

RESEARCH CENTRE

Saclay - Île-de-France

IN PARTNERSHIP WITH:

CNRS, Université Paris-Saclay, Université
Paris-Sud (Paris 11)

2020

ACTIVITY REPORT

DRAFT

Project-Team

TAU

Tackling the Underspecified

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

DOMAIN

Applied Mathematics, Computation and
Simulation

THEME

Optimization, machine learning and
statistical methods

Contents

Project-Team TAU	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	4
2.1 Presentation	4
3 Research program	4
3.1 Toward Good AI	4
3.1.1 Causal modeling and biases	5
3.1.2 Robustness of Learned Models	5
3.2 Hybridizing numerical modeling and learning systems	6
3.3 Learning to learn	8
3.3.1 Auto-*	8
3.3.2 Information theory: adjusting model complexity and data fitting	9
3.3.3 Analyzing and Learning Complex Systems	9
3.4 Organisation of Challenges	10
4 Application domains	11
4.1 Computational Social Sciences	11
4.2 Energy Management	11
4.3 Data-driven Numerical Modeling	11
5 Social and environmental responsibility	12
5.1 Footprint of research activities	12
5.2 Impact of research results	12
6 Highlights of the year	13
6.1 Awards	13
6.2 Selective Fundings	13
7 New software and platforms	13
7.1 New software	13
7.1.1 Codalab	13
7.1.2 Cartolabe	14
7.2 New platforms	14
8 New results	15
8.1 Toward Good AI	15
8.1.1 Causal Modeling	15
8.1.2 Explainability	15
8.1.3 Robustness of AI Systems	16
8.2 Learning to Learn	17
8.2.1 Auto-*	17
8.2.2 Deep Learning: Practical and Theoretical Insights	18
8.2.3 Analyzing and Learning Complex Systems	18
8.3 Computational Social Sciences	19
8.3.1 Labor Studies	19
8.3.2 Health, food, and socio-demographic relationships	20
8.3.3 Scientific Information System and Visual Querying	21
8.3.4 Multi-Agent based simulation framework for social science	21
8.4 Energy Management	22
8.4.1 Power Grids Management	22
8.4.2 Optimization of Local Grids and the Modeling of Worst-case Scenarios	22
8.5 Data-driven Numerical Modelling	23

8.5.1	High Energy Physics	23
8.5.2	Remote Sensing Imagery	23
8.5.3	Space Weather Forecasting	24
8.5.4	Genomic Data and Population Genetics	24
8.5.5	Privacy and synthetic data generation	25
8.5.6	Sampling molecular conformations	26
8.5.7	Earthquake occurrence prediction	26
8.5.8	Reduced order model correction	27
8.5.9	Active Learning for chaotic systems	27
8.5.10	Control of fluid dynamics with reinforcement learning	28
8.5.11	Storm trajectory prediction	28
8.6	Challenges	29
9	Bilateral contracts and grants with industry	30
9.1	Bilateral contracts with industry	30
9.2	Bilateral grants with industry	31
10	Partnerships and cooperations	31
10.1	International initiatives	31
10.1.1	Inria International Labs	31
10.1.2	Inria associate team not involved in an IIL	31
10.1.3	Inria international partners	31
10.1.4	Participation in other international programs	31
10.2	International research visitors	32
10.2.1	Visits of international scientists	32
10.2.2	Visits to international teams	32
10.3	European initiatives	32
10.3.1	FP7 & H2020 Projects	32
10.3.2	Collaborations in European programs, except FP7 and H2020	33
10.3.3	Collaborations with major European organizations	33
10.4	National initiatives	33
10.4.1	ANR	33
10.4.2	Others	34
10.5	Regional initiatives	35
11	Dissemination	35
11.1	Promoting scientific activities	35
11.1.1	Scientific events: organisation	35
11.1.2	Scientific events: selection	35
11.1.3	Journal	35
11.1.4	Invited talks	36
11.1.5	Leadership within the scientific community	36
11.1.6	Scientific expertise	36
11.1.7	Research administration	37
11.2	Teaching - supervision - juries	37
11.2.1	Teaching	37
11.2.2	Supervision	38
11.2.3	Juries	39
11.3	Popularization	40
11.3.1	Internal or external Inria responsibilities	40
11.3.2	Articles and contents	40
11.3.3	Education	40
11.3.4	Interventions	40

12 Scientific production	40
12.1 Major publications	40
12.2 Publications of the year	41
12.3 Other	46
12.4 Cited publications	46

Project-Team TAU

Creation of the Project-Team: 2016 December 01

Keywords

Computer Sciences and Digital Sciences:

- A3.3.3. – Big data analysis
- A3.4. – Machine learning and statistics
- A3.5.2. – Recommendation systems
- A6.2. – Scientific computing, Numerical Analysis & Optimization
- A8.2. – Optimization
- A8.6. – Information theory
- A8.12. – Optimal transport
- 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 members, visitors, external collaborators

Research Scientists

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

Faculty Members

- Philippe Caillou [Univ Paris-Sud, Associate Professor]
- Aurélien Decelle [Univ Paris-Saclay, Associate Professor]
- Cécile Germain [Univ Paris-Saclay, Emeritus, HDR]
- Isabelle Guyon [Université Paris-Saclay and Inria, Professor, HDR]
- Francois Landes [Univ Paris-Saclay, Associate Professor]
- Nicolas Spyrtatos [Univ Paris-Saclay, Emeritus, HDR]

Post-Doctoral Fellows

- Olivier Bui [Inria, from Dec 2020]
- Jean Cury [Univ Paris-Saclay]
- Ksenia Gasnikova [Inria]
- Saumya Jetley [Inria]

PhD Students

- Eleonore Bartenlian [Ministère de l'Enseignement Supérieur et de la Recherche]
- Victor Berger [Inria]
- Guillaume Bied [Univ Paris-Saclay]
- Leonard Blier [Facebook, CIFRE]
- Tony Bonnaire [CNRS]
- Balthazar Donon [Réseau de transport d'électricité, CIFRE]
- Victor Estrade [Univ Paris-Sud]
- Loris Felardos - Saint Jean [Inria]
- Giancarlo Fissore [Univ Paris-Saclay]
- Julien Girard [CEA]

- Jeremy Guez [MNHN]
- Armand Lacombe [Inria]
- Wenzhuo Liu [Institut de recherche technologique System X]
- Zhengying Liu [École polytechnique]
- Nizam Makdoud [Thales, CIFRE]
- Emmanuel Menier [Institut de recherche technologique System X, from Sep 2020]
- Marc Nabhan [Renault, CIFRE, Until Dec. 2020]
- Adrien Pavao [Univ. Paris-Saclay (Région Île-de-France)]
- Adrian Pol [CERN, Until June 2020]
- Herilalaina Rakotoarison [Inria]
- Theophile Sanchez [Univ Paris-Sud]
- Vincenzo Schimmenti [CNRS, from November 2020]
- Nilo Schwencke [Univ Paris-Saclay, from Oct 2020]
- Marion Ullmo [Univ Paris-Sud]
- Elinor Wahal [Université Paris-Saclay]
- Pierre Wolinski [Inria, until June 2020]

Technical Staff

- Victor Alfonso Naya [Inria, Engineer]
- Berna Bakir Batu [Inria, Engineer]
- Laurent Basara [Inria, Engineer, until Jun 2020]

Interns and Apprentices

- Serigne Malick Ba [Inria, from Apr 2020 until Aug 2020]
- Kevin Cardenas Paez [Inria, Mar 2020]
- Martin Cepeda [Inria, from May 2020 until Aug 2020]
- Mirwaisse Djanbaz [Inria, from Jun 2020 until Sep 2020]
- Louis Dumont [Inria, from May 2020 until Oct 2020]
- Adrian El Baz [Inria, from May 2020 until Sep 2020]
- Louis Hernandez [Inria, from May 2020 until Aug 2020]
- Pierre Jobic [Univ Paris-Saclay, from Apr 2020 until Sep 2020]
- Mandie Joulin [Inria, from Sep 2020 until Feb 20]
- Ziheng Li [Inria, from May 2020 until Aug 2020]
- Mathieu Michel [Univ de Rennes I, from Mar 2020 until Jul 2020]
- Benoit Oriol [Inria, from Mar 2020 until Jul 2020]

- Francesco Pezzicoli [Inria, From February 2021 to July 2021]
- Michael Vaccaro [Inria, from Sep 2020]
- Clement Veysiere [Inria, from May 2020 until Sep 2020]
- Dhiaeddine Yousfi [Inria, from May 2020 until Aug 2020]

Administrative Assistants

- Laurence Fontana [Inria, until September 2020]
- Maeva Jeannot [Inria, from Oct 2020]

External Collaborators

- Aurélie Boisbunon [Mydatamodels, from Oct 2020]
- Romain Egele [École polytechnique]
- Burak Yelmen [Institut de génomique de l'Université de Tartu - Estonie]

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 **T**Ackling 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.

3 Research program

3.1 Toward Good AI

As discussed by [\[154\]](#), and in the recent collaborative survey paper [\[24\]](#), 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 [\[148\]](#), 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 legal¹ and societal² issues. While privacy is now recognized as a civil right in Europe, it is feared that the GAFAM, BATX and others can already

¹For 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.

²For instance related to information bubbles and nudge [\[117, 170\]](#).

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

The extraction of causal models, a long goal of AI [152, 130, 153], 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 models³.

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) [131]; 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

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.

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 [171].

³One 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.

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 [140] 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 8.1.3).

3.2 Hybridizing numerical modeling and learning systems

Participants: Michele Alessandro Bucci, Guillaume Charpiat, Cécile Germain, Isabelle Guyon, Marc Schoenauer, Michèle Sebag

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 control⁴ of the phenomena 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.

At the other extreme, machine learning offers the opportunity to model phenomena 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 [134], inverse problem solving [150], and sequential decision making [174, 98], 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* [84], 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.

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

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 8.5 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* [140], can be tackled in terms of domain adaptation [89, 116]. 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 phenomena 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 8.5. Note that the "Cracking the Glass Problem", described in Section 8.2.3 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 [166], supporting by construction *differentiable programming* [136].

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 [164], 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 [107]. Another recent approach uses two deep neural networks, one for the state of the system, the other for the equation itself [158]. 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 [94]. E. Bartenlian's PhD (started Oct. 2018), co-supervised by M. Sebag and F. Pascal (Centrale-Supélec), focusing on musical audio-to-score translation [165], 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 exercises using non verbal cues from its teacher [128]; 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 [122, 116], 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, [145] 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

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 [91], as well as in optimization at large [127], including combinatorial optimization and constraint satisfaction [133, 121] and continuous optimization [86]. 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 [126] (see also Section 3.4), including the AutoDL challenge series [141] launched in 2019⁵ (see also Section 8.6).

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 [162, 121, 87, 86, 120]. Collaborative filtering, considering that a problem instance "likes better" an algorithm configuration yielding a better performance, learns to recommend good algorithms to problem instances [168, 146]. Bayesian optimization proceeds by alternatively building a surrogate model of algorithm performances on *the* problem instance at hand, and tackling it [113]. This last approach currently is the prominent one; as shown in [146], the meta-features developed for AutoML are

⁵<https://autodl.chalearn.org/neurips2019>

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 [106], and still is very active. A more recent approach used in TAU [159] extends the Bayesian Optimization approach with a Multi-Armed Bandit algorithm to generate the full Machine Learning pipeline, competing with the famed AutoSKLearn [113] (see Section 8.2.1).

3.3.2 Information theory: adjusting model complexity and data fitting

Participants: Guillaume Charpiat, Marc Schoenauer, Michèle Sebag

In the 60s, Kolmogorov and Solomonoff provided a well-grounded theory for building (probabilistic) models best explaining the available data [163, 123], 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 [132], directly optimizing a bound on the Kolmogorov complexity of the dataset. More generally variational methods provide quantitative criteria to identify superfluous 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 [172]). Scattering transforms exploit similar principles [94]. 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 [161, 137]. Other approaches might proceed by combining simple models into more powerful ones, e.g. using "Context Tree Weighting" [178] or switch distributions [109]. Another option is to formulate neural architecture design as a reinforcement learning problem [88]; 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

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 [82]. 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) [177]. 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 networks), and make them sustainable through e.g. posterior drop-out [147], weight quantization and probabilistic binary networks [142]. 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 strengthened 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, [114] 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 [179], 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 [115].

A more recent research direction examines the quantity of information in a (deep) neural net along the random matrix theory framework [96]. It is addressed in Giancarlo Fissore's PhD, and is detailed in Section 8.2.3.

Finally, we note the recent surge in using ML to address fundamental physics problems: from turbulence to high-energy physics and soft matter (with amorphous materials at its core) [25] or astrophysics/cosmology as well. 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. In that direction, the PhD thesis of Marion Ullmo and Tony Bonnaire applying statistical method coming either from deep learning or statistical physics to the task of inferring the structure of the cosmic web has show great succes with recents results discussed in Section 8.2.3. 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 8.2.3.

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) [139, 31]. A CNRS grant (80Prime) finances a shared PhD (Vincenzo Schimmenti), 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 [124], to AutoML [125]. The *Higgs challenge* [80], 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 [111] including the *Job Candidate Screening Coopetition* [110]; the *Real Versus Fake Expressed Emotion Challenge* (ICCV 2017) [175]; the *Large-scale Continuous Gesture Recognition Challenge* (ICCV 2017) [175]; the *Large-scale Isolated Gesture Recognition Challenge* (ICCV 2017) [175].

Other challenges have been organized in 2020, or are planned for the near future, detailed in Section 8.6. 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 (EHES); 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 [119] 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 [135]. 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 [167].

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 8.3.

4.2 Energy Management

Participants: Isabelle Guyon, Marc Schoenauer, Michèle Sebag

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).

The daily maintenance 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 (~ 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 8.4 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 8.4.2).

4.3 Data-driven Numerical Modeling

Participants: Michele Alessandro Bucci, Guillaume Charpiat, Cécile Germain, Isabelle Guyon, Flora Jay, Marc Schoenauer, Michèle Sebag

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 [132, 122], 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 [102] 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 Lyapunov coefficients (which express the speed at which simulated trajectories diverge from the true ones) can provide a suitable theoretical framework.

5 Social and environmental responsibility

5.1 Footprint of research activities

Thanks to the pandemia, the impact of our activities regarding carbon footprint have decreased a lot, from our daily commute that have almost completely disappeared as we all switched to tele-working to the transformation of all conferences and workshops into virtual events. We all miss the informal discussions that took place during coffee breaks in the lab as well as during conferences. But when the pandemia vanishes, after the first moments of joy when actually meeting again physically with our colleagues, we will have to think of a new model for the way we work: we were indeed discussing before the pandemia about how to reduce the carbon footprint of the conferences, but now we know that there exist solutions, even though not perfect.

5.2 Impact of research results

All our work on Energy (see Sections 4.2) is ultimately targeted toward optimizing the distribution of electricity, be it in planning the investments in the power network by more accurate previsions of user consumption, or helping the operators of RTE to maintain the French Grid in optimal conditions.

At the outbreak of the covid pandemic in Europe, François Landes got involved in the ICUBAM projet, which aimed at easing the practitioners' job, by providing them with real-time (ICU) beds availability in nearby hospitals. The data was fed by doctors themselves, and they could in return easily picture the ongoing (un)availability of beds in participating hospitals, thus facilitating the task of patient transfer [73].

6 Highlights of the year

6.1 Awards

2020 BBVA Frontiers Frontiers in Research Award : Isabelle Guyon (with Vladimir Vapnik and Bernhard Schölkopf) for contributions to kernel methods (including the invention of SVMs), and causal modeling.

6.2 Selective Fundings

TAU secured the following funded research projects (see Section 10 for more details):

- FET European project **TRUST-AI**, *Transparent, Reliable and Unbiased Smart Tool for AI*, Marc Schoenauer.
- Chaire IA **HUMANIA**, *Democratizing Artificial Intelligence*, Isabelle Guyon
- ANR **HUSH** *HUman Supply cHain behind smart technologies*, Paola Tubaro
- ANR **SPEED** *Simulating Physical PDEs Efficiently with Deep Learning*, Michele Alessandro Bucci and Guillaume Charpiat.
- ANR **RoDAPoG** *Robust Deep learning for Artificial genomics and Population Genetics*, Flora Jay, Cyril Furtlehner and Guillaume Charpiat.
- Inria Challenge **OceanAI**, *AI, Data, Models for a Blue Economy*, Marc Schoenauer and Michèle Sebag
- MITI-contributions **TRIA** *Le TRavail de l'Intelligence Artificielle : éthique et gouvernance de l'automatation*, Paola Tubaro (coordinator).
- DATAIA **ML4CFD**, *Machine Learning for Computational Fluid Dynamics*, Michele Alessandro Bucci.

7 New software and platforms

7.1 New software

7.1.1 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 fifth 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.

News of the Year: Codalab statistics August 2020: Codalab exceeds 50,000 users, 1000 competitions (over 400 in the last year), and 600 submissions per day!

L2RPN July 2020: We launched a new Learning to Run a Power Network competition, in collaboration with ChaLearn and RTE. We have a robustness and an adaptability track. This is an NeurIPS 2020 competition.

Chagrade May 2020: We released a new application to help instructors use challenges in the classroom and grade them called Chagrade.

AutoDL April 2020: The NeurIPS AutoDL challenge ended. But the series of challenges on Automated Deep Learning, in collaboration with ChaLearn, Google Zurich, and 4Paradigm continues with AutoSeries and AutoGraph.

URL: <http://competitions.codalab.org>

Contact: Isabelle Guyon

7.1.2 Cartolabe

Name: 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.

Cartolabe has in particular been 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 sub-map 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. Cartolabe has also been applied in 2020 to the COVID-19 kaggle publication dataset (Cartolabe-COVID project) to explore these publications.

URL: <http://www.cartolabe.fr/>

Publication: [hal-02499006](https://hal.archives-ouvertes.fr/hal-02499006)

Contact: Philippe Caillou

Participants: Philippe Caillou, Jean-Daniel Fekete, Michèle Sebag, Anne-Catherine Letournel

Partners: LRI - Laboratoire de Recherche en Informatique, CNRS

7.2 New platforms

- **CODABENCH:** We developed as part of the COMETH EIT Health EU project a new open-source benchmark platform: (**Codabench**). It expands the possibilities of the **Codalab competitions platform**, an open source project of which we are community lead since 2015, by allowing users to set up benchmarks with a large number of datasets and/or algorithms, to perform systematic studies. A beta version has been released. The code is on **GitHub**.

8 New results

8.1 Toward Good AI

8.1.1 Causal Modeling

Participants: Philippe Caillou, Isabelle Guyon, Michèle Sebag;

PhDs: Armand Lacombe; **Post-doc:** Ksenia Gasnikova, Saumya Jetley **Collaboration:** Olivier Allais (INRAE); Jean-Pierre Nadal & Annick Vignes (CAMS, EHESS); David Lopez-Paz (Facebook).

The causal modelling activity has been continued in 2020 along two directions. The first one concerns the impact of nutrition on health. This study started in the context of the *Initiative de Recherche Stratégique* Nutriperso (2016-2018), headed by Louis-George Soler, INRAE, based on the wealth of data provided by the Kantar panel (170,000 bought products by 10,000 households over Year 2014). The challenges are manifold. Firstly, the number of potential causes is in the thousands, thus larger by an order of magnitude compared to most causal modelling studies. Secondly, a "same" product (e.g. "pizza") has vastly different impacts on health, depending on its composition and (hyper)processing. Lastly, the data is ridden with hidden confounders (e.g. with no information about smoking or sport habits).

On the one hand, the famed *Deconfounder approach* [176, 97, 149, 129] has been investigated and extended to account for the known presence of hidden confounders, as follows. A probabilistic model of the nutritional products based on Latent Dirichlet Allocation has been used, the factors of which are used as substitute confounders (SC) to block the effects of the confounders. On the other hand, the innovative notion of "micro-interventions" has been defined, operating on the basket of products associated to a household to e.g. replace the products with organic products; or increase the amount of alcohol ingested. The average treatment effect of the micro-interventions has been assessed conditionally to each SC, after correction for the biases related to the socio-economic description of the households. Submission in preparation.

Finally, causality is also at the core of TAU participation in the INRIA Challenge *OceanIA*, that will actually start in 2021 [56], and will analyze the ocean data fetched by the Tara expedition. Other motivating applications for causal modeling are described in section 4.1.

8.1.2 Explainability

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

PhD: Marc Nabhan **Collaboration:** MyDataModels

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 [160] 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 [19]. Other directions in this line of research include explaining missing data, with applications in computer vision [20].

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 identification of the border of the failure zone in the parameter space of the autonomous vehicle simulator, main topic of Marc Nabhan's PhD [63], also pertains to explainability (more details in Section 8.1.3).

A completely original approach to DNN explainability might arise from the study of structural glasses (8.2.3), with a parallel to Graph Neural Networks (GNNs), that could become an excellent non-trivial example for developing explainability protocols.

Generic Programming [85] is an Evolutionary Computing technique that evolves models as analytical expressions (Boolean formulae, functions, LISP-like code), that are hopefully easier to understand

than black-box NNs with hundreds of thousands of weights. This idea has been picked up by the European FET project **TRUST-AI** (Transparent, Reliable and Unbiased Smart Tool for AI) that started in October 2020. Alessandro Leite will be joining the project (and the TAU team) on February 1, 2021 on an ARP position. In the mean time, Marc Schoenauer is working together with the startup company *MyDataModels* whose lighthouse product is based on an original variant of Genetic Programming [69]. Both approach are promising just-started or on-going works.

8.1.3 Robustness of AI Systems

Participants: Guillaume Charpiat, Marc Schoenauer, Michèle Sebag

PhDs: Julien Girard, Marc Nabhan, Nizham Makhoud, Roman Bresson

Collaboration: Zakarian Chihani (CEA); Hiba Hage and Yves Tourbier (Renault); Johanne Cohen (LRI-GALAC) and Christophe Labreuche (Thalès); Eyke Hullermeier (U. Paderborn, Germany).

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. Note that several anecdotes are reported in [27], demonstrating that the need for safety bounds also arises in black-box optimization.

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 [46] 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.

To go further and alleviate the computational complexity of formally validating a neural network (naive complexity: exponential in the number of neurons), we explore different strategies to apply solvers to sub-problems that are much simpler. We rely on the fact that ReLU networks (the most common type of modern networks) are actually piecewise-linear, yielding extremely simple problems on each piece [50].

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 fully 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, defended in December 2020 [63]. 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. Several approaches for the identification of failures have been proposed. The thesis then focused on trying to obtain a precise yet simple definition of the border of the failure zone, using different methods, from MILP to Genetic Programming.

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 led 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 [138]. The results demonstrate that LEA is able to learn faster than DDPG if given good-enough policies, and only slightly slower when given lousy advices [51], and Nizam is now completing his PhD, to be defended in September 2021.

Learning Multi-Criteria Decision Aids (Hierarchical Choquet models)

Roman Bresson's PhD is co-supervised with Johanne Cohen (LRI-GALAC), Christophe Labreuche (Thalès) and Eyke Hullermeier (U. Paderborn). The transcription of hierarchical Choquet models (HCI) into a neural architecture enforcing by design the HCI constraints of monotonicity and additivity has been proposed, supporting the end-to-end learning of the HCI with a known hierarchy [42]. A patent (Bresson-Labreuche-Sebag-Cohen) has been filed by Thalès. The approach is being extended to achieve the automatic identification of the hierarchy as well; the unicity of the structure under canonic assumptions is being established.

8.2 Learning to Learn

8.2.1 Auto-*

Participants: Guillaume Charpiat, Isabelle Guyon, Marc Schoenauer, Michèle Sebag

PhDs: Léonard Blier, Guillaume Doquet, Zhengying Liu, Herilalaina Rakotoarison, Pierre Wolinski, Adrien Pavao, Hoazhe Sun

Collaborations: Vincent Renault (SME Artelys); Yann Ollivier (Facebook); Wei-Wei Tu (4Paradigm, Chine); André Elisseeff (Google Zurich); among others (for a full list see <https://autodl.chalearn.org/>)

Auto-★ studies at TAU investigate several research directions.

After proposing MOSAIC [159], that extends and adapts Monte-Carlo Tree Search to explore the structured space of pre-processing + learning algorithm configurations, and performs on par with AutoSklearn, the winner of Auto-★ international competitions in the last few years, Herilalaina Rakotoarison explored in his PhD an original approach in cooperation with Gwendoline de Bie and Gabriel Peyre (ENS). The neural learning from distributions proposed by Gwendoline [93] has been extended to achieve equivariant learning. Formally, the proposed DIDA architecture (Distribution-based Invariant Deep Architecture) learns from set of samples, regardless of the order of the samples and of their descriptive features. Two original tasks have been proposed to train a DIDA [71]: detecting whether two set of samples (with different descriptive features) are extracted from the same overall dataset; ranking two hyper-parameter configurations of a given classification algorithm) w.r.t. their predictive accuracy on the sample set. On both tasks, DIDA significantly outperforms the state of the art. Most interestingly, the main limitation incurred on the latter task (which constitutes a proto-task of AutoML) is the lack of sufficient data. Some augmentation process based on OpenML [173] was required to solve this latter task.

Several works have focused on the automatic adjustment of specific hyper-parameters for neural nets. Pierre Wolinski's PhD (defended in March 2020 [65]) 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 in the variational inference framework [77], which leads to a particular value that the regularizer weight should have in order the network to satisfy some property.

A last direction of investigation concerns the design of challenges, that contribute to the collective advance of research in the Auto- \star direction. The team has been very active in the series of AutoML [169] [59] and AutoDL [74]. An account of the AutoDL challenge series was published following a the NeurIPS 2020 competition track [48]. Post-challenge analyses were conducted on the Jean-Zay super computer and the results will be published in TPAMI paper under final revisions). Two new directions are being pursued: First, a series of challenges on meta-learning has been planned, whose first edition (focusing on few-shot learning) was accepted in conjunction with a workshop on meta-learning at the AAAI 2021 conference (with the sponsorship of Microsoft who provided Azure credits). A scaled up version was submitted to the competition program of NeurIPS 2021. Second, a challenge on Neural Architecture Search (NAS) is in preparation and has been accepted together with a workshop at the CVPR 2021 conference. Preliminary results on NAS have been produced by one of our interns (Romain Egele [72]). More details on challenges are found in Section 8.6).

8.2.2 Deep Learning: Practical and Theoretical Insights

Participants: Guillaume Charpiat, Isabelle Guyon, Marc Schoenauer, Michèle Sebag

PhDs: Léonard Blier, Zhengying Liu

Collaboration: Yann Ollivier (Facebook AI Research, Paris)

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.

During his CIFRE PhD with Facebook AI Research Paris, co-supervised by Yann Ollivier (former TAU member), Léonard Blier has properly formalized the concepts of *successor states* and *multi-goal functions* [68], in particular in the case of continuous state spaces. This allowed him to define unbiased algorithms with finite variance to learn such objects, including the continuous case thanks to approximation functions. In the case of finite environments, new convergence bounds have been obtained for the learning of the value function. These new algorithms capable of learning successor states in turn lead to define and learn new representations for the state space.

The AutoDL challenges, co-organized in TAU (in particular by Isabelle Guyon, and by Zhengying Liu within his PhD), also contribute to a better understanding of Deep Learning. It is interesting to note that no Neural Architecture Search algorithm was proposed to solve the different challenges in AutoDL (corresponding to different data types). See section 8.6 for more details.

The meta-learning setting, which was devised for the Auto \star challenges was analyzed theoretically by Zhengying Liu [55]. Assuming the perfect knowledge of the meta-distribution (i.e., in the limit of a very large number of training tasks), the paper investigates under which conditions algorithm recommendation can benefit from meta-learning, and thus, in some sense, “defeat” the No-Free-Lunch theorem. Four meta-predict strategies are analyzed: Random, Mean, Greedy and Optimal. Conditions of optimality are investigated and experiments conducted on artificial and real data.

8.2.3 Analyzing and Learning Complex Systems

Participants: Cyril Furtlehner, Aurélien Decelle, François Landes

PhDs: Giancarlo Fissore, Tony Bonnaire, Marion Ullmo

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), Nabila Aghanim.

Generative models constitute an important piece of unsupervised ML techniques which is still under rapid development. In this context insights from statistical physics are relevant in particular for energy based models like restricted Boltzmann machines [16, 15, 61]. The information content of a trained restricted Boltzmann machine (RBM) and its learning dynamics can be analyzed precisely with help of

ensemble averaging techniques [99, 100]. 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 leads to a set of mean-field fixed points covering the sample distribution. More insight (work on the verge of being submitted) can be obtained by looking at data of low intrinsic dimension, where exact solutions of the RBM can be obtained thanks to a convex relaxation, along with a Coulomb interpretation of the model, allowing to detect important shortcomings of standard training procedures and their possible resolution in views of concrete applications [41]. Other aspects of generative models are being explored in G. Fissore's PhD, as the connection with independent components analysis (ICA). In this context an efficient learning algorithm has been proposed [47] for the training of deep auto-regressive normalizing flows based on a relative gradient, allowing for density estimators of data embedded in a large dimensional space.

As mentioned earlier, the use of ML to address fundamental physics problems is quickly growing. In that direction two different directions have been taken. On one hand, the PhD thesis of M. Ullmo and T. Bonnaire is focusing on dealing with the characterization of the cosmic web (the baryonic structure taking place at large scale in our universe) in order to track the so-called missing baryons of the standard theory. M. Ullmo demonstrated the feasibility of using Generative Adversarial Network (GAN) on the distributions dark matter at cosmological scale (up to hundreds of Mpc) both using data coming from 2D simulation and 3D simulations [76]. In that setting, she also developed a novel building an encoder capable of inferring the latent structure of the GAN for a given image and showing that many details are recovered. T. Bonnaire on his side worked on designing a new method in order to classifying the structure of the cosmic web into clusters and filaments, directly from the position of the dark matter galaxies. To do so, he developed a method based on the Gaussian mixture model with a prior forcing the centers to "live" on a tree-graph: two centers sharing an edge on this graph benefit from an attractive attraction, forcing the algorithm to adapt the center's position taking into account both the density distribution and the shape of the prior [12, 70].

On the other hand, it 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 [66]. 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 [25, 70], 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, and more recently with a promising M2 internship (Francesco Pezzicoli). Furthermore, this problem is new to the ML community and it provides an original non-trivial example for engineering, testing and benchmarking explainability protocols.

8.3 Computational Social Sciences

Computational Social Sciences (CSS) is making significant progress in the study of social and economic phenomena thanks 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, the labor of human annotators for AI data, 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.

8.3.1 Labor Studies

Participants: Philippe Caillou, Isabelle Guyon, Michèle Sebag, Paola Tubaro

PhDs: Guillaume Bied, Armand Lacombe, Elinor Wahal

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); Francesca Musiani, Mélanie Dulong de Rosnay (CNRS); José Luis Molina (Universitat Autònoma de Barcelona); Antonio Ortega (Universitat de València); Julian Posada (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 human labor behind AI

We look at data "micro-workers" who perform essential, yet marginalized and poorly paid tasks such as labeling objects in a photograph, translating or transcribing short texts, or recording utterances. Formally, micro-workers are independent entrepreneurs, with varying levels of participation and commitment. They are recruited through speacialized intermediaries such as Amazon Mechanical Turk and Microworkers, or suppliers for monopsonist technology giants (such as UHRS for Microsoft).

Micro-workers perform three necessary functions for AI: preparation (data generation and annotation), verification (checking outputs), and impersonation (taking the place of a not-yet-functioning AI) [37]. There are about 260000 micro-workers in France, although many perform these tasks as a side activity, occasionally or intermittently [38]. They mostly include highly-educated people in low-paid jobs, often with family duties [58].

Like many online platform workers, micro-laborers operate outside standard labor and social protection laws [26], contributing to the reinforcement of legacy inequalities [32] and the emergence of new ones [33] that the Covid-19 pandemic has made more apparent [34].

The possibility to use micro-work for research purposes (for example, in online surveys and experiments) raises specific ethical issues that add to the rising number of ethical challenges in today's science [39].

Current work extends this research to map the global production networks of artificial intelligence and their organizational structures [36]. It also looks at how the demand of researchers and companies from (in particular) French and Spanish-speaking countries reaches a chep supply of workers in French-speaking Africa and Spanish-speaking Latin America.

8.3.2 Health, food, and socio-demographic relationships

Participants: Philippe Caillou, Michèle Sebag, Paola Tubaro

PhD: Armand Lacombe; **Post-doc:** Ksenia Gasnikova, Saumya Jetley

Collaboration: Louis-Georges Soler, Olivier Allais (INRA); Jean-Pierre Nadal, Annick Vignes (CAMS, EHESS)

Another area of activity concerns the relationships between eating practices, socio-demographic features and health, and its links with causal learning (see also Section 8.1.1), that has been continued in 2020.

The study about the impact of nutrition on health started in the context of the *Initiative de Recherche Stratégique* Nutriperso (2016-2018), headed by Louis-George Soler, INRAE, based on the wealth of data provided by the Kantar panel (170,000 bought products by 10,000 households over Year 2014). The challenges are manifold. Firstly, the number of potential causes is in the thousands, thus larger by an order of magnitude compared to most causal modelling studies. Secondly, a "same" product (e.g. "pizza") has vastly different impacts on health, depending on its composition and (hyper)processing. Lastly, the data is ridden with hidden confounders (e.g. with no information about smoking or sport habits).

On the one hand, the famed *Deconfounder approach* [176, 97, 149, 129] has been investigated and extended to account for the known presence of hidden confounders, as follows. A probabilistic model of the nutritional products based on Latent Dirichlet Allocation has been used, the factors of which are used as substitute confounders (SC) to block the effects of the confounders. On the other hand, the innovative notion of "micro-interventions" has been defined, operating on the basket of products associated to a household to e.g. replace the products with organic products; or increase the amount of alcohol ingested. The average treatment effect of the micro-interventions has been assessed conditionally to each SC, after correction for the biases related to the socio-economic description of the households. Submission in preparation.

8.3.3 Scientific Information System and Visual Querying

Participants: Philippe Caillou, Michèle Sebag

Engineers: Anne-Catherine Letournel, Victor Alfonso Naya

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 (see also Section 7) 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).

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, Bibliothèque 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. Cartolabe has also been applied in 2020 to the COVID-19 kaggle publication dataset (Cartolabe-COVID project) to explore these publications.

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 has been performed in 2019-2020. 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. A statement of the current state of the project was published in 2020 [13].

8.3.4 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 2020 was on the stability of the platform and on the documentation to provide a stable and well documented framework to the users.

8.4 Energy Management

8.4.1 Power Grids Management

Participants: Isabelle Guyon, Marc Schoenauer

PhDs: Balthazar Donon, Wenzhuo Liu

Collaboration: Rémi Clément, Patrick Panciatici (RTE)

Our collaboration with RTE, during Benjamin Donnot's (2016-2019) [103] and Balthazar Donon's (since 2019) CIFRE PhDs, is centered on the maintainance of the national French Power Grid. In order to maintain the so-called "(n-1) safety" (see Section 4.2), fast simulations of the electrical flows on the grid are mandatory, that the home-brewed simulator HADES is too slow to provide. The main difficulty of using Deep Neural Networks surrogate models is that the topology of the grid (a graph) should be taken into account, and because all topologies cannot be included in the training set, this requires out-of-sample generalization capabilities of the learned models. An original "guided dropout" approach was first proposed [104], and later extended and generalized with the LEAP (Latent Encoding of Atypical Perturbation) architecture [18], 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 tology, 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, 387 lines and 192 substations.

Departing from the LEAP architecture, Balthazar Donon developped a completely different approach based on Graph Neural Networks (GNNs). From a Power Grid perspective, GNNs can be viewed as including the topology in the heart of the structure of the neural network, and learning some generic transfer function amongst nodes that will perform well on any topology. First results [105] [17] 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. Further on-going work [45] completely removed the need to run HADES, thanks to a loss that directly aims to minimize Kirshoff's law on all lines. Theoretical results as well as a generalization of the approach to other optimization problems on graphs are at the heart of Baltazar Donon's PhD (to be defended in Fall 2021).

8.4.2 Optimization of Local Grids 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 Renault (Artelys), Gabriel Peyré and Gwendoline de Bie (ENS).

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 8.2.1); while the Mosaic algorithm is validated on standard AutoML benchmarks [159], 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 [90]. The corresponding approach has been extended to study the trade-off between the optimization of the reconstruction loss and the latent compression of VAEs, both theoretically and numerically [67].

8.5 Data-driven Numerical Modelling

8.5.1 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 [112] 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 [81] and the TRackML challenge [83]. Recent work indicates that the specific issue of extremely poorly separated classes should be addressed through a combination dataset-level inference and iterative refinement of the particle selection.

Anomaly detection (AD) is the subject of A. Pol's PhD [64]. 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. [157] focused 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 [60] 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 [155, 156]. This results in an effective, yet easily trainable and maintainable model.

8.5.2 Remote Sensing Imagery

Participants: Guillaume Charpiat

Collaboration: Yuliya Tarabalka, 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 [143] and, soon after, larger ones by GAFAM).

This year, we concluded a long series of publications with the PhD defense of Nicolas Girard [62].

8.5.3 Space Weather Forecasting

Participants: Cyril Furtlehner, Michèle Sebag

Post-doc: Olivier Bui

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 Multi-scale 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 [95] which has been defended this year in Eindhoven. One 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 [43] 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.

8.5.4 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), UNAM (Mexico), U Brown (USA), U Cornell (USA), TIMC-IMAG (Grenoble), MNHN (Paris), Pasteur Institute (Paris)

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.

⁶After 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.

For the past decades, simulation-based likelihood-free inference methods have enabled researchers to address numerous population genetics problems. As the richness and amount of simulated and real genetic data keep increasing, the field has a strong opportunity to tackle tasks that current methods hardly solve. However, high data dimensionality forces most methods to summarize large genomic datasets into a relatively small number of handcrafted features (summary statistics). In [35], we propose an alternative to summary statistics, based on the automatic extraction of relevant information using deep learning techniques. Specifically, we design artificial neural networks (ANNs) that take as input single nucleotide polymorphic sites (SNPs) found in individuals sampled from a single population and infer the past effective population size history. First, we provide guidelines to construct artificial neural networks that comply with the intrinsic properties of SNP data such as invariance to permutation of haplotypes, long scale interactions between SNPs and variable genomic length. Thanks to a Bayesian hyperparameter optimization procedure, we evaluate the performance of multiple networks and compare them to well established methods like Approximate Bayesian Computation (ABC). Even without the expert knowledge of summary statistics, our approach compares fairly well to an ABC based on handcrafted features. Furthermore we show that combining deep learning and ABC can improve performance while taking advantage of both frameworks. Later, we experimented with other types of permutation invariance, based on similar architectures, and achieved a significant performance gain with respect to the state of the art, including w.r.t. ABC on summary statistics (20% gap), which means that we extract information from raw data that is not present in summary statistics. The question is now how to express this information in a human-friendly way.

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, collab: Pasteur Institute, for ancient DNA: UNAM and U Brown); the longer-term goal is to integrate them in various systems handling genetic data or other biological sequence data. Regarding the bacterial populations, we already implemented a flexible simulator that will allow researchers to investigate complex evolutionary scenarios (e.g. dynamics of antibiotic resistance in 2D space through time) with realistic biological processes (bacterial recombination), which was impossible before (collab. U Cornell, MNHN) [14].

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 [41]. 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 held by companies or public institutes like the Institute of Genomics of Tartu; the latter has a rich diversity within Estonia and informs on population past and recent history as we showed in [30] via population genetics mathematical modelling, inference and simulation). 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.

In collaboration with TIMC-IMAG, we proposed a new factor analysis approach that processes 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 [21]. We also explored other visualization tools, as alternatives to the more classic phylogenetic trees, for present-day populations[54].

8.5.5 Privacy and synthetic data generation

Participants: Isabelle Guyon

PhD: Adrien Pavao

Collaboration: Kristin Bennett (RPI)

Collecting and distributing actual medical data is costly and greatly restrained by laws protecting patients' health and privacy. While beneficial, these laws severely limit access to medical data thus stagnating innovation and limiting research and educational opportunities. The process of obfuscation of medical data is costly and time consuming with high penalties for accidental release. Thus, we have

engaged in developing and using realistic simulated medical data in research and in teaching. In [40] we develop metrics for measuring the quality of synthetic health data for both education and research. We use novel and existing metrics to capture a synthetic dataset's resemblance, privacy, utility and footprint. Using these metrics, we develop an end-to-end workflow based on our generative adversarial network (GAN) method, HealthGAN, that creates privacy preserving synthetic health data. Our workflow meets privacy specifications of our data partner: (1) the HealthGAN is trained inside a secure environment; (2) the HealthGAN model is used outside of the secure environment by external users to generate synthetic data. In [49] we put the HealthGAN methodology that we developed in the previous paper to work in a practical setting. We reproduce the research outcomes obtained on two previously published studies, which used private health data, using synthetic data generated with a method that we developed, called HealthGAN. We demonstrate the value of our methodology for generating and evaluating the quality and privacy of synthetic health data. The dataset are from OptumLabs R Data Warehouse (OLDW). The OLDW is accessed within a secure environment and doesn't allow exporting of patient level data of any type of data, real or synthetic, therefore the HealthGAN exports a privacy-preserving generator model instead. The studies examine questions related to comorbidities of Autism Spectrum Disorder (ASD) using medical records of children with ASD and matched patients without ASD. HealthGAN generates high quality synthetic data that produce similar results while preserving patient privacy. In [44], we extend existing time-series generative models to generate medical data, which is challenging due to this influence of patient covariates. We propose a workflow wherein we leverage existing generative models to generate such data. We demonstrate this approach by generating synthetic versions of several time-series datasets where static covariates influence the temporal values.

8.5.6 Sampling molecular conformations

Participants: Guillaume Charpiat

PhD: Loris Felardos

Collaboration: Jérôme Héning (IBPC), Bruno Raffin (InriaAlpes)

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. We study how to bypass simulations by directly predicting, given a molecule formula, its possible configurations. This is done using Graph Neural Networks [108] in a generative way, producing 3D configurations. The goal is to sample all possible configurations, and with the right probability.

8.5.7 Earthquake occurrence prediction

Participants: François Landes, Marc Schoenauer

PhD: Vincenzo Schimmenti

Collaboration: Alberto Rosso (LPTMS)

Earthquakes occur in brittle regions of the Crust typically located at the depth of 5-15 km and characterized by a solid friction, which is at the origin of the stick-slip behaviour. Their magnitude distribution displays the celebrated Gutenberg-Richter law and a significant increase of the seismic rate is observed after large events (called main shocks). The occurrence of the subsequent earthquakes in the same region, the aftershocks, obeys well established empirical laws that demand to be understood. A change in the seismic rate also happens before a main shock, with an excess of small events compared to the expected rate of aftershocks related to the previous main shock in that region. These additional events are defined as foreshocks of the coming main shock, however they are scarce so that defining them is a very difficult task. For this reason their statistical fingerprint, so important for human security, remains elusive. In this project we combine the techniques of Statistical Physics and Machine Learning to determine the complex spatio-temporal patterns of the events produced by the dynamics of the fault. In

particular we plan to understand the structure of the short sequence of foreshocks, and their potential impact for human applications.

The treatment of rare events by Machine Learning is a challenging yet rapidly evolving domain. At TAU we have a great expertise in data modeling, in particular we are currently working on space weather forecast, a supervised task where, like in seismicity, extreme and rare events are crucial. Bayesian models and Restricted Boltzman Machines (RBMs) have been built to model these weather forecast data. These techniques, inspired from statistical physics, are both based on a probabilistic description of latent variables (i.e. unobserved variables) and have great expressiveness, allowing the modelling of a large span of data correlations. This kind of models can be extended to study spatially resolved earthquakes, the latent variable here being the local stress within the fault and in the ductile regions. Our goal is to characterize the statistical properties of a sequence of events (foreshocks, main shock and aftershocks) and predict its following history. We will first study the sequences obtained from simulations of the physical model [31]. We will answer the following question: given a short sequence of foreshocks, can we predict the future of the sequence? How big will be the main shock? When will it occur? In a second step we will use also the data coming from real sequences, where events are unlabeled. These sequences are public and available (The most accurate catalog is for Southern California, a catalog with 1.81 million earthquakes. It is available at <https://scedc.caltech.edu/research-tools/QTMcatalog.html>). Concretely, the data consists in the earthquakes' magnitude, occurrence time and hypocenter locations.

Two parallel directions are being explored, with our PhD Student, Vincenzo Schimmenti:

- The available data can be used to tune the parameters of the new model to improve its accuracy and generalization properties. We will adjust the parameters of the elastic and friction coefficients in order to produce earthquakes with realistic magnitudes. This will allow us to have information about the physical condition in the fault and in the ductile regions.
- We will use our understanding of foreshocks statistics to perform classification of earthquakes with respect to their nature: foreshock, main shock or after shock, and alignment (assignment of the earthquake to a sequence). These labels are known in the synthetic data and unknown in the catalogs, so this would be an instance of semi-supervised learning. Our final goal is real data completion: presented with an incomplete catalog, the machine is asked to complete it with the missing points.

8.5.8 Reduced order model correction

Participants: Bucci Michele Alessandro, Marc Schoenauer

PhD: Emmanuel Menier

Collaboration: Mouadh Yagoubi (IRT-SystemX)

Numerical simulations of fluid dynamics in industrial applications require the spatial discretization of complex 3D geometries with consequent demanding computational operations for the PDE integration. The computational cost is mitigated by the formulation of Reduced Order Models (ROMs) aiming at describing the flow dynamics in a low dimensional feature space. The Galerkin projection of the driving equations onto a meaningful orthonormal basis speeds up the numerical simulations but introduces numerical errors linked to the underrepresentation of dissipative mechanisms.

Deep Neural Networks can be trained to compensate missing information in the projection basis. By exploiting the projection operation, the ROM correction consists in a forcing term in the reduced dynamical system which has to i) recover the information living in the subspace orthonormal to the projection one ii) ensure that its dynamic is dissipative. A constrained optimization is then employed to minimize the ROM errors but also to ensure the reconstruction and the dissipative nature of the forcing. We tested this solution on benchmarked cases where it is well known that transient dynamics are poorly represented by ROMs. The result shows how the correction term improves the prediction while preserving the guarantees of the ROM.

8.5.9 Active Learning for chaotic systems

Participants: Michele Alessandro Bucci

Collaboration: Lionel Mathelin (LISN), Onofrio Semeraro (LISN), Sergio Chibbaro (UPMC), Alexander Allauzen (ESPCI)

The inference of a data driven model aiming at reproducing chaotic systems is challenging even for the most performing Neural Network architectures. According to the ergodic theory, the amount of data required to converge the invariant measure of a chaotic system goes exponentially with its intrinsic dimension. It follows that for learning the dynamics of a turbulent flow, the computing resources in the world would not be enough to store the necessary data. To circumvent such limitations we generally introduce constraints in the optimization stage in order to preserve physical invariants, when they are known.

In [52] we compare model quality when trained with and without ergodic time series generated by the Lorenz systems (i.e. the chaotic system related to the “butterfly effect”). The ergodic dataset is composed of one long trajectory (27000 time steps), whereas the non ergodic one is composed by 9 short trajectories (9000 time steps each) randomly initialized on the chaotic attractor. Despite the same amount of points, it turns out that the non-ergodic dataset led to biased models. Short trajectories do not ensure statistical knowledge of the phase space. Exploiting the structure of the phase space, 9 trajectories (9000 time steps) emanated from the 3 fix points of the Lorenz systems have been used to generate a new dataset. The fix points and their unstable directions define the skeleton of the phase space. The trajectories emanated from them allow to reduce the entropy of the dataset without introducing bias in the learned models. A dataset incorporating the dynamics around the fix points, not only allows to obtain more robust models with respect to the initialization of the NN parameters but also allows to reduce the size of the dataset by 60% without affecting the quality of the models.

8.5.10 Control of fluid dynamics with reinforcement learning

Participants: Michele Alessandro Bucci

Collaboration: Lionel Mathelin (LISN), Onofrio Semeraro (LISN), Thibaut Guegan (PPrime), Laurent Cordier (PPrime)

The control of fluid dynamics is an active research area given the implications of aerodynamic forces in the transport and energy field. Being able to delay the laminar-to-turbulent transition, stabilize unsteady mechanisms or reduce the pressure forces of an object moving in a fluid, would allow for more ecological vehicles or more efficient wind turbines. For quadratic objective functions and for conditions in which the linearized Navier Stokes equation is a good approximation of the fluid dynamics around the target state, optimal control theory provides the necessary tools (e.g. Riccati equation, direct-adjoint optimization) to recover a robust control policy. In the case of non-linearizable systems, non-quadratic cost functions or in the absence of a model, these tools are no longer valid. Reinforcement learning algorithms allow us to solve the optimal problem even if the model is not available. The control problem, with an infinite time horizon, can be decomposed into local optimal problems if the system is completely observed and its dynamics is Markovian. The solution of the Bellman equation ensure the optimality of the policy if the phase space of the system has been fully explored [57].

We applied actor-critic algorithms (TD3) to control a benchmarked flow configuration: the PinBall case [53]. In the PinBall case, the flow impacting on three cylinders arranged at the vertices of an equilateral triangle generates an unstable wake that causes high aerodynamic forces. Allowing the cylinders to rotate, the RL algorithm provides a control policy capable of reducing the drag by 60% compared to the uncontrolled case. We have also shown how partial observation of the flow velocity field through sensors is not a limiting factor if a temporal state embedding is considered. By reducing the number of sensors and increasing the size of the state with past observations, the efficiency of the policy is not degraded.

8.5.11 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 [101]. 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 [92]. 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 years [118] shows that with deep learning one can outperform the state-of-the-art in many cases [22].

8.6 Challenges

Participants: Cécile Germain, Isabelle Guyon, Adrien Pavao, Anne-Catherine Letournel, Michèle Sebag

PhD: Zhengying Liu, Balthazar Donon, Adrien Pavao, Haozhe Sun

Collaborations: D. Rousseau (LAL), André Elisseff (Google Zurich), Jean-Roch Vilmant (CERN), Antoine Marot and Benjamin Donnot (RTE), Kristin Bennett (RPI), Magali Richard (Université de Grenoble), Wei-Wei Tu (4Paradigm, Chine), Sergio Escalera (U. Barcelona, Espagne).

The TAU group 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 TAU group is community lead of the open-source **Codalab** platform (see Section 7), 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 [125] – PKDD 2018 [126]), a series of challenges on the theme of **AutoDL** [141] 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 [28]. 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. The challenge series is continuing beyond AutoDL, with the **AutoGraph** challenge that was organized for KDD 2020 <https://www.automl.ai/competitions/3> and the newly started Meta-Learning challenge series <https://metalearning.chalearn.org/>, whose first edition took place in conjunction with AAAI 2021. A new challenge on automated reinforcement learning (AutoRL) is currently under design.

A new challenge series in Reinforcement Learning was started with the company RTE France, one the theme “Learning to run a power network” [144] (**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 launched 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 accepted as part of the official program of NeurIPS 2020 [29]. While first competitions aimed at demonstrating the feasibility of applying Reinforcement Learning for controlling electrical flows on a power grid, the NeurIPS competition introduced a realistically-sized grid environment along with two fundamental real-life properties of power grid systems to reconsider while shifting towards a sustainable world: robustness and adaptability. The analysis paper is under review.

The **COMETH** 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 the newly developed Codabench platform (<https://www.codabench.org/>). The challenge gathered transdisciplinary instructors (researchers and professors), students, and health professionals (clinicians). The COMETH project aimed at creating benchmarks permitting practitioners to gain access to advanced algorithms provided by machine learning researchers. We developed a WebApp, which interfaces Codabench with a simplified interface designed for Physicians and which makes robot submissions to Codabench on their behalf. 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.

We have also shared our expertise (and made our challenge platform Codalab available) to support two other NeurIPS challenges: The Black Box Optimization for Machine Learning challenge <https://bbochallenge.com/>, which was then used as part of a class of optimization of the M2 AI of universit  Paris-Saclay; and the Predicting Generalization in Deep Learning challenge <https://sites.google.com/view/pgdl2020>. The latter case is remarkable: Google research selected our platform Codalab to run their challenge, despite the fact that they bought a competing commercial platform (Kaggle). Codalab was also chosen for the second phase of the TrackML challenge, in collaboration with LHC experiments. The goal was to build an algorithm that quickly reconstructs particle tracks from 3D points left in the silicon detectors [11].

The **Paris Ile-de-France project** also took off this year. Codalab and the TAU team were selected to organize the industry machine learning challenge series of the Paris REgion. Adrien pavao, who was the project leader, organized with Dassault aviation a project of “jumeau numerique”, aiming at performing predictive maintenance on airplanes. The Paris Region is offering 500K Euros to the winner, a startup, which would then collaborate with Dassault to productize the solution. The challenge was launched in February 2021 and we are awaiting the results. With this first challenge, Codalab has demonstrated that it is “industry grade”.

It is important to introduce **challenges in ML teaching**. This has been done (and is on-going) in I. Guyon’s Master courses [151] : some assignments to Master students are to **design small challenges**, which are then given to other 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 Mich le Sebag. Along similar line, F. Landes proposed a **challenge** in the context of S. Mallat’s course, at Coll ge 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>).

In terms of dissemination, a collaborative book “AI competitions and benchmarks: The science behind the contests ” written by expert challenge organizers is under way and will appear in the Springer series on challenges in machine learning, see <http://www.chalearn.org/books.html>.

9 Bilateral contracts and grants with industry

9.1 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 led 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.

- **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)
- **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 (72 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 (Alessandro Bucci) to work in all topics mentioned in Section 3.2 relevant to IFPEN activity.
Coordinator: Marc Schoenauer
Participants: Alessandro Bucci, Guillaume Charpiat

9.2 Bilateral grants with industry

10 Partnerships and cooperations

10.1 International initiatives

10.1.1 Inria International Labs

10.1.2 Inria associate team not involved in an IIL

10.1.3 Inria international partners

- Collaboration with CWI Amsterdam on Space weather forecasting (Inria funded post-doc Olivier Bui).

Declared Inria international partners

Informal international partners

10.1.4 Participation in other international programs

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

- **HFSP RGY0075/2019** 2019-2022 (1 M\$) Evolutionary changes in human hosts and their pathogens during first contact in the New World, funded by Human Frontier Science Program
Coordinator: Emilia Huerta-Sanchez (U Brown, USA)
Participants: Flora Jay
Collaboration: M Avila-Arcos (UNAM, Mexico)

10.2 International research visitors

10.2.1 Visits of international scientists

10.2.2 Visits to international teams

Sabbatical programme

Explorer programme

Research stays abroad

- Flora Jay: invited researcher at U Brown, CCMB (1 week, Feb 2020)
- Paola Tubaro: invited researcher at U Greenwich, London (1 week, Feb 2020)

10.3 European initiatives

10.3.1 FP7 & H2020 Projects

TRUST-AI

Title: Transparent, Reliable and Unbiased Smart Tool for AI

Duration: Oct. 2020 – Sept. 2023

Coordinator: INESC TEC - INSTITUTO DE ENGENHARIA DE SISTEMAS E COMPUTADORES, TECNOLOGIA E CIENCIA, Portugal (Gonçalo Figueira)

Partners: 6 partners, see <https://cordis.europa.eu/project/id/952060/fr>

Inria contact: Marc Schoenauer (TAU)

Summary: Due to their black-box nature, existing artificial intelligence (AI) models are difficult to interpret, and hence trust. Practical, real-world solutions to this issue cannot come only from the computer science world. The EU-funded TRUST-AI project is involving human intelligence in the discovery process. It will employ 'explainable-by-design' symbolic models and learning algorithms and adopt a human-centric, 'guided empirical' learning process that integrates cognition. The project will design TRUST, a trustworthy and collaborative AI platform, ensure its adequacy to tackle predictive and prescriptive problems and create an innovation ecosystem in which academics and companies can work independently or together.

TAILOR

Title: Foundations of Trustworthy AI - Integrating Learning, Optimization and Reasoning.

Duration: Sept. 2020 – Aug. 2023

Coordinator: Linkopings Universitet, Sweden (Fredrik Heintz).

Partners: 53 partners, see <https://cordis.europa.eu/project/id/952215/fr>

Inria contact: Marc Schoenauer (TAU)

Summary: Maximising opportunities and minimising risks associated with artificial intelligence (AI) requires a focus on human-centred trustworthy AI. This can be achieved by collaborations between research excellence centres with a technical focus on combining expertise in the areas of learning, optimisation and reasoning. Currently, this work is carried out by an isolated scientific community where research groups are working individually or in smaller networks. The EU-funded TAILOR project aims to bring these groups together in a single scientific network on the Foundations of Trustworthy AI, thereby reducing the fragmentation and increasing the joint AI research capacity of Europe, helping it to take the lead and advance the state-of-the-art in trustworthy AI. The four main instruments are a strategic roadmap, a basic research programme to address grand challenges, a connectivity fund for active dissemination, and network collaboration activities.

VISION

Title: Value and Impact through Synergy, Interaction and coOperation of Networks of AI Excellence Centres

Duration: Sept. 2020 – Aug. 2023

Coordinator: Leiden University, Holland (Holger Hoos)

Partners: 8 partners, see <https://cordis.europa.eu/project/id/952070/fr>.

Inria contact: Joost Geurts (DPE)

Summary: Artificial intelligence (AI) is an area of strategic importance and a key driver of economic development, bringing solutions to many societal challenges ranging from treating diseases to minimising the environmental impact of farming. The EU is focussing on connecting and strengthening AI research centres across Europe and supporting the development of AI applications in key sectors. To ensure Europe stays at the forefront of AI developments, the EU-funded VISION project will build on Europe's world-class community of researchers. The project will also build on the success and organisation of CLAIRE (the Confederation of Laboratories for AI Research in Europe) as well as of AI4EU, which was established to set up the first European Artificial Intelligence On-Demand Platform and Ecosystem.

10.3.2 Collaborations in European programs, except FP7 and H2020

10.3.3 Collaborations with major European organizations

- **C-INDL 2020** (10k euros), *Consolidating the International Network on Digital Labour*, funded by Fondazione Feltrinelli (Milan, Italy) and Independent Social Research Foundation (ISRF).
Coordinator: Paola Tubaro and Antonio A. Casilli (Telecom Paris)

10.4 National initiatives

10.4.1 ANR

- **EPITOME 2017-2020** (225kEuros), *Efficient rePresentatIon TO structure large-scale satellite iMagEs* (Section 8.5.2).
Coordinator: Yuliya Tarabalka (Titane team, INRIA Sophia-Antipolis)
Participant: Guillaume Charpiat
- **Chaire IA HUMANIA 2020-2024** (600kEuros), *Democratizing Artificial Intelligence* (Section 8.1).
Coordinator: Isabelle Guyon (TAU)
Participants: Marc Schoenauer, Michèle Sebag, Anne-Catherine Letournel, François Landes.
- **HUSH 2020-2023** (348k euros), *HUman Supply cHain behind smart technologies*.
Coordinator : Antonio A. Casilli (Telecom Paris)
Participants: Paola Tubaro

- **SPEED** 2021-2024 (49k€) *Simulating Physical PDEs Efficiently with Deep Learning*
Coordinator: Lionel Mathelin (LIMSI)
Participants: Michele Alessandro Bucci, Guillaume Charpiat, Marc Schoenauer.
- **RoDAPoG** 2021-2025 (302k€) *Robust Deep learning for Artificial genomics and Population Genetics*
Coordinator: Flora Jay,
Participants: Cyril Furtlehner, Guillaume Charpiat.

10.4.2 Others

- **ADEME 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: SME ARTELYS
Participants Isabelle Guyon, Marc Schoenauer, Michèle Sebag, Victor Berger (PhD), Herilalaina Rakotoarison (PhD), Berna Bakir Batu (Post-doc)
- **PIA JobAgile** 2018-2021 (379 kEuros) *Evidence-based Recommendation pour l'Emploi et la Formation* (Section 8.3.1).
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 8.6).
Coordinator: Magali Richard (Inria Grenoble)
Participants: Isabelle Guyon
- **BOBCAT** The new B-to-B work intermediaries: comparing business models in the "Collaborative" digital economy, 2018-2021 (100k euros), funded by DARES (French Ministry of Labor).
Coordinator : Odile Chagny (IRES)
Participants: Paola Tubaro
- **IPL HPC-BigData** 2018-2022 (100 kEuros) High Performance Computing and Big Data (Section 8.5.6)
Coordinator: Bruno Raffin (Inria Grenoble)
Participants: Guillaume Charpiat, Loris Felardos (PhD)
- **Inria Challenge** (formerly IPL) **HYAIAI**, 2019-2023, *HYbrid Approaches for Interpretable Artificial Intelligence*
Coordinator: Elisa Fromont (Lacodam, Inria Rennes)
Participants: Marc Schoenauer and Michèle Sebag
- **COMETH** 2020 (60 kEuros), within EIT Health, for the organization of challenges toward personalized medicine (Section 8.6).
Coordinator: Magali Richard (Inria Grenoble)
Participants: Isabelle Guyon
- **TRIA** Le TRavail de l'Intelligence Artificielle : éthique et gouvernance de l'automation, 2020-2021 (131k euros), funded by MITI-CNRS (CNRS mission for interdisciplinary and transverse initiatives).
Coordinator : Paola Tubaro
Participants: A.A. Casilli (Telecom Paris); I. Vasilescu, L. Lamel, Gilles Adda (CNRS-Limsi); N. Seghouani (LRI); T. Allard, David Gross-Amblard (Irisa); J.L. Molina (UAB Barcelona); J.A. Ortega (Univ. València); J. Posada (Univ. Toronto)
- **Inria Challenge OceanAI** 2021-2025, *AI, Data, Models for a Blue Economy*
Coordinator: Nayat Sanchez Pi (Inria Chile)
Participants: Marc Schoenauer and Michèle Sebag

10.5 Regional initiatives

- IRS⁷ **Nutriperso** 2017-2020, 122 kEuros. Personalized recommendations toward healthier eating practices (Section 8.3.2).
Partners: INRA (coordinator), INSERM, Agro Paristech, Mines Telecom
Participants: Philippe Caillou, Flora Jay, Michèle Sebag, Paola Tubaro
- **DATAIA Vadore** 2018-2020 (105 kEuros) VAlorizations of Data to imprOve matching in the labor markEt, with CREST (ENSAE) and Pôle Emploi (Section 8.3.1).
Coordinator: Michèle Sebag
Participants: Philippe Caillou, Isabelle Guyon
- **DATAIA ML4CFD** 2020-2022 (105 kEuros) *Machine Learning for Computational Fluid Dynamics*.
Coordinator: Michele Alessandro Bucci
Participants: Guillaume Charpiat, Marc Schoenauer
Collaboration: IFPEN (Jean-Marc Gratien and Thibault Faney)

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

General chair, scientific chair

- Paola Tubaro: Sunbelt 2020 (eventually canceled due to the pandemic)

Member of the organizing committees

- Paola Tubaro: Unboxing AI online conference, Nov. 2020
- Marc Schoenauer - Steering Committee, Parallel Problem Solving from Nature (PPSN); Steering Committee, Learning and Intelligent OptimizatioN (LION).
- Cecile Germain - Steering committee of the Learning to Discover program of Institut Pascal (originally 2020, postponed to 2022)

11.1.2 Scientific events: selection

Chair of conference program committees

- Michele Sebag: Area Chair, ICLR 2020; Area Chair, NeurIPS 2020.

Member of the conference program committees

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

11.1.3 Journal

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); Editorial Board, ACM Transaction on Evolutionary Learning and Optimization (TELO).

⁷U. Paris-Saclay *Initiative de Recherche Stratégique*

- Michèle Sebag - Editorial Board, Machine Learning, Springer Verlag.
- Paola Tubaro: Sociology, Revue française de sociologie, Journal of Economic Methodology, Lecturas de Economia.

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

11.1.4 Invited talks

- Michele Alessandro Bucci: Keynote at: IMT-Atlantique & RIKEN Joint Workshop and NAFEMS20 France.
- Guillaume Charpiat: keynote at "Journées Recherche de l'IGN", October 5th; Invited talk in the INRIA Willow team, "Input similarity from the neural network's perspective"
- Isabelle Guyon: invited talk, workshop on automated machine learning, ICLR 2020.
- Flora Jay: keynote at JOBIM2020 (FR); invited speaker at workshops: SMPGD20, Scidolyse2020 (FR); invited talk at online seminar series, worldwide: OneWorld ABC 2020, G-BiKE workshop 2021, UHPaleogenomics 2021; local: Ecology&Evolution Imperial College London 2021,
- Paola Tubaro: invited talks at University of Sussex, UK (online), Dec 2020; University of Bologna, IT, Oct 2020; CREST-ENSAE, Sep. 2020; Science and Technology Studies Association of Taipei (online), May 2020; Université de Toulouse Jean Jaurès, Feb 2020.
- Marc Schoenauer: Club DS&AI, Atelier IA Hybride, IRT SystemX, 5/2/2020; Remise de la médaille d'argent CNRS à Paco Chinesta, 20/2/2020; Séminaire EDF R&D, 6/11/2020; CONACYT day, 10/12/2020;
- Michèle Sebag: Symposium ERCIM, May 2020, cancelled; Journées OQaido (Optimisation et QUantification d'Incertitudes pour les Données Onéreuses), 2/6/2020; Séminaire EHESS, 13/11/2020;

11.1.5 Leadership within the scientific community

- Isabelle Guyon: Member of the board, [NeurIPS](#); Member of the Board, [JEDI, Joint European Disruptive Initiative](#); President and co-founder, [ChaLearn, non-for-profit organization](#) dedicated to the organization of challenges.
- Marc Schoenauer: Advisory Board, [ACM-SIGEVO, Special Interest Group on Evolutionary Computation](#); Founding President (since 2015), [SPECIES, Society for the Promotion of Evolutionary Computation In Europe and Surroundings](#), that organizes the yearly series of conferences EvoStar.
- Michèle Sebag: Executive Committee, [Institut de Convergence DataIA](#); Member of IRSN Scientific Council.

11.1.6 Scientific expertise

- Guillaume Charpiat: took part in a engineer position hiring committee for IDRIS (Jean Zay)
- Aurélien Decelle: MdC hiring committee (ex-LIMSI lab in the NLP team).
- Cecile Germain: member of the DFG Review Panel on National Research Data Infrastructures and the review boards for DFG's Strategic Funding Initiative on Artificial Intelligence. Member of the evaluation panel for an H2020 Big Data Supporting Public Health Policies project.
- Flora Jay: MdC hiring committee (BioInfo team, LISN); expert review for ANRT CIFRE, DATAIA and INCEPTION (Institut Convergence des Investissements d'Avenir) grants.
- François Landes: took part in a MdC hiring committee (BioInfo team, LISN).

- Marc Schoenauer, Scientific Advisory Board, BCAM, Bilbao, Spain; "Conseil Scientifique", IFPEN; "Conseil Scientifique", Mines Paritech.
- Michèle Sebag: Pr hiring committee UPSaclay; Chaire UCA-Mines ParisTech; Directeur d'étude EHESS; MdC Centrale Supélec; FNRS-FIB Belgique; Univ. Rutgers.
- Paola Tubaro: MdC hiring committee at Université Technologique de Compiègne

11.1.7 Research administration

- Guillaume Charpiat: part of "Commission Scientifique" (CS) INRIA Saclay
- Cyril Furtlehner, Flora Jay, Marc Schoenauer, Michèle Sebag: members of CCSU27
- 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; External member of the "Commission Recherche" of the Paris-Diderot University.
- Michele Sebag: Member of the Scientific Council of Labex AMIES, Applications des Mathématiques ds l'Industrie, l'Entreprise et la Société; Member of the Lab. Council.

11.2 Teaching - supervision - juries

11.2.1 Teaching

- Licence : Philippe Caillou, Computer Science for students in Accounting and Management, 192h, L1, IUT Sceaux, Univ. Paris Sud.
- Licence : François Landes, Mathematics for Computer Scientists, 51h, L2, Univ. Paris-Sud.
- Licence : François Landes, Introduction to Statistical Learning, 88h, L2, Univ. Paris-Sud.
- 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 : François Landes, A first look inside the ML black box, 25h, M1 Recherche (AI track), U. Paris-Sud.
- Master : Machine Learning, 28h, M2 Univ. Paris-sud, physics department
- Master : Guillaume Charpiat, Deep Learning in Practice, 21h, M2 Recherche, Centrale-Supelec + MVA.
- Master : Guillaume Charpiat, Graphical Models: Discrete Inference and Learning, 12h, M2 Recherche, Centrale-Supelec + MVA.
- Diplôme universitaire: Guillaume Charpiat, Introduction au Deep Learning, 2h30, DU IA, CHU Lille.
- 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 : Paola Tubaro, Sociology of social networks, 24h, M2, EHESS/ENS.
- Master : Paola Tubaro, Social and economic network science, 24h, M2, ENSAE.
- Master : Flora Jay, Population genetics inference, 11h, M2, U PSaclay.
- Master : Isabelle Guyon, Coordination du M1 et M2 [AI], U PSaclay.
- Master : Isabelle Guyon, M1 [AI] project A class (challenge organization)
- Master : Isabelle Guyon, M2 [AI] Advanced Optimization and Automated Machine Learning.

11.2.2 Supervision

Remark: Several PhD students who were supposed to defend in Fall 2020 have been delayed because of the coronavirus situation, and will defend in the first semester 2021.

- PhD - Nicolas GIRARD, *Satellite image vectorization using neural networks*, 16/10/2020, Yuliya Tarabalka & Pierre Alliez (INRIA Sophia-Antipolis) and Guillaume Charpiat, Univ. Nice.
- PhD - Marc NABHAN, *Sûreté de fonctionnement d'un véhicule autonome - évaluation des fausses détections au travers d'un profil de mission réduit*, 23/12/2020, Marc Schoenauer and Yves Tourbier (Renault), Univ. Paris-Saclay.
- PhD - Adrian POL *Machine Learning Anomaly Detection, with application to CMS Data Quality Monitoring*, 8/6/2020, Cécile Germain, Univ. Paris-Saclay.
- PhD - Pierre WOLINSKI, *Learning the Architecture of Neural Networks*, 6/3/2020, Yann Ollivier (Facebook AI Research, Paris) and Guillaume Charpiat, Univ. Paris-Saclay.
- 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 ESTRADÉ *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 - Jérémy GUEZ, *Statistical inference of cultural transmission of reproductive success*, 1/10/2019, Evelyne Heyer (MNHN) and Flora Jay

- 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 - Wenzhuo LIU, *Machine Learning for Numerical Simulation of PDEs*, from 1/11/2019, Mouadh Yagoubi (IRT SystemX) and Marc Schoenauer
- 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 Kodjacobian (Thalès ThereSIS, Palaiseau).
- PhD in progress - Emmanuel MENIER, *Complementary Deep Reduced Order Model*, from 1/9/2020, Michele Alessandro Bucci and Marc Schoenauer
- PhD in progress - Adrien PAVAO, *Theory and practice of challenge organization*, from 1/03/2020, Isabelle Guyon.
- 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 - Vincenzo SCHIMMENTI, *Eartquake Predictions: Machine Learned Features using Expert Models Simulations*, from 1/11/2020, François Landes and Alberto Rosso (LPTMS)
- PhD in progress - Marion ULLMO, *Reconstruction de la toile cosmique*, from 1/10/2018, Nabila Aghanim (Institut d'Astrophysique Spatiale) and Aurélien Decelle
- PhD in progress - Elinor WAHAL, *Micro-work for AI in health applications*, from 1/1/2020, Paola Tubaro

11.2.3 Juries

- Guillaume Charpiat: PhD committee for Rodrigo Dautt (ONERA) + a certain number of half-way PhD committees ("à mi-parcours") + and of MVA master 2 internship defenses
- Isabelle Guyon: PhD jury Justine Falque, U. Paris-Saclay (29/11/2019).
- Flora Jay: half-way PhD committees (R Menegaux, P Guarino-Vignon)
- François Landes: PhD rapporteur of Martina Teruzzi (Condensed matter and Machine Learning PhD, at SISSA, Trieste, Italy).
- Marc Schoenauer: Examiner, HDR Pablo Piantanida, Univ. Paris-Saclay, 11/05/2020; Reviewer, HDR Sylvain Cussat-Blanc, Univ. Toulouse, 23/6/2020; Examiner, PhD Luca Mossina, Univ. Toulouse, 3/12/2020; Examiner, HDR Carola Doerr, Univ. Paris Sorbonne, 18/12/2020 + some half-way juries of Paris-Saclay PhDs.
- Michele Sebag: Reviewer HdR Emmanuel Rachelson, Univ. Toulouse; PhD Stéphane Rivaud, Univ. Strasbourg;
- Paola Tubaro: PhD rapporteur of Mattias Mano, Ecole Polytechnique, 14/12/2020.

11.3 Popularization

11.3.1 Internal or external Inria responsibilities

11.3.2 Articles and contents

- Guillaume Charpiat: interview for the Swiss media Heidi News, on the topic of machine learning for remote sensing imagery, in the context of detecting fishery boats with slave labor
- Aurélien Decelle: dissemination article in *Ithaca: Viaggio nella Scienza (2020)* [15]
- Aurélien Decelle, Cyril Furtlehner, Flora Jay: [41] has been the subject of blog posts and articles in e.g. Science Avenir and Science&Vie, with comments from the authors.
- Flora Jay: Le Journal du CNRS published an online article about our factor analysis method for ancient DNA [21].
- Paola Tubaro: dissemination articles in *Global Dialogue* [79], *La Vie de la Recherche Scientifique* [78]; interview in *Lundi Matin*.

11.3.3 Education

- Paola Tubaro: invited panelist at *Université Populaire de l'Eau et du Développement Durable (UPEDD)* of Val-de-Marne

11.3.4 Interventions

- Guillaume Charpiat: Presentation at "**DIMS 2020: l'intelligence artificielle et les nouvelles technologies au service de l'entreprise augmentée**" about the use of deep learning in remote sensing imagery.
- Paola Tubaro: invited presentations at Geneva Municipal Libraries, CH (online), Nov 2020; union UGICT, Nantes, Sep 2020; AI ethics event, University of Lyon, Feb 2020.

12 Scientific production

12.1 Major publications

- [1] C. Adam-Bourdarios, G. Cowan, C. Germain-Renaud, I. Guyon, B. Kégl and D. Rousseau. 'The Higgs Machine Learning Challenge'. In: *Journal of Physics: Conference Series* 664.7 (Dec. 2015). DOI: [10.1088/1742-6596/664/7/072015](https://doi.org/10.1088/1742-6596/664/7/072015). URL: <https://hal.inria.fr/hal-01745998>.
- [2] L. Da Costa, Á. Fialho, M. Schoenauer and M. Sebag. 'Adaptive Operator Selection with Dynamic Multi-Armed Bandits'. In: *Proc. Genetic and Evolutionary Computation Conference (GECCO)*. ACM-SIGEVO 10-years Impact Award. ACM. 2008, pp. 913–920. DOI: [10.1145/1389095.1389272](https://doi.org/10.1145/1389095.1389272). URL: <https://hal.inria.fr/inria-00278542>.
- [3] C. Furtlehner and A. Decelle. 'Cycle-based Cluster Variational Method for Direct and Inverse Inference'. In: *Journal of Statistical Physics* 164.3 (Aug. 2016), pp. 531–574. URL: <https://hal.inria.fr/hal-01214155>.
- [4] S. Gelly, M. Schoenauer, M. Sebag, O. Teytaud, L. Kocsis, D. Silver and C. Szepesvari. 'The Grand Challenge of Computer Go: Monte Carlo Tree Search and Extensions'. In: *Communications-ACM* 55.3 (2012), pp. 106–113. URL: <https://hal.inria.fr/hal-00695370>.
- [5] O. Goudet, D. Kalainathan, P. Caillou, D. Lopez-Paz, I. Guyon and 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. <https://arxiv.org/abs/1709.05321>. Springer International Publishing, 2018. DOI: [10.1007/978-3-319-98131-4](https://doi.org/10.1007/978-3-319-98131-4). URL: <https://hal.archives-ouvertes.fr/hal-01649153>.

- [6] T. Lucas, C. Tallec, J. Verbeek and Y. Ollivier. ‘Mixed batches and symmetric discriminators for GAN training’. In: *ICML - 35th International Conference on Machine Learning*. Stockholm, Sweden, July 2018. URL: <https://hal.inria.fr/hal-01791126>.
- [7] E. Maggiori, Y. Tarabalka, G. Charpiat and P. Alliez. ‘Convolutional Neural Networks for Large-Scale Remote Sensing Image Classification’. In: *IEEE Transactions on Geoscience and Remote Sensing* 55.2 (2017), pp. 645–657. URL: <https://hal.inria.fr/hal-01369906>.
- [8] M. Mısır and M. Sebag. ‘Alors: An algorithm recommender system’. In: *Artificial Intelligence* 244 (2017). Published on-line Dec. 2016, pp. 291–314. URL: <https://hal.inria.fr/hal-01419874>.
- [9] Y. Ollivier, L. Arnold, A. Auger and N. Hansen. ‘Information-Geometric Optimization Algorithms: A Unifying Picture via Invariance Principles’. In: *Journal of Machine Learning Research* 18.18 (2017), pp. 1–65. URL: <https://hal.inria.fr/hal-01515898>.
- [10] X. Zhang, C. Furtlehner, C. Germain-Renaud and M. Sebag. ‘Data Stream Clustering with Affinity Propagation’. In: *IEEE Transactions on Knowledge and Data Engineering* 26.7 (2014), p. 1. URL: <https://hal.inria.fr/hal-00862941>.

12.2 Publications of the year

International journals

- [11] S. Amrouche, L. Basara, P. Calafiura, D. Emeliyanov, V. Estrade, S. Farrell, C. Germain, V. Vava Gligorov, T. Golling, S. Gorbunov, H. Gray, I. Guyon, M. Hushchyn, V. Innocente, M. Kiehn, M. Kunze, E. Moysé, D. Rousseau, A. Salzburger, A. Ustyuzhanin and J.-R. Vlimant. ‘The Tracking Machine Learning challenge : Throughput phase’. In: *Computing and Software for Big Science* (2021). URL: <https://hal.inria.fr/hal-03159824>.
- [12] T. Bonnaire, N. Aghanim, A. Decelle and M. Douspis. ‘T-ReX: a graph-based filament detection method’. In: *Astronomy and Astrophysics - A&A* 637 (May 2020), A18. DOI: [10.1051/0004-6361/201936859](https://doi.org/10.1051/0004-6361/201936859). URL: <https://hal.archives-ouvertes.fr/hal-02903856>.
- [13] P. Caillou, J. Renault, J.-D. Fekete, A.-C. Letournel and M. Sebag. ‘Cartolabe: A Web-Based Scalable Visualization of Large Document Collections’. In: *IEEE Computer Graphics and Applications* (2020). DOI: [10.1109/MCG.2020.3033401](https://doi.org/10.1109/MCG.2020.3033401). URL: <https://hal.inria.fr/hal-02499006>.
- [14] J. Cury, B. C. Haller, G. Achaz and F. Jay. ‘Simulation of bacterial populations with SLiM’. In: *Peer Community in Evolutionary Biology* (3rd Mar. 2021). DOI: [10.1101/2020.09.28.316869](https://doi.org/10.1101/2020.09.28.316869). URL: <https://hal.archives-ouvertes.fr/hal-03152153>.
- [15] A. Decelle. ‘The Boltzmann Machine: when the Ising model meet Machine Learning’. In: *Ithaca: Viaggio nella Scienza* (29th Oct. 2020). URL: <https://hal.archives-ouvertes.fr/hal-03143397>.
- [16] A. Decelle and C. Furtlehner. ‘Restricted Boltzmann Machine, recent advances and mean-field theory*’. In: *Chinese Physics B* (Nov. 2020). DOI: [10.1088/1674-1056/abd160](https://doi.org/10.1088/1674-1056/abd160). URL: <https://hal.archives-ouvertes.fr/hal-03143314>.
- [17] B. Donon, R. Clément, B. Donnot, A. Marot, I. Guyon and M. Schoenauer. ‘Neural Networks for Power Flow : Graph Neural Solver’. In: *Electric Power Systems Research* 189 (14th Dec. 2020), p. 106547. URL: <https://hal.archives-ouvertes.fr/hal-02372741>.
- [18] B. Donon, B. Donnot, I. Guyon, Z. Liu, A. Marot, P. Panciatici and M. Schoenauer. ‘LEAP Nets for System Identification and Application to Power Systems’. In: *Neurocomputing* 416 (2020), pp. 316–327. URL: <https://hal.inria.fr/hal-02422708>.
- [19] H. J. Escalante, H. Kaya, A. A. Salah, S. Escalera, Y. Gucluturk, U. Güçlü, X. Baró, I. Guyon, J. J. Junior, M. Madadi, S. Ayache, E. Viegas, F. Gurpinar, A. S. Wicaksana, C. C. S. Liem, M. A. J. van Gerven and R. van Lier. ‘Explaining First Impressions: Modeling, Recognizing, and Explaining Apparent Personality from Videos’. In: *IEEE Transactions on Affective Computing* (2020). URL: <https://hal.inria.fr/hal-01991652>.

- [20] S. Escalera, H. J. Escalante, X. Baró, I. Guyon, M. Madadi, J. Wan, S. Ayache, Y. Güçlütürk and U. Güçlü. ‘Guest Editorial: Image and Video Inpainting and Denoising’. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42.5 (5th May 2020), pp. 1021–1024. DOI: [10.1109/TPAMI.2020.2971291](https://doi.org/10.1109/TPAMI.2020.2971291). URL: <https://hal.inria.fr/hal-03159859>.
- [21] O. François and F. Jay. ‘Factor analysis of ancient population genomic samples’. In: *Nature Communications* 11.1 (2020). DOI: [10.1038/s41467-020-18335-6](https://doi.org/10.1038/s41467-020-18335-6). URL: <https://hal.archives-ouvertes.fr/hal-02942333>.
- [22] S. Giffard-Roisin, M. Yang, G. Charpiat, C. K. Bonfanti, B. Kégl and C. Monteleoni. ‘Tropical Cyclone Track Forecasting using Fused Deep Learning from Aligned Reanalysis Data’. In: *Frontiers in Big Data* 3 (10th Jan. 2020), p. 1. DOI: [10.3389/fdata.2020.00001](https://doi.org/10.3389/fdata.2020.00001). URL: <https://hal.archives-ouvertes.fr/hal-02329437>.
- [23] D. Kalainathan, O. Goudet and R. Dutta. ‘Causal Discovery Toolbox: Uncovering causal relationships in Python’. In: *Journal of Machine Learning Research* 21 (Mar. 2020), pp. 1–5. URL: <https://hal.univ-angers.fr/hal-02945539>.
- [24] R. Kusters, D. Misevic, H. Berry, A. Cully, Y. Le Cunff, L. Dandoy, N. Díaz-Rodríguez, M. Ficher, J. Grizou, A. Othmani, T. Palpanas, M. Komorowski, P. Loiseau, C. Moulin-Frier, S. Nanini, D. Quercia, M. Sebag, F. Soulié Fogelman, S. Taleb, L. Tupikina, V. Sahu, J.-J. Vie and F. Wehbi. ‘Interdisciplinary Research in Artificial Intelligence: Challenges and Opportunities’. In: *Frontiers in Big Data* 3 (23rd Nov. 2020). DOI: [10.3389/fdata.2020.577974](https://doi.org/10.3389/fdata.2020.577974). URL: <https://hal.inria.fr/hal-03111148>.
- [25] F. P. Landes, G. Biroli, O. Dauchot, A. Liu and D. Reichman. ‘Attractive versus truncated repulsive supercooled liquids: The dynamics is encoded in the pair correlation function’. In: *Physical Review E* 101.1 (13th Jan. 2020). DOI: [10.1103/PhysRevE.101.010602](https://doi.org/10.1103/PhysRevE.101.010602). URL: <https://hal.inria.fr/hal-02439295>.
- [26] C. Le Ludec, E. Wahal, A. A. Casilli and P. Tubaro. ‘Quel statut pour les petits doigts de l’intelligence artificielle ? Présent et perspectives du micro-travail en France’. In: *Les Mondes du travail* (Nov. 2020). URL: <https://hal.archives-ouvertes.fr/hal-02995186>.
- [27] J. Lehman, J. Clune, D. Misevic, C. Adami, J. Beaulieu, P. J. Bentley, S. Bernard, G. Beslon, D. M. Bryson, N. Cheney, A. Cully, S. Doncieux, F. C. Dyer, K. O. Ellefsen, R. Feldt, S. Fischer, S. Forrest, A. Frenoy, C. Gagneé, L. Le Goff, L. M. Grabowski, B. Hodjat, L. Keller, C. Knibbe, P. Krcah, R. E. Lenski, H. Lipson, R. MacCurdy, C. Maestre, R. Miikkulainen, S. Mitri, D. E. Moriarty, J.-B. Mouret, A. D. Nguyen, C. Ofria, M. Parizeau, D. Parsons, R. T. Pennock, W. F. Punch, T. S. Ray, M. Schoenauer, E. Shulte, K. Sims, K. O. Stanley, F. Taddei, D. Tarapore, S. Thibault, W. Weimer, R. Watson and J. Yosinski. ‘The Surprising Creativity of Digital Evolution: A Collection of Anecdotes from the Evolutionary Computation and Artificial Life Research Communities’. In: *Artificial Life* 26.2 (3rd June 2020), pp. 274–306. DOI: [10.1162/artl_a_00319](https://doi.org/10.1162/artl_a_00319). URL: <https://hal.inria.fr/hal-01735473>.
- [28] Z. Liu, Z. Xu, S. Escalera, I. Guyon, J. C. S. Jacques Junior, M. Madadi, A. Pavao, S. Treguer and W.-W. Tu. ‘Towards Automated Computer Vision: Analysis of the AutoCV Challenges 2019’. In: *Pattern Recognition Letters* 135 (1st July 2020), pp. 196–203. URL: <https://hal.archives-ouvertes.fr/hal-02386805>.
- [29] A. Marot, B. Donnot, C. Romero, L. Veyrin-Forrer, M. Lerousseau, B. Donon and I. Guyon. ‘Learning to run a power network challenge for training topology controllers’. In: *Electric Power Systems Research* 189 (21st Dec. 2020), p. 106635. DOI: [10.1016/j.epsr.2020.106635](https://doi.org/10.1016/j.epsr.2020.106635). URL: <https://hal.inria.fr/hal-03159692>.
- [30] V. Pankratov, F. Montinaro, A. Kushniarevich, G. Hudjashov, F. Jay, L. Saag, R. Flores, D. Marnetto, M. Seppel, M. Kals, U. Võsa, C. Taccioli, M. Möls, L. Milani, A. Aasa, D. J. Lawson, T. Esko, R. Mägi, L. Pagani, A. Metspalu and M. Metspalu. ‘Differences in local population history at the finest level: the case of the Estonian population’. In: *European Journal of Human Genetics* (2020). DOI: [10.1038/s41431-020-0699-4](https://doi.org/10.1038/s41431-020-0699-4). URL: <https://hal.archives-ouvertes.fr/hal-02942330>.

- [31] G. Petrillo, E. Lippiello, F. P. Landes and A. Rosso. ‘The influence of the brittle-ductile transition zone on aftershock and foreshock occurrence’. In: *Nature Communications* 11.1 (15th June 2020). DOI: [10.1038/s41467-020-16811-7](https://doi.org/10.1038/s41467-020-16811-7). URL: <https://hal.inria.fr/hal-02942365>.
- [32] L. Robinson, J. Schulz, G. Blank, M. Ragnedda, H. Ono, B. Hogan, G. Mesch, S. R. Cotten, S. B. Kretchmer, T. M. Hale, T. Drabowicz, P. Yan, B. Wellman, M.-G. Harper, A. Quan-Haase, H. S. Dunn, A. A. Casilli, P. Tubaro, R. Carveth, W. Chen, J. B. Wiest, M. Dodel, M. Stern, C. Ball, K.-T. Huang and A. Khilnani. ‘Digital inequalities 2.0: Legacy inequalities in the information age’. In: *First Monday* 25.7 (July 2020). DOI: [10.5210/fm.v25i7.10842](https://doi.org/10.5210/fm.v25i7.10842). URL: <https://halshs.archives-ouvertes.fr/halshs-02889893>.
- [33] L. Robinson, J. Schulz, H. S. Dunn, A. A. Casilli, P. Tubaro, R. Carveth, W. Chen, J. B. Wiest, M. Dodel, M. Stern, C. Ball, K.-T. Huang, G. Blank, M. Ragnedda, H. Ono, B. Hogan, G. Mesch, S. R. Cotten, S. B. Kretchmer, T. M. Hale, T. Drabowicz, P. Yan, B. Wellman, M.-G. Harper, A. Quan-Haase and A. Khilnani. ‘Digital inequalities 3.0: Emergent inequalities in the information age’. In: *First Monday* 25 (July 2020). DOI: [10.5210/fm.v25i7.10844](https://doi.org/10.5210/fm.v25i7.10844). URL: <https://halshs.archives-ouvertes.fr/halshs-02889891>.
- [34] L. Robinson, J. Schulz, A. Khilnani, H. Ono, S. Cotten, N. McClain, L. Levine, W. Chen, G. Huang, A. A. Casilli, P. Tubaro, M. Dodel, A. Quan-Haase, M. L. Ruiu, M. Ragnedda, D. Aikat and N. Tolentino. ‘Digital inequalities in time of pandemic: COVID-19 exposure risk profiles and new forms of vulnerability’. In: *First Monday* 25.7 (1st July 2020). DOI: [10.5210/fm.v25i7.10845](https://doi.org/10.5210/fm.v25i7.10845). URL: <https://halshs.archives-ouvertes.fr/halshs-02888223>.
- [35] T. Sanchez, J. Cury, G. Charpiat and F. Jay. ‘Deep learning for population size history inference: Design, comparison and combination with approximate Bayesian computation’. In: *Molecular Ecology Resources* (2020). DOI: [10.1111/1755-0998.13224](https://doi.org/10.1111/1755-0998.13224). URL: <https://hal.archives-ouvertes.fr/hal-02942328>.
- [36] P. Tubaro. ‘Disembedded or Deeply Embedded? A Multi-Level Network Analysis of Online Labour Platforms’. In: *Sociology* (31st Jan. 2021), p. 003803852098608. DOI: [10.1177/0038038520986082](https://doi.org/10.1177/0038038520986082). URL: <https://hal.archives-ouvertes.fr/hal-03127861>.
- [37] P. Tubaro, A. A. Casilli and M. Coville. ‘The trainer, the verifier, the imitator: Three ways in which human platform workers support artificial intelligence’. In: *Big Data & Society* 7.1 (Apr. 2020), p. 205395172091977. DOI: [10.1177/2053951720919776](https://doi.org/10.1177/2053951720919776). URL: <https://hal.archives-ouvertes.fr/hal-02554196>.
- [38] P. Tubaro, C. Le Ludec and A. A. Casilli. ‘Counting ‘micro-workers’: societal and methodological challenges around new forms of labour’. In: *Work Organisation, Labour & Globalisation* 14.1 (2020), pp. 67–82. DOI: [10.13169/workorglaboglob.14.1.0067](https://doi.org/10.13169/workorglaboglob.14.1.0067). URL: <https://hal.archives-ouvertes.fr/hal-02898905>.
- [39] P. Tubaro, L. Ryan, A. A. Casilli and A. D’Angelo. ‘Social network analysis: New ethical approaches through collective reflexivity. Introduction to the special issue of Social Networks’. In: *Social Networks* (25th Dec. 2020). DOI: [10.1016/j.socnet.2020.12.001](https://doi.org/10.1016/j.socnet.2020.12.001). URL: <https://hal.archives-ouvertes.fr/hal-03090287>.
- [40] A. Yale, S. Dash, R. Dutta, I. Guyon, A. Pavao and K. P. Bennett. ‘Generation and Evaluation of Privacy Preserving Synthetic Health Data’. In: *Neurocomputing* 416 (27th Nov. 2020), pp. 244–255. DOI: [10.1016/j.neucom.2019.12.136](https://doi.org/10.1016/j.neucom.2019.12.136). URL: <https://hal.inria.fr/hal-03158544>.
- [41] B. Yelmen, A. Decelle, L. Ongaro, D. Marnetto, C. Tallec, F. Montinaro, C. Furtlehner, L. Pagani and F. Jay. ‘Creating artificial human genomes using generative neural networks’. In: *PLoS Genetics* (4th Feb. 2021). DOI: [10.1371/journal.pgen.1009303](https://doi.org/10.1371/journal.pgen.1009303). URL: <https://hal.archives-ouvertes.fr/hal-03149930>.

International peer-reviewed conferences

- [42] R. Bresson, J. Cohen, E. Hüllermeier, C. Labreuche and M. Sebag. ‘Neural Representation and Learning of Hierarchical 2-additive Choquet Integrals’. In: IJCAI-PRICAI-20 - Twenty-Ninth International Joint Conference on Artificial Intelligence and Seventeenth Pacific Rim International Conference on Artificial Intelligence. Yokohama, France, 11th July 2020, pp. 1984–1991. DOI: [10.24963/ijcai.2020/275](https://doi.org/10.24963/ijcai.2020/275). URL: <https://hal.archives-ouvertes.fr/hal-03017078>.
- [43] M. Chandorkar, C. Furtlehner, B. Poduval, E. Camporeale and M. Sebag. ‘Dynamic Time Lag Regression: Predicting What and When’. In: ICLR 2020 - 8th International Conference on Learning Representations. Addis Abeba, Ethiopia, 26th Apr. 2020. URL: <https://hal.inria.fr/hal-02422148>.
- [44] S. Dash, A. Yale, I. Guyon and K. P. Bennett. ‘Medical Time-Series Data Generation using Generative Adversarial Networks’. In: AIME 2020 - International Conference on Artificial Intelligence in Medicine. Minneapolis, United States: https://link.springer.com/chapter/10.1007/978-3-030-59137-3_34, 26th Sept. 2020, pp. 382–391. URL: <https://hal.inria.fr/hal-03158549>.
- [45] B. Donon, W. Liu, A. Marot, Z. Liu, I. Guyon and M. Schoenauer. ‘Deep Statistical Solvers’. In: NeurIPS 2020 - 34th Conference on Neural Information Processing Systems. Vancouver / Virtuel, Canada, 6th Dec. 2020. URL: <https://hal.inria.fr/hal-02974541>.
- [46] J. Girard-Satabin, G. Charpiat, Z. Chihani and M. Schoenauer. ‘CAMUS: A Framework to Build Formal Specifications for Deep Perception Systems Using Simulators’. In: *ECAI 2020, Proceedings of the 24th European Conference on Artificial Intelligence, Santiago de Compostela, Spain, June 10-12, 2020*. ECAI 2020 - 24th European Conference on Artificial Intelligence. Santiago de Compostela, Spain: <http://ecai2020.eu/>, 10th June 2020. URL: <https://hal.inria.fr/hal-02440520>.
- [47] L. Gresele, G. Fissore, A. Javaloy, B. Schölkopf and A. Hyvärinen. ‘Relative gradient optimization of the Jacobian term in unsupervised deep learning’. In: NeurIPS 2020 - 34th Conference on Neural Information Processing Systems. Vancouver / Virtuel, Canada, 6th Dec. 2020. URL: <https://hal.archives-ouvertes.fr/hal-02978662>.
- [48] Z. Liu, Z. Xu, S. Rajaa, M. Madadi, J. C. S. Jacques, S. Escalera, A. Pavao, S. Treguer, W.-W. Tu and I. Guyon. ‘Towards Automated Deep Learning: Analysis of the AutoDL challenge series 2019’. In: NeurIPS 2019 - Thirty-third Conference on Neural Information Processing Systems, Competition and Demonstration Track. Vol. 123. Virtual, United States, 6th Dec. 2020, pp. 242–252. URL: <https://hal.inria.fr/hal-03159795>.
- [49] A. Yale, S. Dash, K. Bhanot, I. Guyon, J. S. Erickson and K. P. Bennett. ‘Synthesizing Quality Open Data Assets from Private Health Research Studies’. In: BIS 2020 - International Conference on Business Information Systems. Colorado Springs, United States: https://link.springer.com/chapter/10.1007/978-3-030-61146-0_26, 8th June 2020, pp. 324–335. URL: <https://hal.inria.fr/hal-03158556>.

National peer-reviewed Conferences

- [50] J. Girard-Satabin, A. Varasse, G. Charpiat, Z. Chihani and M. Schoenauer. ‘Partitionnement en régions linéaires pour la vérification formelle de réseaux de neurones’. In: Journées Francophones des Langages Applicatifs. Saint Médard d’Excideuil, France: <http://jfla.inria.fr/jfla2021.html>, 1st Apr. 2021. URL: <https://hal.inria.fr/hal-03127853>.
- [51] N. Makdoud, J. Kodjabachian and M. Schoenauer. ‘Deep Reinforcement Learning Guided by a Library of Possibly Unreliable Advice’. In: Proc. CAP’2020 (Conférence d’Apprentissage). Vannes, France, 23rd June 2020. URL: <https://hal.inria.fr/hal-03146143>.

Conferences without proceedings

- [52] M. A. Bucci, O. Semeraro, S. Chibbaro, A. Allauzen and L. Mathelin. ‘When deep learning meets ergodic theory’. In: 73rd Annual APS/DFD Meeting. Chicago / Virtual, United States, 22nd Nov. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03101431>.
- [53] T. Guégan, O. Semeraro, M. A. Bucci, L. Cordier and L. Mathelin. ‘Control by Deep Reinforcement Learning of a separated flow’. In: 73rd Annual Meeting of the APS Division of Fluid Dynamics. Chicago, United States, 22nd Nov. 2020. URL: <https://hal.archives-ouvertes.fr/hal-0308845>.
- [54] J.-B. Lamy and F. Jay. ‘Combiner arbres phylogénétiques et visualisation d’ensembles’. In: Atelier Visualisation d’informations, interaction et fouille de données. Bruxelles, Belgium, 28th Jan. 2020. URL: <https://hal.archives-ouvertes.fr/hal-02968296>.
- [55] Z. Liu and I. Guyon. ‘Asymptotic Analysis of Meta-learning as a Recommendation Problem’. In: Meta-learning Workshop @ AAAI 2021. Virtual, Canada, 2nd Feb. 2021. URL: <https://hal.archives-ouvertes.fr/hal-03098180>.
- [56] N. Sanchez-Pi, L. Marti, A. Abreu, O. Bernard, C. de Vargas, D. Eveillard, A. Maass, P. A. Marquet, J. Sainte-Marie, J. Salomon, M. Schoenauer and M. Sebag. ‘Artificial Intelligence, Machine Learning and Modeling for Understanding the Oceans and Climate Change’. In: NeurIPS 2020 Workshop - Tackling Climate Change with Machine Learning. Santiago / Virtual, Chile, Dec. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03138712>.
- [57] O. Semeraro, M. A. Bucci and L. Mathelin. ‘Closed-loop optimal control for shear flows using reinforcement learning’. In: 73rd Annual APS/DFD Meeting. Chicago, United States, 22nd Nov. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03101416>.
- [58] P. Tubaro and A. A. Casilli. ‘Portraits of micro-workers: The real people behind AI in France’. In: 2nd Crowdfunding Symposium 2020. Paderborn, Germany: <https://pace.uni-paderborn.de/pace-phd-programs/digital-future/tagung-2020>, 8th Oct. 2020. URL: <https://hal.archives-ouvertes.fr/hal-02960775>.

Scientific book chapters

- [59] H. J. Escalante, W.-W. Tu, I. Guyon, D. L. Silver, E. Viegas, Y. Chen, W. Dai and Q. Yang. ‘AutoML @ NeurIPS 2018 challenge: Design and Results’. In: *The NeurIPS ’18 Competition*. 11th Mar. 2020, pp. 209–229. URL: <https://hal.inria.fr/hal-03159764>.
- [60] A. A. Pol, C. Germain, M. Pierini and G. Carminara. ‘Data Quality Monitoring Anomaly Detection’. In: *Artificial Intelligence for High Energy Physics*. 2021. DOI: [10.1142/12200](https://doi.org/10.1142/12200). URL: <https://hal.inria.fr/hal-03159873>.

Doctoral dissertations and habilitation theses

- [61] C. Furtlehner. ‘Statistical physics methods for machine learning and traffic forecasting’. Université Paris Saclay, 18th June 2020. URL: <https://hal.inria.fr/tel-02917159>.
- [62] N. Girard. ‘Learning and Geometric Approaches for Automatic Extraction of Objects from Remote Sensing Images’. Université Côte d’Azur, Nice, France, 16th Oct. 2020. URL: <https://hal.inria.fr/tel-03111628>.
- [63] M. Nabhan. ‘Models and algorithms for the exploration of the space of scenarios : toward the validation of the autonomous vehicle’. Université Paris-Saclay, 21st Dec. 2020. URL: <https://tel.archives-ouvertes.fr/tel-03144001>.
- [64] A. A. Pol. ‘Machine Learning Anomaly Detection Applications to Compact Muon Solenoid Data Quality Monitoring’. Université Paris-Saclay, 8th June 2020. URL: <https://tel.archives-ouvertes.fr/tel-02924477>.
- [65] P. Wolinski. ‘Structural Learning of Neural Networks’. Université Paris-Saclay, 6th Mar. 2020. URL: <https://tel.archives-ouvertes.fr/tel-02888604>.

Reports & preprints

- [66] F. Arceri, F. P. Landes, L. Berthier and G. Biroli. *Glasses and aging: A Statistical Mechanics Perspective*. 17th Sept. 2020. URL: <https://hal.inria.fr/hal-02942375>.
- [67] V. Berger and M. Sebag. *Variational Auto-Encoder: not all failures are equal*. 13th Feb. 2020. URL: <https://hal.inria.fr/hal-02497248>.
- [68] L. Blier, C. Tallec and Y. Ollivier. *Learning Successor States and Goal-Dependent Values: A Mathematical Viewpoint*. 19th Jan. 2021. URL: <https://hal.inria.fr/hal-03151901>.
- [69] A. Boisbunon, C. Fanara, I. Grenet, J. Daeden, A. Vighi and M. Schoenauer. *Zoetrope Genetic Programming for Regression*. 2nd Mar. 2021. URL: <https://hal.archives-ouvertes.fr/hal-03155694>.
- [70] T. Bonnaire, A. Decelle and N. Aghanim. *Cascade of Phase Transitions for Multi-Scale Clustering*. 23rd Oct. 2020. URL: <https://hal.archives-ouvertes.fr/hal-02976130>.
- [71] G. De Bie, H. Rakotoarison, G. Peyré and M. Sebag. *Distribution-Based Invariant Deep Networks for Learning Meta-Features*. 26th Feb. 2021. URL: <https://hal.inria.fr/hal-03153200>.
- [72] R. Egele, P. Balaprakash, V. Vishwanath, I. Guyon and Z. Liu. *AgEBO-Tabular: Joint Neural Architecture and Hyperparameter Search with Autotuned Data-Parallel Training for Tabular Data*. 21st Oct. 2020. URL: <https://hal.archives-ouvertes.fr/hal-02973288>.
- [73] C. ICUBAM, L. Bonnasse-Gahot, M. Dénès, G. Dulac-Arnold, S. Girgin, F. Husson, V. Iovene, J. Josse, A. Kimmoun, F. Landes, J.-P. Nadal, R. Primet, F. Quintao, P. G. Raverdy, V. Rouvreau, O. Teboul and R. Yurchak. *ICU Bed Availability Monitoring and analysis in the Grand Est region of France during the COVID-19 epidemic*. 25th May 2020. URL: <https://hal.archives-ouvertes.fr/hal-02620018>.
- [74] Z. Liu, A. Pavao, Z. Xu, S. Escalera, F. Ferreira, I. Guyon, S. Hong, F. Hutter, R. Ji, J. C. S. Jacques Junior, G. Li, M. Lindauer, Z. Luo, M. Madadi, T. Nierhoff, K. Niu, C. Pan, D. Stoll, S. Treguer, J. Wang, P. Wang, C. Wu, Y. Xiong, A. Zela and Y. Zhang. *Winning solutions and post-challenge analyses of the ChaLearn AutoDL challenge 2019*. 30th Nov. 2020. URL: <https://hal.archives-ouvertes.fr/hal-02957135>.
- [75] L. Meunier, H. Rakotoarison, P. K. Wong, B. Roziere, J. Rapin, O. Teytaud, A. Moreau and C. Dorr. *Black-Box Optimization Revisited: Improving Algorithm Selection Wizards through Massive Benchmarking*. 26th Feb. 2021. URL: <https://hal.inria.fr/hal-03154019>.
- [76] M. Ullmo, A. Decelle and N. Aghanim. *Encoding large scale cosmological structure with Generative Adversarial Networks*. 1st Dec. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03034838>.
- [77] P. Wolinski, G. Charpiat and Y. Ollivier. *Interpreting a Penalty as the Influence of a Bayesian Prior*. 4th Feb. 2020. URL: <https://hal.archives-ouvertes.fr/hal-02466702>.

12.3 Other

Scientific popularization

- [78] P. Tubaro. ‘Les « petites mains » derrière le mirage de l’automatisation’. In: *La vie de la recherche scientifique* (July 2020), pp. 76–77. URL: <https://hal.archives-ouvertes.fr/hal-03029739>.
- [79] P. Tubaro. ‘Whose intelligence is artificial intelligence?’ In: *Global Dialogue* (Apr. 2020), pp. 38–39. URL: <https://hal.archives-ouvertes.fr/hal-03029735>.

12.4 Cited publications

- [80] C. Adam-Bourdarios, G. Cowan, C. Germain, I. Guyon, B. Kégl and D. Rousseau. ‘How Machine Learning won the Higgs Boson Challenge’. In: *Proc. European Symposium on ANN, CI and ML*. 2016.

- [81] C. Adam-Bourdarios, G. Cowan, C. Germain-Renaud, I. Guyon, B. Kégl and D. Rousseau. ‘The Higgs Machine Learning Challenge’. In: *Journal of Physics: Conference Series* 664.7 (2015). DOI: [10.1088/1742-6596/664/7/072015](https://doi.org/10.1088/1742-6596/664/7/072015). URL: <https://hal.inria.fr/hal-01745998>.
- [82] D. J. Amit, H. Gutfreund and H. Sompolinsky. ‘Statistical Mechanics of Neural Networks Near Saturation’. In: *Annals of Physics* 173 (1987), pp. 30–67.
- [83] 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 and Y. Yilmaz. ‘The Tracking Machine Learning Challenge: Accuracy Phase’. In: *The NeurIPS ’18 Competition*. Vol. 8. The Springer Series on Challenges in Machine Learning. Springer, Nov. 2019, pp. 231–264. DOI: [10.1007/978-3-030-29135-8_9](https://doi.org/10.1007/978-3-030-29135-8_9). URL: <https://hal.inria.fr/hal-02427989>.
- [84] C. Anderson. *The End of Theory: The Data Deluge Makes the Scientific Method Obsolete*. Wired Magazine, 2008. URL: <https://www.wired.com/2008/06/pb-theory/>.
- [85] W. Banzhaf, F. D. Francone, R. E. Keller and P. Nordin. *Genetic Programming: An Introduction: On the Automatic Evolution of Computer Programs and Its Applications*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1998.
- [86] N. Belkhir. ‘Per Instance Algorithm Configuration for Continuous Black Box Optimization’. Theses. Université Paris-Saclay, Nov. 2017.
- [87] N. Belkhir, J. Dréo, P. Savéant and M. Schoenauer. ‘Per instance algorithm configuration of CMA-ES with limited budget’. In: *Proc. ACM-GECCO*. 2017, pp. 681–688.
- [88] I. Bello, B. Zoph, V. Vasudevan and Q. V. Le. ‘Neural Optimizer Search with Reinforcement Learning’. In: *34th ICML*. 2017, pp. 459–468.
- [89] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira and J. W. Vaughan. ‘A theory of learning from different domains’. In: *Machine Learning* 79.1 (2018), pp. 151–175. DOI: [10.1007/s10994-009-5152-4](https://doi.org/10.1007/s10994-009-5152-4).
- [90] V. Berger and 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, Sept. 2019. URL: <https://hal.inria.fr/hal-02431955>.
- [91] J. Bergstra, R. Bardenet, Y. Bengio and B. Kégl. ‘Algorithms for Hyper-Parameter Optimization’. In: *NIPS* 25. 2011, pp. 2546–2554.
- [92] M. de Bezenac, A. Pajot and P. Gallinari. ‘Deep Learning for Physical Processes: Incorporating Prior Scientific Knowledge’. In: *arXiv preprint arXiv:1711.07970* (2017).
- [93] G. de Bie, G. Peyré and M. Cuturi. ‘Stochastic Deep Networks’. In: *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*. Ed. by K. Chaudhuri and R. Salakhutdinov. Vol. 97. Proceedings of Machine Learning Research. PMLR, 2019, pp. 1556–1565. URL: <http://proceedings.mlr.press/v97/de-bie19a.html>.
- [94] J. Bruna and S. Mallat. ‘Invariant Scattering Convolution Networks’. In: *IEEE Trans. Pattern Anal. Mach. Intell.* 35.8 (2013), pp. 1872–1886.
- [95] M. Chandorkar. ‘Machine Learning in Space Weather’. Theses. Université of Eindhoven, Nov. 2019. URL: <https://hal.inria.fr/tel-02430788>.
- [96] R. Couillet and M. Debbah. *Random matrix methods for wireless communications*. Cambridge University Press, 2011.
- [97] A. D’Amour. *On Multi-Cause Causal Inference with Unobserved Confounding: Counterexamples, Impossibility, and Alternatives*. 2019. arXiv: [1902.10286](https://arxiv.org/abs/1902.10286) [stat.ML].
- [98] David Silver et al. ‘Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm’. In: *CoRR* abs/1712.01815 (2017). URL: <http://arxiv.org/abs/1712.01815>.

- [99] A. Decelle, G. Fissore and C. Furtlehner. ‘Spectral dynamics of learning in restricted Boltzmann machines’. In: *EPL (Europhysics Letters)* 119.6 (2017), p. 60001.
- [100] A. Decelle, G. Fissore and C. Furtlehner. ‘Thermodynamics of Restricted Boltzmann Machines and Related Learning Dynamics’. In: *J. Stat. Phys.* 172 (2018), pp. 1576–1608.
- [101] M. DeMaria, M. Mainelli, L. K. Shay, J. A. Knaff and J. Kaplan. ‘Further improvements to the statistical hurricane intensity prediction scheme (SHIPS)’. In: *Weather and Forecasting* 20.4 (2005), pp. 531–543.
- [102] L. Dinh, J. Sohl-Dickstein and S. Bengio. ‘Density estimation using Real NVP’. In: *Int. Conf. on Learning Representations (ICLR)*. 2017.
- [103] B. Donnot. ‘Deep learning methods for predicting flows in power grids : novel architectures and algorithms’. Theses. Université Paris Saclay (COMUE), Feb. 2019. URL: <https://tel.archives-ouvertes.fr/tel-02045873>.
- [104] B. Donnot, I. Guyon, M. Schoenauer, A. Marot and P. Panciatici. ‘Fast Power system security analysis with Guided Dropout’. In: *26th European Symposium on Artificial Neural Networks*. Electronic Proceedings ESANN 2018. <https://arxiv.org/abs/1801.09870>. Bruges, Belgium, Apr. 2018. URL: <https://hal.archives-ouvertes.fr/hal-01695793>.
- [105] B. Donon, B. Donnot, I. Guyon and A. Marot. ‘Graph Neural Solver for Power Systems’. In: *IJCNN 2019 - International Joint Conference on Neural Networks*. Budapest, Hungary, July 2019. URL: <https://hal.archives-ouvertes.fr/hal-02175989>.
- [106] G. F. Doquet and 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, Sept. 2019. URL: <https://hal.archives-ouvertes.fr/hal-02436824>.
- [107] T. Duriez, S. L. Brunton and B. R. Noack. *Machine Learning Control – Taming Nonlinear Dynamics and Turbulence*. Springer International Publishing, 2017.
- [108] D. K. Duvenaud, D. Maclaurin, J. Aguilera-Iparraguirre, R. Gómez-Bombarelli, T. Hirzel, A. Aspuru-Guzik and R. P. Adams. ‘Convolutional Networks on Graphs for Learning Molecular Fingerprints’. In: *NIPS*. 2015, pp. 2224–2232.
- [109] T. v. Erven, P. Grunwald and 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* 74.3 (2012), pp. 361–417. DOI: [10.1111/j.1467-9868.2011.01025.x](https://doi.org/10.1111/j.1467-9868.2011.01025.x). (Visited on 08/11/2018).
- [110] 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 and R. van Lier. ‘Design of an Explainable Machine Learning Challenge for Video Interviews’. In: *IJCNN 2017 - 30th International Joint Conference on Neural Networks*. Neural Networks (IJCNN), 2017 International Joint Conference on. Anchorage, AK, United States: IEEE, 2017, pp. 1–8. URL: <https://hal.inria.fr/hal-01668386>.
- [111] S. Escalera, X. Baró, H. J. Escalante and 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.
- [112] V. Estrade, C. Germain, I. Guyon and D. Rousseau. ‘Systematics aware learning: a case study in High Energy Physics’. In: *ESANN 2018 - 26th European Symposium on Artificial Neural Networks*. Bruges, Belgium, Apr. 2018. URL: <https://hal.inria.fr/hal-01715155>.
- [113] M. Feurer, A. Klein, K. Eggensperger, J. Springenberg, M. Blum and F. Hutter. ‘Efficient and Robust Automated Machine Learning’. In: *NIPS 28*. 2015, pp. 2962–2970.
- [114] C. Furtlehner and A. Decelle. ‘Cycle-Based Cluster Variational Method for Direct and Inverse Inference’. In: *Journal of Statistical Physics* 164.3 (2016), pp. 531–574.
- [115] C. Furtlehner, M. Sebag and X. Zhang. ‘Scaling analysis of affinity propagation’. In: *Physical Review E* 81.6 (2010), p. 066102.

- [116] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand and V. Lempitsky. ‘Domain-Adversarial Training of Neural Networks’. In: *Journal of Machine Learning Research* 17.59 (2016), pp. 1–35.
- [117] K. Garimella, G. D. F. Morales, A. Gionis and M. Mathioudakis. ‘Political Discourse on Social Media: Echo Chambers, Gatekeepers, and the Price of Bipartisanship’. In: *WWW. ACM*, 2018, pp. 913–922.
- [118] S. Giffard-Roisin, M. Yang, G. Charpiat, B. Kégl and C. Monteleoni. ‘Deep Learning for Hurricane Track Forecasting from Aligned Spatio-temporal Climate Datasets’. In: *Modeling and decision-making in the spatiotemporal domain NIPS workshop*. Montréal, Canada, Dec. 2018. URL: <https://hal.archives-ouvertes.fr/hal-01905408>.
- [119] J. Giles. ‘Computational social science: Making the links’. In: *Nature - News* 488.7412 (2012), pp. 448–450.
- [120] F. Gonard. ‘Cold-start recommendation : from Algorithm Portfolios to Job Applicant Matching’. Theses. Université Paris-Saclay, May 2018. URL: <https://tel.archives-ouvertes.fr/tel-01825220>.
- [121] F. Gonard, M. Schoenauer and M. Sebag. ‘ASAPV2 and ASAPV3: Sequential optimization of an Algorithm Selector and a Scheduler’. In: *Open Algorithm Selection Challenge 2017*. Vol. 79. Proceedings of Machine Learning Research. 2017, pp. 8–11.
- [122] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville and Y. Bengio. ‘Generative Adversarial Nets’. In: *NIPS 27*. Ed. by Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence and K. Q. Weinberger. Curran Associates, Inc., 2014, pp. 2672–2680.
- [123] P. Grunwald. *The Minimum Description Length Principle*. MIT Press, 2007.
- [124] I. Guyon, C. F. Aliferis, G. F. Cooper, A. Elisseeff, J.-P. Pellet, P. Spirtes and 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.
- [125] I. Guyon, K. Bennett, G. Cawley, H. J. Escalante, S. Escalera, T. K. Ho, N. Macia, B. Ray, M. Saeed and A. Statnikov. ‘Design of the 2015 ChaLearn AutoML challenge’. In: *Proc. IJCNN*. IEEE. 2015, pp. 1–8.
- [126] 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 and E. Viegas. ‘Analysis of the AutoML Challenge series 2015-2018’. In: *AutoML: Methods, Systems, Challenges*. Ed. by F. Hutter, L. Kotthoff and J. Vanschore. The Springer Series on Challenges in Machine Learning. Springer Verlag, 2018. URL: <https://hal.archives-ouvertes.fr/hal-01906197>.
- [127] H. H. Hoos. ‘Programming by Optimization’. In: *Commun. ACM* 55.2 (2012), pp. 70–80.
- [128] D. Hothersall. *History of Psychology*. McGraw-Hill, 2004.
- [129] K. Imai and Z. Jiang. *Discussion of "The Blessings of Multiple Causes" by Wang and Blei*. 2019. arXiv: 1910.06991 [stat.ME].
- [130] G. W. Imbens and D. B. Rubin. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press, 2015.
- [131] D. Kalainathan, O. Goudet, P. Caillou, M. Sebag and P. Tubaro. *Conditions objectives de travail et ressenti des individus : le rôle du management*. Synthèse n. 14 de "La Fabrique de l'Industrie". 2017. URL: <https://hal.inria.fr/hal-01742592>.
- [132] D. Kingma and M. Welling. ‘Auto-encoding variational Bayes.’ In: *Int. Conf. on Learning Representations (ICLR)*. 2014.
- [133] L. Kotthoff. ‘Algorithm Selection for Combinatorial Search Problems: A Survey’. In: *Data Mining and Constraint Programming: Foundations of a Cross-Disciplinary Approach*. Ed. by Ch. Bessiere et al. Cham: Springer Verlag, 2016, pp. 149–190.
- [134] A. Krizhevsky, I. Sutskever and 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.

- [135] 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 and M. Van Alstyne. ‘Life in the network: the coming age of computational social science’. In: *Science* 323.5915 (2009), pp. 721–723.
- [136] Y. LeCun. *Deep Learning est mort. Vive Differentiable Programming!* Facebook Post. 2018. URL: <https://www.facebook.com/yann.lecun/posts/10155003011462143>.
- [137] J. Z. Liang, E. Meyerson and R. Miikkulainen. ‘Evolutionary architecture search for deep multitask networks’. In: *Proc. ACM-GECCO*. Ed. by H. E. Aguirre and K. Takadama. ACM, 2018, pp. 466–473. DOI: [10.1145/3205455.3205489](https://doi.org/10.1145/3205455.3205489). URL: <http://doi.acm.org/10.1145/3205455.3205489>.
- [138] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver and D. Wierstra. *Continuous control with deep reinforcement learning*. 2015. arXiv: [1509.02971](https://arxiv.org/abs/1509.02971) [cs.LG].
- [139] E. Lippiello, G. Petrillo, F. Landes and A. Rosso. ‘Fault Heterogeneity and the Connection between Aftershocks and Afterslip’. In: *Bulletin of the Seismological Society of America* 109.3 (Apr. 2019), pp. 1156–1163. DOI: [10.1785/0120180244](https://doi.org/10.1785/0120180244). URL: <https://hal.inria.fr/hal-02156407>.
- [140] H. Lipson and J. B. Pollack. ‘Automatic design and manufacture of robotic lifeforms’. In: *Nature Letters* 406 (2000), pp. 974–978.
- [141] Z. Liu, O. Bousquet, A. Elisseff, S. Escalera, I. Guyon, J. Jacques, A. Pavao, D. Silver, L. Sun-Hosoya, S. Treguer, W.-W. Tu, J. Wang and Q. Yao. ‘AutoDL Challenge Design and Beta Tests-Towards automatic deep learning’. In: *CiML workshop @ NIPS2018*. Montreal, Canada, Dec. 2018. URL: <https://hal.archives-ouvertes.fr/hal-01906226>.
- [142] C. Louizos, M. Reisser, T. Blankevoort, E. Gavves and M. Welling. ‘Relaxed Quantization for Discretized Neural Networks’. In: *ICLR*. 2019.
- [143] E. Maggiori, Y. Tarabalka, G. Charpiat and 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. URL: <https://hal.inria.fr/hal-01468452>.
- [144] A. Marot, B. Donon, I. Guyon and B. Donnot. ‘Learning To Run a Power Network Competition’. In: *CiML Workshop, NeurIPS*. Montréal, Canada, Dec. 2018. URL: <https://hal.archives-ouvertes.fr/hal-01968295>.
- [145] L. Mescheder, A. Geiger and S. Nowozin. ‘Which Training Methods for GANs do actually Converge?’ In: *35th ICML. International Conference on Machine Learning*. 2018, pp. 3481–3490.
- [146] M. Mısır and M. Sebag. ‘Alors: An algorithm recommender system’. In: *Artificial Intelligence* 244 (2017). Published on-line Dec. 2016, pp. 291–314. URL: <https://hal.inria.fr/hal-01419874>.
- [147] D. Molchanov, A. Ashukha and D. Vetrov. ‘Variational Dropout Sparsifies Deep Neural Networks’. In: *ArXiv e-prints* (Jan. 2017). arXiv: [1701.05369](https://arxiv.org/abs/1701.05369) [stat.ML].
- [148] C. O’Neill. *Weapons of Math Destruction*. Crown Books, 2016.
- [149] E. L. Ogburn, I. Shpitser and E. J. T. Tchetgen. *Comment on "Blessings of Multiple Causes"*. 2019. arXiv: [1910.05438](https://arxiv.org/abs/1910.05438) [stat.ME].
- [150] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior and K. Kavukcuoglu. ‘WaveNet: A Generative Model for Raw Audio’. In: *CoRR* abs/1609.03499 (2016). URL: <http://arxiv.org/abs/1609.03499>.
- [151] A. Pavao, D. Kalainathan, L. Sun-Hosoya, K. Bennett and 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, Dec. 2019. URL: <https://hal.inria.fr/hal-02415639>.
- [152] J. Pearl. *Causality: Models, Reasoning, and Inference (2nd edition)*. Cambridge University Press, 2009.
- [153] J. Pearl. ‘Theoretical Impediments to Machine Learning With Seven Sparks from the Causal Revolution’. In: *11th ACM WSDM*. 2018, p. 3.

- [154] G. Pleiss, M. Raghavan, F. Wu, J. M. Kleinberg and K. Q. Weinberger. ‘On Fairness and Calibration’. In: *NIPS*. 2017, pp. 5684–5693.
- [155] A. A. Pol, V. Berger, G. Cerminara, C. Germain and M. Pierini. ‘Anomaly Detection With Conditional Variational Autoencoders’. In: *ICMLA 2019 - 18th IEEE International Conference on Machine Learning and Applications*. 18th International Conference on Machine Learning Applications. Boca Raton, United States, Dec. 2019. URL: <https://hal.inria.fr/hal-02396279>.
- [156] A. A. Pol, V. Berger, G. Cerminara, C. Germain and M. Pierini. *Trigger Rate Anomaly Detection with Conditional Variational Autoencoders at the CMS Experiment*. Machine Learning and the Physical Sciences Workshop at the 33rd Conference on Neural Information Processing Systems (NeurIPS). Poster. Dec. 2019. URL: <https://hal.inria.fr/hal-02428005>.
- [157] A. A. Pol, G. Cerminara, C. Germain, M. Pierini and 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 3.1* (Jan. 2019). DOI: [10.1007/s41781-018-0020-1](https://doi.org/10.1007/s41781-018-0020-1). URL: <https://hal.inria.fr/hal-02427998>.
- [158] M. Raissi. ‘Deep Hidden Physics Models: Deep Learning of Nonlinear Partial Differential Equations’. In: *JMLR* 19 (2018), pp. 1–24.
- [159] H. Rakotoarison, M. Schoenauer and 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, Aug. 2019, pp. 3296–3303. DOI: [10.24963/ijcai.2019/457](https://doi.org/10.24963/ijcai.2019/457). URL: <https://hal.inria.fr/hal-02300884>.
- [160] G. Ramírez, E. Villatoro, B. Ionescu, H. J. Escalante, S. Escalera, M. Larson, H. Müller and I. Guyon. ‘Overview of the Multimedia Information Processing for Personality and Social Networks Analysis Contest’. In: *International Conference on Pattern Recognition (ICPR)*. Ed. by T. Tan, J. Kittler and A. Jai. IEEE, 2018, pp. 127–139.
- [161] E. Real, S. Moore, A. Selle, S. Saxena, Y. L. Suematsu, J. T. Q. Le and A. Kurakin. ‘Large-Scale Evolution of Image Classifiers’. In: *34th ICML*. 2017. arXiv: [1703.01041](https://arxiv.org/abs/1703.01041).
- [162] J. R. Rice. ‘The Algorithm Selection Problem’. In: *Advances in Computers* 15 (1976). Ed. by M. Rubinoff and M. C. Yovits, pp. 65–118. DOI: [10.1016/S0065-2458\(08\)60520-3](https://doi.org/10.1016/S0065-2458(08)60520-3). URL: <http://www.sciencedirect.com/science/article/pii/S0065245808605203>.
- [163] J. Rissanen. *Information and Complexity in Statistical Modeling*. Information Science and Statistics. Springer-Verlag, 2007.
- [164] M. Schmidt and H. Lipson. ‘Distilling free-form natural laws from experimental data’. In: *Science* 324.5923 (2009), pp. 81–85.
- [165] S. Sigtia, E. Benetos and S. Dixon. ‘An End-to-End Neural Network for Polyphonic Piano Music Transcription’. In: *IEEE/ACM Trans. Audio, Speech & Language Processing* 24.5 (2016), pp. 927–939.
- [166] J. Sirignano and K. Spiliopoulos. ‘DGM: A deep learning algorithm for solving partial differential equations’. In: *Journal of Computational Physics* 375 (2018), pp. 1339–1364. DOI: <https://doi.org/10.1016/j.jcp.2018.08.029>.
- [167] I. Stengers. *Another Science Is Possible*. Open Humanities Press, 2013.
- [168] D. H. Stern, R. Herbrich and T. Graepel. ‘Matchbox: large scale online bayesian recommendations’. In: *WWW*. Madrid, Spain: ACM Press, 2010, p. 111. DOI: [10.1145/1526709.1526725](https://doi.org/10.1145/1526709.1526725).
- [169] L. Sun-Hosoya, I. Guyon and M. Sebag. ‘Lessons learned from the AutoML challenge’. In: *Conférence sur l’Apprentissage Automatique 2018*. Rouen, France, June 2018. URL: <https://hal.inria.fr/hal-01811454>.
- [170] C. Sunstein and R. Thaler. *Nudge: Improving Decisions about Health, Wealth, and Happiness*. Yale University Press, 2008.
- [171] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow and R. Fergus. ‘Intriguing properties of neural networks’. In: *CoRR* abs/1312.6199 (2013). URL: <http://arxiv.org/abs/1312.6199>.

- [172] C. Tallec and Y. Ollivier. ‘Unbiased online recurrent optimization’. In: *International Conference On Learning Representation*. Vancouver, Canada, Apr. 2018. URL: <https://hal.archives-ouvertes.fr/hal-01972587>.
- [173] J. Vanschoren, J. N. van Rijn, B. Bischl and L. Torgo. ‘OpenML: networked science in machine learning’. In: *SIGKDD Explorations* 15.2 (2013), pp. 49–60. URL: <https://arxiv.org/abs/1407.7722>.
- [174] Volodymyr Mnih et al. ‘Human-level control through deep reinforcement learning’. In: *Nature* 518.7540 (2015), pp. 529–533. URL: <https://doi.org/10.1038/nature14236>.
- [175] J. Wan, S. Escalera, G. Anbarjafari, H. J. Escalante, X. Baró, I. Guyon, M. Madadi, J. Allik, C. Lin and 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. URL: <https://hal.inria.fr/hal-01677974>.
- [176] Y. Wang and D. M. Blei. *The Blessings of Multiple Causes: A Reply to Ogburn et al.* (2019). 2019. arXiv: 1910.07320 [stat.ML].
- [177] M. Welling. *Intelligence per Kilowatthour*. Keynote talk @ICML. 2018. URL: <https://icml.cc/Conferences/2018/Schedule?showEvent=1866>.
- [178] F. M. J. Willems, Y. M. Shtarkov and T. J. Tjalkens. ‘The context-tree weighting method: basic properties’. In: *IEEE Transactions on Information Theory* 41.3 (1995), pp. 653–664.
- [179] X. Zhang, C. Furtlehner, C. Germain-Renaud and M. Sebag. ‘Data Stream Clustering with Affinity Propagation’. In: *IEEE Transactions on Knowledge and Data Engineering* 26.7 (2014), p. 1. URL: <https://hal.inria.fr/hal-00862941>.