

RESEARCH CENTRE

**Inria Centre
at Université Côte d'Azur**

IN PARTNERSHIP WITH:
Université Côte d'Azur

2023

ACTIVITY REPORT

Project-Team
MAASAI

Models and Algorithms for Artificial Intelligence

IN COLLABORATION WITH: Laboratoire informatique, signaux systèmes
de Sophia Antipolis (I3S), Laboratoire Jean-Alexandre Dieudonné (JAD)

DOMAIN

**Applied Mathematics, Computation and
Simulation**

THEME

**Optimization, machine learning and
statistical methods**

Inria

Contents

Project-Team MAASAI	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	4
3 Research program	4
4 Application domains	6
5 Highlights of the year	7
5.1 Publications	7
5.2 Important scientific events	7
6 New software, platforms, open data	7
6.1 R and Python packages	7
6.2 SAAS platforms	11
7 New results	12
7.1 Unsupervised learning	12
7.1.1 Generalized Mutual Information: A Framework for Discriminative Clustering	12
7.1.2 Sparse GEMINI for Joint Discriminative Clustering and Feature Selection	13
7.1.3 A Deep Dynamic Latent Block Model for the Co-clustering of Zero-Inflated Data Matrices	13
7.1.4 Dimension-Grouped Mixed Membership Models for Multivariate Categorical Data	14
7.1.5 Embedded Topics in the Stochastic Block Model	14
7.1.6 Deep Latent Position Topic Model for Clustering and Representation of Networks with Textual Edges	15
7.1.7 Clustering: from modeling to visualizing Mapping clusters as spherical Gaussians	16
7.1.8 A Partition-Based EM Stopping Rule for Avoiding Degeneracy in Gaussian Mixtures with Missing Data	16
7.2 Understanding (deep) learning models	17
7.2.1 Explainability as statistical inference	17
7.2.2 Kernel-Matrix Determinant Estimates from stopped Cholesky Decomposition	17
7.2.3 The Risks of Recourse in Binary Classification	18
7.2.4 Logic Explained Networks	18
7.2.5 Forward Approximate Solution for Linear Quadratic Tracking	20
7.2.6 A Sea of Words: An In-Depth Analysis of Anchors for Text Data	20
7.2.7 Interpretable Neural-Symbolic Concept Reasoning	21
7.3 Adaptive and robust learning	21
7.3.1 On the Robustness of Text Vectorizers	21
7.3.2 Knowledge-Driven Active Learning	21
7.3.3 Taming the Diversity of Computational Notebooks	22
7.3.4 Toward Novel Optimizers: A Moreau-Yosida View of Gradient-Based Learning	22
7.4 Learning with heterogeneous and corrupted data	23
7.4.1 Mind the map! Accounting for existing map information when estimating online HDMaps from sensor data	23
7.4.2 Exploring the Road Graph in Trajectory Forecasting for Autonomous Driving	23
7.4.3 Machinery Anomaly Detection using artificial neural networks and signature feature extraction	23
7.4.4 Are labels informative in semi-supervised learning? Estimating and leveraging the missing-data mechanism.	24
7.4.5 AnoRand: Deep Learning-Based Semi-Supervised Anomaly Detection with Synthetic Labels	25
7.4.6 Fed-MIWAE: Federated Imputation of Incomplete Data via Deep Generative Models	26

7.4.7	Don't fear the unlabelled: safe deep semi-supervised learning via simple debiasing	26
7.4.8	The graph embedded topic model	27
7.4.9	Model-based clustering with Missing Not At Random Data	28
7.4.10	Continuous Latent Position Models for Instantaneous Interactions	28
7.4.11	DeepWILD: wildlife identification, localization and population estimation from camera trap videos in the Parc National du Mercantour	29
7.4.12	Another Point of View on Visual Speech Recognition	30
8	Bilateral contracts and grants with industry	31
9	Partnerships and cooperations	32
9.1	International initiatives	32
9.1.1	Participation in other International Programs	33
9.2	European initiatives	33
9.2.1	FP7 & H2020 Projects	33
9.3	National initiatives	34
9.4	Regional initiatives	34
10	Dissemination	35
10.1	Promoting scientific activities	35
10.1.1	Scientific events: organisation	35
10.1.2	Scientific publishing	35
10.1.3	Leadership within the scientific community	36
10.1.4	Scientific expertise	36
10.2	Teaching - Supervision - Juries	36
10.3	Popularization	37
10.4	Interventions	37
11	Scientific production	37
11.1	Major publications	37
11.2	Publications of the year	38
11.3	Other	42

Project-Team MAASAI

Creation of the Project-Team: 2020 February 01

Keywords

Computer sciences and digital sciences

- A3.1. – Data
 - A3.1.10. – Heterogeneous data
 - A3.1.11. – Structured data
- A3.4. – Machine learning and statistics
 - A3.4.1. – Supervised learning
 - A3.4.2. – Unsupervised learning
 - A3.4.6. – Neural networks
 - A3.4.7. – Kernel methods
 - A3.4.8. – Deep learning
- A9. – Artificial intelligence
 - A9.2. – Machine learning

Other research topics and application domains

- B3.6. – Ecology
- B6.3.4. – Social Networks
- B7.2.1. – Smart vehicles
- B8.2. – Connected city
- B9.6. – Humanities

1 Team members, visitors, external collaborators

Research Scientist

- Pierre-Alexandre Mattei [INRIA, Researcher]

Faculty Members

- Charles Bouveyron [Team leader, UNIV COTE D'AZUR, Professor, HDR]
- Marco Corneli [UNIV COTE D'AZUR, Chaire de Professeur Junior]
- Damien Garreau [UNIV COTE D'AZUR, Associate Professor]
- Diane Lingrand [UNIV COTE D'AZUR, Associate Professor, from Sep 2023]
- Frederic Precioso [UNIV COTE D'AZUR, Professor, HDR]
- Michel Riveill [UNIV COTE D'AZUR, Professor, HDR]
- Aude Sportisse [UNIV COTE D'AZUR, from Oct 2023, Junior Fellow 3IA Côte d'Azur]
- Vincent Vandewalle [UNIV COTE D'AZUR, Professor, HDR]

Post-Doctoral Fellows

- Alessandro Betti [UNIV COTE D'AZUR, Post-Doctoral Fellow]
- Gabriele Ciravegna [UNIV COTE D'AZUR, Post-Doctoral Fellow, until May 2023]
- Aude Sportisse [INRIA, Post-Doctoral Fellow, until Sep 2023]
- Remy Sun [CNRS, Post-Doctoral Fellow, from Feb 2023]

PhD Students

- Davide Adamo [CNRS, from Oct 2023]
- Kilian Burgi [UNIV COTE D'AZUR]
- Gatien Caillet [UNIV COTE D'AZUR]
- Antoine Collin [UNIV COTE D'AZUR]
- Célia Dacruz [UNIV COTE D'AZUR]
- Mariam Grigoryan [UNIV COTE D'AZUR, from Oct 2023]
- Gianluigi Lopardo [UNIV COTE D'AZUR]
- Giulia Marchello [INRIA, from Sep 2023]
- Giulia Marchello [UNIV COTE D'AZUR, until Aug 2023]
- Hugo Miralles [ORANGE, until Nov 2023]
- Kevin Mottin [UNIV COTE D'AZUR]
- Seydina Ousmane Niang [UNIV COTE D'AZUR, from Oct 2023]
- Louis Ohl [UNIV COTE D'AZUR]
- Baptiste Pouthier [INRIA, from Jun 2023 until Aug 2023]

- Baptiste Pouthier [NXP, until May 2023]
- Hugo Schmutz [INRIA, from Oct 2023]
- Hugo Schmutz [UNIV COTE D'AZUR, until Sep 2023]
- Hugo Senetaire [UNIV DTU]
- Julie Tores [UNIV COTE D'AZUR]
- Cédric Vincent-Cuaz [INRIA, from Feb 2023 until Mar 2023]
- Cédric Vincent-Cuaz [UNIV COTE D'AZUR, until Jan 2023]
- Xuchun Zhang [INRIA, from Sep 2023 until Nov 2023]
- Xuchun Zhang [UNIV COTE D'AZUR, until Aug 2023]

Technical Staff

- Lucas Boiteau [INRIA, Engineer]
- Leonie Borne [INRIA, Engineer, from Sep 2023]
- Stephane Petiot [INRIA, Engineer, until Feb 2023]
- Li Yang [CNRS, Engineer]
- Mansour Zoubeirou A Mayaki [PRO BTP, Engineer]

Interns and Apprentices

- Davide Adamo [INRIA, Intern, until Feb 2023]
- Seydina Ousmane Niang [INRIA, Intern, from Apr 2023 until Sep 2023]
- Mathieu Occhipinti [UNIV COTE D'AZUR, Intern, from May 2023 until Aug 2023]
- Kenza Roche [INRIA, Intern, from Mar 2023 until Sep 2023]
- Andela Todorovic [INRIA, Intern, from Apr 2023 until Sep 2023]
- Ali Youness [INRIA, Intern, from Apr 2023 until Aug 2023]

Administrative Assistant

- Claire Senica [INRIA]

Visiting Scientists

- Marco Gori [UNIV SIENA, HDR]
- Nicolo Navarin [UNIV PADOVA, from May 2023 until Jul 2023]

External Collaborators

- Alexandre Destere [CHU NICE, from Mar 2023]
- Amosse Edouard [INSTANT SYSTEM]
- Pierre Latouche [UNIV CLERMONT AUVERG, HDR]
- Hans Ottosson [IBM]

2 Overall objectives

Artificial intelligence has become a key element in most scientific fields and is now part of everyone's life thanks to the digital revolution. Statistical, machine and deep learning methods are involved in most scientific applications where a decision has to be made, such as medical diagnosis, autonomous vehicles or text analysis. The recent and highly publicized results of artificial intelligence should not hide the remaining and new problems posed by modern data. Indeed, despite the recent improvements due to deep learning, the nature of modern data has brought new specific issues. For instance, learning with high-dimensional, atypical (networks, functions, ...), dynamic, or heterogeneous data remains difficult for theoretical and algorithmic reasons. The recent establishment of deep learning has also opened new questions such as: How to learn in an unsupervised or weakly-supervised context with deep architectures? How to design a deep architecture for a given situation? How to learn with evolving and corrupted data?

To address these questions, the Maasai team focuses on topics such as unsupervised learning, theory of deep learning, adaptive and robust learning, and learning with high-dimensional or heterogeneous data. The Maasai team conducts a research that links practical problems, that may come from industry or other scientific fields, with the theoretical aspects of Mathematics and Computer Science. In this spirit, the Maasai project-team is totally aligned with the "Core elements of AI" axis of the Institut 3IA Côte d'Azur. It is worth noticing that the team hosts three 3IA chairs of the Institut 3IA Côte d'Azur, as well as several PhD students funded by the Institut.

3 Research program

Within the research strategy explained above, the Maasai project-team aims at developing statistical, machine and deep learning methodologies and algorithms to address the following four axes.

Unsupervised learning The first research axis is about the development of models and algorithms designed for unsupervised learning with modern data. Let us recall that unsupervised learning — the task of learning without annotations — is one of the most challenging learning challenges. Indeed, if supervised learning has seen emerging powerful methods in the last decade, their requirement for huge annotated data sets remains an obstacle for their extension to new domains. In addition, the nature of modern data significantly differs from usual quantitative or categorical data. We ambition in this axis to propose models and methods explicitly designed for unsupervised learning on data such as high-dimensional, functional, dynamic or network data. All these types of data are massively available nowadays in everyday life (omics data, smart cities, ...) and they remain unfortunately difficult to handle efficiently for theoretical and algorithmic reasons. The dynamic nature of the studied phenomena is also a key point in the design of reliable algorithms.

On the one hand, we direct our efforts towards the development of unsupervised learning methods (clustering, dimension reduction) designed for specific data types: high-dimensional, functional, dynamic, text or network data. Indeed, even though those kinds of data are more and more present in every scientific and industrial domains, there is a lack of sound models and algorithms to learn in an unsupervised context from such data. To this end, we have to face problems that are specific to each data type: How to overcome the curse of dimensionality for high-dimensional data? How to handle multivariate functional data / time series? How to handle the activity length of dynamic networks? On the basis of our recent results, we ambition to develop generative models for such situations, allowing the modeling and the unsupervised learning from such modern data.

On the other hand, we focus on deep generative models (statistical models based on neural networks) for clustering and semi-supervised classification. Neural network approaches have demonstrated their efficiency in many supervised learning situations and it is of great interest to be able to use them in unsupervised situations. Unfortunately, the transfer of neural network approaches to the unsupervised context is made difficult by the huge amount of model parameters to fit and the absence of objective quantity to optimize in this case. We therefore study and design model-based deep learning methods that can handle unsupervised or semi-supervised problems in a statistically grounded way.

Finally, we also aim at developing explainable unsupervised models that can ease the interaction with the practitioners and their understanding of the results. There is an important need for such models,

in particular when working with high-dimensional or text data. Indeed, unsupervised methods, such as clustering or dimension reduction, are widely used in application fields such as medicine, biology or digital humanities. In all these contexts, practitioners are in demand of efficient learning methods which can help them to make good decisions while understanding the studied phenomenon. To this end, we aim at proposing generative and deep models that encode parsimonious priors, allowing in turn an improved understanding of the results.

Understanding (deep) learning models The second research axis is more theoretical, and aims at improving our understanding of the behavior of modern machine learning models (including, but not limited to, deep neural networks). Although deep learning methods and other complex machine learning models are obviously at the heart of artificial intelligence, they clearly suffer from an overall weak knowledge of their behavior, leading to a general lack of understanding of their properties. These issues are barriers to the wide acceptance of the use of AI in sensitive applications, such as medicine, transportation, or defense. We aim at combining statistical (generative) models with deep learning algorithms to justify existing results, and allow a better understanding of their performances and their limitations.

We particularly focus on researching ways to understand, interpret, and possibly explain the predictions of modern, complex machine learning models. We both aim at studying the empirical and theoretical properties of existing techniques (like the popular LIME), and at developing new frameworks for interpretable machine learning (for example based on deconvolutions or generative models). Among the relevant application domains in this context, we focus notably on text and biological data.

Another question of interest is: what are the statistical properties of deep learning models and algorithms? Our goal is to provide a statistical perspective on the architectures, algorithms, loss functions and heuristics used in deep learning. Such a perspective can reveal potential issues in existing deep learning techniques, such as biases or miscalibration. Consequently, we are also interested in developing statistically principled deep learning architectures and algorithms, which can be particularly useful in situations where limited supervision is available, and when accurate modeling of uncertainties is desirable.

Adaptive and Robust Learning The third research axis aims at designing new learning algorithms which can learn incrementally, adapt to new data and/or new context, while providing predictions robust to biases even if the training set is small.

For instance, we have designed an innovative method of so-called cumulative learning, which allows to learn a convolutional representation of data when the learning set is (very) small. The principle is to extend the principle of Transfer Learning, by not only training a model on one domain to transfer it once to another domain (possibly with a fine-tuning phase), but to repeat this process for as many domains as available. We have evaluated our method on mass spectrometry data for cancer detection. The difficulty of acquiring spectra does not allow to produce sufficient volumes of data to benefit from the power of deep learning. Thanks to cumulative learning, small numbers of spectra acquired for different types of cancer, on different organs of different species, all together contribute to the learning of a deep representation that allows to obtain unequalled results from the available data on the detection of the targeted cancers. This extension of the well-known Transfer Learning technique can be applied to any kind of data.

We also investigate active learning techniques. We have for example proposed an active learning method for deep networks based on adversarial attacks. An unlabelled sample which becomes an adversarial example under the smallest perturbations is selected as a good candidate by our active learning strategy. This does not only allow to train incrementally the network but also makes it robust to the attacks chosen for the active learning process.

Finally, we address the problem of biases for deep networks by combining domain adaptation approaches with Out-Of-Distribution detection techniques.

Learning with heterogeneous and corrupted data The last research axis is devoted to making machine learning models more suitable for real-world, "dirty" data. Real-world data rarely consist in a single kind of Euclidean features, and are generally heterogeneous. Moreover, it is common to find some form of

corruption in real-world data sets: for example missing values, outliers, label noise, or even adversarial examples.

Heterogeneous and non-Euclidean data are indeed part of the most important and sensitive applications of artificial intelligence. As a concrete example, in medicine, the data recorded on a patient in an hospital range from images to functional data and networks. It is obviously of great interest to be able to account for all data available on the patients to propose a diagnostic and an appropriate treatment. Notice that this also applies to autonomous cars, digital humanities and biology. Proposing unified models for heterogeneous data is an ambitious task, but first attempts (e.g. the Linkage¹ project) on combination of two data types have shown that more general models are feasible and significantly improve the performances. We also address the problem of conciliating structured and non-structured data, as well as data of different levels (individual and contextual data).

On the basis of our previous works (notably on the modeling of networks and texts), we first intend to continue to propose generative models for (at least two) different types of data. Among the target data types for which we would like to propose generative models, we can cite images and biological data, networks and images, images and texts, and texts and ordinal data. To this end, we explore modelings through common latent spaces or by hybridizing several generative models within a global framework. We are also interested in including potential corruption processes into these heterogeneous generative models. For example, we are developing new models that can handle missing values, under various sorts of missingness assumptions.

Besides the modeling point of view, we are also interested in making existing algorithms and implementations more fit for "dirty data". We study in particular ways to robustify algorithms, or to improve heuristics that handle missing/corrupted values or non-Euclidean features.

4 Application domains

The Maasai research team has the following major application domains:

Medicine Most of team members apply their research work to Medicine or extract theoretical AI problems from medical situations. In particular, our main applications to Medicine are concerned with pharmacovigilance, medical imaging, and omics. It is worth noticing that medical applications cover all research axes of the team due to the high diversity of data types and AI questions. It is therefore a preferential field of application of the models and algorithms developed by the team.

Digital humanities Another important application field for Maasai is the increasingly dynamic one of digital humanities. It is an extremely motivating field due to the very original questions that are addressed. Indeed, linguists, sociologists, geographers and historians have questions that are quite different than the usual ones in AI. This allows the team to formalize original AI problems that can be generalized to other fields, allowing to indirectly contribute to the general theory and methodology of AI.

Multimedia The last main application domain for Maasai is multimedia. With the revolution brought to computer vision field by deep learning techniques, new questions have appeared such as combining subsymbolic and symbolic approaches for complex semantic and perception problems, or as edge AI to embed machine learning approaches for multimedia solutions preserving privacy. This domain brings new AI problems which require to bridge the gap between different views of AI.

Other application domains Other topics of interest of the team include astrophysics, bioinformatics and ecology.

¹The Linkage project: linkage.fr

5 Highlights of the year

5.1 Publications

Among the numerous publications of the year, we can highlight the publications of 4 papers at ICML (International Conference on Machine Learning), 2 at JMLR (Journal of Machine Learning Research), 1 at ICLR (International Conference on Learning Representations), 1 at AISTATS (International Conference on Artificial Intelligence and Statistics), 2 at ECML (European Conference on Machine Learning), and 1 at IJCNN (International Joint Conference on Artificial Neural Networks). We also had a best paper award at AIXIA 2023 (Advances in Artificial Intelligence: XXIIInd International Conference of the Italian Association for Artificial Intelligence, Rome, Italy).

5.2 Important scientific events

- Pierre-Alexandre Mattei co-organized and taught the first Generative Models Summer School (GeMSS), held in Copenhagen (June 26th to 30th, 2023). More details at gemss.ai/2023/. Several PhD students from Maasai were among the hundred participants.
- Pierre-Alexandre Mattei co-organized the Generative Models and Uncertainty quantification workshop (GenU), held in Copenhagen (September 20-21, 2023). More details at genu.ai/2023/.
- Damien Garreau and Frédéric Precioso co-organized the 2nd Edition of the Nice Workshop on Interpretability (NWI), held in Nice (November 30 - December 1, 2023). More details at sites.google.com/view/damien-garreau/home.
- Charles Bouveyron, Marco Corneli and Pierre-Alexandre Mattei were member of the scientific committee of the Workshop Statlearn, Montpellier, 5-7 April 2023. More details at statlearn.sciencesconf.org.

6 New software, platforms, open data

For the Maasai research team, the main objective of the software implementations is to experimentally validate the results obtained and ease the transfer of the developed methodologies to industry. Most of the software will be released as R or Python packages that requires only a light maintaining, allowing a relative longevity of the codes. Some platforms are also proposed to ease the use of the developed methodologies by users without a strong background in Machine Learning, such as scientists from other fields.

6.1 R and Python packages

The team maintains several R and Python packages, among which the following ones have been released or updated in 2023:

GEMCLUS Web site: <https://github.com/gemini-clustering/GemClus>.

- Software Family : vehicle;
- Audience: community;
- Evolution and maintenance: basic;
- Duration of the Development (Duration): 1 year;
- Free Description: The GemClus package [59] is a python software for using GEMNI with various models, from logistic regression to multi-layered perceptrons, intended for small scale datasets. The package aims at minimal dependencies and incorporates some methods from related discriminative clustering models such as regularized information maximization (RIM). The package documentation is available here: <https://gemini-clustering.github.io> and the open-source code is available here: <https://github.com/gemini-clustering/GemClus>, see also Figure 1.

Welcome to GemClus documentation!

Welcome and thank you for checking GEMINI out, this really pleases us.

GEMINI is a clustering-dedicated function derived from information theory that allows to do clustering without hypotheses on the data distributions. We try to provide overall some implementations of small-scale discriminative clustering methods in this package.

Contents

- [Getting started.](#)
- [User Guide.](#)
- [API.](#)
- [Examples](#)

 [downloads](#) 3k

Figure 1: Home page of the GemClus package documentation.

SMACE. Web site: <https://github.com/gianluigilopardo/smace>.

- Software Family : vehicle;
- Audience: community;
- Evolution and maintenance: basic;
- Duration of the Development (Duration): 1 year;
- Free Description: this **Python** package implements SMACE, the first *Semi-Model-Agnostic Contextual Explainer*. The code is available on Github as well as on pypi at <https://pypi.org/project/smace>, distributed under the MIT License.

POT. Web site: <https://PythonOT.github.io/>.

- Software Family : vehicle;
- Audience: community;
- Evolution and maintenance: lts, long term support.
- Duration of the Development (Duration): 23 Releases since April 2016. MAASAI contribution: since release 0.8.0 In November 2021.
- Free Description: Open source **Python** library that provides several solvers for optimization problems related to Optimal Transport for signal, image processing and machine learning. Distribution: PyPI distribution, Anaconda distribution. The library has been tested on Linux, MacOSX and Windows. It requires a C++ compiler for building/installing. License: MIT license. Website and documentation: <https://PythonOT.github.io/> Source Code (MIT): <https://github.com/PythonOT/POT> The software contains implementations of more than 40 research papers providing new solvers for Optimal Transport problems.

CLPM. Web site: <https://github.com/marcogenni/CLPM>.

- Software Family : vehicle;
- Audience: community;
- Evolution and maintenance: basic;
- Duration of the Development (Duration): 2 years;
- Free Description: this **Python** software that implements CLPM, a continuous time extension of the Latent Position Model for graphs embedding. The code is available on Github and distributed under the MIT License.

ordinalLBM. Web site: <https://cran.r-project.org/web/packages/ordinalLBM/index.html>.

- Software Family: vehicle;
- Audience: community;
- Evolution and maintenance: basic;
- Duration of the Development (Duration): 3 years;
- Free Description: this **R** package implements the inference for the ordinal latent block model for not missing at random data. The code is available on the CRAN repository and distributed under the GPL-2 | GPL-3 licence.

R-miss-tastic. Web site: <https://rmissstastic.netlify.app/>.

- Software Family: vehicle.
- Audience: community.
- Evolution and maintenance: basic.
- Duration of the Development (Duration): 2 years.
- Free Description: “R-miss-tastic” platform aims to provide an overview of standard missing values problems, methods, and relevant implementations of methodologies. Beyond gathering and organizing a large majority of the material on missing data (bibliography, courses, tutorials, implementations), “R-miss-tastic” covers the development of standardized analysis workflows. Several pipelines are developed in **R** and **Python** to allow for hands-on illustration of and recommendations on missing values handling in various statistical tasks such as matrix completion, estimation and prediction, while ensuring reproducibility of the analyses. Finally, the platform is dedicated to users who analyze incomplete data, researchers who want to compare their methods and search for an up-to-date bibliography, and also teachers who are looking for didactic materials (notebooks, video, slides). The platform takes the form of a reference website: <https://rmissstastic.netlify.app/>.

GEMINI Web site: <https://github.com/oshillou/GEMINI>.

- Software Family: vehicle;
- Audience: community;
- Evolution and maintenance: basic;
- Duration of the Development (Duration): 1 year
- Free Description: a **Python** software that allows users to manipulate GEMINI objectives functions on their own data. By specifying a configuration file, users may plug their own data to GEMINI clustering as well as some custom models. The core of the software essentially lies in the file entitled *losses.py* which contains all of the core objective functions for clustering. The software is currently under no licence, but we are discussing about setting it under a GPL v3 licence.

FunHDDC. Web site: <https://cran.r-project.org/web/packages/funHDDC/index.html>

- Software Family: vehicle;
- Audience: community;
- Evolution and maintenance: basic;
- Duration of the Development (Duration): 2 years;
- Free Description: this **R** package implements the inference for Clustering multivariate functional data in group-specific functional subspaces. The code is available on the CRAN repository and distributed under the GPL-2 | GPL-3 licence.

FunFEM. Web site: <https://cran.r-project.org/web/packages/funFEM/index.html>

- Software Family : `vehicle`;
- Audience: `community`;
- Evolution and maintenance: `basic`;
- Duration of the Development (Duration): 2 years;
- Free Description: released in 2021, this **R** package implements the inference for the clustering of functional data by modeling the curves within a common and discriminating functional subspace. The code is available on the CRAN repository and distributed under the GPL-2 | GPL-3 licence.

FunLBM. Web site: <https://cran.r-project.org/web/packages/funLBM/index.html>

- Software Family : `vehicle`;
- Audience: `community`;
- Evolution and maintenance: `basic`;
- Duration of the Development (Duration): 1 years;
- Free Description: released in 2022, this **R** package implements the inference for the co-clustering of functional data (time series) with application to the air pollution data in the South of France. The code is available on the CRAN repository and distributed under the GPL-2 | GPL-3 licence.

MIWAE. Web Site: <https://github.com/pamattei/miwae>

- Software Family: `vehicle`;
- Audience: `community`;
- Evolution and maintenance: `basic`;
- Free Description: this is the implementations of the MIWAE method for handling missing data with deep generative modeling, as described in previous works of PA. Mattei. The **Python** code is available on Github and freely distributed.

not-MIWAE. Web Site: <https://github.com/nbip/notMIWAE>

- Software Family: `vehicle`;
- Audience: `community`;
- Evolution and maintenance: `basic`;
- Free Description: this is the implementations of the not-MIWAE method for handling missing not-at-random data with deep generative modeling. The **Python** code is available on Github and freely distributed.

supMIWAE. Web Site: <https://github.com/nbip/suptMIWAE>

- Software Family: `vehicle`;
- Audience: `community`;
- Evolution and maintenance: `basic`;
- Free Description: this is the implementations of the supMIWAE method for supervised deep learning with missing values. The **Python** code is available on Github and freely distributed.

fisher-EM. Web Site: <https://cran.r-project.org/web/packages/FisherEM/index.html>

- Software Family: vehicle;
- Audience: community;
- Evolution and maintenance: basic;
- Free Description: The FisherEM algorithm, proposed by Bouveyron in previous works is an efficient method for the clustering of high-dimensional data. FisherEM models and clusters the data in a discriminative and low-dimensional latent subspace. It also provides a low-dimensional representation of the clustered data. A sparse version of Fisher-EM algorithm is also provided in this package created in 2020. Distributed under the GPL-2 licence.

6.2 SAAS platforms

The team is also proposing some SAAS (software as a service) platforms in order to allow scientists from other fields or companies to use our technologies. The team developed the following platforms:

DiagnoseNET: Automatic Framework to Scale Neural Networks on Heterogeneous Systems. Web Site: <https://diagnosenet.github.io/>.

- Software Family: Transfer;
- Audience: partners;
- Evolution and maintenance: basic;
- Free Description: DiagnoseNET is a platform oriented to design a green intelligence medical workflow for deploying medical diagnostic tools with minimal infrastructure requirements and low power consumption. The first application built was to automate the unsupervised patient phenotype representation workflow trained on a mini-cluster of Nvidia Jetson TX2. The **Python** code is available on Github and freely distributed.

Indago. Web site: <http://indago.inria.fr>. (Inria internal)

- Software Family: transfer.
- Audience: partners
- Evolution and maintenance: lts: long term support.
- Duration of the Development (Duration): 1.8 years
- Free Description: Indago implements a textual graph clustering method based on a joint analysis of the graph structure and the content exchanged between each nodes. This allows to reach a better segmentation than what could be obtained with traditional methods. Indago's main applications are built around communication network analysis, including social networks. However, Indago can be applied on any graph-structured textual network. Thus, Indago have been tested on various data, such as tweet corpus, mail networks, scientific paper co-publication network, etc.

The software is used as a fully autonomous SaaS platform with 2 parts :

1. A Python kernel that is responsible for the actual data processing.
2. A web application that handles collecting, pre-processing and saving the data, such as providing a set of visualization for the interpretation of the results.

Indago is deployed internally on the Inria network and used mainly by the development team for testing and research purposes. We also build tailored versions for industrial or academic partners that use the software externally (with contractual agreements).

Topix. Web site: <https://topix.mi.parisdescartes.fr>

- Software Family: research;
- Audience: universe;
- Evolution and maintenance: lts;
- Free Description: Topix is an innovative AI-based solution allowing to summarize massive and possibly extremely sparse data bases involving text. Topix is a versatile technology that can be applied in a large variety of situations where large matrices of texts / comments / reviews are written by users on products or addressed to other individuals (bi-partite networks). The typical use case consists in an e-commerce company interested in understanding the relationship between its users and the sold products thanks to the analysis of user comments. A simultaneous clustering (co-clustering) of users and products is produced by the Topix software, based on the key emerging topics from the reviews and by the underlying model. The Topix demonstration platform allows you to upload your own data on the website, in a totally secured framework, and let the AI-based software analyze them for you. The platform also provides some typical use cases to give a better idea of what Topix can do.

7 New results

7.1 Unsupervised learning

7.1.1 Generalized Mutual Information: A Framework for Discriminative Clustering

Participants: Louis Ohl, Pierre-Alexandre Mattei, Frédéric Precioso.

Keywords: Clustering, Deep learning, Information Theory, Mutual Information

Collaborations: Mickael Leclercq, Arnaud Droit (Centre de recherche du CHU de Québec-Université, Université Laval), Warith Harchaoui (Jellysmack)

In the last decade, recent successes in deep clustering majorly involved the mutual information (MI) as an unsupervised objective for training neural networks with increasing regularizations. While the quality of the regularizations has been largely discussed for improvements, little attention has been dedicated to the relevance of MI as a clustering objective. In this paper, we first highlight how the maximization of MI does not lead to satisfying clusters. We identified the Kullback-Leibler divergence as the main reason of this behavior. Hence, we generalize the mutual information by changing its core distance, introducing the generalized mutual information (GEMINI): a set of metrics for unsupervised neural network training [51]. Unlike MI, some GEMINIs do not require regularizations when training. Some of these metrics are geometry-aware thanks to distances or kernels in the data space. Finally, we highlight that GEMINIs can automatically select a relevant number of clusters, a property that has been little studied in deep clustering context where the number of clusters is a priori unknown.




Figure 2: Entropy maps of the predictions of each MLP trained using a GEMINI or the MI. For more details, see 7.1.1.

7.1.2 Sparse GEMINI for Joint Discriminative Clustering and Feature Selection

Participants: Louis Ohl, Pierre-Alexandre Mattei, Charles Bouveyron, Mickael Leclercq, Arnaud Droit, Frédéric Precioso.

Keywords: Clustering, Deep learning, Sparsity, Feature Selection

Collaborations: Centre de recherche du CHU de Québec-Université, Université Laval

Feature selection in clustering is a hard task which involves simultaneously the discovery of relevant clusters as well as relevant variables with respect to these clusters. While feature selection algorithms are often model-based through optimized model selection or strong assumptions on the data distribution, we introduce a discriminative clustering model trying to maximize a geometry-aware generalization of the mutual information called GEMINI with a simple ℓ_1 penalty: the Sparse GEMINI [52]. This algorithm avoids the burden of combinatorial feature subset exploration and is easily scalable to high-dimensional data and large amounts of samples while only designing a discriminative clustering model (see Figure 3). We demonstrate the performances of Sparse GEMINI on synthetic datasets and large-scale datasets. Our results show that Sparse GEMINI is a competitive algorithm and has the ability to select relevant subsets of variables with respect to the clustering without using relevance criteria or prior hypotheses.

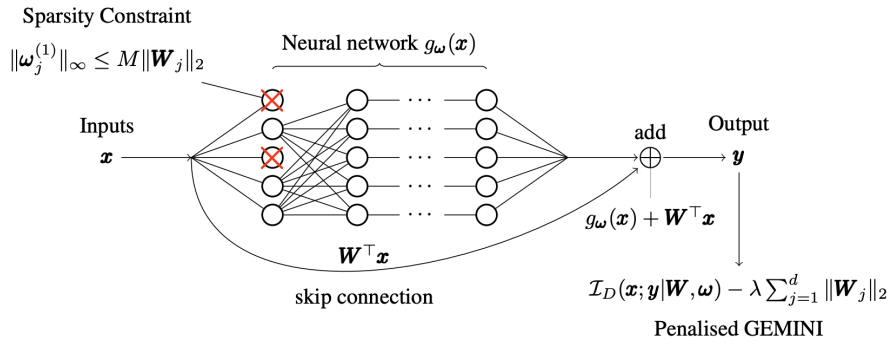


Figure 3: Description of the complete Sparse GEMINI model. For more details, see 7.1.2.

7.1.3 A Deep Dynamic Latent Block Model for the Co-clustering of Zero-Inflated Data Matrices

Participants: Giulia Marchello, Marco Corneli, Charles Bouveyron.

Keywords: Co-clustering, Latent Block Model, zero-inflated distributions, dynamic systems, VEM algorithm.

Collaborations: Regional Center of Pharmacovigilance (RCPV) of Nice.

The simultaneous clustering of observations and features of data sets (known as co-clustering) has recently emerged as a central machine learning application to summarize massive data sets. However, most existing models focus on continuous and dense data in stationary scenarios, where cluster assignments do not evolve over time. In [36] we introduce a novel latent block model for the dynamic co-clustering of data matrices with high sparsity. To properly model this type of data, we assume that the observations follow a time and block dependent mixture of zero-inflated distributions, thus combining stochastic processes with the time-varying sparsity modeling. To detect abrupt changes in the dynamics of both cluster memberships and data sparsity, the mixing and sparsity proportions are modeled through systems of ordinary differential equations. The inference relies on an original variational procedure whose maximization step trains fully connected neural networks in order to solve the dynamical systems. Numerical

experiments on simulated data sets demonstrate the effectiveness of the proposed methodology in the context of count data. The proposed method, called $ZI_{\mathcal{D}}$ -dLBM, was then applied to two real data sets. The first is the data set on the London Bike sharing system while the second is a pharmacovigilance data set, on adverse drug reaction (ADR) reported to the Regional Center of Pharmacovigilance (RCPV) in Nice, France. Figure 4 shows some of the main results obtained through the application of $ZI_{\mathcal{D}}$ -dLBM on the pharmacovigilance data set.

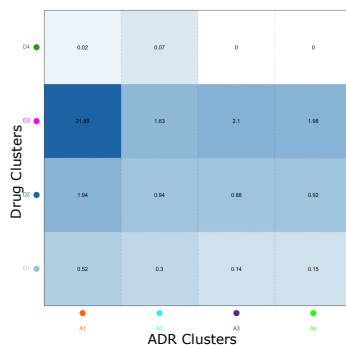


Figure 4: Estimated Poisson intensities, each color represents a different drug (ADR) cluster. For more details see Section 7.1.3.

7.1.4 Dimension-Grouped Mixed Membership Models for Multivariate Categorical Data

Participants: Elena Erosheva.

Keywords: Bayesian estimation, grant peer review, inter-rater reliability, maximum likelihood estimation, measurement, mixed-effects models

Collaborations: Yuqi Gu (Columbia University), Gongjun Xu (University of Michigan), David B. Dunson (Duke University)

Mixed Membership Models (MMMs) are a popular family of latent structure models for complex multivariate data. Instead of forcing each subject to belong to a single cluster, MMMs incorporate a vector of subject-specific weights characterizing partial membership across clusters. With this flexibility come challenges in uniquely identifying, estimating, and interpreting the parameters. In [17], we propose a new class of Dimension-Grouped MMMs (Gro-M³s) for multivariate categorical data, which improves parsimony and interpretability. In Gro-M³s, observed variables are partitioned into groups such that the latent membership is constant for variables within a group but can differ across groups. Traditional latent class models are obtained when all variables are in one group, while traditional MMMs are obtained when each variable is in its own group. The new model corresponds to a novel decomposition of probability tensors. Theoretically, we derive transparent identifiability conditions for both the unknown grouping structure and model parameters in general settings. Methodologically, we propose a Bayesian approach for Dirichlet Gro-M³s to infer the variable grouping structure and estimate model parameters. Simulation results demonstrate good computational performance and empirically confirm the identifiability results. We illustrate the new methodology through an application to a functional disability dataset.

7.1.5 Embedded Topics in the Stochastic Block Model

Participants: Charles Bouveyron, Rémi Boutin, Pierre Latouche.




Figure 5: Clustering on a simulated network with SBM (left), SBM+ETM (center) and ETSBM (right). See Section 7.1.5 for more details.

Keywords: generative models, clustering, networks, text, topic modeling

Collaborations: service politique du journal Le Monde

Communication networks such as emails or social networks are now ubiquitous and their analysis has become a strategic field. In many applications, the goal is to automatically extract relevant information by looking at the nodes and their connections. Unfortunately, most of the existing methods focus on analyzing the presence or absence of edges and textual data is often discarded. However, all communication networks actually come with textual data on the edges. In order to take into account this specificity, we consider in [12] networks for which two nodes are linked if and only if they share textual data. We introduce a deep latent variable model allowing embedded topics to be handled called ETSBM to simultaneously perform clustering on the nodes while modeling the topics used between the different clusters (see Figure 5). ETSBM extends both the stochastic block model (SBM) and the embedded topic model (ETM) which are core models for studying networks and corpora, respectively. The inference is done using a variational-Bayes expectation-maximization algorithm combined with a stochastic gradient descent. The methodology is evaluated on synthetic data and on a real world dataset.

7.1.6 Deep Latent Position Topic Model for Clustering and Representation of Networks with Textual Edges

Participants: Charles Bouveyron, Rémi Boutin, Pierre Latouche.

Keywords: generative models, clustering, networks, text, topic modeling

Numerical interactions leading to users sharing textual content published by others are naturally represented by a network where the individuals are associated with the nodes and the exchanged texts with the edges. To understand those heterogeneous and complex data structures, clustering nodes into homogeneous groups as well as rendering a comprehensible visualization of the data is mandatory. To address both issues, we introduced in [47] Deep-LPTM, a model-based clustering strategy relying on a variational graph auto-encoder approach as well as a probabilistic model to characterize the topics of discussion. Deep-LPTM allows to build a joint representation of the nodes and of the edges in two embeddings spaces. The parameters are inferred using a variational inference algorithm. We also introduce IC2L, a model selection criterion specifically designed to choose models with relevant clustering and visualization properties. An extensive benchmark study on synthetic data is provided. In particular, we find that Deep-LPTM better recovers the partitions of the nodes than the state-of-the-art ETSBM and STBM (see Figure 6). Eventually, the emails of the Enron company are analyzed and visualizations of the results are presented, with meaningful highlights of the graph structure.

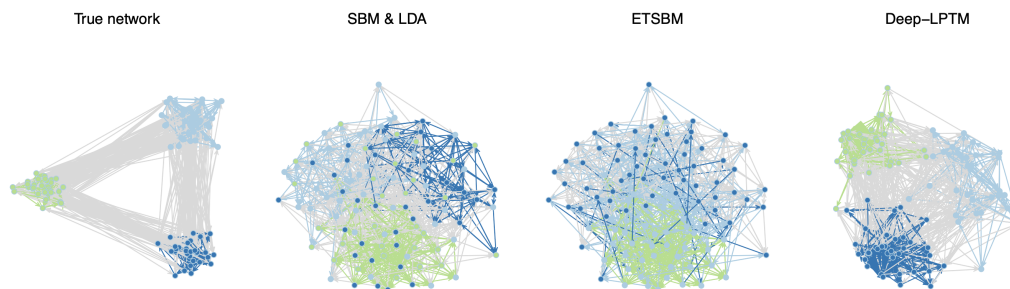


Figure 6: Illustration of Deep-LPTM main contributions on a synthetic network.

7.1.7 Clustering: from modeling to visualizing Mapping clusters as spherical Gaussians

Participants: Vincent Vandewalle.

Collaborations: Christophe Biernacki, Matthieu Marbac

Keywords: model-based clustering, visualization

A generic method is introduced to visualize in a "Gaussian-like way," and onto \mathbb{R}^2 , results of Gaussian or non-Gaussian-based clustering. The key point is to explicitly force a visualization based on a spherical Gaussian mixture to inherit from the within cluster overlap that is present in the initial clustering mixture. The result is a particularly user-friendly drawing of the clusters, displayed for instance in Figure 7, providing any practitioner with an overview of the potentially complex clustering result. An entropic measure provides information about the quality of the drawn overlap compared with the true one in the initial space. This work has been presented in two international conferences [43, 30].

7.1.8 A Partition-Based EM Stopping Rule for Avoiding Degeneracy in Gaussian Mixtures with Missing Data

Participants: Vincent Vandewalle.

Collaborations: Christophe Biernacki

Keywords: model-based clustering, Gaussian mixtures, degeneracy

The missing data problem is well-known, but its frequency increases with the growing size of multivariate modern datasets. In Gaussian model-based clustering, the EM algorithm easily takes into account such data but the degeneracy problem in Gaussian mixtures is dramatically aggravated during the EM runs. Indeed, numerical experiments clearly reveal that parameter degeneracy is quite slow and also more frequent than with complete data. Consequently, parameter degenerated solutions may be confused with valuable parameter solutions and, in addition, computing time may be wasted through wrong runs. A theoretical and practical study of the parameter degeneracy reveals that, in practice, due to its very specific dynamic, degenerated situations are difficult to detect or to avoid efficiently with traditional parameter-based methods used in the no missing data case. However, a simple and low informational condition on the latent partition, by opposition to more classical conditions on the mixture parameters themselves, produces versions of the EM algorithm which can efficiently prevent degeneracy. In particular, we present and experiment a very simple partition-based stopping rule of EM which shows good behavior on real and simulated data. This work has been presented in an international conference [29].




Figure 7: Illustration of the Gaussian visualization of the clustering of categorical data.

7.2 Understanding (deep) learning models

7.2.1 Explainability as statistical inference

Participants: Hugo Senetaire, Damien Garreau, Pierre-Alexandre Mattei.

Keywords: Interpretability, Human and AI, Explainability, latent variable models

Collaborations: Jes Frellsen (Technical University of Denmark)

A wide variety of model explanation approaches have been proposed in recent years, all guided by very different rationales and heuristics. In [40], we take a new route and cast interpretability as a statistical inference problem. We propose a general deep probabilistic model designed to produce interpretable predictions (see Fig. 8). The model's parameters can be learned via maximum likelihood, and the method can be adapted to any predictor network architecture, and any type of prediction problem. Our method is a case of amortized interpretability models, where a neural network is used as a selector to allow for fast interpretation at inference time. Several popular interpretability methods are shown to be particular cases of regularized maximum likelihood for our general model. We propose new datasets with ground truth selection which allow for the evaluation of the features importance map. Using these datasets, we show experimentally that using multiple imputation provides more reasonable interpretation.

7.2.2 Kernel-Matrix Determinant Estimates from stopped Cholesky Decomposition

Participants: Damien Garreau.

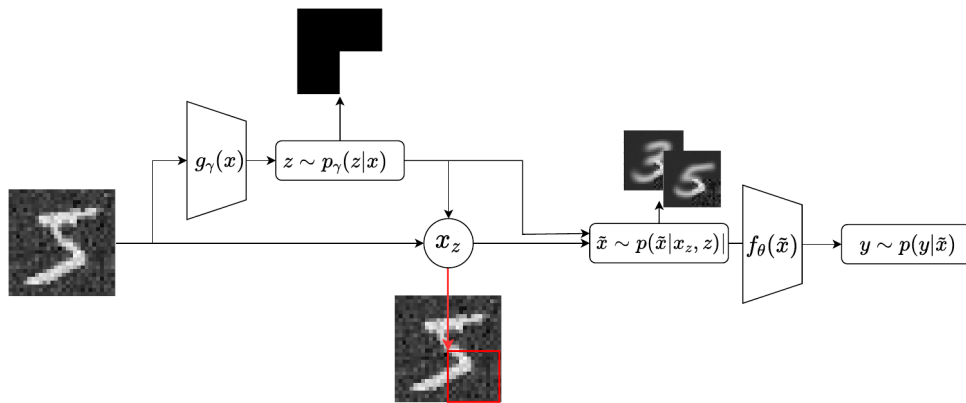


Figure 8: The LEX pipeline allows us to transform any prediction model into an explainable one. For more details, see [7.2.1](#).

Keywords: Gaussian processes

Collaborations: Simon Bartels, Wouter Boomsma, Jes Frelsen (Technical University of Denmark)

Algorithms involving Gaussian processes or determinantal point processes typically require computing the determinant of a kernel matrix. Frequently, the latter is computed from the Cholesky decomposition, an algorithm of cubic complexity in the size of the matrix. We showed in [11] that, under mild assumptions, it is possible to estimate the determinant from only a sub-matrix, with probabilistic guarantee on the relative error. We present an augmentation of the Cholesky decomposition that stops under certain conditions before processing the whole matrix. Experiments demonstrate that this can save a considerable amount of time while having an overhead of less than 5% when not stopping early. More generally, we present a probabilistic stopping strategy for the approximation of a sum of known length where addends are revealed sequentially. We do not assume independence between addends, only that they are bounded from below and decrease in conditional expectation.

7.2.3 The Risks of Recourse in Binary Classification

Participants: Damien Garreau.

Keywords: Interpretability, Recourse, Machine Learning Theory

Collaborations: Hidde Fokkema, Tim van Erven

Algorithmic recourse provides explanations that help users overturn an unfavorable decision by a machine learning system. But so far very little attention has been paid to whether providing recourse is beneficial or not. We introduce in [48] an abstract learning-theoretic framework that compares the risks (i.e. expected losses) for classification with and without algorithmic recourse. This allows us to answer the question of when providing recourse is beneficial or harmful at the population level. Surprisingly, we find that there are many plausible scenarios in which providing recourse turns out to be harmful, because it pushes users to regions of higher class uncertainty and therefore leads to more mistakes. We further study whether the party deploying the classifier has an incentive to strategize in anticipation of having to provide recourse, and we find that sometimes they do, to the detriment of their users. Providing algorithmic recourse may therefore also be harmful at the systemic level. We confirm our theoretical findings in experiments on simulated and real-world data. All in all, we conclude that the current concept of algorithmic recourse is not reliably beneficial, and therefore requires rethinking.

7.2.4 Logic Explained Networks

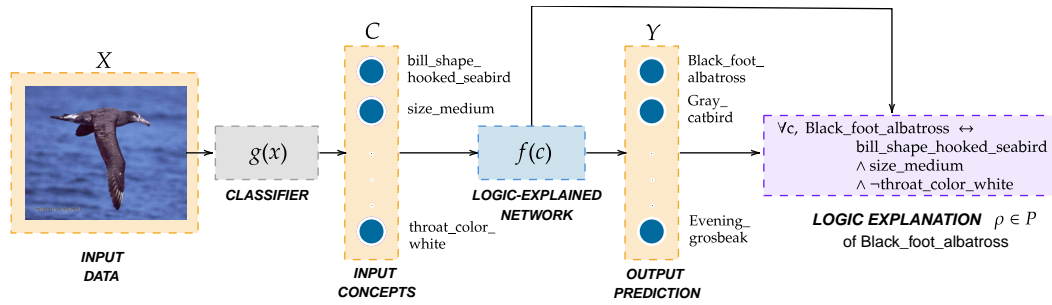


Figure 9: An instantiation of the Logic Explained Network framework of Section 7.2.4, inspired by the CUB 200-2011 fine-grained classification dataset. Classes are divided into a two-level hierarchy. A LEN is placed on top of a convolutional neural network $g(\cdot)$ in order to (i) classify the species of the bird in input and (ii) provide an explanation on why it belongs to this class. The logic explanation in the example showcases the predicted output class.

Participants: Gabriele Ciravegna, Marco Gori.

Keywords: XAI, Explainability-by-design, Concept-based Explanations, Human and AI

Collaborations: Pietro Barbiero, Pietro Lió (University of Cambridge), Francesco Giannini, Marco Maggini, Stefano Melacci (Università di Siena)

In [13] we present a unified framework for XAI allowing the design of a family of neural models, the Logic Explained Networks (LENs, see Fig. 9), which are trained to solve-and-explain a categorical learning problem integrating elements from deep learning and logic. Differently from vanilla neural architectures, LENs can be directly interpreted by means of a set of first order logic (FOL) formulas. To implement such a property, LENs require their inputs to represent the activation scores of human-understandable concepts. Then, specifically designed learning objectives allow LENs to make predictions in a way that is well suited for providing FOL-based explanations that involve the input concepts. To reach this goal, LENs leverage parsimony criteria aimed at keeping their structure simple. There are several different computational pipelines in which a LEN can be configured, depending on the properties of the considered problem and on other potential experimental constraints. For example, LENs can be used to directly classify data in an explainable manner, or to explain another black-box neural classifier. Moreover, according to the user expectations, different kinds of logic rules may be provided.

We investigate three different use-cases comparing different ways of implementing the LEN models. While most of the emphasis of this paper is on supervised classification, we also show how LEN can be leveraged in fully unsupervised settings. Additional human priors could be eventually incorporated into the learning process, in the architecture, and, following previous works, what we propose can be trivially extended to semi-supervised learning. Our work contributes to the XAI research field in the following ways: (1) It generalizes existing neural methods for solving and explaining categorical learning problems into a broad family of neural networks, i.e., the Logic Explained Networks (LENs). In particular, we extend the use of ψ networks also to directly provide interpretable classifications, and we introduce other two main instances of LENs, i.e. ReLU networks and μ networks. (2) It describes how users may interconnect LENs in the classification task under investigation, and how to express a set of preferences to get one or more customized explanations. (3) It shows how to get a wide range of logic-based explanations, and how logic formulas can be restricted in their scope, working at different levels of granularity (explaining a single sample, a subset of the available data, etc). (4) It reports experimental results using three out-of-the-box preset LENs showing how they may generalize better in terms of model accuracy than established white-box models such as decision trees on complex Boolean tasks. (5) It advertises our public implementation of LENs in a GitHub repository with an extensive documentation about LENs models, implementing different trade-offs between interpretability, explainability and accuracy.

7.2.5 Forward Approximate Solution for Linear Quadratic Tracking

Participants: Alessandro Betti, Marco Gori.

Collaborations: Michele Casoni

Keywords: Linear Quadratic Problem, Forward Approximation, Optimal Control.

In [46], we discuss an approximation strategy for solving the Linear Quadratic Tracking that is both forward and local in time. We exploit the known form of the value function along with a time reversal transformation that nicely addresses the boundary condition consistency. We provide the results of an experimental investigation with the aim of showing how the proposed solution performs with respect to the optimal solution. Finally, we also show that the proposed solution turns out to be a valid alternative to model predictive control strategies, whose computational burden is dramatically reduced.

7.2.6 A Sea of Words: An In-Depth Analysis of Anchors for Text Data

Participants: Gianluigi Lopardo, Damien Garreau, Frédéric Precioso.

Keywords: Interpretability, Explainable Artificial Intelligence, Natural Language Processing

Anchors (Ribeiro et al., 2018) is a post-hoc, rule-based interpretability method. For text data, it proposes to explain a decision by highlighting a small set of words (an anchor) such that the model to explain has similar outputs when they are present in a document. In [35], we present the first theoretical analysis of Anchors, considering that the search for the best anchor is exhaustive. After formalizing the algorithm for text classification, illustrated in Figure 10, we present explicit results on different classes of models when the preprocessing step is term frequency-inverse document frequency (TF-IDF) vectorization, including elementary if-then rules and linear classifiers. We then leverage this analysis to gain insights on the behavior of Anchors for any differentiable classifiers. For neural networks, we empirically show that the words corresponding to the highest partial derivatives of the model with respect to the input, reweighted by the inverse document frequencies, are selected by Anchors.




Figure 10: An illustration of the algorithm with evaluation function $p = \text{Prec}$. Each blue dot is an anchor, with x coordinate its length and y coordinate its value for p . Here, $\epsilon = 0.2$ and the maximal length of an anchor is $b = 10$ (the length of ζ). In the end, the anchor A such that $|A| = 3$ and $p(A) = 0.9$ is selected (red circle).

7.2.7 Interpretable Neural-Symbolic Concept Reasoning

Participants: Gabriele Ciravegna, Frédéric Precioso.

Collaborations: Pietro Barbiero, Francesco Giannini, Mateo Espinosa Zarlenga, Lucie Charlotte Magister, Alberto Tonda, Pietro Lio, Mateja Jamnik, Giuseppe Marra

Keywords: Artificial Intelligence, Machine Learning, Neural and Evolutionary Computing, FOS: Computer and information sciences

Deep learning methods are highly accurate, yet their opaque decision process prevents them from earning full human trust. Concept-based models aim to address this issue by learning tasks based on a set of human-understandable concepts. However, state-of-the-art concept-based models rely on high-dimensional concept embedding representations which lack a clear semantic meaning, thus questioning the interpretability of their decision process. To overcome this limitation, we propose, in [26, 27], the Deep Concept Reasoner (DCR), the first interpretable concept-based model that builds upon concept embeddings. In DCR, neural networks do not make task predictions directly, but they build syntactic rule structures using concept embeddings. DCR then executes these rules on meaningful concept truth degrees to provide a final interpretable and semantically-consistent prediction in a differentiable manner. Our experiments show that DCR: (i) improves up to +25% w.r.t. state-of-the-art interpretable concept-based models on challenging benchmarks (ii) discovers meaningful logic rules matching known ground truths even in the absence of concept supervision during training, and (iii), facilitates the generation of counterfactual examples providing the learnt rules as guidance.

7.3 Adaptive and robust learning

7.3.1 On the Robustness of Text Vectorizers

Participants: Damien Garreau.

Keywords: Natural Language Processing, Robustness

Collaborations: Rémi Catellier, Samuel Vaïter

A fundamental issue in machine learning is the robustness of the model with respect to changes in the input. In natural language processing, models typically contain a first embedding layer, transforming a sequence of tokens into vector representations. While the robustness with respect to changes of continuous inputs is well-understood, the situation is less clear when considering discrete changes, for instance replacing a word by another in an input sentence. Our work [32] formally proves that popular embedding schemes, such as concatenation, TF-IDF, and Paragraph Vector (a.k.a. doc2vec), exhibit robustness in the Hölder or Lipschitz sense with respect to the Hamming distance. We provide quantitative bounds for these schemes and demonstrate how the constants involved are affected by the length of the document. These findings are exemplified through a series of numerical examples.

7.3.2 Knowledge-Driven Active Learning

Participants: Gabriele Ciravegna, Alessandro Betti, Kévin Mottin, Marco Gori, Frédéric Precioso.

Keywords: Active Learning, Knowledge Representation, Deep Learning

Deep Learning (DL) methods have achieved impressive results over the last few years in fields ranging from computer vision to machine translation. Most of the research, however, focused on improving model performances, while little attention has been paid to overcome the intrinsic limits of DL algorithms. In particular, in this work [33] we will focus on the amount of data problem. Indeed, deep neural networks need large amounts of labelled data to be properly trained. With the advent of Big Data, sample collection

does not represent an issue any more. Nonetheless, the number of supervised data in some contexts is limited, and manual labelling can be expensive and time-consuming. Therefore, a common situation is the unlabelled pool scenario, where many data are available, but only some are annotated. Historically, two strategies have been devised to tackle this situation: semi-supervised learning which focus on improving feature representations by processing the unlabelled data with unsupervised techniques; active learning in which the training algorithm indicates which data should be annotated to improve the most its performances. The main assumption behind active learning strategies is that there exists a subset of samples that allows to train a model with a similar accuracy as when fed with all training data. Iteratively, the model indicates the optimal samples to be annotated from the unlabelled pool. This is generally done by ranking the unlabelled samples w.r.t. a given measure and by selecting the samples associated to the highest scores. In this paper, we propose an active learning strategy that compares the predictions over the unsupervised data with the available domain knowledge and exploits the inconsistencies as an index for selecting the data to be annotated. Domain knowledge can be generally expressed as First-Order Logic (FOL) clauses and translated into real-valued logic constraints by means of T-Norms. This formulation has been employed in the semi-supervised learning scenario to improve classifier performance by enforcing the constraints on the unsupervised data. More recently, constraints violation has been effectively used also as a metric to detect adversarial attacks. To the best of our knowledge, however, domain-knowledge (in the form of logic constraints) violation has never been used as an index in the selection process of an active learning strategy. We show that the proposed strategy is more data-efficient than the standard uncertain sample selection method, particularly in those contexts where domain-knowledge is rich. We empirically demonstrate that this is mainly due to the fact that the proposed strategy allows discovering data distributions lying far from training data, unlike uncertainty-based approaches. Neural networks, indeed, are known to be over-confident of their prediction, and they are generally unable to recognize samples lying far from the training data distribution. This issue, beyond exposing them to adversarial attacks, prevents uncertainty-based strategies from detecting these samples as points that would require an annotation. On the contrary, even though a neural network may be confident of its predictions, the interaction between the predicted classes may still offer a way to spot out-of-the-distribution samples. Finally, the Knowledge-driven Active Learning (KAL) strategy can be also employed in the object-detection context where standard uncertainty-based ones are difficult to apply.

7.3.3 Taming the Diversity of Computational Notebooks

Participants: Frédéric Precioso.

Collaborations: Yann Brault, Yassine El Amraoui, Mireille Blay-Fornarino, Philippe Collet, Florent Jaillet.
Keywords: Computing methodologies, Machine learning, Software development techniques, Reusability, Software product lines

In many applications of Computational Science and especially Data Science, notebooks are the cornerstone of knowledge and experiment sharing. Their diversity is multiple (problem addressed, input data, algorithm used, overall quality) and is not made explicit at all. As they are heavily reused through a clone-and-own approach, the tailoring process from an existing notebook to a specific problem is cumbersome, error-prone, and particularly uncertain. In [31], we propose a tooling approach that captures the different dimensions of variability in computational notebooks. It allows one to seek an existing notebook that suits her requirements, or to generate most parts of a new one.

7.3.4 Toward Novel Optimizers: A Moreau-Yosida View of Gradient-Based Learning

Participants: Alessandro Betti, Gabriele Ciravegna, Marco Gori, Kevin Mottin, Frédéric Precioso.

Collaborations: Stefano Melacci.

Keywords: Artificial Intelligence, Machine Learning, Optimization

Best Paper Award at AIXIA 2023 – Advances in Artificial Intelligence: XXIInd International Conference of the Italian Association for Artificial Intelligence, Rome, Italy!

Machine Learning (ML) strongly relies on optimization procedures that are based on gradient descent. Several gradient-based update schemes have been proposed in the scientific literature, especially in the context of neural networks, that have become common optimizers in software libraries for ML. In [28], we re-frame gradient-based update strategies under the unifying lens of a Moreau-Yosida (MY) approximation of the loss function. By means of a first-order Taylor expansion, we make the MY approximation concretely exploitable to generalize the model update. In turn, this makes it easy to evaluate and compare the regularization properties that underlie the most common optimizers, such as gradient descent with momentum, ADAGRAD, RMSprop, and ADAM. The MY-based unifying view opens to the possibility of designing novel update schemes with customizable regularization properties. As case-study we propose to use the network outputs to deform the notion of closeness in the parameter space.

7.4 Learning with heterogeneous and corrupted data

7.4.1 Mind the map! Accounting for existing map information when estimating online HDMaps from sensor data

Participants: Rémy Sun, Li Yang, Diane Lingrand, Frédéric Precioso.

Keywords: Autonomous Driving, HDMaps, Online HDMap estimation

Collaborations: ANR Project MultiTrans

Online High Definition Map (HDMap) estimation from sensors offers a low-cost alternative to manually acquired HDMaps. As such, it promises to lighten costs for already HDMap-reliant Autonomous Driving systems, and potentially even spread their use to new systems. We propose in [55] to improve online HDMap estimation by accounting for already existing maps. We identify 3 reasonable types of useful existing maps (minimalist, noisy, and outdated). We also introduce MapEX (see Fig. 11), a novel online HDMap estimation framework that accounts for existing maps. MapEX achieves this by encoding map elements into query tokens and by refining the matching algorithm used to train classic query based map estimation models. We demonstrate that MapEX brings significant improvements on the nuScenes dataset. For instance, given noisy maps, MapEX improves by 38% over the MapTRv2 detector it is based on and by 16% over the current SOTA (state-of-the-art).

7.4.2 Exploring the Road Graph in Trajectory Forecasting for Autonomous Driving

Participants: Rémy Sun, Diane Lingrand, Frédéric Precioso.

Keywords: Autonomous Driving, HDMaps, Trajectory Forecasting, Graphs

Collaborations: ANR Project MultiTrans

As Deep Learning tackles complex tasks like trajectory forecasting in autonomous vehicles, a number of new challenges emerge. In particular, autonomous driving requires accounting for vast a priori knowledge in the form of HDMaps. Graph representations have emerged as the most convenient representation for this complex information. Nevertheless, little work has gone into studying how this road graph should be constructed and how it influences forecasting solutions. In [42], we explore the impact of spatial resolution, the graph's relation to trajectory outputs and how knowledge can be embedded into the graph. To this end, we propose thorough experiments for 2 graph-based frameworks (PGP, LAformer) over the nuScenes dataset, with additional experiments on the LaneGCN framework and Argoverse 1.1 dataset.

7.4.3 Machinery Anomaly Detection using artificial neural networks and signature feature extraction

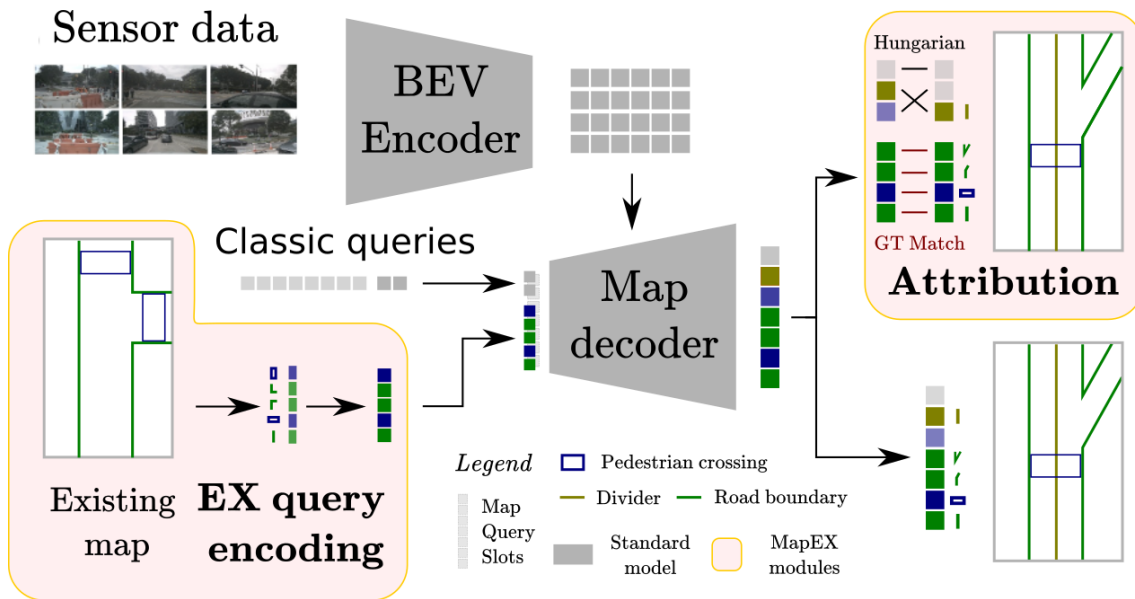


Figure 11: Overview of our MapEX method (see Sec. 7.4.1). We add two modules (EX query encoding, Attribution) to the standard query based map estimation pipeline (in gray on the figure). Map elements are encoded into EX queries, then decoded with a standard decoder.

Participants: Mansour Zoubeirou A Mayaki, Michel Riveill.

Keywords: Fault diagnosis, Anomaly detection, Predictive maintenance, Concept drift detection, Data streams, Signature, Machine learning

Machine learning models are increasingly being used in predictive maintenance. However, due to the complexity of vibration and audio signals used in fault diagnosis, some pre-processing is required before feeding them into the machine learning algorithm. Fast Fourier Transform (FFT) and the Hilbert transform (HT) envelope spectrum are mostly used in the literature for pre-processing. However, these frequency domain transforms are not very effective when applied to rotating systems (e.g. bearings) fault detection. In fact, in these applications, the fault signal patterns are usually very weak relative to background noise and other interference in the early damage stage. In this paper [37], we propose to use signature coefficients to feed machine learning models for fault detection. Our experimental results show that this method outperforms most state-of-the-art methods on fault diagnosis datasets. For example, in the Case Western Reserve University (CWRU) data set, the accuracy of the proposed method ranges from 96.59 % to 100%. Moreover, the results show that this method is particularly well suited for high-dimensional time series. The results also show that compared to Fast Fourier Transform (FFT), the signature method requires fewer data points to detect failure. This means that in a situation where the two methods have similar performances, the signature method detects failure faster than FFT.

The architecture and dataflow of the proposed approach. is shown in Figure 12.

7.4.4 Are labels informative in semi-supervised learning? Estimating and leveraging the missing-data mechanism.

Participants: Aude Sportisse, Hugo Schmutz, Charles Bouveyron, Pierre-Alexandre Mattei.

Keywords: missing data, semi-supervised learning, deep learning

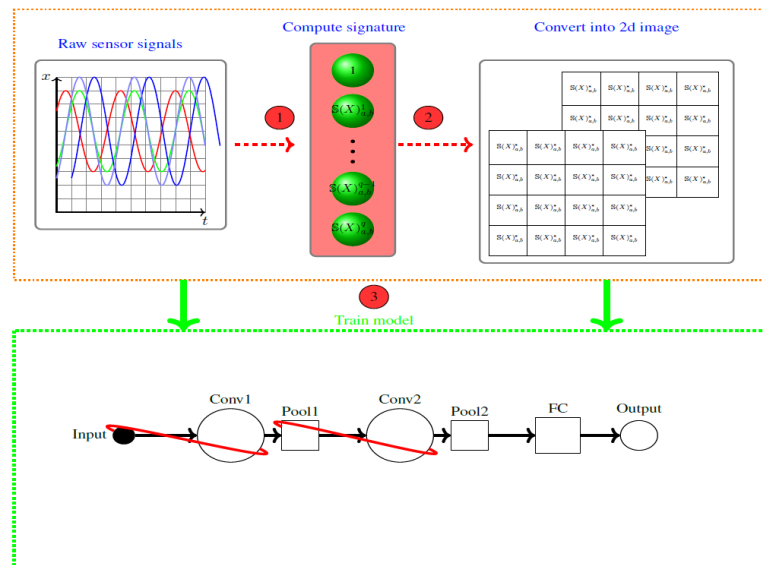


Figure 12: Architecture and dataflow of the proposed approach. For more details, see Section 7.4.3.

Collaborations: Olivier Humbert (PU-PH, Centre Antoine Lacassagne)

Semi-supervised learning (SSL) is a powerful technique for leveraging unlabeled data to improve machine learning models, but it can be affected by the presence of “informative” labels, which occur when some classes are more likely to be labeled than others. In the missing data literature, such labels are called missing not at random. In this paper, we propose a novel approach to address this issue by estimating the missing-data mechanism and using inverse propensity weighting to debias any SSL algorithm, including those using data augmentation. We also propose a likelihood ratio test to assess whether or not labels are indeed informative. Finally, we demonstrate the performance of the proposed methods on different datasets, in particular on two medical datasets for which we design pseudo-realistic missing data scenarios. This work has been accepted as an oral to ICML 2023 (International Conference on Machine Learning) [41].

7.4.5 AnoRand: Deep Learning-Based Semi-Supervised Anomaly Detection with Synthetic Labels

Participants: Mansour Zoubeirou A Mayaki, Michel Riveill.

Keywords: Medicare fraud, Anomaly detection, Deep learning, Auto encoder, Machine learning

Anomaly detection, or more generally outlier detection, is one of the most popular and challenging topics in theoretical and applied machine learning. The main challenge is that in general we have access to very few labeled data or no labels at all. In this paper [56], we present a new semi-supervised anomaly detection method called **AnoRand** by combining a deep learning architecture with random synthetic label generation. The proposed architecture has two building blocks: (1) a noise detection (ND) block composed of feed forward perceptron and (2) an autoencoder (AE) block. The main idea of this new architecture is to learn one class (e.g. the majority class in case of anomaly detection) as well as possible by taking advantage of the ability of auto encoders to represent data in a latent space and the ability of Feed Forward Perceptron (FFP) to learn one class when the data is highly imbalanced. First, we create synthetic anomalies by randomly disturbing a few samples (e.g., 2%) from the training set. Second, we use the normal and synthetic samples as input to our model. We compared the performance of the proposed method to 17 state-of-the-art unsupervised anomaly detection methods on synthetic datasets and 57 real-world datasets. Our results show that this new method generally outperforms most of the state-of-the-art methods and has the best performance (AUC ROC and AUC PR) on the vast majority of

reference datasets.

The full network architecture is described in Figure 13.

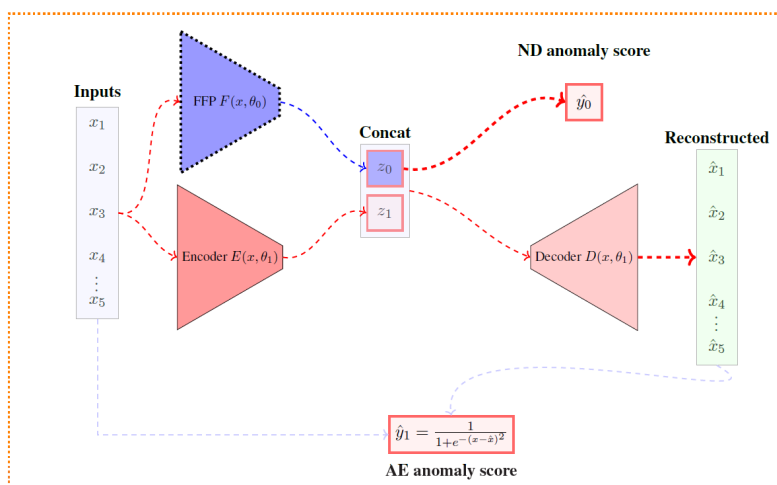


Figure 13: Overall architecture of AnoRand. For more details, see 7.4.5.

7.4.6 Fed-MIWAE: Federated Imputation of Incomplete Data via Deep Generative Models

Participants: Aude Sportisse, Pierre-Alexandre Mattei.

Keywords: missing data, federated learning, federated pre-processing, variational autoencoders, deep learning

Collaborations: Irene Balelli (Epione, Inria Center at Université Côte d’Azur), Marco Lorenzi (Epione, Inria Center at Université Côte d’Azur), Francesco Cremonesi (Epione, Inria Center at Université Côte d’Azur)

Federated learning allows for the training of machine learning models on multiple decentralized local datasets without requiring explicit data exchange. However, data pre-processing, including strategies for handling missing data, remains a major bottleneck in real-world federated learning deployment, and is typically performed locally. This approach may be biased, since the subpopulations locally observed at each center may not be representative of the overall one. To address this issue, this paper first proposes a more consistent approach to data standardization through a federated model. Additionally, we propose Fed- MIWAE, a federated version of the state-of-the-art imputation method MIWAE, a deep latent variable model for missing data imputation based on variational autoencoders. MIWAE has the great advantage of being easily trainable with classical federated aggregators. Furthermore, it is able to deal with MAR (Missing At Random) data, a more challenging missing-data mechanism than MCAR (Missing Completely At Random), where the missingness of a variable can depend on the observed ones. We evaluate our method on multi-modal medical imaging data and clinical scores from a simulated federated scenario with the ADNI dataset. We compare Fed-MIWAE with respect to classical imputation methods, either performed locally or in a centralized fashion. Fed-MIWAE allows to achieve imputation accuracy comparable with the best centralized method, even when local data distributions are highly heterogeneous. In addition, thanks to the variational nature of Fed-MIWAE, our method is designed to perform multiple imputation, allowing for the quantification of the imputation uncertainty in the federated scenario. The working document is available in HAL [45].

7.4.7 Don’t fear the unlabelled: safe deep semi-supervised learning via simple debiasing

Participants: Hugo Schmutz, Pierre-Alexandre Mattei.

Collaborations: Olivier Humbert

Keywords: Semi-supervised learning, safeness, debiasing, control variates, asymptotic statistics, proper scoring rules

Semi-supervised learning (SSL) provides an effective means of leveraging unlabelled data to improve a model’s performance. Even though the domain has received a considerable amount of attention in the past years, most methods present the common drawback of lacking theoretical guarantees. In [39], our starting point is to notice that the estimate of the risk that most discriminative SSL methods minimize is biased, even asymptotically. This bias impedes the use of standard statistical learning theory and can hurt empirical performance. We propose a simple way of removing the bias, as seen in the blue part of the following equation.

$$\hat{\mathcal{R}}_{DeSSL}(\theta) = \frac{1}{n_l} \sum_{i=1}^{n_l} L(\theta; x_i, y_i) + \frac{\lambda}{n} \sum_{i=1}^n H(\theta; x_i) - \frac{\lambda}{n_l} \sum_{i=1}^{n_l} H(\theta; x_i) \quad (1)$$

Our debiasing approach is straightforward to implement and applicable to most deep SSL methods. We provide simple theoretical guarantees on the trustworthiness of these modified methods, without having to rely on the strong assumptions on the data distribution that SSL theory usually requires. In particular, we provide generalization error bounds for the proposed methods by deriving Rademacher complexity. We evaluate debiased versions of different existing SSL methods, such as the Pseudo-label method and Fixmatch, and show that debiasing can compete with classic deep SSL techniques in various settings by providing better calibrated models. For instance, in Figure 14, we show that the classic PseudoLabel method fails to learn correctly the minority classes in an unbalanced dataset setting. Additionally, we provide a theoretical explanation of the intuition of the popular SSL methods.




Figure 14: Class accuracies (without the majority class) on DermaMNIST trained with $n_l = 1000$ labelled data on five folds. (Left) CompleteCase (B-Acc: $26.88 \pm 2.26\%$); (Middle) PseudoLabel (B-Acc: $22.03 \pm 1.45\%$); (Right) DePseudoLabel (B-Acc: $28.84 \pm 1.02\%$), with 95% CI.

7.4.8 The graph embedded topic model

Participants: Dingge Liang, Marco Corneli, Charles Bouveyron, Pierre Latouche.

Keywords: Graph neural networks, Topic modeling, Deep latent variable models, Clustering, Network analysis

Most of existing graph neural networks (GNNs) developed for the prevalent text-rich networks typically treat texts as node attributes. This kind of approach unavoidably results in the loss of important semantic structures and restricts the representational power of GNNs. In [18], we introduce a document similarity-based graph convolutional network (DS-GCN) encoder to combine graph convolutional networks and embedded topic models for text-rich network representation. Then, a latent position-based decoder is used to reconstruct the graph while preserving the network topology. Similarly, the document matrix

is rebuilt using a decoder that takes both topic and word embeddings into account. By including a cluster membership variable for each node in the network, we thus develop an end-to-end clustering technique relying on a new deep probabilistic model called the graph embedded topic model (GETM), see Figure 15. Numerical experiments on three simulated scenarios emphasize the ability of GETM in fusing the graph topology structure and the document embeddings, and highlight its node clustering performance. Moreover, an application on the Cora-enrich citation network is conducted to demonstrate the effectiveness and interest of GETM in practice.

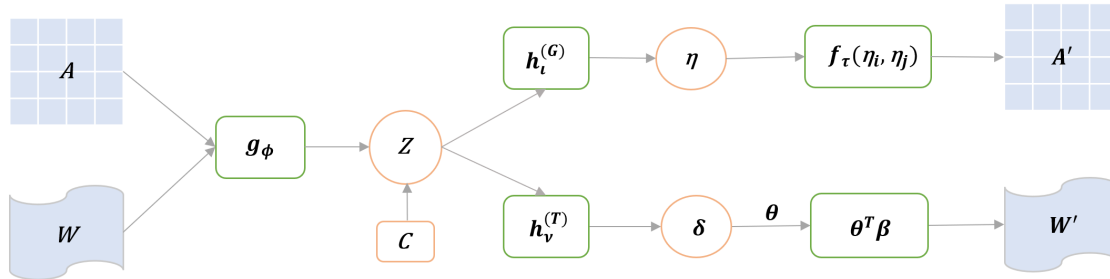


Figure 15: Model architecture of GETM. For more details, see Section 7.4.8.

7.4.9 Model-based clustering with Missing Not At Random Data

Participants: Aude Sportisse.

Keywords: model-based clustering, generative models, missing data

Collaborations: Christophe Biernacki (Inria Lille), Claire Boyer (Sorbonne Université), Julie Josse (Inria Montpellier) Matthieu Marbac (Ensaï Rennes)

Model-based unsupervised learning, as any learning task, stalls as soon as missing data occurs. This is even more true when the missing data are informative, or said missing not at random (MNAR). In this paper, we propose model-based clustering algorithms designed to handle very general types of missing data, including MNAR data. To do so, we introduce a mixture model for different types of data (continuous, count, categorical and mixed) to jointly model the data distribution and the MNAR mechanism, remaining vigilant to the relative degrees of freedom of each. Several MNAR models are discussed, for which the cause of the missingness can depend on both the values of the missing variable themselves and on the class membership. However, we focus on a specific MNAR model, called MNARz, for which the missingness only depends on the class membership. We first underline its ease of estimation, by showing that the statistical inference can be carried out on the data matrix concatenated with the missing mask considering finally a standard MAR mechanism. Consequently, we propose to perform clustering using the Expectation Maximization algorithm, specially developed for this simplified reinterpretation. Finally, we assess the numerical performances of the proposed methods on synthetic data and on the real medical registry TraumaBase as well. This work is available in HAL [54], as well as its accompanying note [53].

7.4.10 Continuous Latent Position Models for Instantaneous Interactions

Participants: Marco Corneli.

Keywords: Latent Position Models, Dynamic Networks, Non-Homogeneous Poisson Process, Spatial Embeddings, Statistical Network Analysis

Collaborations: Riccardo Rastelli (UCD, Dublin)

In [22] we create a framework to analyze the timing and frequency of instantaneous interactions between pairs of entities. This type of interaction data is especially common nowadays, and easily available. Examples of instantaneous interactions include email networks, phone call networks and some common types of technological and transportation networks. Our framework relies on a novel extension of the latent position network model: we assume that the entities are embedded in a latent Euclidean space, and that they move along individual trajectories which are continuous over time. These trajectories are used to characterize the timing and frequency of the pairwise interactions. We discuss an inferential framework where we estimate the individual trajectories from the observed interaction data, and propose applications on artificial and real data. Figure 16 shows the evolving latent positions of a dynamic graph.

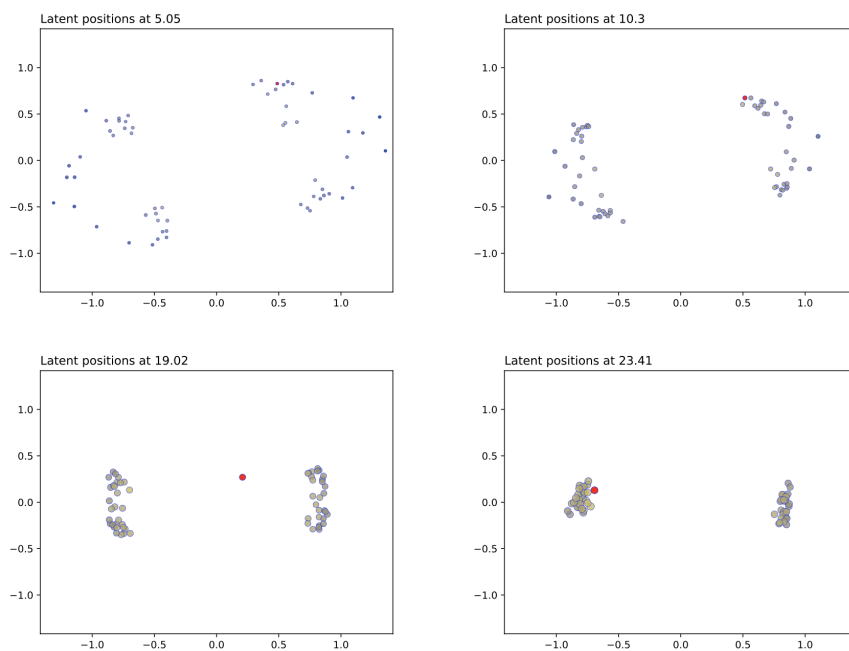


Figure 16: Snapshots of the evolving latent positions of a dynamic graph. Two communities emerge and a node (red) migrates between them. For more details, see Section 7.4.10.

7.4.11 DeepWILD: wildlife identification, localization and population estimation from camera trap videos in the Parc National du Mercantour

Participants: Charles Bouveyron, Frédéric Precioso.

Keywords: image analysis,

Collaborations: Fanny Simoes (Institut 3IA Côte d’Azur), Nathalie Siefert (Parc National du Mercantour)

Videos and images from camera traps are more and more used by ecologists to estimate the population of species on a territory. Most of the time, it is a laborious work since the experts analyze manually all this data. It takes also a lot of time to filter these videos when there are plenty of empty videos or with humans presence. Fortunately, deep learning algorithms for object detection could help ecologists to identify multiple relevant species on their data and to estimate their population. In [23], we propose to go even further by using object detection model to detect, classify and count species on camera traps videos (see Figure 17). We developed a 3-parts process to analyze camera trap videos. At the first stage, after splitting videos into images, we annotate images by associating bounding boxes to each label thanks to MegaDetector algorithm. Then, we extend MegaDetector based on Faster R-CNN architecture

with backbone Inception-ResNet-v2 in order to not only detect the 13 species considered but also to classify them. Finally, we define a method to count species based on maximum number of bounding boxes detected, it included only detection results and an evolved version of this method included both, detection and classification results. The results obtained during the evaluation of our model on the test dataset are: (i) 73,92% mAP for classification, (ii) 96,88% mAP for detection with a ratio Intersection-Over-Union (IoU) of 0.5 (overlapping ratio between groundtruth bounding box and the detected one), and (iii) 89,24% mAP for detection at IoU=0.75. Big species highly represented, like human, have highest values of mAP around 81% whereas species less represented in the train dataset, such as dog, have lowest values of mAP around 66%. As regards to our method of counting, we predicted a count either exact or ± 1 unit for 87% with detection results and 48% with detection and classification results of our video sample. Our model is also able to detect empty videos. To the best of our knowledge, this is the first study in France about the use of object detection model on a French national park to locate, identify and estimate the population of species from camera trap videos.

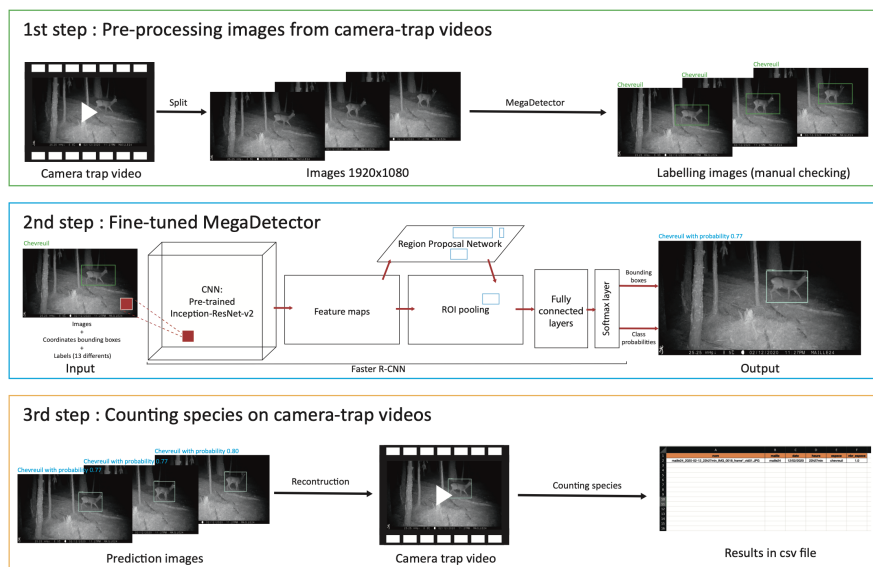


Figure 17: The DeepWILD process for the wildlife identification, localization and population estimation from camera trap videos in the Parc National du Mercantour

7.4.12 Another Point of View on Visual Speech Recognition

Participants: Baptiste Pouthier, Charles Bouveyron, Frédéric Precioso.

Keywords: visual speech recognition, graph convolutional network, point cloud definition

Collaborations: Laurent Pilati, Giacomo Valenti (NXP)

Standard Visual Speech Recognition (VSR) systems directly process images as input features without any a-priori link between raw pixel data and facial traits. Pixel information is smartly sieved when facial landmarks are extracted from pictures and repurposed as graph nodes. Their evolution through time is thus modeled by a Graph Convolutional Network. However, with graph-based VSR being in its infancy, the selection of points and their correlation are still ill-defined and often bound to a-prioristic knowledge and handcrafted techniques. In [38], we investigate the graph approach for VSR and its ability to learn the correlation between points beyond the mouth region (see Figure 18). We also study the different contributions that each facial region brings to the system accuracy, proving that more scattered but better connected graphs can be both computationally light and accurate.

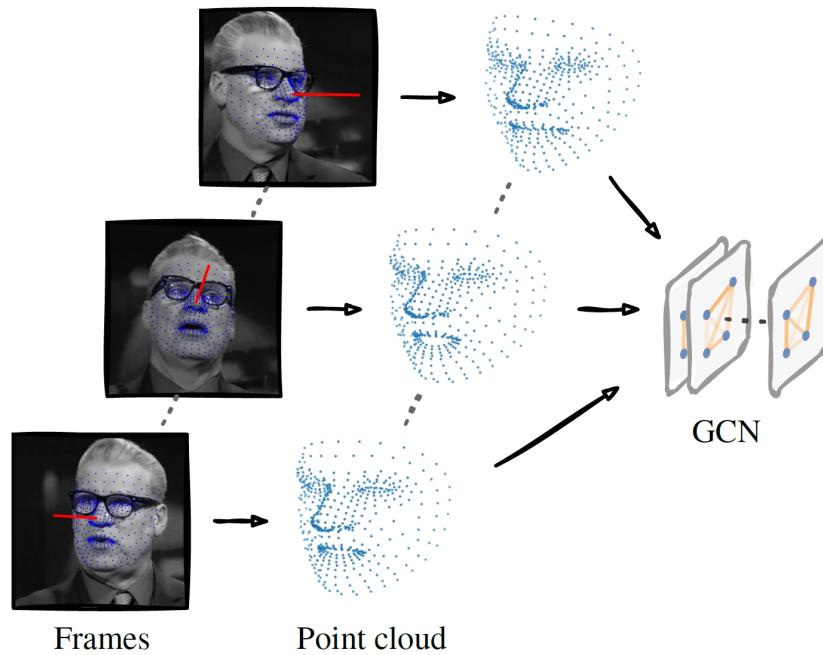


Figure 18: Graph-based VSR system pipeline: landmark detection on video frames is applied to extract the point cloud. Following head pose estimation (red line) and normalization, the point cloud is processed by a GCN model outputting speech transcription.

8 Bilateral contracts and grants with industry

The team is particularly active in the development of research contracts with private companies. The following contracts were active during 2022:

- **Pulse Audition** This contract was the fruit of the "start it up" program of the 3IA Côte d'Azur. The goal is to work on semi-supervised learning for hearing glasses. A research engineer (Léonie Borne) was recruited via the "start it up" program. Amount: 15 000€.

Participants: Pierre-Alexandre Mattei, Léonie Borne.

- **NXP:** This collaboration contract is a France Relance contract. Drift detection and predictive maintenance. Amount: 45 000€.

Participants: Mansour Mayaki Zoubeirou, Michel Riveill.

- **Orange:** it is a CIFRE build upon the PhD of Gatien Caillet on decentralized and efficient federated AutoML learning for heterogeneous embedded devices. External participants: Tamara Tosic (Orange), Frédéric Guyard (Orange). Amount: 30 000€.

Participants: Vincent Vandewalle.

- **Naval Group:** The goal of this project is the development of an open-source Python library for semi-supervised learning, via the hiring of a research engineer, Lucas Boiteau. External participants: Alexandre Gense, Quentin Oliveau (Naval Group). Amount: 125 000€.

Participants: Pierre-Alexandre Mattei, Lucas Boiteau, Hugo Schmutz, Aude Sportisse.

- **Orange:** it is a CIFRE contract built upon the PhD of Hugo Miralles on Distributed device-embedded classification and prediction in near-to-real time. External participants: Tamara Tosic (Orange), Thierry Nagellen (Orange). Amount: 45 000€.

Participants: Michel Riveill, Hugo Miralles.

- **NXP:** This collaboration contract is a CIFRE contract built upon the PhD of Baptiste Pouthier on Deep Learning and Statistical Learning on audio-visual data for embedded systems. Participants: Frederic Precioso, Charles Bouveyron, Baptiste Pouthier. External participants: Laurent Pilati (NXP). Amount: 45 000€.

Participants: Frédéric Precioso, Charles Bouveyron, Baptiste Pouthier.

- **Instant System:** This collaboration contract is a France Relance contract. The objective is to design new recommendation systems based on deep learning for multimodal public transport recommendations (e.g. combining on a same trip: bike, bus, e-scooter, metro, then bike again). Amount: 45 000€.

Participants: Frédéric Precioso, Michel Riveill, Amosse Edouard.

- **EDF:** In this project, we developed model-based clustering and co-clustering methods to summarize massive and multivariate functional data of electricity consumption. The data are coming from Linky meters, enriched by meteorological and spatial data. The developed algorithms were released as open source R packages. External participants: F. Simoes, J. Jacques. Amount: 50 000€.

Participants: Charles Bouveyron.

9 Partnerships and cooperations

9.1 International initiatives

The Maasai team has informal relationships with the following international teams:

- Department of Statistics of the University of Washington, Seattle (USA) through collaborations with Elena Erosheva and Adrian Raftery,
- SAILAB team at Università di Siena, Siena (Italy) through collaborations with Marco Gori,
- School of Mathematics and Statistics, University College Dublin (Ireland) through the collaborations with Brendan Murphy, Riccardo Rastelli and Michael Fop,

- Department of Computer Science, University of Tübingen (Germany) through the collaboration with Ulrike von Luxburg,
- Université Laval, Québec (Canada) through the Research Program DEEL (DEpendable and EXplainable Learning) with François Laviolette and Christian Gagné, and through a FFCR funding with Arnaud Droit (including the planned supervision of two PhD students in 2022),
- DTU Compute, Technical University of Denmark, Copenhagen (Denmark), through collaborations with Jes Frellsen and his team (including the co-supervision of a PhD student in Denmark: Hugo Sénétaire).

9.1.1 Participation in other International Programs

DEpendable EXplainable Learning Program (DEEL), Québec, Canada

Participants: Frédéric Precioso.

Collaborations: François Laviolette (Prof. U. Laval), Christian Gagné (Prof. U. Laval)

The DEEL Project involves academic and industrial partners in the development of dependable, robust, explainable and certifiable artificial intelligence technological bricks applied to critical systems. We are involved in the Workpackage Robustness and the Workpackage Interpretability, in the co-supervision of several PhD thesis, Post-docs, and Master internships.

CHU Québec–Laval University Research Centre, Québec, Canada

Participants: Frédéric Precioso, Pierre-Alexandre Mattei, Louis Ohl.

Collaborations: Arnaud Droit (Prof. U. Laval), Mickael Leclercq (Chercheur U. Laval), Khawla Seddiki (doctorante, U. Laval)

This collaboration framework covers several research projects: one project is related to the PhD thesis of Khawla Seddiki who works on Machine Learning/Deep Learning methods for classification and analysis of mass spectrometry data; another project is related to the France Canada Research Fund (FCRF) which provides the PhD funding of Louis Ohl, co-supervised by all the collaborators. This project investigates Machine Learning solutions for Aortic Stenosis (AS) diagnosis.

SAILAB: Lifelong learning in computer vision

Participants: Frédéric Precioso.

Keywords: computer vision, lifelong learning, focus of attention in vision, virtual video environments.

Collaborations: Lucile Sassatelli, Dario (Universität Erlangen-Nürnberg), Alessandro Betti (UNISI), Stefano Melacci (UNISI), Matteo Tiezzi (UNISI), Enrico Meloni (UNISI), Simone Marullo (UNISI).

This collaboration concerns the current hot machine learning topics of Lifelong Learning, “on developing versatile systems that accumulate and refine their knowledge over time”), or continuous learning which targets tackling catastrophic forgetting via model adaptation. The most important expectations of this research is that of achieving object recognition visual skills by a little supervision, thus overcoming the need for the expensive accumulation of huge labelled image databases.

9.2 European initiatives

9.2.1 FP7 & H2020 Projects

Maasai is one of the 3IA-UCA research teams of **AI4Media**, one of the 4 ICT-48 Center of Excellence in Artificial Intelligence which has started in September 2020. There are 30 partners (Universities and companies), and 3IA-UCA received about 325k€.

9.3 National initiatives

Institut 3IA Côte d'Azur

Following the call of President Macron to found several national institutes in AI, we presented in front of an international jury our project for the Institut 3IA Côte d'Azur in April 2019. The project was selected for funding (50 M€ for the first 4 years, including 16 M€ from the PIA program) and started in september 2019. Charles Bouveyron and Marco Gori are two of the 29 3IA chairs which were selected *ab initio* by the international jury and Pierre-Alexandre Mattei was awarded a 3IA chair in 2021. Charles Bouveyron is also the Director of the institute since January 2021, after being the Deputy Scientific Director on 2019-2020. The research of the institute is organized around 4 thematic axes: Core elements of AI, Computational Medicine, AI for Biology and Smart territories. The Maasai reserch team is totally aligned with the first axis of the Institut 3IA Côte d'Azur and also contributes to the 3 other axes through applied collaborations. The team has 7 Ph.D. students and postdocs who are directly funded by the institute.

Web site: 3ia.univ-cotedazur.eu

Projet ANR MultiTrans

In MultiTrans project, we propose to tackle autonomous driving algorithms development and deployment jointly. The idea is to enable data, experience and knowledge to be transferable across the different systems (simulation, robotic models, and real-word cars), thus potentially accelerating the rate at which an embedded intelligent system can gradually learn to operate at each deployment stage. Existing autonomous vehicles are able to learn how to react and operate in known domains autonomously but research is needed to help these systems during the perception stage, allowing them to be operational and safer in a wider range of situations. MultiTrans proposes to address this issue by developing an intermediate environment that allows to deploy algorithms in a physical world model, by re-creating more realistic use cases that would contribute to a better and faster transfer of perception algorithms to and from a real autonomous vehicle test-bed and between multiple domains.

Web site: anr-multitrans.github.io

9.4 Regional initiatives

Parc National du Mercantour

Participants: Frédéric Precioso, Charles Bouveyron.

Keywords: Deep learning, image recognition,

Collaborators: Fanny Simoes (3IA Tech pool), Nathalie Siefert and Stéphane Combeau (Parc National du Mercantour)

The team started in 2021 a collaboration with the Parc National du Mercantour to exploit the camera-traps installed in the Park to monitor and conserve wild species. We developed, in collaboration with the engineer team of Institut 3IA Côte d'Azur, an AI pipeline allowing to automatically detect, classify and count specific endangered wild species in camera-trap videos. A demonstrator of the methodology has been presented to the general public at *Le Fête des Sciences* in Antibes in October 2021.

Centre de pharmacovigilance, CHU Nice

Participants: Charles Bouveyron, Marco Corneli, Giulia Marchello, Michel Riveill, Xuchun Zhang

Keywords: Pharmacovigilance, co-clustering, count data, text data

Collaborateurs: Milou-Daniel Drici, Audrey Freysse, Fanny Serena Romani

The team works very closely with the Regional Pharmacovigilance Center of the University Hospital Center of Nice (CHU) through several projects. The first project concerns the construction of a dashboard to classify spontaneous patient and professional reports, but above all to report temporal breaks. To this end, we are studying the use of dynamic co-classification techniques to both detect significant ADR patterns and identify temporal breaks in the dynamics of the phenomenon. The second project focuses on the analysis of medical reports in order to extract, when present, the adverse events for characterization. After studying a supervised approach, we are studying techniques requiring fewer annotations.

Interpretability for automated decision services

Participants: Frédéric Precioso, Damien Garreau.

Keywords: interpretability, deep learning

Collaborators: Greger Ottosson (IBM)

Businesses rely more and more frequently on machine learning to make automated decisions. In addition to the complexity of these models, a decision is rarely by using only one model. Instead, the crude reality of business decision services is that of a jungle of models, each predicting key quantities for the problem at hand, that are then agglomerated to produce the final decision, for instance by a decision tree. In collaboration with IBM, we want to provide principled methods to obtain interpretability of these automated decision processes.

10 Dissemination

10.1 Promoting scientific activities

10.1.1 Scientific events: organisation

- The 2nd Nice Workshop on Interpretability, organized by Damien Garreau, Frédéric Precioso and Gianluigi Lopardo. The workshop aims to create links between researchers working on interpretability of machine learning models, in a broad sense. With the objective of animating fruitful discussions and facilitating valuable knowledge sharing, on topics such as Logic-Based Explainability in Machine Learning, Consistent Sufficient Explanations and Minimal Local Rules for explaining regression and classification models, On the Trade-off between Actionable Explanations and the Right to be Forgotten, Explainability of a Model under stress, Learning interpretable scoring rules...The workshop took place on November 30 - December 1, 2023 in Nice, and counted 6 senior research talks, and 11 young research talks, with about 40 participants. Web : <https://sites.google.com/view/nwi2023/home>
- Statlearn 2023: the workshop Statlearn is a scientific workshop held every year since 2010, which focuses on current and upcoming trends in Statistical Learning. Statlearn is a scientific event of the French Society of Statistics (SFdS). Conferences and tutorials are organized alternatively every other year. The 2023 edition was the 12th edition of the Statlearn series and welcomed about 50 participants in Montpellier, 5-7 April 2023. The Statlearn conference was founded by Charles Bouveyron in 2010. Since 2019, Marco Corneli and Pierre-Alexandre Mattei are members of the scientific committee of the conference. Web : <https://statlearn.sciencesconf.org>
- Pierre-Alexandre Mattei co-founded and co-organized the Generative Models and Uncertainty quantification workshop (GenU), held in Copenhagen (September 20-21, 2023). More details at <https://genu.ai/2023/>.
- GenU 2023: Pierre-Alexandre Mattei co-founded in 2019 the on Generative Models and Uncertainty Quantification (GenU) workshop. This small-scale workshop has been held physically in Copenhagen in the Fall. The 2023 edition was on September 20-21, 2023 (Web: <https://genu.ai/2023/>).
- SophIA Summit: AI conference that brings together researchers and companies doing AI, held every Fall in Sophia Antipolis. C. Bouveyron was a member of the scientific committee in 2020. P.-A. Mattei was a member of the scientific committee in 2022 and 2023. Web: <https://univ-cotedazur.eu/events/sophia-summit>.

10.1.2 Scientific publishing

Member of the editorial boards

- Charles Bouveyron is Associate Editor for the Annals of Applied Statistics since 2016.

Reviewer - reviewing activities All permanent members of the team are serving as reviewers for the most important journals and conferences in statistical and machine learning, such as (non exhaustive list):

- International journals:
 - Annals of Applied Statistics,
 - Statistics and Computing,
 - Journal of the Royal Statistical Society, Series C,
 - Journal of Computational and Graphical Statistics,
 - Journal of Machine Learning Research
 - Transactions on Machine Learning Research
- International conferences:
 - Neural Information Processing Systems (Neurips),
 - International Conference on Machine Learning (ICML),
 - International Conference on Learning Representations (ICLR),
 - International Joint Conference on Artificial Intelligence (IJCAI),
 - International Conference on Artificial Intelligence and Statistics (AISTATS),
 - International Conference on Computer Vision and Pattern Recognition

10.1.3 Leadership within the scientific community

- Charles Bouveyron is the Director of the Institut 3IA Côte d’Azur since January 2021 and of the EFELIA Côte d’Azur education program since September 2022.
- Vincent Vandewalle is the Deputy Scientific director of the EFELIA Côte d’Azur education program since September 2022.

10.1.4 Scientific expertise

- Frédéric Precioso was the Scientific Responsible and Program Officer for AI at the French Research Agency (ANR) from September 2019 to August 2023. He was thus in charge of all the programs related to the National Plan IA, and of the new French Priority Equipment and Research Programme (PEPR) on AI, French Priority Equipment and Research Programme (PEPR) on Digital Health, and Programs for Platforms in AI (DeepGreen for embedded AI, and Platform DATA for open source AI libraries, Interoperability, AI Cloud).
- Charles Bouveyron is member of the Scientific Orientation Council of Centre Antoine Lacassagne, Unicancer center of Nice.

10.2 Teaching - Supervision - Juries

C. Bouveyron, M. Riveill and V. Vandewalle are full professors, D. Garreau is assistant-professor at Université Côte d’Azur and therefore handle usual teaching duties at the university. F. Precioso is full professor at Université Côte d’Azur but he is detached to ANR for 60% of his time, his teaching duties are thus 40% of standard ones. M. Corneli and P.-A. Mattei are also teaching around 60h per year at Université Côte d’Azur. P.-A. Mattei is also teaching a graphical models course at the MVA masters from ENS Paris Saclay. M. Corneli has been hired in September 2022 on a “Chaire de Professeur Junior” on AI for Archeology and Historical Sciences.

M. Riveill is the current director of the Master of Science “Data Sciences and Artificial Intelligence” at Université Côte d’Azur, since September 2020. C. Bouveyron was the founder and first responsible (Sept. 2018 - Aug. 2020) of that MSc.

Since September 2022, C. Bouveyron and V. Vandewalle are respectively the Director and Deputy Scientific Director of the EFELIA Côte d'Azur program (<https://univ-cotedazur.fr/efelia-cote-dazur>), funded by the French national plan "France 2030", through the "Compétences et Métiers d'Avenir" initiative (8M€ for 5 years). This program aims at enlarging the teaching capacities in AI of the Institut 3IA Côte d'Azur and developing new education programs for specialists and non-specialists.

All members of the team are also actively involved in the supervision of postdocs, Ph.D. students, interns and participate frequently to Ph.D. and HDR defenses. They are also frequently part of juries for the recruitment of research scientists, assistant-professors or professors.

10.3 Popularization

- F. Precioso, C. Bouveyron, F. Simoes and J. Torres Sanchez have developed a demonstrator for general public on the recognition and monitoring of wild species in the French National Park of Mercantour. This demonstrator has been exhibited during the "Fête des Sciences" in Antibes in October 2023. Web: <https://3ia-demos.inria.fr/mercantour/>
- F. Precioso has developed an experimental platform both for research projects and scientific mediation on the topic of autonomous cars. This platform is currently installed in the "Maison de l'Intelligence Artificielle" where high school students have already experimented coding autonomous remote control cars. Web: <https://maison-intelligence-artificielle.com/experimenter-projets-ia/>
- C. Bouveyron, F. Simoes and S. Bottini have developed an interactive software allowing to visualize the relationships between pollution and a health disease (dispnea) in the Région Sud. This platform is currently installed at the "Maison de l'Intelligence Artificielle".

10.4 Interventions

- S. Petiot presented the Indago platform (<http://indago.inria.fr>) for the analysis of communication networks during the 1st World AI Cannes Festival (WAICF) in Cannes, France, in February 2023.

11 Scientific production

11.1 Major publications

- [1] P. Barbiero, G. Ciravegna, F. Giannini, M. E. Zarlenga, L. C. Magister, A. Tonda, P. Lio, F. Precioso, M. Jamnik and G. Marra. 'Interpretable Neural-Symbolic Concept Reasoning'. In: *Proceedings of the 2023 International Conference on Machine Learning*. ICML 2023 - 40th International Conference on Machine Learning. Honolulu, Hawaii, United States: arXiv, 2023. DOI: [10.48550/arXiv.2304.14068](https://doi.org/10.48550/arXiv.2304.14068). URL: <https://hal.science/hal-04244362>.
- [2] S. Bartels, W. Boomsma, J. Frellsen and D. Garreau. 'Kernel-Matrix Determinant Estimates from stopped Cholesky Decomposition'. In: *Journal of Machine Learning Research* (Jan. 2023). URL: <https://hal.science/hal-03359274>.
- [3] R. Boutin, C. Bouveyron and P. Latouche. 'Embedded Topics in the Stochastic Block Model'. In: *Statistics and Computing* 33.5 (14th Apr. 2023), p. 95. DOI: [10.1007/s11222-023-10265-9](https://doi.org/10.1007/s11222-023-10265-9). URL: <https://hal.science/hal-03782528>.
- [4] D. Liang, M. Corneli, C. Bouveyron and P. Latouche. 'The graph embedded topic model'. In: *Neurocomputing* 562 (17th Jan. 2023), p. 126900. DOI: [10.1016/j.neucom.2023.126900](https://doi.org/10.1016/j.neucom.2023.126900). URL: <https://hal.science/hal-03942487>.
- [5] G. Lopardo, F. Precioso and D. Garreau. 'A Sea of Words: An In-Depth Analysis of Anchors for Text Data'. In: International Conference on Artificial Intelligence and Statistics (AISTATS) 2023. Valencia, Spain, 25th Apr. 2023. URL: <https://hal.science/hal-03696324>.

- [6] M. Z. A. Mayaki and M. Riveill. ‘Machinery Anomaly Detection using artificial neural networks and signature feature extraction’. In: 2023 International Joint Conference on Neural Networks (IJCNN). Gold Coast, Australia: IEEE, 18th June 2023, pp. 1–10. DOI: [10.1109/IJCNN54540.2023.10191814](https://doi.org/10.1109/IJCNN54540.2023.10191814). URL: <https://hal.science/hal-03967924>.
- [7] H. Schmutz, O. Humbert and P.-A. Mattei, eds. *Don't fear the unlabelled: safe deep semi-supervised learning via simple debiasing*. International Conference on Learning Representations. 2023. URL: <https://hal.science/hal-03610272>.
- [8] H. H. J. Senetaire, D. Garreau, J. Frellsen and P.-A. Mattei, eds. *Explainability as statistical inference*. Proceedings of the 40th International Conference on Machine Learning. 2023. URL: <https://hal.science/hal-03888528>.
- [9] F. Simões, C. Bouveyron and F. Precioso. ‘DeepWILD : Wildlife Identification, Localisation and estimation on camera trap videos using Deep learning’. In: *Ecological Informatics* 75 (July 2023). DOI: [10.1016/j.ecoinf.2023.102095](https://doi.org/10.1016/j.ecoinf.2023.102095). URL: <https://hal.science/hal-03797530>.
- [10] A. Sportisse, H. Schmutz, O. Humbert, C. Bouveyron and P.-A. Mattei. *Are labels informative in semi-supervised learning? Estimating and leveraging the missing-data mechanism*. 2023. URL: <https://hal.science/hal-03982898>.

11.2 Publications of the year

International journals

- [11] S. Bartels, W. Boomsma, J. Frellsen and D. Garreau. ‘Kernel-Matrix Determinant Estimates from stopped Cholesky Decomposition’. In: *Journal of Machine Learning Research* (Jan. 2023). URL: <https://hal.science/hal-03359274>.
- [12] R. Boutin, C. Bouveyron and P. Latouche. ‘Embedded Topics in the Stochastic Block Model’. In: *Statistics and Computing* 33.5 (14th Apr. 2023), p. 95. DOI: [10.1007/s11222-023-10265-9](https://doi.org/10.1007/s11222-023-10265-9). URL: <https://hal.science/hal-03782528>.
- [13] G. Ciravegna, P. Barbiero, F. Giannini, M. Gori, P. Liò, M. Maggini and S. Melacci. ‘Logic explained networks’. In: *Artificial Intelligence* 314 (Jan. 2023), p. 103822. DOI: [10.1016/j.artint.2022.103822](https://doi.org/10.1016/j.artint.2022.103822). URL: <https://hal.science/hal-03862456>.
- [14] A. Destere, P. Marquet, M. Labriffe, M.-D. Drici and J.-B. Woillard. ‘A Hybrid Algorithm Combining Population Pharmacokinetic and Machine Learning for Isavuconazole Exposure Prediction’. In: *Pharmaceutical Research* 40.4 (29th Mar. 2023), pp. 951–959. DOI: [10.1007/s11095-023-03507-y](https://doi.org/10.1007/s11095-023-03507-y). URL: <https://unilim.hal.science/hal-04414322>.
- [15] A. Destere, D. Merino, T. Lavrut, F. Rocher, D. Viard, M.-D. Drici and A. Gérard. ‘Drug-induced cardiac toxicity and adverse drug reactions, a narrative review’. In: *Therapies* (Oct. 2023). DOI: [10.1016/j.therap.2023.10.008](https://doi.org/10.1016/j.therap.2023.10.008). URL: <https://hal.science/hal-04321284>.
- [16] A. Destere, M. Teisseyre, D. Merino, M. Cremoni, A. Gérard, T. Crepin, N. Jourde-Chiche, D. Graça, K. Zorzi, C. Fernandez, V. Brglez, S. Benzaken, V. L. Esnault, S. Benito, M.-D. Drici and B. Seitz-Polski. ‘Optimization of Rituximab Therapy in Adult Patients With PLA2R1-Associated Membranous Nephropathy With Artificial Intelligence’. In: *Kidney International Reports* (2023). DOI: [10.1016/j.ekir.2023.10.023](https://doi.org/10.1016/j.ekir.2023.10.023). URL: <https://hal.science/hal-04321277>.
- [17] Y. Gu, E. Erosheva, G. Xu and D. Dunson. ‘Dimension-Grouped Mixed Membership Models for Multivariate Categorical Data’. In: *Journal of Machine Learning Research* 24.88 (2023), p. 49. URL: <https://hal.science/hal-03522275>.
- [18] D. Liang, M. Corneli, C. Bouveyron and P. Latouche. ‘The graph embedded topic model’. In: *Neurocomputing* 562 (17th Jan. 2023), p. 126900. DOI: [10.1016/j.neucom.2023.126900](https://doi.org/10.1016/j.neucom.2023.126900). URL: <https://hal.science/hal-03942487>.
- [19] D. Merino, A. Gérard, E. van Obberghen, A. Destere, M. Lanteri-Minet and M.-D. Drici. ‘The Neuropsychiatric Safety Profile of Lasmiditan: A Comparative Disproportionality Analysis with Triptans’. In: *Neurotherapeutics* (12th July 2023). DOI: [10.1007/s13311-023-01404-1](https://doi.org/10.1007/s13311-023-01404-1). URL: <https://hal.science/hal-04189133>.

- [20] D. Merino, A. O. Gérard, A. Destere, F. Askenazy, E. Dor, M. Benoit, F. Cherikh and M.-D. Drici. ‘Iatrogenic triggers for anorexia nervosa and bulimia nervosa: A WHO safety database disproportionality analysis’. In: *Psychiatry Research* 327 (Sept. 2023), p. 115415. DOI: [10.1016/j.psychres.2023.115415](https://doi.org/10.1016/j.psychres.2023.115415). URL: <https://hal.science/hal-04189150>.
- [21] A. Paschier, A. Destere, C. Monchaud, M. Labriffe, P. Marquet and J.-b. Woillard. ‘Tacrolimus population pharmacokinetics in adult heart transplant patients’. In: *British Journal of Clinical Pharmacology* (Dec. 2023). DOI: [10.1111/bcp.15857](https://doi.org/10.1111/bcp.15857). URL: <https://hal.science/hal-04189142>.
- [22] R. Rastelli and M. Corneli. ‘Continuous Latent Position Models for Instantaneous Interactions’. In: *Network Science* 11.4 (24th July 2023), pp. 560–588. DOI: [10.1017/nws.2023.14](https://doi.org/10.1017/nws.2023.14). URL: <https://hal.science/hal-03517392>.
- [23] F. Simões, C. Bouveyron and F. Precioso. ‘DeepWILD : Wildlife Identification, Localisation and estimation on camera trap videos using Deep learning’. In: *Ecological Informatics* 75 (July 2023). DOI: [10.1016/j.ecoinf.2023.102095](https://doi.org/10.1016/j.ecoinf.2023.102095). URL: <https://hal.science/hal-03797530>.

International peer-reviewed conferences

- [24] G. Babykina and V. Vandewalle. ‘Classification non supervisée des processus d’événements récurrents’. In: *JdS 2023 - 54es Journées de Statistique de la SFdS. Bruxelles (BEL), France, 3rd July 2023*. URL: <https://inria.hal.science/hal-04411145>.
- [25] P. Barbiero, G. Ciravegna, F. Giannini, M. Espinosa Zarlenga, L. C. Magister, A. Tonda, P. Liò, F. Precioso, M. Jamnik and G. Marra. ‘Deep Concept Reasoning: Beyond the Accuracy-Interpretability Trade-Off’. In: *Workshop on Neuro-Explicit AI and Expert-informed Machine Learning for Engineering and Physical Sciences, of the European Conference on Machine Learning and Principles and Practices of Knowledge Discovery in Databases (ECML-PKDD) 2023. Torino, Italy, 18th Sept. 2023*. URL: <https://hal.science/hal-04244408>.
- [26] P. Barbiero, G. Ciravegna, F. Giannini, M. E. Zarlenga, L. C. Magister, A. Tonda, P. Lio, F. Precioso, M. Jamnik and G. Marra. ‘Interpretable Neural-Symbolic Concept Reasoning’. In: *International Conference on Machine Learning (ICML). Hawaii, United States: arXiv, 2023*. DOI: [10.48550/arXiv.2304.14068](https://doi.org/10.48550/arXiv.2304.14068). URL: <https://hal.science/hal-04244362>.
- [27] P. Barbiero, G. Ciravegna, F. Giannini, M. E. Zarlenga, L. C. Magister, A. Tonda, P. Lió, F. Precioso, M. Jamnik and G. Marra. ‘Interpretable Neural-Symbolic Concept Reasoning ★’. In: *CEUR workshop proceedings. NeSy 2023 - 17th International Workshop on Neural-Symbolic Learning and Reasoning. Vol. 3432. Certosa di Pontignano, Siena, Italy, 3rd July 2023*. DOI: [10.48550/arXiv.2304.14068](https://doi.org/10.48550/arXiv.2304.14068). URL: <https://hal.science/hal-04244373>.
- [28] A. Betti, G. Ciravegna, M. Gori, S. Melacci, K. Mottin and F. Precioso. ‘Toward Novel Optimizers: A Moreau-Yosida View of Gradient-Based Learning’. In: *Lecture Notes in Computer Science. Lecture Notes in Artificial Intelligence. AIXIA 2023 – XXIIInd International Conference of the Italian Association for Artificial Intelligence. Vol. LNCS. LNAI-14318. AIXIA 2023 – Advances in Artificial Intelligence. Rome, France: Springer Nature Switzerland, 2nd Nov. 2023, pp. 218–230*. DOI: [10.1007/978-3-031-47546-7_15](https://doi.org/10.1007/978-3-031-47546-7_15). URL: <https://hal.science/hal-04410778>.
- [29] C. Biernacki, C. Boyer, G. Celeux, J. Josse, F. Laporte, M. Marbac, A. Sportisse and V. Vandewalle. ‘Impact of missing data on mixtures and clustering with illustrations in Biology and Medicine’. In: *SPSR 2023 - The 24th annual Conference of the Romanian Society of Probability and Statistics. Bucharest, Romania, 21st Apr. 2023*. URL: <https://inria.hal.science/hal-04370870>.
- [30] C. Biernacki, V. Vandewalle and M. Marbac. ‘Clustering: from modeling to visualizing Mapping clusters as spherical Gaussians’. In: *SFC 2023 - Rencontres de la Société Francophone de Classification. Rencontres de la Société Francophone de Classification. Strasbourg, France, 6th July 2023*. URL: <https://inria.hal.science/hal-04370886>.

- [31] Y. Brault, Y. El Amraoui, M. Blay-Fornarino, P. Collet, F. Jaillet and F. Precioso. 'Taming the Diversity of Computational Notebooks'. In: *ACM digital library*. SPLC 2023 - 27th ACM International Systems and Software Product Line Conference. SPLC '23: Proceedings of the 27th ACM International Systems and Software Product Line Conference - Volume A. Tokyo, Japan: ACM, 28th Aug. 2023, pp. 27–33. DOI: [10.1145/3579027.3608974](https://doi.org/10.1145/3579027.3608974). URL: <https://hal.science/hal-04247860>.
- [32] R. Catellier, S. Vaiter and D. Garreau. 'On the Robustness of Text Vectorizers'. In: *PMLR. Proceedings of Machine Learning Research*. ICML 2023 - Fortieth International Conference on Machine Learning. Honolulu, United States, July 2023. URL: <https://hal.science/hal-04403681>.
- [33] G. Ciravegna, F. Precioso, A. Betti, K. Mottin and M. Gori. 'Knowledge-driven Active Learning'. In: *Lecture notes in computer science*. ECML PKDD 2023 - European Conference on Machine Learning and Knowledge Discovery in Databases. Vol. LNCS-14169. Joint European Conference on Machine Learning and Knowledge Discovery in Databases. Torino, Italy, 17th Sept. 2023. DOI: [10.1007/978-3-031-43412-9_3](https://doi.org/10.1007/978-3-031-43412-9_3). URL: <https://hal.science/hal-03527128>.
- [34] G. Lopardo, D. Garreau, F. Precioso and G. Ottosson. 'SMACE: A New Method for the Interpretability of Composite Decision Systems'. In: *ECML PKDD 2022 - European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*. Vol. LNCS. LNAI - 13713. LNCS. Lecture Notes in Computer Science. Grenoble, France: Springer International Publishing; Springer International Publishing, 17th Mar. 2023, pp. 325–339. DOI: [10.1007/978-3-031-26387-3_20](https://doi.org/10.1007/978-3-031-26387-3_20). URL: <https://hal.science/hal-03527129>.
- [35] G. Lopardo, F. Precioso and D. Garreau. 'A Sea of Words: An In-Depth Analysis of Anchors for Text Data'. In: *International Conference on Artificial Intelligence and Statistics (AISTATS) 2023*. Valencia, Spain, 25th Apr. 2023. URL: <https://hal.science/hal-03696324>.
- [36] G. Marchello, M. Corneli and C. Bouveyron. 'A Deep Dynamic Latent Block Model for the Co-clustering of Zero-Inflated Data Matrices'. In: *ECML-PKDD 2023 - European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*. Turin, Italy, Italy, 18th Sept. 2023. URL: <https://hal.science/hal-04150292>.
- [37] M. Z. A. Mayaki and M. Riveill. 'Machinery Anomaly Detection using artificial neural networks and signature feature extraction'. In: *2023 International Joint Conference on Neural Networks (IJCNN)*. Gold Coast, Australia: IEEE, 18th June 2023, pp. 1–10. DOI: [10.1109/IJCNN54540.2023.10191814](https://doi.org/10.1109/IJCNN54540.2023.10191814). URL: <https://hal.science/hal-03967924>.
- [38] B. Pouthier, L. Pilati, G. Valenti, C. Bouveyron and F. Precioso. 'Another Point of View on Visual Speech Recognition'. In: *Proceedings INTERSPEECH 2023*. 24th INTERSPEECH Conference. Vol. 2023. Dublin, Ireland: ISCA, 20th Aug. 2023, pp. 4089–4093. DOI: [10.21437/Interspeech.2023-969](https://doi.org/10.21437/Interspeech.2023-969). URL: <https://hal.science/hal-04250885>.
- [39] H. Schmutz, O. Humbert and P.-A. Mattei. 'Don't fear the unlabelled: safe deep semi-supervised learning via simple debiasing'. In: *International Conference on Learning Representations*. Kigali, Rwanda, 2023. URL: <https://hal.science/hal-03610272>.
- [40] H. H. J. Senetaire, D. Garreau, J. Frellsen and P.-A. Mattei. 'Explainability as statistical inference'. In: *International Conference on Machine Learning (ICML)*. Hawaii, United States, 2023. URL: <https://hal.science/hal-03888528>.
- [41] A. Sportisse, H. Schmutz, O. Humbert, C. Bouveyron and P.-A. Mattei. 'Are labels informative in semi-supervised learning? Estimating and leveraging the missing-data mechanism'. In: *International Conference on Machine Learning (ICML)*. Hawaii, United States, 2023. URL: <https://hal.science/hal-03982898>.
- [42] R. Sun, D. Lingrand and F. Precioso. 'Exploring the Road Graph in Trajectory Forecasting for Autonomous Driving'. In: *IEEE Xplore. ICCV 2023 - 2023 IEEE/CVF International Conference on Computer Vision Workshops - 1st Workshop on Scene Graphs and Graph Representation Learning*. 2023 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW). Paris, France: IEEE, 2nd Oct. 2023, pp. 71–80. DOI: [10.1109/ICCVW60793.2023.00014](https://doi.org/10.1109/ICCVW60793.2023.00014). URL: <https://hal.science/hal-04385182>.

Conferences without proceedings

- [43] C. Biernacki, M. Marbac and V. Vandewalle. ‘Clustering: from modeling to visualizing Mapping clusters as spherical Gaussians’. In: Working Group on Model-Based Clustering. Pittsburg, Etats-Unis, United States, 17th July 2023. URL: <https://inria.hal.science/hal-04370893>.

Doctoral dissertations and habilitation theses

- [44] C. Vincent-Cuaz. ‘Optimal transport for graph representation learning’. Université Côte d’Azur, 29th Mar. 2023. URL: <https://theses.hal.science/tel-04146481>.

Reports & preprints

- [45] I. Balelli, A. Sportisse, F. Cremonesi, P.-A. Mattei and M. Lorenzi. *Fed-MIWAE: Federated Imputation of Incomplete Data via Deep Generative Models*. 2023. URL: <https://hal.science/hal-04069795>.
- [46] A. Betti, M. Casoni and M. Gori. *Forward Approximate Solution for Linear Quadratic Tracking*. 30th Dec. 2022. URL: <https://hal.science/hal-03939954>.
- [47] R. Boutin, P. Latouche and C. Bouveyron. *The Deep Latent Position Topic Model for Clustering and Representation of Networks with Textual Edges*. 14th Apr. 2023. URL: <https://hal.science/hal-04068665>.
- [48] H. Fokkema, D. Garreau and T. van Erven. *The Risks of Recourse in Binary Classification*. 1st June 2023. URL: <https://hal.science/hal-04403713>.
- [49] G. Lopardo, F. Precioso and D. Garreau. *Faithful and Robust Local Interpretability for Textual Predictions*. 30th Oct. 2023. URL: <https://inria.hal.science/hal-04394149>.
- [50] G. Lopardo, F. Precioso and D. Garreau. *Understanding Post-hoc Explainers: The Case of Anchors*. 15th Mar. 2023. URL: <https://hal.science/hal-04038665>.
- [51] L. Ohl, P.-A. Mattei, C. Bouveyron, W. Harchaoui, M. Leclercq, A. Droit and F. Precioso. *Generalised Mutual Information: a Framework for Discriminative Clustering*. 6th Sept. 2023. URL: <https://hal.science/hal-04198445>.
- [52] L. Ohl, P.-A. Mattei, C. Bouveyron, M. Leclercq, A. Droit and F. Precioso. *Sparse GEMINI for joint discriminative clustering and feature selection*. 7th Feb. 2023. URL: <https://hal.science/hal-04195208>.
- [53] A. Sportisse, M. Marbac, F. Laporte, G. Celeux, C. Boyer, C. Biernacki and J. Josse. *Accompanying note : Model-based Clustering with Missing Not At Random Data*. 21st Dec. 2023. URL: <https://hal.science/hal-04358192>.
- [54] A. Sportisse, M. Marbac, F. Laporte, G. Celeux, C. Boyer, J. Josse and C. Biernacki. *Model-based Clustering with Missing Not At Random Data*. 21st Dec. 2023. URL: <https://hal.science/hal-03494674>.
- [55] R. Sun, L. Yang, D. Lingrand and F. Precioso. *Mind the map! Accounting for existing map information when estimating online HDMaps from sensor data*. 2023. DOI: [10.48550/arXiv.2311.10517](https://doi.org/10.48550/arXiv.2311.10517). URL: <https://hal.science/hal-04385135>.
- [56] M. Zoubeirou a Mayaki and M. Riveill. *AnoRand: A Semi Supervised Deep Learning Anomaly Detection Method by Random Labeling*. 28th May 2023. URL: <https://hal.science/hal-04116457>.

Other scientific publications

- [57] A. Barbe, M. T. Dhoub, C. Faron, M. Corneli and A. Zucker. ‘Zoo-KG: A knowledge graph of zoological data’. In: AIUCD 2023 - Associazione per l’Informatica Umanistica e la Cultura Digitale. Sienna (Italy), Italy, 5th June 2023. URL: <https://hal.science/hal-04235805>.

- [58] X. Zhang, M. Riveill and M.-D. Drici. 'Contrastive Clustering of medical reports to characterise Adverse Drug Reactions'. In: Statlearn 2023 workshop : Challenging problems in Statistical Learning. Montpellier, France, 5th Apr. 2023. URL: <https://hal.science/hal-04282263>.

11.3 Other

Softwares

- [59] [SW] L. Ohl, P.-A. Mattei and F. Precioso, *GemClus* version 0.2.0, 27th Apr. 2023. Inria & Université Cote d'Azur, CNRS, I3S, Sophia Antipolis, France. LIC: GNU General Public License v3.0 or later. HAL: [hal-04209625](https://hal.science/hal-04209625), URL: <https://hal.science/hal-04209625>, VCS: <https://github.com/gemini-clustering/GemClus>.