

RESEARCH CENTRE

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

IN PARTNERSHIP WITH:
Université Côte d'Azur

2022

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**

The Inria logo is a stylized, cursive script in red, positioned in the bottom right corner of the page.

Contents

Project-Team MAASAI	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	3
3 Research program	3
4 Application domains	6
5 Highlights of the year	6
5.1 Recruitments and promotions	6
5.2 Fundings	6
5.3 Awards	6
5.4 Conferences co-organised by team members	7
5.5 Innovation and transfer	7
5.6 Nominations	7
6 New software and platforms	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 Generalised Mutual Information for Discriminative Clustering	12
7.1.2 Continual Unsupervised Learning for Optical Flow Estimation with Deep Networks	12
7.1.3 A Deep Dynamic Latent Block Model for the Co-clustering of Zero-Inflated Data Matrices	12
7.1.4 Co-Clustering of Multivariate Functional Data for Air Pollution Analysis	13
7.1.5 Semi-supervised Consensus Clustering Based on Closed Patterns	14
7.1.6 Dimension-Grouped Mixed Membership Models for Multivariate Categorical Data	14
7.1.7 Tensor decomposition for learning Gaussian mixtures from moments	15
7.1.8 Dynamic Co-Clustering for PharmaCovigilance	15
7.1.9 Embedded Topics in the Stochastic Block Model	16
7.2 Understanding (deep) learning models	17
7.2.1 Explainability as statistical inference	17
7.2.2 Concept Embedding Models	17
7.2.3 How to scale hyperparameters for quickshift image segmentation	19
7.2.4 Interpretable Prediction of Post-Infarct Ventricular Arrhythmia using Graph Convolutional Network	19
7.2.5 Logic Explained Networks	20
7.2.6 Extending Logic Explained Network to Text Classification	21
7.2.7 Foveated Neural Computation	22
7.2.8 Continual Learning through Hamilton Equations	22
7.2.9 A free boundary singular transport equation as a formal limit of a discrete dynamical system	22
7.2.10 Forward Approximate Solution for Linear Quadratic Tracking	22
7.2.11 Comparing Feature Importance and Rule Extraction for Interpretability on Text Data	23
7.2.12 A Sea of Words: An In-Depth Analysis of Anchors for Text Data	23
7.2.13 Learning and Reasoning for Cultural Metadata Quality	23
7.2.14 SMACE: A New Method for the Interpretability of Composite Decision Systems	24
7.3 Adaptive and robust learning	24
7.3.1 Model-agnostic out-of-distribution detection using combined statistical tests	24
7.3.2 Autoregressive based Drift Detection Method	26
7.3.3 PARTIME: Scalable and Parallel Processing Over Time with Deep Neural Networks	26

7.3.4	Unobserved classes and extra variables detection in high-dimensional discriminant analysis	27
7.3.5	Knowledge-Driven Active Learning	27
7.4	Learning with heterogeneous and corrupted data	28
7.4.1	Don't fear the unlabelled: safe deep semi-supervised learning via simple debiasing	28
7.4.2	¹⁸ FDG PET/CT and Machine Learning for the prediction of lung cancer response to immunotherapy	29
7.4.3	Stochastic Coherence Over Attention Trajectory For Continuous Learning In Video Streams	30
7.4.4	Robust Lasso-Zero for sparse corruption and model selection with missing covariates	30
7.4.5	R-miss-tastic: a unified platform for missing values methods and workflows	30
7.4.6	Deep latent position model for node clustering in graphs	31
7.4.7	The graph embedded topic model	31
7.4.8	Graph Neural Network for Graph Drawing	32
7.4.9	Semi-relaxed Gromov-Wasserstein divergence for graphs classification	33
7.4.10	Template based Graph Neural Network with Optimal Transport Distances	33
7.4.11	Semi-relaxed Gromov-Wasserstein divergence with applications on graphs	34
7.4.12	How to deal with missing data in supervised deep learning?	35
7.4.13	Model-based clustering with Missing Not At Random Data	36
7.4.14	Unsupervised Text Clusterisation to characterize Adverse Drug Reactions from hospitalization reports	36
7.4.15	Auto-encoder Based Medicare Fraud Detection	38
7.4.16	Multiple Inputs Neural Networks for Fraud Detection	38
7.4.17	Continuous Latent Position Models for Instantaneous Interactions	38
7.4.18	DeepWILD: wildlife identification, localisation and population estimation from camera trap videos in the Parc National du Mercantour	39
8	Bilateral contracts and grants with industry	40
9	Partnerships and cooperations	41
9.1	International initiatives	41
9.1.1	Participation in other International Programs	42
9.2	European initiatives	42
9.2.1	FP7 & H2020 Projects	42
9.3	National initiatives	42
9.4	Regional initiatives	42
10	Dissemination	43
10.1	Promoting scientific activities	43
10.1.1	Scientific events: organisation	43
10.1.2	Journal	44
10.1.3	Invited talks	44
10.1.4	Leadership within the scientific community	45
10.1.5	Scientific expertise	45
10.2	Teaching - Supervision - Juries	45
10.3	Popularization	45
10.4	Interventions	46
11	Scientific production	46
11.1	Major publications	46
11.2	Publications of the year	47
11.3	Cited publications	52

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]
- Damien Garreau [UNIV COTE D'AZUR, Associate Professor]
- Frederic Precioso [UNIV COTE D'AZUR, Professor]
- Michel Riveill [UNIV COTE D'AZUR, Professor]
- Vincent Vandewalle [UNIV COTE D'AZUR, Professor, from Sep 2022, HDR]

Post-Doctoral Fellows

- Alessandro Betti [UNIV COTE D'AZUR]
- Gabriele Ciravegna [UNIV COTE D'AZUR]
- Aude Sportisse [INRIA]

PhD Students

- Kilian Burgi [UNIV COTE D'AZUR, from Sep 2022]
- Gatien Caillet [ORANGE, from Nov 2022]
- Antoine Collin [UNIV COTE D'AZUR, from Jun 2022]
- Célia Dacruz [UNIV COTE D'AZUR]
- Kevin Dsouza [UNIV COTE D'AZUR]
- Dingge Liang [UNIV COTE D'AZUR]
- Gianluigi Lopardo [UNIV COTE D'AZUR]
- Giulia Marchello [UNIV COTE D'AZUR]
- Hugo Miralles [ORANGE]
- Kevin Mottin [UNIV COTE D'AZUR]
- Louis Ohl [UNIV COTE D'AZUR]
- Baptiste Pouthier [NXP]
- Hugo Schmutz [UNIV COTE D'AZUR]
- Hugo Senetaire [UNIV DTU]
- Julie Tores [UNIV COTE D'AZUR, from Nov 2022]
- Cédric Vincent-Cuaz [UNIV COTE D'AZUR]
- Xuchun Zhang [UNIV COTE D'AZUR]

Technical Staff

- Lucas Boiteau [INRIA, Engineer, from Aug 2022]
- Marco Corneli [UNIV COTE D'AZUR, Engineer, until Aug 2022]
- Amosse Edouard [INSTANT SYSTEM , Engineer, from Feb 2022]
- Stephane Petiot [INRIA, Engineer]
- Li Yang [CNRS, Engineer, from Feb 2022]
- Mansour Zoubairou A Mayaki [PRO BTP, Engineer]

Interns and Apprentices

- Davide Adamo [INRIA, from Sep 2022]

Administrative Assistant

- Claire Senica [INRIA]

External Collaborators

- Marco Corneli [Université Côte d'Azur, from Sep 2022, Chaire de Professeur Junior]
- Marco Gori [UNIV FLORENCE]
- Pierre Latouche [UNIV CLERMONT AUVERG, from Nov 2022, HDR]
- Hans Ottosson [IBM]

2 Overall objectives

Artificial intelligence has become a key element in most scientific fields and is now part of everyone 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 behaviour 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 behaviour, 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 modelling 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 modelling 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.

¹The Linkage project: linkage.fr

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 astronomy, bioinformatics, recommender systems and ecology.

5 Highlights of the year

5.1 Recruitments and promotions

- The team benefited from the recruitment in 2022 of Vincent Vandewalle (coming from the Modal project-team of Inria Lille) as a Full Professor with Université Côte d'Azur. He joined the Maasai team on September, 1st, 2022.
- In the meantime, Marco Corneli (who was research engineer with Université Côte d'Azur) has been hired on a Chaire de Professeur Junior on AI & Archeology, effective on September, 1st, 2022.

5.2 Fundings

- Vincent Vandewalle was granted a 3IA chair from Institut 3IA Côte d'Azur.
- Marco Corneli was granted a chair of "Professor Junior" from Université Côte d'Azur.

5.3 Awards

- Hugo Schmutz was awarded a "highlight lecture" at the 35th Annual Congress of the European Association of Nuclear Medicine in Barcelona (2022), for his work in collaboration with P.-A. Mattei and O. Humbert.
- Cédric Vincent-Cuaz received a "NeurIPS Top Reviewer" award at the Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS'22)
- Pierre-Alexandre Mattei was ranked among the top 10% reviewers of the International Conference on Machine Learning (ICML) in 2022.
- Louis Ohl received a "NeurIPS Scholar Award" award at the Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS'22).

5.4 Conferences co-organised by team members

- The 1st Nice Workshop on Interpretability, organised by Damien Garreau, Frédéric Precioso and Gianluigi Lopardo. The workshop took place on November 17-18, 2022 in Nice, and counted 6 senior research talks, and 11 young research talks, with about 40 participants. Web : <https://sites.google.com/view/nwi2022/home>
- Statlearn 2022: 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). The 2022 edition was the 11th edition of the Statlearn series and welcomed about 50 participants. Web : <https://statlearn.sciencesconf.org>
- GenU 2022: this small-scale workshop has been held physically in Copenhagen in the Fall. The 2022 edition was on September 14-15, 2022 (Web: <https://genu.ai/2022/>).
- SophIA Summit: AI conference that brings together researchers and companies doing AI, held every Fall in Sophia Antipolis. The 2022 edition was held on 23-25th November 2022. Web: <https://univ-cotedazur.eu/events/sophia-summit>.

5.5 Innovation and transfer

- A contract has been signed with the company Naval Group for the development of an open-source Python library for semi-supervised learning, via the hiring of a research engineer. Lucas Boiteau started on August 1st, 2022.

5.6 Nominations

- Vincent Vandewalle has been nominated Deputy Scientific Director of the EFELIA Côte d'Azur program, effective October 1st, 2022. The EFELIA Côte d'Azur program is funded by the AMI Compétences et Métiers d'Avenir to develop education in AI in France.

6 New software and platforms

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 2022:

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://rmisstastic.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.r>

- 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:** realized 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: realized 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 modelling, as described in previous works of P.A. 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 modelling. 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 visualisation 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 a 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 proposes some typical use cases to give a better idea of what Topix can do.

7 New results

7.1 Unsupervised learning

7.1.1 Generalised Mutual Information for Discriminative Clustering

Participants: Louis Ohl, Pierre-Alexandre Mattei, Charles Bouveyron, 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 regularisations. While the quality of the regularisations have 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 maximisation of MI does not lead to satisfying clusters. We identified the Kullback-Leibler divergence as the main reason of this behaviour. Hence, we generalise in [27] the mutual information by changing its core distance, introducing the generalised mutual information (GEMINI): a set of metrics for unsupervised neural network training. Unlike MI, some GEMINIs do not require regularisations 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.

7.1.2 Continual Unsupervised Learning for Optical Flow Estimation with Deep Networks

Participants: Alessandro Betti

Collaborations: Simone Marullo, Matteo Tiezzi, Lapo Faggi, Enrico Meloni, Stefano Melacci

Keywords: Continual Learning, Optical Flow, Online Learning.

In [41] we present an extensive study on how neural networks can learn to estimate optical flow in a continual manner while observing a long video stream and reacting online to the streamed information without any further data buffering. To this end, we rely on photo-realistic video streams that we specifically created using 3D virtual environments, as well as on a real-world movie. Our analysis considers important model selection issues that might be easily overlooked at a first glance, comparing different neural architectures and also state-of-the-art models pretrained in an offline manner. Our results not only show the feasibility of continual unsupervised learning in optical flow estimation, but also indicate that the learned models, in several situations, are comparable to state-of-the-art offline-pretrained networks. Moreover, we show how common issues in continual learning, such as catastrophic forgetting, do not affect the proposed models in a disruptive manner, given the task at hand.

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

Participants: G. Marchello, M. Corneli, C. 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 [64], 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_{φ} -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. Fig. 1 shows some of the main results obtained through the application of ZI_{φ} -dLBM on the pharmacovigilance data set.



Figure 1: Estimated Poisson intensities, each color represents a different drug (ADR) cluster.

7.1.4 Co-Clustering of Multivariate Functional Data for Air Pollution Analysis

Participants: Charles Bouveyron.

Keywords: generative models, model-based co-clustering, functional data, air pollution, public health
Collaborations: J. Jacques and A. Schmutz (Univ. de Lyon), Fanny Simoes and Silvia Bottini (MDIab, MSI, Univ. Côte d’Azur)

In [11], we focused on Air pollution, which is nowadays a major treat for public health, with clear links with many diseases, especially cardiovascular ones. The spatio-temporal study of pollution is of great interest for governments and local authorities when deciding for public alerts or new city policies against pollution raise. The aim of this work is to study spatio-temporal profiles of environmental data collected in the south of France (Région Sud) by the public agency AtmoSud. The idea is to better understand the exposition to pollutants of inhabitants on a large territory with important differences in term of geography and urbanism. The data gather the recording of daily measurements of five environmental variables, namely three pollutants (PM10, NO2, O3) and two meteorological factors (pressure and temperature) over six years. Those data can be seen as multivariate functional data: quantitative entities evolving along time, for which there is a growing need of methods to summarize and understand them. For this purpose, a novel co-clustering model for multivariate functional data is defined. The model is based on a functional latent block model which assumes for each co-cluster a probabilistic distribution for multivariate functional principal component scores. A Stochastic EM algorithm, embedding a Gibbs sampler, is proposed for model inference, as well as a model selection criteria for choosing the number of co-clusters. The application of the proposed co-clustering algorithm on environmental data of the Région

Sud allowed to divide the region composed by 357 zones in six macro-areas with common exposure to pollution. We showed that pollution profiles vary accordingly to the seasons and the patterns are conserved during the 6 years studied. These results can be used by local authorities to develop specific programs to reduce pollution at the macro-area level and to identify specific periods of the year with high pollution peaks in order to set up specific prevention programs for health. Overall, the proposed co-clustering approach is a powerful resource to analyse multivariate functional data in order to identify intrinsic data structure and summarize variables profiles over long periods of time. Figure 2 illustrates the spatial and temporal clustering results.

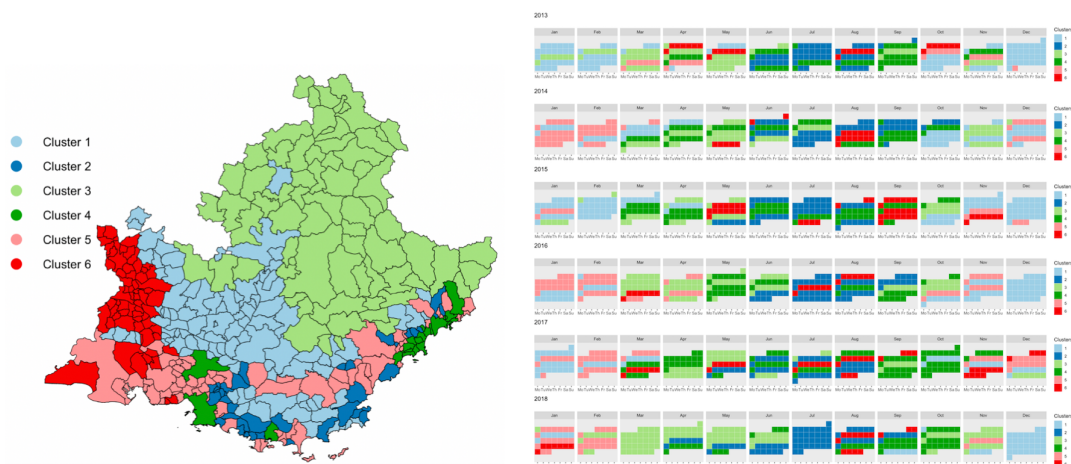


Figure 2: Spatial clustering of the area zones according to the air pollution dynamic over the studied period (left panel) and temporal segmentation of the time (right panel). Those tools may offer meaningful summaries on such massive pollution data to experts or local authorities.

7.1.5 Semi-supervised Consensus Clustering Based on Closed Patterns

Participants: Frédéric Precioso.

Keywords: Clustering; Semi-supervised learning; Semi-supervised consensus clustering; Frequent closed itemsets

Collaborations: Tianshu Yang (Université Côte d’Azur, Amadeus), Nicolas Pasquier (Université Côte d’Azur), Luca Marchetti (Amadeus), Michael Defoin Pratel (Amadeus), in a CIFRE PhD project with Amadeus

Semi-supervised consensus clustering, also called semi-supervised ensemble clustering, is a recently emerged technique that integrates prior knowledge into consensus clustering in order to improve the quality of the clustering result. In this article [23], we propose a novel semi-supervised consensus clustering algorithm extending the previous work on the MultiCons multiple consensus clustering approach. By using closed pattern mining technique, the proposed Semi-MultiCons algorithm manages to generate a recommended consensus solution with a relevant inferred number of clusters based on ensemble members with different and pairwise constraints. Compared with other semi-supervised and/or consensus clustering approaches, Semi-MultiCons does not require the number of generated clusters as an input parameter, and is able to alleviate the widely reported negative effect related to the integration of constraints into clustering. The experimental results demonstrate that the proposed method outperforms state of the art semi-supervised consensus clustering algorithms.

7.1.6 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 [61], we propose a new class of Dimension-Grouped MMMs (Gro-M³s) for multivariate categorical data, which improve 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.

7.1.7 Tensor decomposition for learning Gaussian mixtures from moments

Participants: Pierre-Alexandre Mattei.

Keywords: model-based clustering, tensor decomposition, method of moments

Collaborations: Rima Khouja, Bernard Mourrain (Inria Sophia-Antipolis, AROMATH team)

In [16] consider the problem of estimation of Gaussian mixture models. As an alternative to maximum-likelihood, our focus is on the method of moments. More specifically, we investigate symmetric tensor decomposition methods, where the tensor is built from empirical moments of the data distribution. We consider identifiable tensors, which have a unique decomposition, showing that moment tensors built from spherical Gaussian mixtures have this property. We prove that symmetric tensors with interpolation degree strictly less than half their order are identifiable and we present an algorithm, based on simple linear algebra operations, to compute their decomposition. Illustrative experimentations show the impact of the tensor decomposition method for recovering Gaussian mixtures, in comparison with other state-of-the-art approaches.

7.1.8 Dynamic Co-Clustering for PharmaCovigilance

Participants: Charles Bouveyron, Marco Corneli, Giulia Marchello.

Keywords: generative models, dynamic co-clustering, count data, pharmacovigilance

Collaborations: Audrey Fresse (Centre de Pharmacovigilance, CHU de Nice)

We consider in [17] the problem of co-clustering count matrices with a high level of missing values that may evolve in time. We introduce a generative model, named dynamic latent block model (dLBM), which extends the classical binary latent block model (LBM) to the dynamic case. The time dependent counting data are modeled via non-homogeneous Poisson processes (HHPPs). The continuous time is handled by a partition of the whole considered time period, with the interaction counts being aggregated on the time




Figure 3: Evolution of the relation between the drug clusters and the ADR clusters over time. Each color corresponds to a different cluster of adverse drug reaction.

intervals of such partition. In this way, a sequence of static matrices that allows us to identify meaningful time clusters is obtained. The model inference is done using a SEM-Gibbs algorithm and the ICL criterion is used for model selection. Numerical experiments on simulated data highlight the main features of the proposed approach and show the interest of dLBM with respect to related works. An application to adverse drug reaction (ADR) dataset, obtained thanks to the collaboration with the Regional Center of Pharmacovigilance (RCPV) of Nice (France), is also proposed. One of the missions of RCPVs is safety signal detection. However, the current expert detection of safety signals, despite being unavoidable, has the disadvantage of being incomplete due to the workload it represents. For this reason, developing automatized methods of safety signal detection is currently a major issue in pharmacovigilance. The application of dLBM on this dataset allowed us to extract meaningful patterns for medical authorities. In particular, dLBM identifies 7 drug clusters, 10 ADRs clusters and 6 time clusters. The clusters identified by the algorithm are coherent with previous knowledge and adequately represent the variety of drugs present in the dataset. Moreover, an in depth analysis of the clusters found by the model revealed that dLBM correctly detected the three drugs that gave rise to the health scandals that took place between 2010 and 2020, demonstrating its potential as a routine tool in pharmacovigilance. Figure 3 illustrates this work.

7.1.9 Embedded Topics in the Stochastic Block Model

Participants: Charles Bouveyron, Rémi Boutin, Pierre Latouche.

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 analysing 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 [56] networks for which two nodes are linked if and only if they share textual

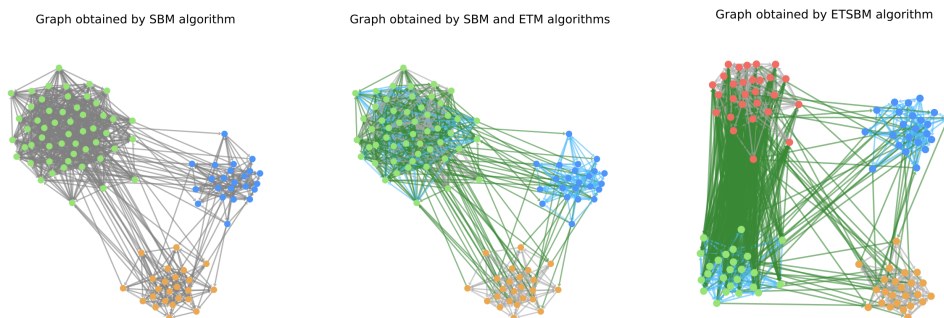


Figure 4: Clustering on a simulated network with SBM (left), SBM+ETM (center) and ETSBM (right).

data. We introduce a deep latent variable model allowing embedded topics to be handled called ETSBM to simultaneously perform clustering on the nodes while modelling the topics used between the different clusters. 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-maximisation algorithm combined with a stochastic gradient descent. The methodology is evaluated on synthetic data and on a real world dataset.

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 Frelsen (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 [69], 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. 5). 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 Concept Embedding Models

Participants: Gabriele Ciravegna, Frederic Precioso

Keywords: Deep Learning, Interpretability, Human and AI, Concept-based Explanations

Collaborations: Mateo Espinosa Zarlenga, Pietro Barbiero, Zohreh Shams, Adrian Weller, Pietro Lio, Mateja Jamnik (University of Cambridge), Francesco Giannini, Michelangelo Diligenti, Stefano Melacci (Università di Siena), Giuseppe Marra, (Katholieke Universiteit Leuven)

While any child can explain what an “apple” is by enumerating its characteristics, deep neural networks (DNNs) fail to explain what they learn in human-understandable terms despite their high prediction accuracy. This accuracy-vs-interpretability trade-off has become a major concern as high-performing DNNs become commonplace in practice, thus questioning the ethical and legal ramifications of their deployment. Concept bottleneck models (CBMs) aim at replacing “black-box” DNNs by first learning to predict a set of concepts, that is, “interpretable” high-level units of information (e.g., “colour”




Figure 5: The LEX pipeline allows us to transform any prediction model into an explainable one.

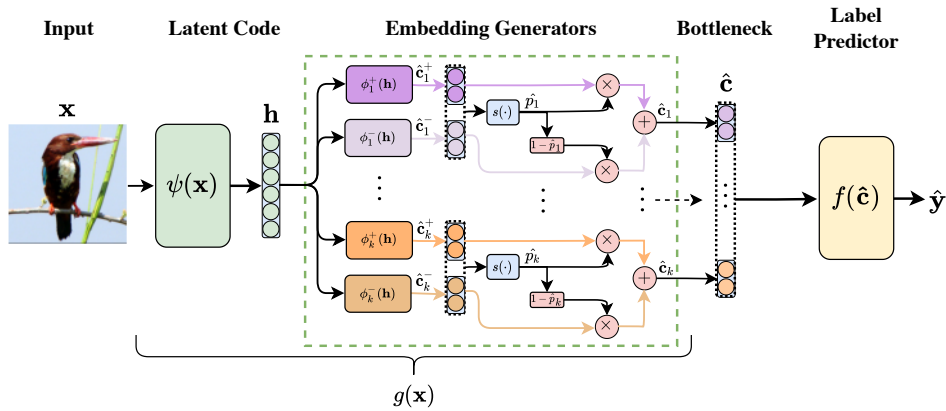


Figure 6: **Concept Embedding Model:** from an intermediate latent code \mathbf{h} , we learn two embeddings per concept, one for when it is active (i.e., $\hat{\mathbf{c}}_i^+$), and another when it is inactive (i.e., $\hat{\mathbf{c}}_i^-$). Each concept embedding (shown in this example as a vector with $m = 2$ activations) is then aligned to its corresponding ground truth concept through the scoring function $s(\cdot)$, which learns to assign activation probabilities \hat{p}_i for each concept. These probabilities are used to output an embedding for each concept via a weighted mixture of each concept's positive and negative embedding.

or “shape”) provided at training time, and then using these concepts to learn a downstream classification task. Predicting tasks as a function of concepts engenders user trust by allowing predictions to be explained in terms of concepts and by supporting human interventions, where at test-time an expert can correct a mis-predicted concept, possibly changing the CBM’s output. That said, concept bottlenecks may impair task accuracy, especially when concept labels do not contain all the necessary information for accurately predicting a downstream task (i.e., they form an “incomplete” representation of the task). In principle, extending a CBM’s bottleneck with a set of unsupervised neurons may improve task accuracy. However, such a hybrid approach not only significantly hinders the performance of concept interventions, but it also affects the interpretability of the learnt bottleneck, thus undermining user trust.

In [45], we propose Concept Embedding Models (CEMs, see fig. 6), a novel concept bottleneck model which overcomes the current accuracy-vs-interpretability trade-off found in concept-incomplete settings. Furthermore, we introduce two new metrics for evaluating concept representations and use them to help understand why our approach circumvents the limits found in the current state-of-the-art CBMs. Our experiments provide significant evidence in favour of CEM’s accuracy/interpretability and, consequently, in favour of its real-world deployment. In particular, CEMs offer: (1) state-of-the-art task accuracy, (2) interpretable concept representations aligned with human ground truths, (3) effective interventions on

learnt concepts, and (4) robustness to incorrect concept interventions. While in practice CBMs require carefully selected concept annotations during training, which can be as expensive as task labels to obtain, our results suggest that CEM is more efficient in concept-incomplete settings, requiring less concept annotations and being more applicable to real-world tasks. While there is room for improvement in both concept alignment and task accuracy in challenging benchmarks such as CUB or CelebA, as well as in resource utilization during inference/training, our results indicate that CEM advances the state-of-the-art for the accuracy-vs-interpretability trade-off, making progress on a crucial concern in explainable AI.

7.2.3 How to scale hyperparameters for quickshift image segmentation

Participants: Damien Garreau

Keywords: Computer vision, clustering

Quickshift is a popular algorithm for image segmentation, used as a preprocessing step in many applications. Unfortunately, it is quite challenging to understand the hyperparameters' influence on the number and shape of superpixels produced by the method. In [60], we study theoretically a slightly modified version of the quickshift algorithm, with a particular emphasis on homogeneous image patches with i.i.d. pixel noise and sharp boundaries between such patches. Leveraging this analysis, we derive a simple heuristic to scale quickshift hyperparameters with respect to the image size, which we check empirically (see fig. 7).




Figure 7: Using our heuristic, we can scale hyperparameters with respect to the size of the image. *Left panel:* original image size, 15 superpixels obtained with kernel size 5.0 and infinite maximal distance; *Middle panel:* same hyperparameters, upscaled image by a factor 2 yield approximately 4 times more superpixels; *Right panel:* multiplying kernel size by 2, we end up with roughly the same numbers of superpixels as before.

7.2.4 Interpretable Prediction of Post-Infarct Ventricular Arrhythmia using Graph Convolutional Network

Participants: Damien Garreau

Collaborations: Buntheng Ly, Sonny Finsterbach, Marta Nuñez-Garcia, Pierre Jaïs, Hubert Cochet, Maxime Sermesant

Keywords: Interpretability, graph neural networks, Ventricular Arrhythmia

Heterogeneity of left ventricular (LV) myocardium infarction scar plays an important role as anatomical substrate in ventricular arrhythmia (VA) mechanism. LV myocardium thinning, as observed on cardiac computed tomography (CT), has been shown to correlate with LV myocardial scar and with abnormal electrical activity. In [26], we propose an automatic pipeline for VA prediction, based on CT images, using a Graph Convolutional Network (GCN). The pipeline includes the segmentation of LV masks from the input CT image, the short-axis orientation reformatting, LV myocardium thickness computation and mid-wall surface mesh generation. An average LV mesh was computed and fitted to every patient in order to use the same number of vertices with point-to-point correspondence. The GCN model was trained using the thickness value as the node feature and the atlas edges as the adjacency matrix. This allows the model to process the data on the 3D patient anatomy and bypass the “grid” structure limitation of the traditional convolutional neural network. The model was trained and evaluated on a dataset of 600 patients (27% VA), using 451 (3/4) and 149 (1/4) patients as training and testing data, respectively. The evaluation results showed that the graph model (81% accuracy) outperformed the clinical baseline (67%), the left ventricular ejection fraction, and the scar size (73%). We further studied the interpretability of

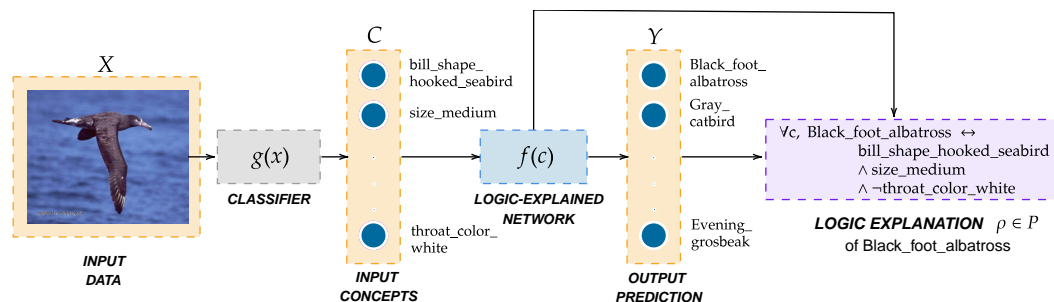


Figure 8: An instantiation of the Logic Explained Network framework inspired by the CUB 200-2011 fine-grained classification dataset. Classes are divided into a two-level hierarchy. A LEN is placed on top of a convolutional neural network $g(\cdot)$ in order to (i) classify the species of the bird in input and (ii) provide an explanation on why it belongs to this class. The logic explanation in the example showcases the predicted output class.

the trained model using LIME and integrated gradients and found promising results on the personalised discovering of the specific regions within the infarct area related to the arrhythmogenesis.

7.2.5 Logic Explained Networks

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 [12] we present a unified framework for XAI allowing the design of a family of neural models, the Logic Explained Networks (LENs, see fig. 8), 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 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

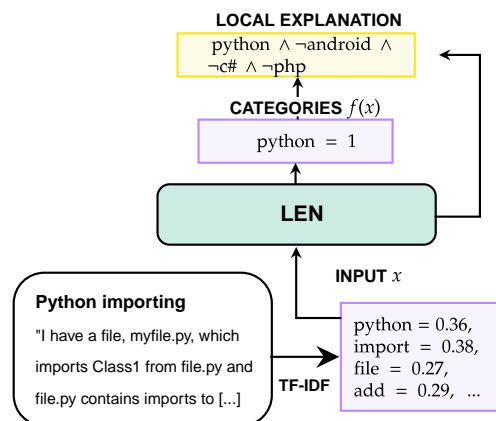


Figure 9: Example of $LENO^p$ local explanations for a text predicted as 'python' tag.

of LENs in a GitHub repository with an extensive documentation about LENs models, implementing different trade-offs between interpretability, explainability and accuracy.

7.2.6 Extending Logic Explained Network to Text Classification

Participants: Gabriele Ciravegna

Keywords: XAI, Logic Explanation, Text Classification

Collaborations: Rishabh Jain, Pietro Barbiero, Pietro Lio (University of Cambridge), Francesco Giannini (Università di Siena), Davide Buffelli (Università di Padova)

The majority of the data found in an organization tends to be unstructured (with some estimates being over 80%). Unstructured data tends to be text heavy. Sifting and sorting this data by hand require a lot of effort and time. Text classification is a useful way of automating this process, with applications ranging from small tasks (e.g., spam-email classification), to safety-critical ones (e.g., legal-document risk assessment). The development of Deep Neural Networks has enabled the creation of high accuracy text classifiers with state-of-the-art models leveraging different forms of architectures, like RNNs (GRU, LSTM) or Transformer models. However, these architectures are considered as black-box models, since their decision processes are not easy to explain and depend on a very large set of parameters. In order to shed light on neural models' decision processes, eXplainable Artificial Intelligence (XAI) techniques attempt to understand text attribution to certain classes, for instance by using white-box models. Interpretable-by-design models engender higher trust in human users with respect to explanation methods for black-boxes, at the cost, however, of lower prediction performance. Previous works introduced the Logic Explained Network (LEN), an explainable-by-design neural network combining interpretability of white-box models with high performance of neural networks. However, the authors only compared LENs with white-box models and on tabular/computer vision tasks.

For these reasons, in [37] we apply an improved version of the $LENO^p$ to the text classification problem (see fig. 9, and we compare it with LIME a standard and very-well known explanation method. LEN and LIME provide different kind of explanations, respectively FOL formulae and feature-importance vectors, and we assess their user-friendliness by means of a user-study. As an evaluation benchmark, we considered Multi-Label Text Classification for the tag classification task on the "StackSample: 10% of Stack Overflow Q&A" dataset. The paper aims to apply LENs to the text classification problem and to test the generated explanations. More specifically, its purpose are to: (1) improve LEN explanation algorithm with $LENO^p$ (2) confirm the small performance drop when employing LENs, w.r.t. using a black-box model; (3) compare the faithfulness and the sensitivity of the explanations provided by LENs and LIME; (4) assess the user-friendliness of the two kinds of explanations.

7.2.7 Foveated Neural Computation

Participants: Alessandro Betti, Marco Gori

Collaborations: Matteo Tiezzi, Simone Marullo, Enrico Meloni, Lapo Faggi, Stefano Melacci

Keywords: Foveated Convolutional Layers, Convolutional Neural Networks, Visual Attention.

In [43] this paper we introduce the notion of Foveated Convolutional Layer (FCL), that formalizes the idea of location-dependent convolutions with foveated processing, i.e., fine-grained processing in a given-focused area and coarser processing in the peripheral regions. We show how the idea of foveated computations can be exploited not only as a filtering mechanism, but also as a mean to speed-up inference with respect to classic convolutional layers, allowing the user to select the appropriate trade-off between level of detail and computational burden. FCLs can be stacked into neural architectures and we evaluate them in several tasks, showing how they efficiently handle the information in the peripheral regions, eventually avoiding the development of misleading biases. When integrated with a model of human attention, FCL-based networks naturally implement a foveated visual system that guides the attention toward the locations of interest, as we experimentally analyze on a stream of visual stimuli.

7.2.8 Continual Learning through Hamilton Equations

Participants: Alessandro Betti, Marco Gori

Collaborations: Lapo Faggi, Matteo Tiezzi, Simone Marullo, Enrico Meloni, Stefano Melacci

Keywords: Continual Learning, Optimal Control, Hamilton-Jacobi.

In [35] we consider a fully new perspective, rethinking the methodologies to be used to tackle continual learning, instead of re-adapting offline-oriented optimization. In particular, we propose a novel method to frame continual and online learning within the framework of optimal control. The proposed formulation leads to a novel interpretation of learning dynamics in terms of Hamilton equations. As a case study for the theory, we consider the problem of unsupervised optical flow estimation from a video stream. An experimental proof of concept for this learning task is discussed with the purpose of illustrating the soundness of the proposed approach, and opening to further research in this direction.

7.2.9 A free boundary singular transport equation as a formal limit of a discrete dynamical system

Participants: Alessandro Betti

Collaborations: Giovanni Bellettini, Maurizio Paolini

Keywords: PDE, Continuous Open Mancala, Transport Equation.

In [53] we study the continuous version of a hyperbolic rescaling of a discrete game, called open mancala. The resulting PDE turns out to be a singular transport equation, with a forcing term taking values in $\{0, 1\}$, and discontinuous in the solution itself. We prove existence and uniqueness of a certain formulation of the problem, based on a nonlocal equation satisfied by the free boundary dividing the region where the forcing is one (active region) and the region where there is no forcing (tail region). Several examples, most notably the Riemann problem, are provided, related to singularity formation. Interestingly, the solution can be obtained by a suitable vertical rearrangement of a multi-function. Furthermore, the PDE admits a Lyapunov functional.

7.2.10 Forward Approximate Solution for Linear Quadratic Tracking

Participants: Alessandro Betti, Marco Gori

Collaborations: Michele Casoni

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

In [54], 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.

7.2.11 Comparing Feature Importance and Rule Extraction for Interpretability on Text Data

Participants: Gianluigi Lopardo, Damien Garreau

Keywords: Interpretability, Explainable Artificial Intelligence, Natural Language Processing

Complex machine learning algorithms are used more and more often in critical tasks involving text data, leading to the development of interpretability methods. Among local methods, two families have emerged: those computing importance scores for each feature and those extracting simple logical rules. In [39] we show that using different methods can lead to unexpectedly different explanations, even when applied to simple models for which we would expect qualitative coincidence, as in Figure 10. To quantify this effect, we propose a new approach to compare explanations produced by different methods.

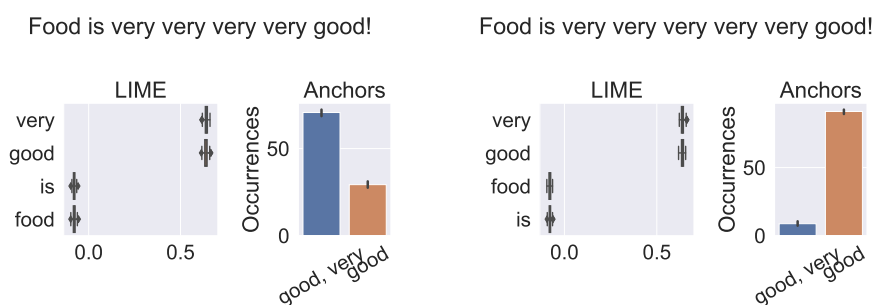


Figure 10: Making a word disappear from the explanation by adding one occurrence. The classifier $\mathbf{I}_{(very \in z \text{ and } good \in z)}$ is applied when the multiplicities are $m_{very} = 4$ (left) and $m_{very} = 5$ (right).

7.2.12 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 [63], 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 11, we present explicit results on different classes of models when the preprocessing step is 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.

7.2.13 Learning and Reasoning for Cultural Metadata Quality

Participants: Frédéric Precioso.

Keywords: Deep Learning, Image Recognition, Semantic Web, Knowledge Graph

Collaborations: Anna Bobasheva, Fabien Gandon (Inria)

This work [10] combines semantic reasoning and machine learning to create tools that allow curators of the visual art collections to identify and correct the annotations of the artwork as well as to improve the relevance of the content-based search results in these collections. The research is based on the Joconde database maintained by French Ministry of Culture that contains illustrated artwork records from main French public and private museums representing archeological objects, decorative arts, fine

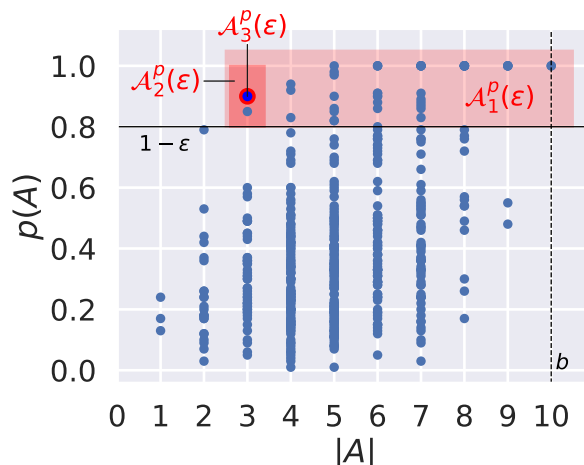


Figure 11: An illustration of the algorithm with evaluation function $p = \text{Prec}$. Each blue dot is an anchor, with x coordinate its length and y coordinate its value for p . Here, $\varepsilon = 0.2$ and the maximal length of an anchor is $b = 10$ (the length of ξ). In the end, the anchor A such that $|A| = 3$ and $p(A) = 0.9$ is selected (red circle).

arts, historical and scientific documents, etc. The Joconde database includes semantic metadata that describes properties of the artworks and their content. The developed methods create a data pipeline that processes metadata, trains a Convolutional Neural Network image classification model, makes prediction for the entire collection and expands the metadata to be the base for the SPARQL search queries. We developed a set of such queries to identify noise and silence in the human annotations and to search image content with results ranked according to the relevance of the objects quantified by the prediction score provided by the deep learning model. We also developed methods to discover new contextual relationships between the concepts in the metadata by analyzing the contrast between the concepts similarities in the Joconde's semantic model and other vocabularies and we tried to improve the model prediction scores based on the semantic relations. Our results show that cross-fertilization between symbolic AI and machine learning can indeed provide the tools to address the challenges of the museum curators work describing the artwork pieces and searching for the relevant images.

7.2.14 SMACE: A New Method for the Interpretability of Composite Decision Systems

Participants: Gianluigi Lopardo, Damien Garreau, Frédéric Precioso, Greger Ottosson

Keywords: Interpretability, Composite AI, Decision-making

Collaborations: IBM France

Interpretability is a pressing issue for decision systems. Many post hoc methods have been proposed to explain the predictions of a single machine learning model. However, business processes and decision systems are rarely centered around a unique model. These systems combine multiple models that produce key predictions, and then apply decision rules to generate the final decision (see Figure 13 for an illustration). To explain such decisions, we propose in [40] the Semi-Model-Agnostic Contextual Explainer (SMACE), a new interpretability method that combines a geometric approach for decision rules with existing interpretability methods for machine learning models to generate an intuitive feature ranking tailored to the end user. We show that established model-agnostic approaches produce poor results on tabular data in this setting, in particular giving the same importance to several features, whereas SMACE can rank them in a meaningful way.

7.3 Adaptive and robust learning

7.3.1 Model-agnostic out-of-distribution detection using combined statistical tests

Participants: Pierre-Alexandre Mattei, Hugo Senetaire, Hugo Schmutz

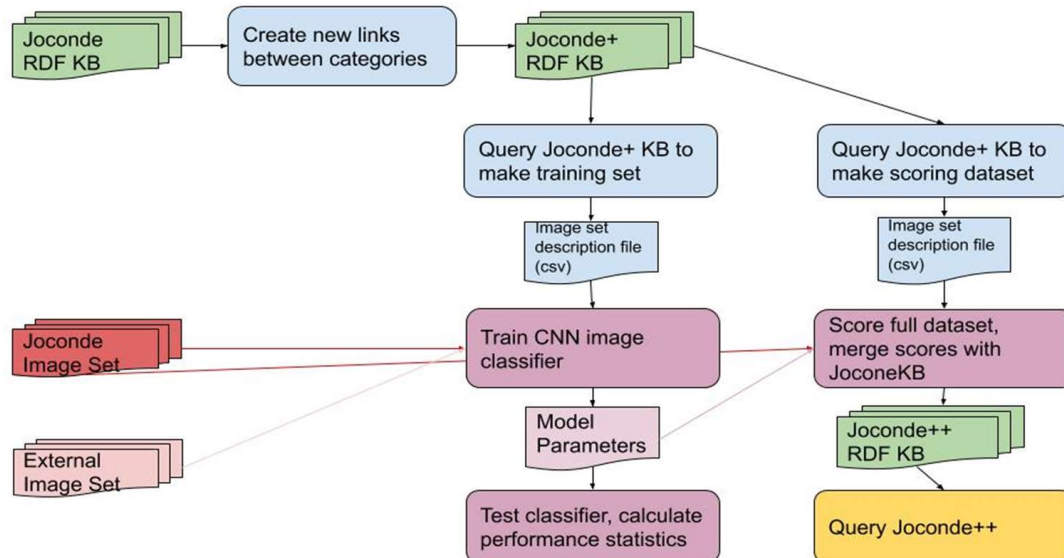


Figure 12: Data processing pipeline combining reasoning and learning.




Figure 13: Structure of a composite decision system with D input features x_1, \dots, x_D , and N models m_1, \dots, m_N . A decision policy P (i.e., a set of decision rules) is finally applied to produce an outcome O . Note that in general both the models and the rules take a subset of input features as input, though not necessarily the same.

Collaborations: Jakob Havtorn, Lars Maaløe, Søren Hauberg, Jes Frelsen

Keywords: Anomaly detection, statistical tests

We present simple methods for out-of-distribution detection using a trained generative model. These techniques, based on classical statistical tests, are model-agnostic in the sense that they can be applied to any differentiable generative model. The idea is to combine a classical parametric test (Rao's score test) with the recently introduced typicality test. These two test statistics are both theoretically well-founded and exploit different sources of information based on the likelihood for the typicality test and its gradient for the score test. We show that combining them using Fisher's method overall leads to a more accurate out-of-distribution test. We also discuss the benefits of casting out-of-distribution detection as a statistical testing problem, noting in particular that false positive rate control can be valuable for practical out-of-distribution detection. Despite their simplicity and generality, these methods can be competitive with model-specific out-of-distribution detection algorithms without any assumptions on the out-distribution.

7.3.2 Autoregressive based Drift Detection Method

Participants: Mansour Zoubeirou A Mayaki, Michel Riveill

Keywords: Concept drift detection, Data streams, Auto-regressive model, Machine learning, Deep neural networks

In the classic machine learning framework, models are trained on historical data and used to predict future values. It is assumed that the data distribution does not change over time (stationarity). However, in real-world scenarios, the data generation process changes over time and the model has to adapt to the new incoming data. This phenomenon is known as concept drift and leads to a decrease in the predictive model's performance. We proposed a new concept drift detection method based on autoregressive models called ADDM [48]. This method can be integrated into any machine learning algorithm from deep neural networks to simple linear regression model. Our results show that this new concept drift detection method outperforms the state-of-the-art drift detection methods, both on synthetic data sets and real-world data sets. Our approach is theoretically guaranteed as well as empirical and effective for the detection of various concept drifts. In addition to the drift detector, we proposed a new method of concept drift adaptation based on the severity of the drift. The architecture and dataflow of ADDM is shown in Figure 14.

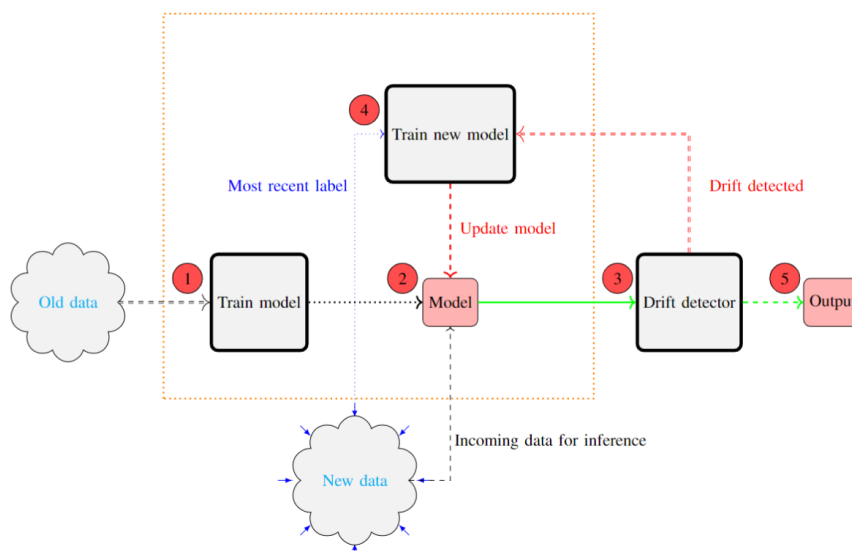


Figure 14: Architecture and Dataflow of ADDM.

7.3.3 PARTIME: Scalable and Parallel Processing Over Time with Deep Neural Networks

Participants: Alessandro Betti, Marco Gori

Collaborations: Enrico Meloni, Lapo Faggi, Simone Marullo, Matteo Tiezzi, Stefano Melacci

Keywords: PyTorch, PARTIME, Software Library Transport Equation.

In [66] this paper, we present PARTIME, a software library written in Python and based on PyTorch, designed specifically to speed up neural networks whenever data is continuously streamed over time, for both learning and inference. Existing libraries are designed to exploit data-level parallelism, assuming that samples are batched, a condition that is not naturally met in applications that are based on streamed data. Differently, PARTIME starts processing each data sample at the time in which it becomes available from the stream. PARTIME wraps the code that implements a feed-forward multi-layer network and it distributes the layer-wise processing among multiple devices, such as Graphics Processing Units (GPUs). Thanks to its pipeline-based computational scheme, PARTIME allows the devices to perform computations in parallel. At inference time this results in scaling capabilities that are theoretically linear

with respect to the number of devices. During the learning stage, PARTIME can leverage the non-i.i.d. nature of the streamed data with samples that are smoothly evolving over time for efficient gradient computations. Experiments are performed in order to empirically compare PARTIME with classic non-parallel neural computations in online learning, distributing operations on up to 8 NVIDIA GPUs, showing significant speedups that are almost linear in the number of devices, mitigating the impact of the data transfer overhead.

7.3.4 Unobserved classes and extra variables detection in high-dimensional discriminant analysis

Participants: Charles Bouveyron, Pierre-Alexandre Mattei.

Keywords: Adaptive supervised classification; conditional estimation; model-based discriminant analysis; unobserved classes; variable selection.

Collaborations: Michael Fop and Brendan Murphy (University College Dublin, Ireland)

In supervised classification problems, the test set may contain data points belonging to classes not observed in the learning phase. Moreover, the same units in the test data may be measured on a set of additional variables recorded at a subsequent stage with respect to when the learning sample was collected. In this situation, the classifier built in the learning phase needs to adapt to handle potential unknown classes and the extra dimensions. We introduce in [15] a model-based discriminant approach, Dimension-Adaptive Mixture Discriminant Analysis (D-AMDA), which can detect unobserved classes and adapt to the increasing dimensionality. Model estimation is carried out via a full inductive approach based on an EM algorithm. The method is then embedded in a more general framework for adaptive variable selection and classification suitable for data of large dimensions. A simulation study and an artificial experiment related to classification of adulterated honey samples are used to validate the ability of the proposed framework to deal with complex situations. Figure 15 illustrates the general framework of the proposed approach.

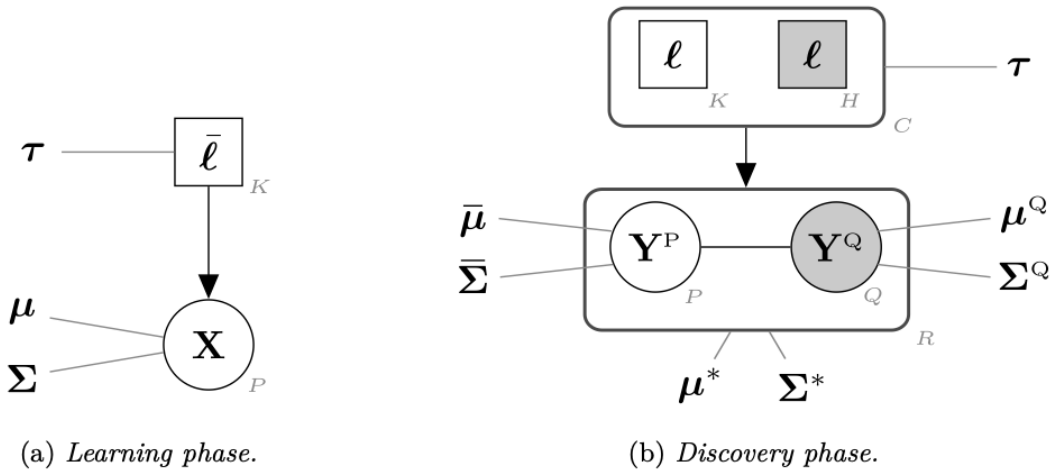


Figure 15: General framework of the inductive approach for Dimension-Adaptive Mixture Discriminant Analysis (DAMDA).

7.3.5 Knowledge-Driven Active Learning

Participants: Gabriele Ciravegna, 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 [75]. 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 [58] 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 outperforms 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.

7.4 Learning with heterogeneous and corrupted data

7.4.1 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 [68], our starting point is to notice that the estimate of the risk that most discriminative SSL methods minimise 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.

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

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 generalisation 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 16, 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 16: 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.

7.4.2 ¹⁸FDG PET/CT and Machine Learning for the prediction of lung cancer response to immunotherapy

Participants: Hugo Schmutz, Pierre-Alexandre Mattei

Collaborations: Sara Contu, David Chardin, Olivier Humbert

Keywords: FDG PET, immunotherapy, lung cancer, adverse events, biomarker, machine learning, heterogeneous data, features selection

In patients with non-small cell lung cancer (NSCLC) treated with immunotherapy, individual biological and PET imaging prognostic biomarkers have been recently identified. However, a combination of biomarkers has not been studied yet. This study [42] aims to combine clinical, biological and ¹⁸FDG PET/CT parameters and use machine-learning algorithms to build more accurate prognostic models of NSCLC response to immunotherapy.

Patients with metastatic NSCLC, treated with either pembrolizumab or nivolumab in monotherapy, were prospectively included in 2 different monocentric prospective trials (NCT03584334; ID-RCB: 2018-A00915-50). For all patients, a total of 28 baseline quantitative features were analyzed (12 clinical, 6 biological and 10 PET/CT parameters), such as the patient’s age, weight, height, ECOG performance status (PS), PD-L1 tumour expression level (PD-L1%), neutrophil to lymphocyte blood ratio, number and metabolism of lesions, metabolic tumour volume (MTV) and spleen to liver metabolic ratio. Patients had a clinical follow-up of at least 12 months, with progression-free survival at 6 months (6M-PFS) and overall survival at 12 months (12M-OS) as endpoints. To evaluate the models, we set aside 20% of this exploratory cohort as the final set. On the remaining 80%, we sorted the features according to their selection frequency by a LASSO logistic regression on extensive cross-validation. The feature selection’s performance was then studied by training a logistic ridge regression with various features for both outcomes. Finally, we built a model using the intersection of the most frequent features of both outcomes. The performance of this model using only 8 features was evaluated on the exploratory cohort but also on two external validation cohorts, using the area under the receiver operating characteristic (AUROC).

117 patients were included (93 for training and 24 for testing). The AUROC performance of the final model for the 6M-PFS (respectively the 12M-OS) was 74.31% (resp. 85.52%) on the test set of the exploratory cohort. On the external validation cohorts, the AUROC were respectively 88.50% and 82.14% for the 6M-FPS and 88.72% and 91.03% for the 12M-OS.

The combination of heterogeneous biomarkers provides a powerful model for predicting the outcome of NSCLC patients treated with immunotherapy. Rigorous feature selection is a critical point in ML approaches to avoid overfitting.

7.4.3 Stochastic Coherence Over Attention Trajectory For Continuous Learning In Video Streams

Participants: Alessandro Betti

Collaborations: Matteo Tiezzi, Simone Marullo, Lapo Faggi, Enrico Meloni, Stefano Melacci

Keywords: Online Learning, Attention, Motion

In [72] the paper we propose a novel neural-network-based approach to progressively and autonomously develop pixel-wise representations in a video stream. The proposed method is based on a human-like attention mechanism that allows the agent to learn by observing what is moving in the attended locations. Spatio-temporal stochastic coherence along the attention trajectory, paired with a contrastive term, leads to an unsuper-vised learning criterion that naturally copes with the considered setting. Differently from most existing works, the learned representations are used in open-set class-incremental classification of each frame pixel, relying on few supervisions. Our experiments leverage 3D virtual environments and they show that the proposed agents can learn to distinguish objects just by observing the video stream.

7.4.4 Robust Lasso-Zero for sparse corruption and model selection with missing covariates

Participants: Aude Sportisse

Keywords: incomplete data, informative missing values, Lasso-Zero, sparse corruptions, support recovery

Collaborations: Pascaline Descloux (Université of Geneva), Claire Boyer (Sorbonne Unviersité), Julie Josse (Inria Montpellier), Sylvain Sardy (University of Geneva)

In a first part, this paper [14] focuses on estimating the support of the parameter vector in the sparse corruption problem. An extension of the Lasso-Zero methodology, initially introduced for standard sparse linear models, is studied in the sparse corruptions problem. Theoretical guarantees on the sign recovery of the parameters are provided for a slightly simplified version of the estimator, called Thresholded Justice Pursuit This paper also handles the case where the matrix of covariates contains missing values, which can be due to manual errors, poor calibration, insufficient resolution, etc. In the high-dimensional setting, note that the naive complete case analysis, which discards all incomplete rows, is not an option, because the missingness of a single entry causes the loss of an entire row, which contains a lot of information when the number of covariates is large. Showing that missing values in the covariates can be reformulated into a sparse corruption problem, the Robust Lasso-Zero can be used for dealing with missing data. Numerical experiments and a medical application underline the relevance of Robust Lasso-Zero in such a context with few available competitors. The method is easy to use and implemented in the R library lass0.

7.4.5 R-miss-tastic: a unified platform for missing values methods and workflows

Participants: Aude Sportisse

Keywords: missing data; state of the art; bibliography; reproducibility; guided workflows; teaching material; statistical analysis community

Collaborations: Imke Mayer (Institute of Public Health, Berlin), Julie Josse (Inria Montpellier), Nicholas Tiernay (Monash University), Nathalie Vialaneix (INRAE)

Missing values are unavoidable when working with data. Their occurrence is exacerbated as more data from different sources become available. However, most statistical models and visualization methods require complete data, and improper handling of missing data results in information loss or biased analyses. Since the seminal work of Rubin, a burgeoning literature on missing values has arisen, with heterogeneous aims and motivations. This led to the development of various methods, formalizations, and tools. For practitioners, it remains nevertheless challenging to decide which method is most suited for their problem, partially due to a lack of systematic covering of this topic in statistics or data science curricula. To help address this challenge, we have launched the "R-miss-tastic" platform [19], which 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. Indeed, we have developed several pipelines 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




Figure 17: Probability of Sign Recovery with automatic tuning, for sparsity levels $s = 3$ and $s = 10$ (subplots columns), proportions of missing values 5% or 20% (subplots rows), and two missing data mechanisms (MCAR vs MNAR). Non-correlated case ($\rho = 0$).

their methods and search for an up-to-date bibliography, and also teachers who are looking for didactic materials (notebooks, video, slides).

	0.1 MCAR	0.3 MCAR	0.5 MCAR	0.1 MAR	0.3 MAR	0.5 MAR	0.1 MNAR	0.3 MNAR	0.5 MNAR
<i>X.pca</i>	0.74	0.76	0.78	0.75	0.78	0.81	0.76	0.78	0.81
<i>X.forest</i>	0.77	0.8	0.86	0.78	0.81	0.87	0.78	0.81	0.88
<i>X.mice</i>	0.82	0.83	0.86	0.83	0.86	0.9	0.84	0.87	0.9
<i>X.soft</i>	0.93	0.86	0.87	0.97	1	1.1	1	1.1	1.1
<i>X.mean</i>	1	0.99	1	1.1	1.1	1.1	1.2	1.1	1.1

Figure 18: Output of the function `how_to_impute` in R. The results are truncated to two digits.

7.4.6 Deep latent position model for node clustering in graphs

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

Keywords: Network Analysis, Graph Clustering, Latent Position Models, Graph Convolutional Networks

With the significant increase of interactions between individuals through numeric means, the clustering of vertex in graphs has become a fundamental approach for analysing large and complex networks. We propose here the deep latent position model (DeepLPM), an end-to-end clustering approach which combines the widely used latent position model (LPM) for network analysis with a graph convolutional network (GCN) encoding strategy. Thus, DeepLPM can automatically assign each node to its group without using any additional algorithms and better preserves the network topology. Numerical experiments on simulated data and an application on the Cora citation network are conducted to demonstrate its effectiveness and interest in performing unsupervised clustering tasks.

7.4.7 The graph embedded topic model

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

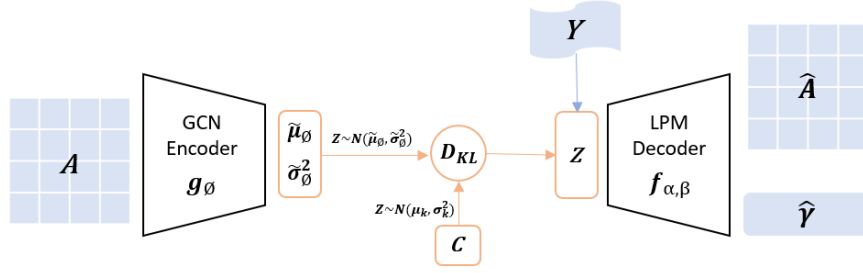


Figure 19: A deep-learning-like model view of DeepLPM.

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 this work, 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). 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.

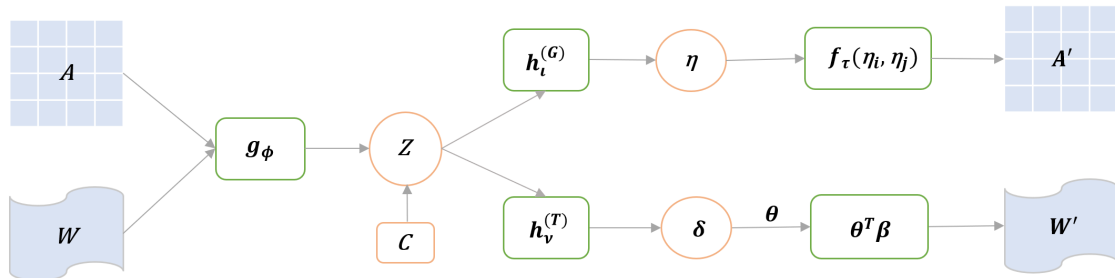


Figure 20: Model architecture of GETM.

7.4.8 Graph Neural Network for Graph Drawing

Participants: Gabriele Ciravegna, Marco Gori

Keywords: Graph Drawing, Graph Representation Learning, Graph Neural Drawers, Graph Neural Networks

Collaborations: Matteo Tiezzi (Università di Siena)

Visualizing complex relations and interaction patterns among entities is a crucial task, given the increasing interest in structured data representations. The Graph Drawing literature aims at developing algorithmic techniques to construct drawings of graphs for example via the node-link paradigm. The readability of graph layouts can be evaluated following some aesthetic criteria such as the number of

crossing edges, minimum crossing angles, community preservation, edge length variance, etc. The final goal is to find suitable coordinates for the node positions, and this often requires to explicitly express and combine these criteria through complicated mathematical formulations. Moreover, effective approaches such as energy-based models or spring-embedders require hands-on expertise and trial and error processes to achieve certain, desired visual properties. Additionally, such methods define loss or energy functions that must be optimized for each new graph to be drawn, often requiring to adapt algorithm-specific parameters. Lately, two interesting directions have emerged in the Graph Drawing community. The former one leverages the power of Gradient Descent to explore the manifold given by pre-defined loss functions or combinations of them. Stochastic Gradient Descent (SGD) can be used to move sub-samples of vertices couples in the direction of the gradient of spring-embedder losses, substituting complicated techniques such as Majorization. The latter novel direction consists in the exploitation of Deep Learning models. Indeed, the flexibility of neural networks and their approximation capability can come in handy also when dealing with the Graph Drawing scenario.

In [21], we propose a framework, Graph Neural Drawers (GND), which embraces both the aforementioned directions. We borrow the representational capability and computational efficiency of neural networks to prove that (1) differentiable loss functions guiding the common Graph Drawing pipeline can be provided directly by a neural network, a Neural Aesthete, even when the required aesthetic criteria cannot be directly optimized. In particular, we propose a proof-of-concept where we focus on the criteria of edge crossing, proving that a neural network can learn to identify if two arcs are crossing or not and provide a differentiable loss function towards non-intersection. Otherwise, in fact, this simple aesthetic criterion cannot be achieved through direct optimization, because it is non-differentiable. Instead, the Neural Aesthete provides a useful and flexible gradient direction that can be exploited by (Stochastic) Gradient Descent methods. Moreover, (2) we prove that GNNs, even in the non-attributed graph scenario if enriched with appropriate node positional features, can be used to process the topology of the input graph with the purpose of mapping the obtained node representation in a 2D layout. We compare various commonly used GNN models, proving that the proposed framework is flexible enough to give these models the ability to learn a wide variety of solutions. In particular, GND is capable to draw graphs (1) from supervised coordinates, i.e. emulating Graph Drawing Packages, (2) minimizing common aesthetic loss functions and, additionally, (3) by descending towards the gradient direction provided by the Neural Aesthete.

7.4.9 Semi-relaxed Gromov-Wasserstein divergence for graphs classification

Participants: Cédric Vincent-Cuaz, Marco Corneli

Keywords: Optimal Transport, Graph Dictionary Learning; Graph classification

Collaborations: Rémi Flamary, Titouan Vayer, Nicolas Courty

Comparing structured objects such as graphs is a fundamental operation involved in many learning tasks. To this end, the Gromov-Wasserstein (GW) distance, based on Optimal Transport (OT), has been successful in providing meaningful comparison between such entities. GW operates on graphs, seen as probability measures over spaces depicted by their nodes connectivity relations. At the core of OT is the idea of mass conservation, which imposes a coupling between all the nodes from the two considered graphs. We argue in [31] that this property can be detrimental for tasks such as graph dictionary learning (DL), and we relax it by proposing a new semi-relaxed Gromov-Wasserstein divergence. The latter leads to immediate computational benefits and naturally induces a new graph DL method, illustrated in Figure 21 shown to be relevant for unsupervised representation learning and classification of graphs.

7.4.10 Template based Graph Neural Network with Optimal Transport Distances

Participants: Cédric Vincent-Cuaz, Marco Corneli

Keywords: Graph Neural Network, Optimal Transport, Supervised Graph Representation Learning

Collaborations: Rémi Flamary, Titouan Vayer, Nicolas Courty

Current Graph Neural Networks (GNN) architectures generally rely on two important components: node features embedding through message passing, and aggregation with a specialized form of pooling. The structural (or topological) information is implicitly taken into account in these two steps. We propose in [30] a novel point of view, which places distances to some learnable graph templates at the core

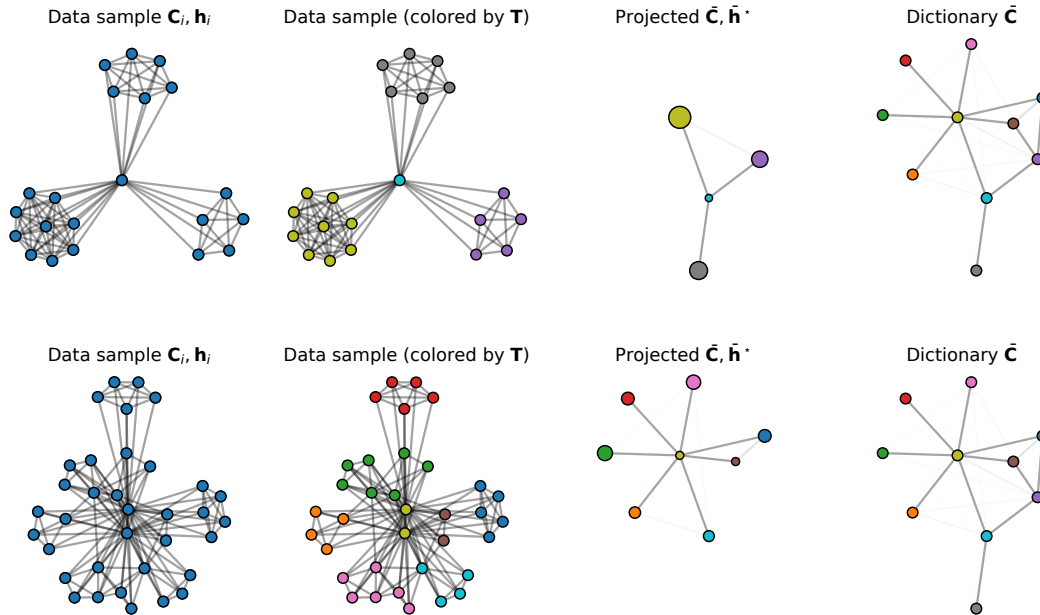


Figure 21: Illustration of the embedding of different graphs from the IMDB dataset on the estimated dictionary $\bar{\mathbf{C}}$. Each row corresponds to one observed graph and we show its graph (left), its graph with nodes colored corresponding to the OT plan (center left), the projected graph on the dictionary with optimal weight \bar{h}^* and the full dictionary with uniform mass (right).

of the graph representation. This distance embedding is constructed thanks to an optimal transport distance: the Fused Gromov-Wasserstein (FGW) distance, which encodes simultaneously feature and structure dissimilarities by solving a soft graph-matching problem. We postulate that the vector of FGW distances to a set of template graphs has a strong discriminative power, which is then fed to a non-linear classifier for final predictions. Distance embedding can be seen as a new layer, and can leverage on existing message passing techniques to promote sensible feature representations. Interestingly enough, in our work the optimal set of template graphs is also learnt in an end-to-end fashion by differentiating through this (TFGW) layer. The complete architecture of the model is illustrated in Figure 22. After describing the corresponding learning procedure, we empirically validate our claim on several synthetic and real life graph classification datasets, where our method is competitive or surpasses kernel and GNN state-of-the-art approaches. We complete our experiments by an ablation study and a sensitivity analysis to parameters.

7.4.11 Semi-relaxed Gromov-Wasserstein divergence with applications on graphs

Participants: Cédric Vincent-Cuaz, Marco Corneli

Keywords: Optimal Transport, Graph Dictionary Learning; Graph Partitioning; Graph Clustering; Graph Completion

Collaborations: Rémi Flamary, Titouan Vayer, Nicolas Courty

Comparing structured objects such as graphs is a fundamental operation involved in many learning tasks. To this end, the Gromov-Wasserstein (GW) distance, based on Optimal Transport (OT), has proven to be successful in handling the specific nature of the associated objects. More specifically, through the nodes connectivity relations, GW operates on graphs, seen as probability measures over specific spaces. At the core of OT is the idea of conservation of mass, which imposes a coupling between all the nodes from the two considered graphs. We argue in [29] that this property can be detrimental for tasks such as graph dictionary or partition learning, and we relax it by proposing a new semi-relaxed Gromov-Wasserstein divergence (Illustrated in Figure 23). Aside from immediate computational benefits, we discuss its properties, and show that it can lead to an efficient graph dictionary learning algorithm.

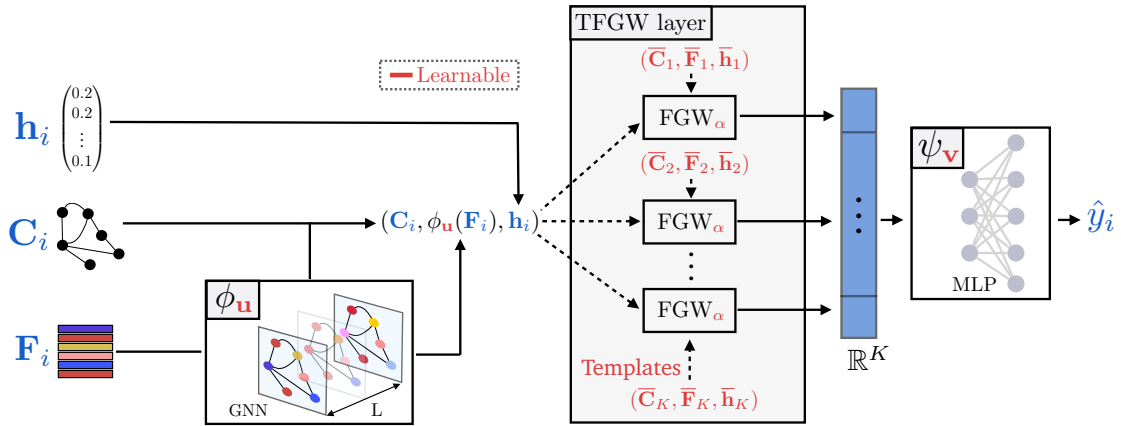


Figure 22: Illustration of the proposed model. **(left)** The input graph is represented as a triplet $(\mathbf{C}_i, \mathbf{F}_i, \mathbf{h}_i)$ where the matrix \mathbf{C}_i encodes the structure, \mathbf{F}_i the features, \mathbf{h}_i the nodes' weights. A GNN $\phi_{\mathbf{u}}$ is applied to the raw features in order to extract a meaningful node representations. **(center)** The TFGW layer is applied to the filtered graph and provides a vector representation as FGW distances to templates. **(right)** a final MLP $\psi_{\mathbf{v}}$ is applied to this vector in order to predict the final output of the model. All objects in red are parameters that are learned from the data.

We empirically demonstrate its relevance for complex tasks on graphs such as partitioning, clustering and completion.




Figure 23: Comparison of the GW matching (left) and asymmetric srGW matchings (middle and right) between graphs \mathbf{C} and $\bar{\mathbf{C}}$ with uniform distributions. Nodes of the source graph are colored based on their clusters. The OT from the source to the target nodes is represented by arcs colored depending on the corresponding source node color. The nodes in the target graph are colored by averaging the (rgb) color of the source nodes, weighted by the entries of the OT plan.

7.4.12 How to deal with missing data in supervised deep learning?

Participants: Pierre-Alexandre Mattei.

Keywords: deep learning, missing data

Collaborations: Niels Bruun Ipsen, Jes Frellsen

The issue of missing data in supervised learning has been largely overlooked, especially in the deep learning community. We investigate strategies to adapt neural architectures for handling missing values. Here, we focus on regression and classification problems where the features are assumed to be missing at

random. Of particular interest are schemes that allow reusing as-is a neural discriminative architecture. To address supervised deep learning with missing values, we propose in [25] to marginalize over missing values in a joint model of covariates and outcomes. Thereby, we leverage both the flexibility of deep generative models to describe the distribution of the covariates and the power of purely discriminative models to make predictions. More precisely, a deep latent variable model can be learned jointly with the discriminative model, using importance-weighted variational inference, essentially using importance sampling to mimic averaging over multiple imputations. In low-capacity regimes, or when the discriminative model has a strong inductive bias, we find that our hybrid generative/discriminative approach generally outperforms single imputations methods.

7.4.13 Model-based clustering with Missing Not At Random Data

Participants: Aude Sportisse.

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

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

With the increase of large datasets, the model-based clustering has become a very popular, flexible and interpretable methodology for data exploration in a well-defined statistical framework. However, in large scale data analysis, the problem of missing data is ubiquitous. We propose a novel approach by embedding missing data directly within model-based clustering algorithms. In particular, we consider the general case of Missing Not At Random (MNAR) values. We introduce in [71] a selection model for the joint distribution of data and missing-data indicator. It corresponds to a mixture model for the data distribution and a general MNAR model for the missing-data mechanism, for which the missingness may depend on the underlying classes (unknown) and/or the values of the missing variables themselves. A large set of meaningful MNAR sub-models is derived and the identifiability of the parameters is studied for each of the sub-models, which is usually a key issue for any MNAR proposals. The EM and Stochastic EM algorithms are considered for estimation. Finally, we perform empirical evaluations for the proposed sub-models on synthetic data (see e.g. Fig. 24) and we illustrate the relevance of our method on a medical register, the TraumaBase[®] dataset.




Figure 24: Relative effect on theoretical ARI of both the mixture component separation strength Δ_μ and the MNAR evidence Δ_{perc} ($\Delta_{perc} = 0$ is the MCAR case, the more $\Delta_{perc} = 0$ increases, the greater the deviation from the MCAR hypothesis.)

7.4.14 Unsupervised Text Clusterisation to characterize Adverse Drug Reactions from hospitalization reports

Participants: Michel Riveill, Xuchun Zhang.

Keywords: clustering, NLP, pharmacovigilance

Collaborations: Milou-Drici Daniel (Centre de Pharmacovigilance, Nice)

The detection of Adverse Drug Reactions (ADRs) in clinical records plays a pivotal role in pharmacovigilance (PhV). Achieving near-ideal practice relies on well-trained health professionals, who are trained to identify, assess, and report to health authorities ADRs occurring after drug marketing approval, including those that are infrequent. Despite being mandatory for health care probationers to report ADRs when suspected, notifications of ADRs amount to a mere 5-10 percent of all ADRs. However, the efficiency to detecting ADRs is limited due to the lack of well-trained professionals, the underreporting and the enormous amount of clinical reports at disposition.

Unsupervised learning can be a powerful resource in post-marketing pharmacovigilance, as it can exploit the big amount of data produced by daily trials of a larger populations and avoiding simultaneously the big cost of annotating data. In [46], we proposed a model (see fig. 25) using unsupervised learning technique to make use of modern text features extraction technique with BERT based models and explored the possibility of clustering ADR-related representations together in semantic space. The structure of our model is shown as in Figure 1. We applied our model on the well-formed electronic health records (EHRs) data in MADE 1.0 challenge dataset and also on the more casual text data of CADEC. The results indicate that with only contextual tokens as input, the model representation, especially those who obtained from domain-specific pretrained model like BioBERT, can be helpful in classifying ADR-related textual blocks with non-ADR blocks, especially for corpus like EHRs.

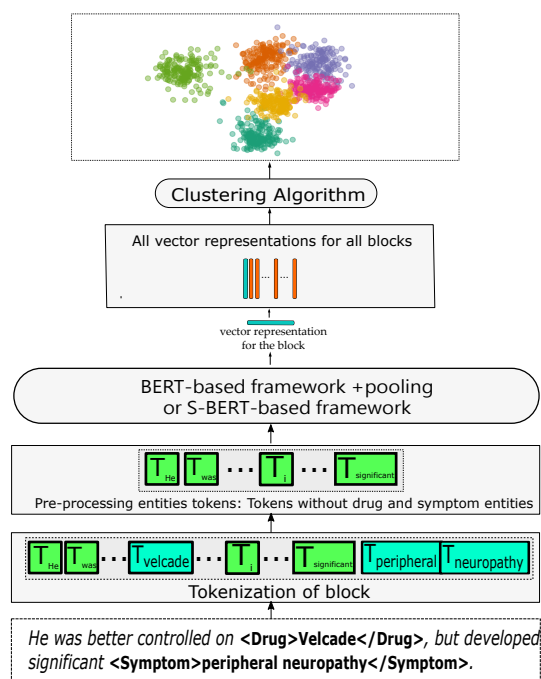


Figure 25: The structure of our model. For each text block, only the contextual tokens around drug/symptom named entities are selected for BERT-related embedding. For BERT-like models we applied a pooling strategy between tokens from the last layer to obtain a single vector representation for the block as the Sentence-BERT like models do, and then the output embedding vectors will go through the clustering algorithm to get the cluster assignment for each one of them.

7.4.15 Auto-encoder Based Medicare Fraud Detection

Participants: Mansour Zoubeirou A Mayaki, Michel Riveill

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

In this study, we used deep learning based multiple inputs classifier with a Long-short Term Memory (LSTM) autoencoder component to detect medicare fraud. The proposed model is made of two separate block: MLP block and auto encoder feature extraction block [47]. The MLP block extracts high level feature from the invoice data and the auto encoder block extracts high level features from data that describes the provider behavior over time. This architecture makes it possible to take into account many sources of data without mixing them. The latent features extracted from the LSTM autoencoder have a strong discriminating power and separate the providers into homogeneous clusters. We use the data sets from the Centers for Medicaid and Medicare Services (CMS) of the US federal government. Our results show that baseline artificial neural network give good performances compared to classical machine learning models but they are outperformed by our model.

7.4.16 Multiple Inputs Neural Networks for Fraud Detection

Participants: Mansour Zoubeirou A Mayaki, Michel Riveill

Keywords: Medicare fraud detection ; Anomaly detection ; Imbalanced data ; Machine learning ; Deep neural networks

This study aims to use artificial neural network based classifiers to predict fraud, particularly that related to health insurance. Medicare fraud results in considerable losses for governments and insurance companies and results in higher premiums from clients. Medicare fraud costs around 13 billion euros in Europe and between 21 billion and 71 billion US dollars per year in the United States. To detect medicare frauds, we propose a multiple inputs deep neural network based classifier with an autoencoder component [49]. This architecture makes it possible to take into account many sources of data without mixing them and makes the classification task easier for the final model. We use the data sets from the Centers for Medicaid and Medicare Services (CMS) of the US federal government and four benchmarks fraud detection data sets. Our results show that although baseline artificial neural network give good performances, they are outperformed by our multiple inputs neural networks. We have shown that using an autoencoder to embed the provider behavior gives better results and makes the classifiers more robust to class imbalance. The proposed method is described in Figure 26.

7.4.17 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 [59] 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 27 shows the evolving latent positions of a dynamic graph.

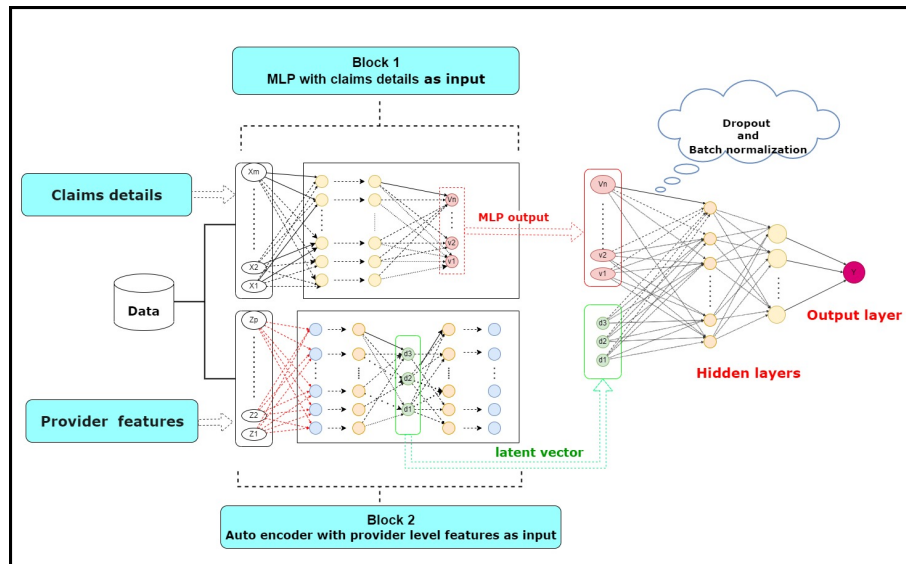


Figure 26: Visualization of the proposed neural network architecture. Block 1 receives features related to the invoice. Block 2 receives features related to the provider behavior and trains an autoencoder. The latent vector of the autoencoder and the output of block 1 are concatenated and used as input for the next hidden layers of the model.

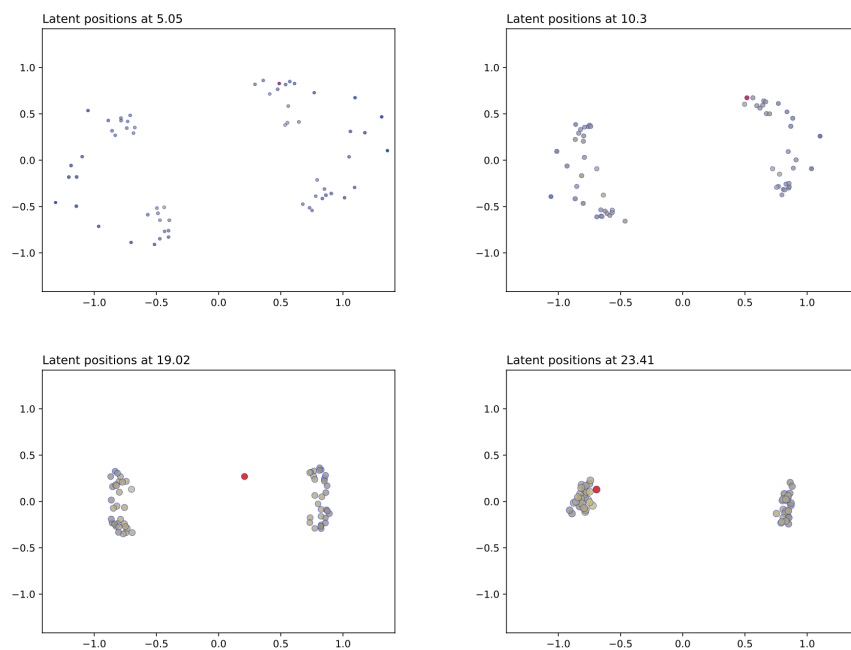


Figure 27: Snapshots of the evolving latent positions of a dynamic graph. Two communities emerge and a node (red) migrates between them.

7.4.18 DeepWILD: wildlife identification, localisation 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 analyse 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 [70], we propose to go even further by using object detection model to detect, classify and count species on camera traps videos. We developed a 3-parts process to analyse 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.

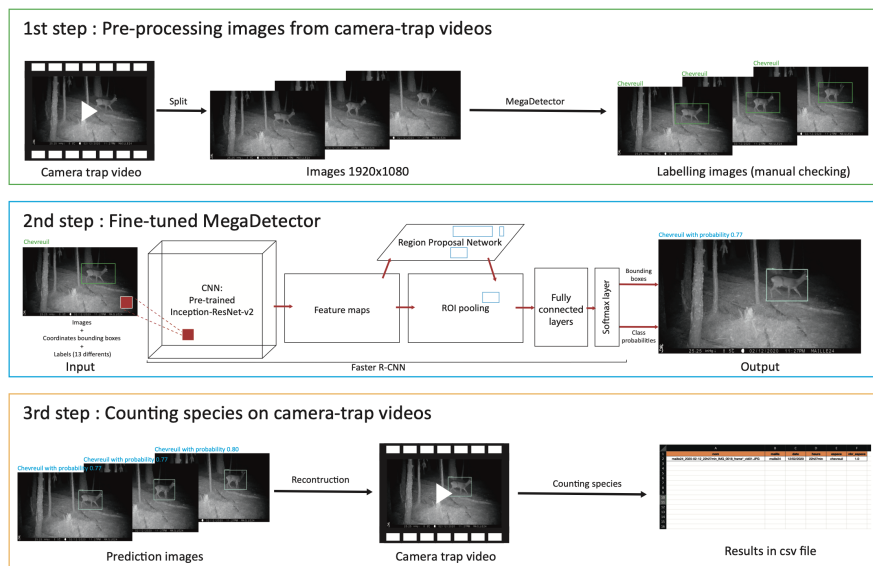


Figure 28: The DeepWILD process for the wildlife identification, localisation and population estimation from camera trap videos in the Parc National du Mercantour

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:

- **NXP:** This collaboration contract is a France Relance contract. Drift detection and predictive maintenance. Participant: Mansour Mayaki Zoubeirou, Michel Riveill. External participants: . Amount: 45 000€.

- **Orange:** it is a CIFRE build upon the PhD of Gatien Caillet on decentralized and efficient federated AutoML learning for heterogeneous embedded devices. Participant: Vincent Vandewalle. External participants: Tamara Tosic (Orange), Frédéric Guyard (Orange). Amount: 30 000€.
- **Naval Group:** The goal of this project will be the development of an open-source Python library for semi-supervised learning, via the hiring of a research engineer, Lucas Boiteau. Participant: Pierre-Alexandre Mattei, Hugo Schmutz, Aude Sportisse. External participants: Alexandre Gense, Quentin Oliveau (Naval Group). Amount: 125 000€.
- **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. Participants: Hugo Miralles, Michel Riveill. External participants: Tamara Tosic (Orange), Thierry Nagellen (Orange). Amount: 45 000€.
- **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€.
- **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). Participant: Michel Riveill, Frédéric Precioso. External participants: Amosse Edouard. Amount: 45 000€.
- **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. Participant: C. Bouveyron. External participants: F. Simoes, J. Jacques. Amount: 50 000€.

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: Frederic 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: Pierre-Alexandre Mattei, Frederic Precioso, Louis Ohl (doctorant)

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: Lucile Sassatelli and Frédéric Precioso (UCA)

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

Collaborations: 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

9.4 Regional initiatives

Parc National du Mercantour

Participants: Charles Bouveyron, Frédéric Precioso and Fanny Simoës

Keywords: Deep learning, image recognition,

Collaborators: 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: Damien Garreau, Frédéric Precioso

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 1st Nice Workshop on Interpretability, organised 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 17-18, 2022 in Nice, and counted 6 senior research talks, and 11 young research talks, with about 40 participants. Web : <https://sites.google.com/view/nwi2022/home>
- Statlearn 2022: 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. In 2020, a one-week spring-school should have been held in Cargèse (March, 23-27), but has been postponed in 2022 (April, 4-7, 2022) due to the pandemic crisis. The 2022 edition was the 11th edition of the Statlearn series and welcomed about 50 participants. 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>

- GenU 2022: 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 2022 edition was on September 14-15, 2022 (Web: <https://genu.ai/2022/>).
- 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. Frédéric Precioso was a member of the scientific committee in 2020 and 2021. P.-A. Mattei was a member of the scientific committee in 2022. Web: <https://univ-cotedazur.eu/events/sophia-summit>.
- The Deep Learning School @UCA, organized by Frédéric Precioso since 2017, took place in July 2020 and July 2021 (with a pause in 2022 after two years of organization under Covid conditions). The summer school is gathering some of the top scientists in Deep Learning for giving lectures of about 3 hours on a specific topic in Deep Learning every day of the week, then the other half of each day is dedicated to three hours of practice session on the topic of the lecture. The practice sessions are made by Maasai PhDs and post-docs under F. Precioso supervision. In 2020, and in 2021, for remote editions about 300 participants from both academia and industry attended each year the lectures. Web: <https://univ-cotedazur.fr/deep-learning-school/homepage>.

10.1.2 Journal

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
- 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 Invited talks

- Charles Bouveyron was invited for a keynote talk at the [17th conference of the International Federation of Classification Societies](#).

10.1.4 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.

10.1.5 Scientific expertise

- Frédéric Precioso is the Scientific Responsible and Program Officer for AI at the French Research Agency (ANR) since September 2019. He is 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

- C. Bouveyron and S. Petiot participated to a series of 4 articles in the journal Le Monde (editions of 1st and 2nd April 2022) about the analysis of the 2022 French Presidential election on Twitter. Web: <https://www.lemonde.fr/politique/article/2022/03/31/presidentielle...>
- 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 2022. 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 visualise 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”.
- F. Precioso gave webinar - Master class on "Intelligence artificielle : tout ce que vous avez toujours voulu savoir, avancées et cas d’usage", October, the 27th, 2022, for BPI Université and INRIA Academy. About 300 companies attended live, and about 800 accessed the replay.
- F. Precioso gave a conference to the Université Nice Inter-Ages (UNIA) on "Les différents systèmes d’IA et leurs applications. Un regard décomplexé." in a series of conferences to explain AI to the elderly.
- F. Precioso gave a conference - Master Class "Tout ce que vous avez toujours voulu savoir sur l’IA sans jamais oser le demander" at the Semaine de l’industrie 2022 event in Grasse BIOTECH for local companies.
- C. Bouveyron participated in several TV (AzurTV, B-SMART TV, AzurIA, ...) and newspaper (Le Monde, AEF, NewsThank, ActuIA, ...) interviews to promote the actions in AI of the team and the Institut 3IA Côte d’Azur, or to provide his view on the recent evolution of AI.
- F. Precioso also participated in TV shows (Lumni TV, l’Entre-Deux, “L’intelligence artificielle - Vidéo Culture générale - Lumni”; France 3 Regional news for “La semaine du cerveau”, 17 March 2022) to promote the actions of the team in AI or to provide his view on the recent and future evolution of AI.

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 April 2022.
- P.-A. Mattei gave (with Serena Villata) an Introductory lecture on AI to the 25th “Ecole Jeunes Chercheurs et Chercheurs en Informatique Mathématique” (<https://ejcim2022.sciencesconf.org/>), a summer school dedicated to graduate students in theoretical computer science. All lectures of the school were published as a book by the CNRS [52].

11 Scientific production

11.1 Major publications

- [1] F. Bergamin, P.-A. Mattei, J. Havtorn, H. Senetaire, H. Schmutz, L. Maaløe, S. Hauberg and J. Frellsen. ‘Model-agnostic out-of-distribution detection using combined statistical tests’. In: AISTATS 2022 - 25th International Conference on Artificial Intelligence and Statistics. Vol. 151. Valence, Spain, 30th Mar. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03600783>.
- [2] C. Bouveyron, J. Jacques, A. Schmutz, F. Simoes and S. Bottini. ‘Co-Clustering of Multivariate Functional Data for the Analysis of Air Pollution in the South of France’. In: *Annals of Applied Statistics* 16.3 (2022), pp. 1400–1422. DOI: [10.1214/21-A0AS1547](https://doi.org/10.1214/21-A0AS1547). URL: <https://hal.archives-ouvertes.fr/hal-02862177>.
- [3] N. B. Ipsen, P.-A. Mattei and J. Frellsen. ‘How to deal with missing data in supervised deep learning?’ In: ICLR 2022 - 10th International Conference on Learning Representations. International Conference on Learning Representations. Virtual conference, France, 25th Apr. 2022. URL: <https://hal.inria.fr/hal-03044144>.
- [4] G. Marchello, A. Fresse, M. Corneli and C. Bouveyron. ‘Co-clustering of evolving count matrices with the dynamic latent block model: application to pharmacovigilances’. In: *Statistics and Computing* 32.41 (2022). DOI: [10.1007/s11222-022-10098-y](https://doi.org/10.1007/s11222-022-10098-y). URL: <https://hal.archives-ouvertes.fr/hal-03146769>.

- [5] L. Ohl, P.-A. Mattei, C. Bouveyron, W. Harchaoui, M. Leclercq, A. Droit and F. Precioso. ‘Generalised Mutual Information for Discriminative Clustering’. In: *Advances in Neural Information Processing Systems*. NeurIPS - Thirty-sixth Conference on Neural Information Processing Systems. Vol. 36. New Orleans, United States, 28th Nov. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03815229>.
- [6] M. Tiezzi, G. Ciravegna and M. Gori. ‘Graph Neural Networks for Graph Drawing’. In: *IEEE Transactions on Neural Networks and Learning Systems* (2022). DOI: [10.1109/TNNLS.2022.3184967](https://doi.org/10.1109/TNNLS.2022.3184967). URL: <https://hal.archives-ouvertes.fr/hal-03854543>.
- [7] C. Vincent-Cuaz, R. Flamary, M. Corneli, T. Vayer and N. Courty. ‘Semi-relaxed Gromov Wasserstein divergence with applications on graphs’. In: *ICLR 2022 - 10th International Conference on Learning Representations*. Virtual, France, 25th Apr. 2022, pp. 1–28. URL: <https://hal.archives-ouvertes.fr/hal-03832589>.
- [8] C. Vincent-Cuaz, R. Flamary, M. Corneli, T. Vayer and N. Courty. ‘Template based Graph Neural Network with Optimal Transport Distances’. In: *NeurIPS 2022 – 36th Conference on Neural Information Processing Systems*. New Orleans, United States, 2022. URL: <https://hal.archives-ouvertes.fr/hal-03839517>.
- [9] M. E. Zarlenga, P. Barbiero, G. Ciravegna, G. Marra, F. Giannini, M. Diligenti, F. Precioso, S. Melacci, A. Weller, P. Lio and M. Jamnik. ‘Concept Embedding Models’. In: *NeurIPS 2022 - 36th Conference on Neural Information Processing Systems*. New Orleans, United States, 28th Nov. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03854550>.

11.2 Publications of the year

International journals

- [10] A. Bobasheva, F. Gandon and F. Precioso. ‘Learning and Reasoning for Cultural Metadata Quality’. In: *Journal on Computing and Cultural Heritage* 15.3 (Sept. 2022). DOI: [10.1145/3485844](https://doi.org/10.1145/3485844). URL: <https://hal.archives-ouvertes.fr/hal-03363442>.
- [11] C. Bouveyron, J. Jacques, A. Schmutz, F. Simoes and S. Bottini. ‘Co-Clustering of Multivariate Functional Data for the Analysis of Air Pollution in the South of France’. In: *Annals of Applied Statistics* 16.3 (2022), pp. 1400–1422. DOI: [10.1214/21-A0AS1547](https://doi.org/10.1214/21-A0AS1547). URL: <https://hal.archives-ouvertes.fr/hal-02862177>.
- [12] 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.archives-ouvertes.fr/hal-03862456>.
- [13] V. Comte, H. Schmutz, D. Chardin, F. Orlhac, J. Darcourt and O. Humbert. ‘Development and validation of a radiomic model for the diagnosis of dopaminergic denervation on [18F]FDOPA PET/CT’. In: *European Journal of Nuclear Medicine and Molecular Imaging* 49.11 (Sept. 2022), pp. 3787–3796. DOI: [10.1007/s00259-022-05816-7](https://doi.org/10.1007/s00259-022-05816-7). URL: <https://hal.archives-ouvertes.fr/hal-03839299>.
- [14] P. Descloux, C. Boyer, J. Josse, A. Sportisse and S. Sardy. ‘Robust Lasso-Zero for sparse corruption and model selection with missing covariates’. In: *Scandinavian Journal of Statistics* (2022). URL: <https://hal.archives-ouvertes.fr/hal-02569696>.
- [15] M. Fop, P.-A. Mattei, C. Bouveyron and T. B. Murphy. ‘Unobserved classes and extra variables in high-dimensional discriminant analysis’. In: *Advances in Data Analysis and Classification* 16 (Mar. 2022), pp. 55–92. DOI: [10.1007/s11634-021-00474-3](https://doi.org/10.1007/s11634-021-00474-3). URL: <https://hal.archives-ouvertes.fr/hal-03132362>.
- [16] R. Khouja, P.-A. Mattei and B. Mourrain. ‘Tensor decomposition for learning Gaussian mixtures from moments’. In: *Journal of Symbolic Computation* 113 (2022), pp. 193–210. DOI: [10.1016/j.jsc.2022.04.002](https://doi.org/10.1016/j.jsc.2022.04.002). URL: <https://hal.archives-ouvertes.fr/hal-03244448>.

- [17] G. Marchello, A. Fresse, M. Corneli and C. Bouveyron. ‘Co-clustering of evolving count matrices with the dynamic latent block model: application to pharmacovigilances’. In: *Statistics and Computing* 32.41 (2022). DOI: [10.1007/s11222-022-10098-y](https://doi.org/10.1007/s11222-022-10098-y). URL: <https://hal.archives-ouvertes.fr/hal-03146769>.
- [18] A. Mathieu, M. Leclercq, M. Sanabria, O. Perin and A. Droit. ‘Machine Learning and Deep Learning Applications in Metagenomic Taxonomy and Functional Annotation’. In: *Frontiers in Microbiology* 13 (14th Mar. 2022), p. 811495. DOI: [10.3389/fmicb.2022.811495](https://doi.org/10.3389/fmicb.2022.811495). URL: <https://hal.inria.fr/hal-03622551>.
- [19] I. Mayer, A. Sportisse, N. Tierney, N. Vialaneix and J. Josse. ‘R-miss-tastic: a unified platform for missing values methods and workflows’. In: *The R Journal* (30th July 2022). URL: <https://hal.archives-ouvertes.fr/hal-02879337>.
- [20] F. Simões, C. Bouveyron, D. Piga, D. Borel, S. Descombes, V. Paquis-Flucklinger, J. Levrant, P. Gibelin and S. Bottini. ‘Cardiac dyspnea risk zones in the South of France identified by geo-pollution trends study’. In: *Scientific Reports* (2022). URL: <https://hal.archives-ouvertes.fr/hal-03519068>.
- [21] M. Tiezzi, G. Ciravegna and M. Gori. ‘Graph Neural Networks for Graph Drawing’. In: *IEEE Transactions on Neural Networks and Learning Systems* (2022). DOI: [10.1109/TNNLS.2022.3184967](https://doi.org/10.1109/TNNLS.2022.3184967). URL: <https://hal.archives-ouvertes.fr/hal-03854543>.
- [22] T. U. Tran, H. T. Thi Hoan, P. H. Dang and M. Riveill. ‘Toward a multitask aspect-based sentiment analysis model using deep learning’. In: *IAES International Journal of Artificial Intelligence (IJ-AI)* 11.2 (1st June 2022), p. 9. DOI: [10.11591/ijai.v11.i2.pp516-524](https://doi.org/10.11591/ijai.v11.i2.pp516-524). URL: <https://hal.inria.fr/hal-03766855>.
- [23] T. Yang, N. Pasquier and F. Precioso. ‘Semi-supervised Consensus Clustering Based on Closed Patterns’. In: *Knowledge-Based Systems*. Article 107599 235 (10th Jan. 2022). DOI: [10.1016/j.knsys.2021.107599](https://doi.org/10.1016/j.knsys.2021.107599). URL: <https://hal.archives-ouvertes.fr/hal-03402517>.

International peer-reviewed conferences

- [24] S. d’Angelo, F. Precioso and F. Gandon. ‘Revisiting Artistic Style Transfer for Data Augmentation in A Real-Case Scenario’. In: *IEEE ICIP 2022 - 29th IEEE International Conference on Image Processing*. Bordeaux, France: IEEE, 16th Oct. 2022, pp. 4178–4182. DOI: [10.1109/ICIP46576.2022.9897728](https://doi.org/10.1109/ICIP46576.2022.9897728). URL: <https://hal.archives-ouvertes.fr/hal-03921565>.
- [25] N. B. Ipsen, P.-A. Mattei and J. Frellsen. ‘How to deal with missing data in supervised deep learning?’ In: *ICLR 2022 - 10th International Conference on Learning Representations*. International Conference on Learning Representations. Virtual conference, France, 25th Apr. 2022. URL: <https://hal.inria.fr/hal-03044144>.
- [26] B. Ly, S. Finsterbach, M. Nuñez-Garcia, P. Jaïs, D. Garreau, H. Cochet and M. Sermesant. ‘Interpretable Prediction of Post-Infarct Ventricular Arrhythmia using Graph Convolutional Network’. In: *STACOM 2022 - 13th Workshop on Statistical Atlases and Computational Modelling of the Heart*. Singapore, Singapore, 18th Sept. 2022. URL: <https://hal.inria.fr/hal-03829609>.
- [27] L. Ohl, P.-A. Mattei, C. Bouveyron, W. Harchaoui, M. Leclercq, A. Droit and F. Precioso. ‘Generalised Mutual Information for Discriminative Clustering’. In: *Advances in Neural Information Processing Systems*. NeurIPS - Thirty-sixth Conference on Neural Information Processing Systems. Vol. 36. New Orleans, United States, 28th Nov. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03815229>.
- [28] T. U. Tran, H. T. T. Hoang, P. H. Dang and M. Riveill. ‘Aspect-Based Sentiment Analysis with Deep Learning: A Multidomain and Multitask Approach’. In: *Intelligence of Things: Technologies and Applications*. ICIT 2022 - The First International Conference on Intelligence of Things. Vol. LNDECT-148. Lecture Notes on Data Engineering and Communications Technologies. Hanoi, Vietnam: Springer International Publishing, 23rd Aug. 2022, pp. 134–145. DOI: [10.1007/978-3-031-15063-0_12](https://doi.org/10.1007/978-3-031-15063-0_12). URL: <https://hal.archives-ouvertes.fr/hal-03844120>.

- [29] C. Vincent-Cuaz, R. Flamary, M. Corneli, T. Vayer and N. Courty. ‘Semi-relaxed Gromov Wasserstein divergence with applications on graphs’. In: ICLR 2022 - 10th International Conference on Learning Representations. Virtual, France, 25th Apr. 2022, pp. 1–28. URL: <https://hal.archives-ouvertes.fr/hal-03832589>.
- [30] C. Vincent-Cuaz, R. Flamary, M. Corneli, T. Vayer and N. Courty. ‘Template based Graph Neural Network with Optimal Transport Distances’. In: NeurIPS 2022 – 36th Conference on Neural Information Processing Systems. New Orleans, United States, 2022. URL: <https://hal.archives-ouvertes.fr/hal-03839517>.

National peer-reviewed Conferences

- [31] C. Vincent-Cuaz, R. Flamary, M. Corneli, T. Vayer and N. Courty. ‘Semi-relaxed Gromov-Wasserstein divergence for graphs classification’. In: Colloque GRETSI 2022 - XXVIIIème Colloque Francophone de Traitement du Signal et des Images. Nancy, France, 6th Sept. 2022. URL: <https://hal.science/hal-03839524>.

Conferences without proceedings

- [32] Y. E. Amraoui, M. Blay-Fornarino, P. Collet, F. Precioso and J. Muller. ‘Evolvable SPL management with partial knowledge: an application to anomaly detection in time series’. In: SPLC 2022 - 26th ACM International Systems and Software Product Line Conference. Vol. A. Graz, Austria, 12th Sept. 2022, pp. 222–233. DOI: [10.1145/3546932.3547008](https://doi.org/10.1145/3546932.3547008). URL: <https://hal.archives-ouvertes.fr/hal-03811038>.
- [33] P. Barbiero, G. Ciravegna, F. Giannini, P. Lió, M. Gori and S. Melacci. ‘Entropy-based Logic Explanations of Neural Networks’. In: AAAI-22. Vancouver, United States, 22nd Feb. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03507746>.
- [34] F. Bergamin, P-A. Mattei, J. Havtorn, H. Senetaire, H. Schmutz, L. Maaløe, S. Hauberg and J. Frellsen. ‘Model-agnostic out-of-distribution detection using combined statistical tests’. In: AISTATS 2022 - 25th International Conference on Artificial Intelligence and Statistics. Vol. 151. Valence, Spain, 30th Mar. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03600783>.
- [35] A. Betti, L. Faggi, M. Gori, M. Tiezzi, S. Marullo, E. Meloni and S. Melacci. ‘Continual Learning through Hamilton Equations’. In: CoLLAs 2022 - Conference on Lifelong Learning Agents. Montreal, Canada, 22nd Aug. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03881404>.
- [36] O. Carrillo, C. J. B. Hernandez, F. Le Mouel, H. Barrera, Y. Denneulin, J. T. Hernandez, F. J. Vargas, L. X. B. Roza, C. Roncancio and M. Riveill. ‘Scalable Architectures to Support Sustainable Advanced Information Technologies’. In: CLUSTER 2022 - IEEE International Conference on Cluster Computing. Heidelberg, Germany: IEEE, 5th Sept. 2022, pp. 512–515. DOI: [10.1109/CLUSTER51413.2022.00065](https://doi.org/10.1109/CLUSTER51413.2022.00065). URL: <https://hal.archives-ouvertes.fr/hal-03844189>.
- [37] R. Jain, G. Ciravegna, P. Barbiero, F. Giannini, D. Buffelli and P. Lio. ‘Extending Logic Explained Networks to Text Classification’. In: Empirical Methods in Natural Language Processing. Abu Dhabi, United Arab Emirates, 7th Dec. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03863362>.
- [38] D. Liang, M. Corneli, C. Bouveyron and P. Latouche. ‘Deep latent position model for node clustering in graphs’. In: The 30th European Symposium on Artificial Neural Networks (ESANN 2022). Bruges, Belgium, 5th Oct. 2022. URL: <https://hal.science/hal-03874698>.
- [39] G. Lopardo and D. Garreau. ‘Comparing Feature Importance and Rule Extraction for Interpretability on Text Data’. In: XAIE 2022 - 2-nd workshop on Explainable and Ethical AI ICPR 2022. Montreal, Canada, 21st Aug. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03828177>.
- [40] G. Lopardo, D. Garreau, F. Precioso and G. Ottosson. ‘SMACE: A New Method for the Interpretability of Composite Decision Systems’. In: European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases 2022. Grenoble, France, 19th Sept. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03527129>.

- [41] S. Marullo, M. Tiezzi, A. Betti, L. Faggi, E. Meloni and S. Melacci. ‘Continual Unsupervised Learning for Optical Flow Estimation with Deep Networks’. In: *CoLLAs 2022 - Conference on Lifelong Learning Agents*. Montreal, Canada, 22nd Aug. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03881291>.
- [42] H. Schmutz, P.-A. Mattei, S. Contu, D. Chardin and O. Humbert. ‘18FDG PET/CT and Machine Learning for the prediction of lung cancer response to immunotherapy’. In: *European Journal of Nuclear Medicine and Molecular Imaging*. he 35th Annual EANM Congres - Annual Congress of the European Association of Nuclear Medicine. Vol. 49. S1. Barcelona, Spain, Sept. 2022, p. 245. DOI: 10.1007/s00259-022-05924-4. URL: <https://hal.archives-ouvertes.fr/hal-03876074>.
- [43] M. Tiezzi, S. Marullo, A. Betti, E. Meloni, L. Faggi, M. Gori and S. Melacci. ‘Foveated Neural Computation’. In: *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases - ECML PKDD 2022*. Grenoble, France, 19th Sept. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03881258>.
- [44] V. Vandewalle, G. Babykina, J.-B. Beuscart, J. A. Carretero Bravo and F. Visade. ‘Clustering of recurrent events data applied to the re-admission of elderly people at hospital’. In: *European Conference on Data Analysis 2022*. Naples, Italy, 14th Sept. 2022. URL: <https://hal.inria.fr/hal-03899727>.
- [45] M. E. Zarlenga, P. Barbiero, G. Ciravegna, G. Marra, F. Giannini, M. Diligenti, F. Precioso, S. Melacci, A. Weller, P. Lio and M. Jamnik. ‘Concept Embedding Models’. In: *NeurIPS 2022 - 36th Conference on Neural Information Processing Systems*. New Orleans, United States, 28th Nov. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03854550>.
- [46] X. Zhang, M.-D. Drici and M. Riveill. ‘Unsupervised Text Clusterisation to characterize Adverse Drug Reactions from hospitalization reports’. In: *4th International Conference on Advances in Signal Processing and Artificial Intelligence (ASPAI 2022)*. Corfu, France, 19th Oct. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03847229>.
- [47] M. Zoubeirou a Mayaki and M. Riveill. ‘Auto-encoder Based Medicare Fraud Detection’. In: *ASPAI 2022 - 4th International Conference on Advances in Signal Processing and Artificial Intelligence*. CORFOU, Greece, 19th Oct. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03847301>.
- [48] M. Zoubeirou a Mayaki and M. Riveill. ‘Autoregressive based Drift Detection Method’. In: *IEEE WCCI 2022 - IEEE world congress on computational intelligenceWORLD CONGRESS ON COMPUTATIONAL INTELLIGENCE*. Padoue, Italy, 21st July 2022. URL: <https://hal.archives-ouvertes.fr/hal-03740180>.
- [49] M. Zoubeirou a Mayaki and M. Riveill. ‘Multiple Inputs Neural Networks for Fraud Detection’. In: *The 2022 International Conference on Machine Learning, Control, and Robotics (MLCR 2022)*. Suzhou, China, 29th Oct. 2022. URL: <https://hal.science/hal-03607722>.

Scientific books

- [50] A. Betti, M. Gori and S. Melacci. *Deep Learning to See: Towards New Foundations of Computer Vision*. SpringerBriefs in Computer Science. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03528028>.

Scientific book chapters

- [51] R. G. Hussain, M. Tiezzi, G. Ciravegna and M. Gori. ‘Minimizing Cross Intersections in Graph Drawing via Linear Splines’. In: *Artificial Neural Networks in Pattern Recognition*. Vol. 13739. Lecture Notes in Computer Science. Springer International Publishing, 11th Nov. 2022, pp. 28–39. DOI: 10.1007/978-3-031-20650-4_3. URL: <https://hal.archives-ouvertes.fr/hal-03854553>.
- [52] P.-A. Mattei and S. Villata. ‘Introduction à l’intelligence artificielle et aux modèles génératifs’. In: *Informatique Mathématique: Une photographie en 2022*. CNRS Editions, 2022. URL: <https://hal.archives-ouvertes.fr/hal-03849387>.

Reports & preprints

- [53] G. Bellettini, A. Betti and M. Paolini. *A free boundary singular transport equation as a formal limit of a discrete dynamical system*. 28th Nov. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03874836>.
- [54] A. Betti, M. Casoni and M. Gori. *Forward Approximate Solution for Linear Quadratic Tracking*. 30th Dec. 2022. URL: <https://hal.science/hal-03939954>.
- [55] A. Betti, M. Gori, S. Melacci, M. Pelillo and F. Roli. *Can machines learn to see without visual databases?* 28th Nov. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03874787>.
- [56] R. Boutin, P. Latouche and C. Bouveyron. *Embedded Topics in the Stochastic Block Model*. 21st Sept. 2022. DOI: [10.48550/arXiv.2209.10097](https://doi.org/10.48550/arXiv.2209.10097). URL: <https://hal.science/hal-03782528>.
- [57] G. Ciravegna, F. Precioso and M. Gori. *KNOWLEDGE-DRIVEN ACTIVE LEARNING*. 14th Jan. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03526510>.
- [58] G. Ciravegna, F. Precioso and M. Gori. *Knowledge-driven Active Learning*. 15th Jan. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03527128>.
- [59] M. Corneli and R. Rastelli. *Continuous Latent Position Models for Instantaneous Interactions*. 7th Jan. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03517392>.
- [60] D. Garreau. *How to scale hyperparameters for quickshift image segmentation*. 7th Nov. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03842452>.
- [61] Y. Gu, E. Erosheva, G. Xu and D. Dunson. *Dimension-Grouped Mixed Membership Models for Multivariate Categorical Data*. 12th Jan. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03522275>.
- [62] D. Liang, M. Corneli, C. Bouveyron and P. Latouche. *Clustering by Deep Latent Position Model with Graph Convolutional Network*. 4th Apr. 2022. URL: <https://hal.science/hal-03629104>.
- [63] G. Lopardo, D. Garreau and F. Precioso. *A Sea of Words: An In-Depth Analysis of Anchors for Text Data*. 15th June 2022. URL: <https://hal.archives-ouvertes.fr/hal-03696324>.
- [64] G. Marchello, M. Corneli and C. Bouveyron. *A Deep Dynamic Latent Block Model for the Co-clustering of Zero-Inflated Data Matrices*. 6th Oct. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03800210>.
- [65] P.-A. Mattei and J. Frellsen. *Uphill Roads to Variational Tightness: Monotonicity and Monte Carlo Objectives*. 17th Oct. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03817920>.
- [66] E. Meloni, L. Faggi, S. Marullo, A. Betti, M. Tiezzi, M. Gori and S. Melacci. *PARTIME: Scalable and Parallel Processing Over Time with Deep Neural Networks*. 28th Nov. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03874858>.
- [67] M. Sanabria, T. Menguy, F. Precioso and P.-A. Mattei. *A Multi-stage deep architecture for summary generation of soccer videos*. 9th May 2022. URL: <https://hal.inria.fr/hal-03662328>.
- [68] H. Schmutz, O. Humbert and P.-A. Mattei. *Don't fear the unlabelled: safe deep semi-supervised learning via simple debiasing*. 16th Mar. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03610272>.
- [69] H. H. J. Senetaire, D. Garreau, J. Frellsen and P.-A. Mattei. *Explainability as statistical inference*. 7th Dec. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03888528>.
- [70] F. Simões, C. Bouveyron and F. Precioso. *DeepWILD: Wildlife Identification, Localisation and estimation on camera trap videos using Deep learning*. 1st Oct. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03797530>.
- [71] A. Sportisse, M. Marbac, C. Biernacki, C. Boyer, G. Celeux, F. Laporte and J. Josse. *Model-based Clustering with Missing Not At Random Data*. 2022. URL: <https://hal.science/hal-03494674>.
- [72] M. Tiezzi, S. Marullo, L. Faggi, E. Meloni, A. Betti and S. Melacci. *Stochastic Coherence Over Attention Trajectory For Continuous Learning In Video Streams*. 28th Nov. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03874819>.

- [73] M. Zoubeirou a Mayaki and M. Riveill. *Autoregressive based Drift Detection Method*. 14th Mar. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03607737>.

Other scientific publications

- [74] X. Zhang, M. D. Drici and M. Riveill. ‘Unsupervised Text Clusterisation to characterize Adverse Drug Reactions from hospitalization reports’. In: ECML PKDD 2022 - European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases. Grenoble, France, 19th Sept. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03794158>.

11.3 Cited publications

- [75] Y. LeCun, Y. Bengio and G. Hinton. ‘Deep learning’. In: *nature* 521.7553 (2015), pp. 436–444.