Neural Networks", Bruges, Belgium, April 2018, https://hal.inria.fr/hal-01715155

[94] S. FALKNER, A. KLEIN, F. HUTTER. *BOHB: Robust and Efficient Hyperparameter Optimization at Scale*, in "Proc. 35th ICML", 2018, pp. 1436–1445

[95] M. FEURER, A. KLEIN, K. EGGENSPERGER, J. SPRINGENBERG, M. BLUM, F. HUTTER. *Efficient and Robust Automated Machine Learning*, in "NIPS 28", 2015, pp. 2962–2970

[96] L. FRANCESCHI, P. FRASCONI, S. SALZO, R. GRAZZI, M. PONTIL. *Bilevel Programming for Hyperparameter Optimization and Meta-Learning*, in "Proc. 35th ICML", J. G. DY, A. KRAUSE (editors), 2018, pp. 1563–1572

[97] C. FURTLEHNER, A. DECELLE. *Cycle-Based Cluster Variational Method for Direct and Inverse Inference*, in "Journal of Statistical Physics", 2016, vol. 164, n^o 3, pp. 531–574

[98] C. FURTLEHNER, M. SEBAG, X. ZHANG. *Scaling analysis of affinity propagation*, in "Physical Review E", 2010, vol. 81, n^o 6, 066102 p.

[99] Y. GANIN, E. USTINOVA, H. AJAKAN, P. GERMAIN, H. LAROCHELLE, F. LAVIOLETTE, M. MARCHAND, V. LEMPITSKY. *Domain-Adversarial Training of Neural Networks*, in "Journal of Machine Learning Research", 2016, vol. 17, n^o 59, pp. 1-35

[100] K. GARIMELLA, G. D. F. MORALES, A. GIONIS, M. MATHIOUDAKIS. *Political Discourse on Social Media: Echo Chambers, Gatekeepers, and the Price of Bipartisanship*, in "WWW", ACM, 2018, pp. 913–922

[101] S. GIFFARD-ROISIN, M. YANG, G. CHARPIAT, B. KÉGL, C. MONTELEONI. *Deep Learning for Hurricane Track Forecasting from Aligned Spatio-temporal Climate Datasets*, in "Modeling and decision-making in the spatiotemporal domain NIPS workhop", Montréal, Canada, December 2018, https://hal.archives-ouvertes.fr/hal-01905408

[102] J. GILES. *Computational social science: Making the links*, in "Nature - News", 2012, vol. 488, n^o 7412, pp. 448-450

[103] F. GONARD. *Cold-start recommendation : from Algorithm Portfolios to Job Applicant Matching*, Université Paris-Saclay, May 2018, https://tel.archives-ouvertes.fr/tel-01825220

[104] F. GONARD, M. SCHOENAUER, M. SEBAG. *ASAP.V2 and ASAP.V3: Sequential optimization of an Algorithm Selector and a Scheduler*, in "Open Algorithm Selection Challenge 2017 ", Proceedings of Machine Learning Research, 2017, vol. 79, pp. 8-11

[105] I. GOODFELLOW, J. POUGET-ABADIE, M. MIRZA, B. XU, D. WARDE-FARLEY, S. OZAIR, A. COURVILLE, Y. BENGIO. *Generative Adversarial Nets*, in "NIPS 27", 2014, pp. 2672–2680

[106] I. GOODFELLOW, J. POUGET-ABADIE, M. MIRZA, B. XU, D. WARDE-FARLEY, S. OZAIR, A. COURVILLE, Y. BENGIO. *Generative Adversarial Nets*, in "NIPS 27", Z. GHAHRAMANI, M. WELLING, C. CORTES, N. D. LAWRENCE, K. Q. WEINBERGER (editors), Curran Associates, Inc., 2014, pp. 2672–2680

[107] O. GOUDET, D. KALAINATHAN, P. CAILLOU, D. LOPEZ-PAZ, I. GUYON, M. SEBAG. *Learning Functional Causal Models with Generative Neural Networks*, in "Explainable and Interpretable Models in Computer Vision and Machine Learning", Springer Series on Challenges in Machine Learning, Springer International Publishing, 2018 [*DOI : 10.1007/978-3-319-98131-4*], https://hal.archives-ouvertes.fr/hal-01649153

[108] P. GRUNWALD. *The Minimum Description Length Principle*, MIT Press, 2007

[109] I. GUYON, C. F. ALIFERIS, G. F. COOPER, A. ELISSEEFF, J.-P. PELLET, P. SPIRTES, A. R. STATNIKOV. *Design and Analysis of the Causation and Prediction Challenge*, in "WCCI Causation and Prediction Challenge", JMLR W-CP, 2008, pp. 1–33

[110] I. GUYON, K. BENNETT, G. CAWLEY, H. J. ESCALANTE, S. ESCALERA, T. K. HO, N. MACIA, B. RAY, M. SAEED, A. STATNIKOV. *Design of the 2015 ChaLearn AutoML challenge*, in "Proc. IJCNN", IEEE, 2015, pp. 1–8

[111] I. GUYON, K. BENNETT, G. CAWLEY, H. J. ESCALANTE, S. ESCALERA, T. K. HO, N. MACIA, B. RAY, M. SAEED, A. STATNIKOV. *Design of the 2015 ChaLearn AutoML challenge*, in "Proc. IJCNN", IEEE, 2015, pp. 1–8

[112] I. GUYON, L. BOTTOU, B. SCHÖLKOPF, A. STATNIKOV, E. VIEGAS, J. M. ROBINS. *NIPS workshop on causality*, 2013, http://clopinet.com/isabelle/Projects/NIPS2013/

[113] I. GUYON, L. SUN-HOSOYA, M. BOULLÉ, H. J. ESCALANTE, S. ESCALERA, Z. LIU, D. JAJETIC, B. RAY, M. SAEED, M. SEBAG, A. STATNIKOV, W.-W. TU, E. VIEGAS. *Analysis of the AutoML Challenge series 2015-2018*, in "AutoML: Methods, Systems, Challenges", F. HUTTER, L. KOTTHOFF, J. VANSCHORE (editors), The Springer Series on Challenges in Machine Learning, Springer Verlag, 2018, forthcoming, https://hal.archives-ouvertes.fr/hal-01906197

[114] D. HERNÁNDEZ-LOBATO, P. MORALES-MOMBIELA, D. LOPEZ-PAZ, A. SUÁREZ. *Non-linear Causal Inference using Gaussianity Measures*, in "Journal of Machine Learning Research", 2016, vol. 17, n⁰ 28, pp. 1-39, http://jmlr.org/papers/v17/14-375.html

[115] H. H. HOOS. *Programming by Optimization*, in "Commun. ACM", 2012, vol. 55, n⁰ 2, pp. 70–80

[116] D. HOTHERSALL. *History of Psychology*, McGraw-Hill, 2004

[117] G. W. IMBENS, D. B. RUBIN. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*, Cambridge University Press, 2015

[118] A. JACOT, F. GABRIEL, C. HONGLER. *Neural Tangent Kernel: Convergence and Generalization in Neural Networks*, in "Advances in Neural Information Processing Systems 31", 2018, pp. 8571–8580

[119] D. KALAINATHAN, O. GOUDET, P. CAILLOU, M. SEBAG, P. TUBARO. *Conditions objectives de travail et ressenti des individus : le rôle du management*, 2017, 12 p. , Synthèse n. 14 de "La Fabrique de l'Industrie", https://hal.inria.fr/hal-01742592

[120] D. KALAINATHAN, O. GOUDET, I. GUYON, D. LOPEZ-PAZ, M. SEBAG. *SAM: Structural Agnostic Model, Causal Discovery and Penalized Adversarial Learning*, August 2018, working paper or preprint, https://hal.archives-ouvertes.fr/hal-01864239

[121] D. KINGMA, M. WELLING. *Auto-encoding variational Bayes.*, in "Int. Conf. on Learning Representations (ICLR)", 2014

[122] L. KOTTHOFF. *Algorithm Selection for Combinatorial Search Problems: A Survey*, in "Data Mining and Constraint Programming: Foundations of a Cross-Disciplinary Approach", CH. BESSIERE ET AL. (editor), Springer Verlag, 2016, pp. 149–190

[123] A. KRIZHEVSKY, I. SUTSKEVER, G. E. HINTON. *ImageNet Classification with Deep Convolutional Neural Networks*, in "Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1", NIPS'12, 2012, pp. 1097–1105

[124] D. LAZER, A. PENTLAND, L. ADAMIC, S. ARAL, A. L. BARABASI, D. BREWER, N. CHRISTAKIS, N. CONTRACTOR, J. FOWLER, M. GUTMANN, T. JEBARA, G. KING, M. MACY, D. ROY, M. VAN ALSTYNE. *Life in the network: the coming age of computational social science*, in "Science", 2009, vol. 323, n⁰ 5915, pp. 721–723

[125] Y. LECUN. *Deep Learning est mort. Vive Differentiable Programming!*, 2018, Facebook Post, https://www.facebook.com/yann.lecun/posts/10155003011462143

[126] J. Z. LIANG, E. MEYERSON, R. MIIKKULAINEN. *Evolutionary architecture search for deep multitask networks*, in "Proc. ACM-GECCO", H. E. AGUIRRE, K. TAKADAMA (editors), ACM, 2018, pp. 466–473, http://doi.acm.org/10.1145/3205455.3205489

[127] T. P. LILLICRAP, J. J. HUNT, A. PRITZEL, N. HEESS, T. EREZ, Y. TASSA, D. SILVER, D. WIERSTRA. *Continuous control with deep reinforcement learning*, 2015

[128] H. LIPSON, J. B. POLLACK. *Automatic design and manufacture of robotic lifeforms*, in "Nature Letters", 2000, vol. 406, pp. 974–978

[129] Z. LIU, O. BOUSQUET, A. ELISSEEFF, S. ESCALERA, I. GUYON, J. JACQUES, A. PAVAO, D. SILVER, L. SUN-HOSOYA, S. TREGUER, W.-W. TU, J. WANG, Q. YAO. *AutoDL Challenge Design and Beta Tests-*

Towards automatic deep learning, in "CiML workshop @ NIPS2018", Montreal, Canada, December 2018, https://hal.archives-ouvertes.fr/hal-01906226

[130] D. LOPEZ-PAZ, K. MUANDET, B. SCHÖLKOPF, I. O. TOLSTIKHIN. *Towards a Learning Theory of Cause-Effect Inference*, in "Proceedings of the 32nd International Conference on Machine Learning, ICML", F. R. BACH, D. M. BLEI (editors), 2015, pp. 1452–1461, http://jmlr.org/proceedings/papers/v37/lopez-paz15.html

[131] C. LOUIZOS, M. REISSER, T. BLANKEVOORT, E. GAVVES, M. WELLING. *Relaxed Quantization for Discretized Neural Networks*, in "ICLR", 2019

[132] T. LUCAS, C. TALLEC, J. VERBEEK, Y. OLLIVIER. *Mixed batches and symmetric discriminators for GAN training*, in "ICML - 35th International Conference on Machine Learning", Stockholm, Sweden, July 2018, https://hal.inria.fr/hal-01791126

[133] E. MAGGIORI, Y. TARABALKA, G. CHARPIAT, P. ALLIEZ. *Can Semantic Labeling Methods Generalize to Any City? The Inria Aerial Image Labeling Benchmark*, in "IEEE International Symposium on Geoscience and Remote Sensing (IGARSS)", Fort Worth, United States, July 2017, https://hal.inria.fr/hal-01468452

[134] A. MAROT, B. DONON, I. GUYON, B. DONNOT. *Learning To Run a Power Network Competition*, in "CiML Workshop, NeurIPS", Montréal, Canada, December 2018, https://hal.archives-ouvertes.fr/hal-01968295

[135] L. MESCHEDER, A. GEIGER, S. NOWOZIN. *Which Training Methods for GANs do actually Converge?*, in "35th ICML", 2018, pp. 3481–3490

[136] D. MOLCHANOV, A. ASHUKHA, D. VETROV. *Variational Dropout Sparsifies Deep Neural Networks*, in "ArXiv e-prints", January 2017

[137] M. MISIR, M. SEBAG. *Alors: An algorithm recommender system*, in "Artificial Intelligence", 2017, vol. 244, pp. 291-314, Published on-line Dec. 2016, https://hal.inria.fr/hal-01419874

[138] C. O'NEILL. *Weapons of Math Destruction*, Crown Books, 2016

[139] J. PEARL. *Causality: Models, Reasoning, and Inference (2nd edition)*, Cambridge University Press, 2009

[140] J. PEARL. *Theoretical Impediments to Machine Learning With Seven Sparks from the Causal Revolution*, in "11th ACM WSDM", 2018, 3 p.

[141] G. PLEISS, M. RAGHAVAN, F. WU, J. M. KLEINBERG, K. Q. WEINBERGER. *On Fairness and Calibration*, in "NIPS", 2017, pp. 5684–5693

[142] M. RAISSI. *Deep Hidden Physics Models: Deep Learning of Nonlinear Partial Differential Equations*, in "JMLR", 2018, vol. 19, pp. 1–24

[143] G. RAMÍREZ, E. VILLATORO, B. IONESCU, H. J. ESCALANTE, S. ESCALERA, M. LARSON, H. MÜLLER, I. GUYON. *Overview of the Multimedia Information Processing for Personality and Social Networks Analysis Contest*, in "International Conference on Pattern Recognition (ICPR)", T. TAN, J. KITTLER, A. JAI (editors), IEEE, 2018, pp. 127-139

[144] S. RAVANBAKHSH, J. SCHNEIDER, B. PÓCZOS. *Equivariance Through Parameter-Sharing*, in "International Conference on Machine Learning", 2017, pp. 2892–2901, http://proceedings.mlr.press/v70/ravanbakhsh17a.html

[145] E. REAL, S. MOORE, A. SELLE, S. SAXENA, Y. L. SUEMATSU, J. T. Q. LE, A. KURAKIN. *Large-Scale Evolution of Image Classifiers*, in "34th ICML", 2017

[146] J. R. RICE. *The Algorithm Selection Problem*, in "Advances in Computers", 1976, vol. 15, pp. 65 – 118 [*DOI : 10.1016/S0065-2458(08)60520-3*], http://www.sciencedirect.com/science/article/pii/S0065245808605203

[147] J. RISSANEN. *Information and Complexity in Statistical Modeling*, Information Science and Statistics, Springer-Verlag, 2007

[148] T. SANCHEZ, G. CHARPIAT, F. JAY. *SPI-DNA: End-to-end Deep Learning Approach for Demographic History Inference*, in "Paris-Saclay Junior Conference on Data Science and Engineering", Orsay, France, September 2017, https://hal.archives-ouvertes.fr/hal-01679385

[149] M. SCHMIDT, H. LIPSON. *Distilling free-form natural laws from experimental data*, in "Science", 2009, vol. 324, n^o 5923, pp. 81–85

[150] S. SIGTIA, E. BENETOS, S. DIXON. *An End-to-End Neural Network for Polyphonic Piano Music Transcription*, in "IEEE/ACM Trans. Audio, Speech & Language Processing", 2016, vol. 24, n^o 5, pp. 927–939

[151] J. SIRIGNANO, K. SPILIOPOULOS. *DGM: A deep learning algorithm for solving partial differential equations*, in "Journal of Computational Physics", 2018, vol. 375, pp. 1339 – 1364 [*DOI : 10.1016/J.JCP.2018.08.029*]

[152] I. STENGERS. *Another Science Is Possible*, Open Humanities Press, 2013

[153] D. H. STERN, R. HERBRICH, T. GRAEPEL. *Matchbox: large scale online bayesian recommendations*, in "WWW", ACM Press, 2010, 111 p. [*DOI : 10.1145/1526709.1526725*]

[154] L. SUN-HOSOYA, I. GUYON, M. SEBAG. *Lessons learned from the AutoML challenge*, in "Conférence sur l'Apprentissage Automatique 2018", Rouen, France, June 2018, https://hal.inria.fr/hal-01811454

[155] C. SUNSTEIN, R. THALER. *Nudge: Improving Decisions about Health, Wealth, and Happiness*, Yale University Press, 2008

[156] C. SZEGEDY, W. ZAREMBA, I. SUTSKEVER, J. BRUNA, D. ERHAN, I. J. GOODFELLOW, R. FERGUS. *Intriguing properties of neural networks*, in "CoRR", 2013, vol. abs/1312.6199, http://arxiv.org/abs/1312.6199

[157] C. TALLEC, Y. OLLIVIER. *Can recurrent neural networks warp time?*, in "ICLR", 2018, https://openreview.net/forum?id=SJcKhk-Ab

[158] C. TALLEC, Y. OLLIVIER. *Can recurrent neural networks warp time?*, in "International Conference on Learning Representation 2018", Vancouver, France, April 2018, https://hal.inria.fr/hal-01812064

[159] C. Tallec, Y. Ollivier. *Unbiased online recurrent optimization*, in "International Conference On Learning Representation", Vancouver, Canada, April 2018, https://hal.archives-ouvertes.fr/hal-01972587

[160] S. Triantafillou, I. Tsamardinos. *Constraint-based Causal Discovery from Multiple Interventions over Overlapping Variable Sets*, in "Journal of Machine Learning Research", 2015, vol. 16, pp. 2147-2205

[161] J. Vanschoren, J. N. van Rijn, B. Bischl, L. Torgo. *OpenML: networked science in machine learning*, in "SIGKDD Explorations", 2013, vol. 15, n^o 2, pp. 49–60, https://arxiv.org/abs/1407.7722

[162] Volodymyr Mnih et al.. *Human-level control through deep reinforcement learning*, in "Nature", 2015, vol. 518, n^o 7540, pp. 529–533, https://doi.org/10.1038/nature14236

[163] J. Wan, S. Escalera, G. Anbarjafari, H. J. Escalante, X. Baró, I. Guyon, M. Madadi, J. Allik, C. Lin, Y. Xie. *Results and Analysis of ChaLearn LAP Multi-modal Isolated and Continuous Gesture Recognition, and Real versus Fake Expressed Emotions Challenges*, in "International Conference on Computer Vision - ICCV 2017", 2017, https://hal.inria.fr/hal-01677974

[164] M. Weimer, A. Karatzoglou, Q. Le, A. Smola. *CoFiRank - Maximum Margin Matrix Factorization for Collaborative Ranking*, in "21th NIPS", 2007, pp. 222-230

[165] M. Welling. *Intelligence per Kilowatthour*, 2018, Keynote talk @ICML, https://icml.cc/Conferences/2018/Schedule?showEvent=1866

[166] F. M. J. Willems, Y. M. Shtarkov, T. J. Tjalkens. *The context-tree weighting method: basic properties*, in "IEEE Transactions on Information Theory", 1995, vol. 41, n^o 3, pp. 653-664

[167] A. Zampieri, G. Charpiat, N. Girard, Y. Tarabalka. *Multimodal image alignment through a multiscale chain of neural networks with application to remote sensing*, in "European Conference on Computer Vision (ECCV)", Munich, Germany, September 2018, https://hal.inria.fr/hal-01849389

[168] X. Zhang, C. Furtlehner, C. Germain-Renaud, M. Sebag. *Data Stream Clustering with Affinity Propagation*, in "IEEE Transactions on Knowledge and Data Engineering", 2014, vol. 26, n^o 7, 1 p. , https://hal.inria.fr/hal-00862941

[169] M. de Bezenac, A. Pajot, P. Gallinari. *Deep Learning for Physical Processes: Incorporating Prior Scientific Knowledge*, in "arXiv preprint arXiv:1711.07970", 2017

[170] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, K. Kavukcuoglu. *WaveNet: A Generative Model for Raw Audio*, in "CoRR", 2016, vol. abs/1609.03499, http://arxiv.org/abs/1609.03499