

Université de Bretagne Occidentale

École Doctorale MathSTIC

Sciences et Technologies de l'Information et de la Communication

Habilitation to Supervise Research

Presented by:

Bastien PASDELOUP

TRANSFERRING ADVANCES FROM MODERN AI TO HEALTHCARE

Defended publicly on **November 12th, 2026** before the jury composed of:

Reviewers

Rémi BARDENET, Senior Researcher, HDR
CRISAL, CNRS & University of Lille

David PICARD, Senior Researcher, HDR
LIGM, CNRS

Dorina THANOU, Senior Researcher
LTS4, EPFL

Examiners

Nicolas KERIVEN, Researcher, HDR
CNRS, IRISA

Gwenolé QUELLEC, Senior Researcher, HDR
LaTIM, Inserm & UBO

François ROUSSEAU, Professor, HDR
LaTIM, Inserm & IMT Atlantique

Serge TIMSIT, PU-PH, HDR
Inserm & UBO

TABLE OF CONTENTS



Foreword	5
TL;DR	5
Formatting of This Manuscript	5
Global Vision of My Research	6
Other Things You Will Find in This Manuscript	7
Generative AI Usage	7
TRANSFERRING ADVANCES FROM MODERN AI TO HEALTHCARE	9
I Introduction	11
I.1 How I Got Here	12
I.2 What Interests Me	14
I.3 The People I Work With	14
I.4 Bringing People to Science	15
I.5 Organization of This Manuscript	16
II Overview of Past Activities	17
II.1 Graph Signal Processing	18
II.1.1 The Domain in a Nutshell	18
II.1.2 Inferring Graphs from Signals	24
II.1.3 The Uncertainty Principle Applied to Signals on Graphs	26
II.1.4 From GSP and CNNs to GCNs	28
II.1.5 Application: GSP for Model Analysis & Improvement	30
II.1.6 Application: GSP for Brain Signals	31
II.2 Artificial Intelligence	33
II.2.1 The Domain in a Coconutshell	33
II.2.2 Few-Shot Learning	34
II.2.3 From FSL to LLMs	40
II.2.4 REVE: A Foundation Model for Brain Signals	45
II.3 Applications in Healthcare	50
II.3.1 Works in Neurology	51
II.3.2 Works in Cardiology	53

II.3.3	Works in Oncology	57
II.4	Other Collaborations	60
II.4.1	Operational Research & Decision Aid	60
II.4.2	Physics & Time Series	64
III	Future Work & Directions	73
III.1	Short-Term Perspectives	74
III.1.1	Projects on Fundamental AI	74
III.1.2	Applications of AI	79
III.2	Future Research Directions	82
III.2.1	Text as the Pivot of Multimodality	82
III.2.2	Agentic AI & Trust	83
III.2.3	Genomics as the Next Modality	84
III.2.4	A Cross-Service Agentic Environment	85
IV	Conclusion	87
APPENDICES		89
A	Teaching Activities	91
A.1	Focus on a Few Courses	92
A.1.1	PyRat	92
A.1.2	How to Lie with Statistics	101
A.1.3	DS Track	104
A.2	Q&A on Specific Topics	105
B	Curriculum Vitae	109
B.1	General Information	110
B.1.1	Personal Details	110
B.1.2	Higher Education	110
B.1.3	Academic Positions	110
B.1.4	Details on My PhD	111
B.2	Supervisions	111
B.2.1	PhD Students	111
B.2.2	MD Students	113
B.2.3	Postdoctoral Fellows	114
B.2.4	Research Engineers	115
B.2.5	Interns	115
B.2.6	Academic Student Projects	120
B.2.7	Industrial Student Projects	127
B.2.8	Tutored Students	128
B.3	Funded Projects	130
B.3.1	Research Projects I Wrote or Co-Wrote	130

	B.3.2	Research Projects I Contribute(d) To	133
	B.3.3	Other Funding	135
B.4		Teaching Experience	136
	B.4.1	Responsibilities	136
	B.4.2	In-Class Teaching Activities	137
	B.4.3	Certifications	141
	B.4.4	Organization of Programming Challenges	142
B.5		Other Activities	143
	B.5.1	Research Stays	143
	B.5.2	Invited Talks	143
	B.5.3	Organization of Scientific Conferences & Symposiums	145
	B.5.4	Scientific Dissemination	146
	B.5.5	Reviewing Activities	148
C		Data Analytics	151
	C.1	Teaching	152
		C.1.1 In-Class Teaching Hours Per Year	152
		C.1.2 Supervised Internships & Student Projects Per Year	152
	C.2	Research	153
		C.2.1 Publications Per Year	153
		C.2.2 Citations Profile	153
		C.2.3 Graph of Co-Authors	154
RESOURCES			155
		Bibliography	157
		References I Contributed To	157
		PhD & MD Theses	157
		Book Chapters	158
		Journals	158
		Conferences	159
		Workshops	162
		Seminars & Invited Talks	164
		Preprints	165
		Under Review	165
		Internship & Project Reports	166
		Teaching Materials	169
		Animations	169
		Other References	170
		Online Resources	190
Acronyms			195

FOREWORD



TL;DR

This manuscript, entitled “*Transferring Advances from Modern AI to Healthcare*”, compiles my previous research, and offers my perspective on what I believe are important considerations for the not-so-distant future in AI, especially in the context of healthcare.

I obviously do not claim to describe in this manuscript what the perfect AI system for healthcare will be, but rather to give an overview of the directions of my future research, and how I plan to contribute to this broader objective. In particular, main directions include integration of multiple modalities in a single prediction, use of text as a central modality, and agentic AI.

I hope you enjoy the read!

Formatting of This Manuscript

To make this document reader-friendly, I gave special attention to its formatting. In particular, I strongly recommend that you read its PDF version – using a viewer that can handle tooltips (e.g., [ACROBAT READER 🌐]; EVINCE 🌐; OVERLEAF 🌐; VSCODE PDF VIEWER 🌐)) – which offers the following features:

- **Acronyms** — In addition to the usual list of [Acronyms](#) at the end of this document, you can hover any acronym with the mouse. In most viewers, it should make a tooltip appear, showing the acronym meaning.
- **References** — Similarly, in-text references are formatted to be easily identifiable. The documents I contributed to are identified by a single-letter code indicating the type of publication (e.g., [J_i] for a **J**ournal, [B_i] for a **B**ook chapter, [W_i] for a **W**orkshop, etc.). External references appear like [BOYLE, 1665 📄] for articles, and [WORLD WIDE WEB 🌐] for websites. All references are associated with a tooltip that will appear when hovering them, for a quick glimpse at the article or website name. The associated icon also provides quick access to the full resource online.
- **Bibliography** — Entries in the [Bibliography](#) are sorted by publication type. In addition, for the documents I contributed to, I put significant work into retrieving all related public materials, such as articles, slides, posters, code repositories, etc. These resources can be directly accessed through the icons under the bibliographic entries.



The [MANUSCRIPT TEMPLATE 🌐] I designed is available publicly under the [GNU GPL v3 🌐] license. In addition to the features described above, it provides a few environments, and simplifies the management of $\LaTeX$ labels. It can also be used for a PhD manuscript, either in French or in English. Feel free to share it!

Global Vision of My Research

My research revolves around modern AI technologies, with the purpose of exploring how to get the best out of data by exploiting its structure or variability rather than its sheer amount. I am interested both in the fundamental aspects of the question and in the use of models for applied purposes. I believe this is particularly important in healthcare.

To reconcile these two facets of my research, I am part of the BRAIN team (formerly BRAIn) within the Lab-STICC laboratory at IMT Atlantique. This environment, composed of six permanent members and (on average) about a dozen PhD students, postdoctoral researchers, and research engineers, provides an ideal setting to work on the fundamental aspects of models, which form a common ground for all its members.

Staying relevant when conducting research in AI in 2026 is not straightforward within a small team with limited financial resources. It is therefore necessary to adopt an efficient and distinctive strategy. Mine, aligned with that of the team, is to identify the major research trends in AI and to draw inspiration from them, in order to advance domains that are not priorities for major industrial actors, particularly those related to healthcare. To achieve this, a significant part of my research focuses on flagship domains of AI research, *i.e.*, vision and language. Working directly on the associated technologies allows me to stay up to date in a rapidly evolving field and to identify gaps in other domains. A practical example of this approach is the article REVE [C₃₁], which draws inspiration from recent models for vision and text to propose a foundation model for EEG signals.

Focusing part of my research on core AI allows me to bring a relevant perspective to my external collaborations. Whenever possible, I try to orient these collaborations toward medical applications, where I believe there is still a great deal to accomplish with tangible societal impact. In particular, I have worked with oncologists, cardiologists, ophthalmologists, and neurologists, and envisage extending my interests to omics.

One conclusion I draw from these projects with healthcare researchers and practitioners is that there is a strong need for multimodal integration, in order to aggregate data from scans, questionnaires, electrocardiograms, clinical records, and more. Like many researchers in both AI and healthcare, I believe that one of the most promising challenges for modern AI is to identify, within heterogeneous data, complementary information that can enable more precise diagnoses or better predictions [MOHSEN *et al.*, 2022 [1]; SUN *et al.*, 2024 [2]; KRONES *et al.*, 2025 [3]]. Recent techniques in AI provide elements of a solution, both from a theoretical perspective [SHWARTZ ZIV *et al.*, 2024 [4]] and a practical one [JAEGLE *et al.*, 2021 [5]]. I have also chosen to explore this question through the prism of diversity – within data or representations – in the ANR JCJC project “Encouraging Diversity in FSL” (ENDIVE) I conduct.

The regular emergence of any-to-any models [HUGGINGFACE 

Another – in my opinion, more promising – avenue to multimodal integration of diverse sources of information in a single prediction is the rise of multi-agent systems [DU *et al.*, 2024 [10]; TANG *et al.*, 2024 [11]]. Indeed, specialized agents can be used to analyze independent modalities and improve results over a single model [KLANG *et al.*, 2026 [12]], using text as a central modality for communication between agents. This approach sort of mimics real life, in the sense that a central orchestrator agent could act as a general practitioner, and ask specialized agents to provide a description for an MRI scan, a clinical report, a blood analysis, etc. [ZHOU *et al.*, 2026 [13]].

This approach highlights the importance of text as a versatile modality, as was already suggested numerous times in the literature [HUANG *et al.*, 2022 ; SINGHAL *et al.*, 2023 ; TANG *et al.*, 2024 ]. In addition, some works have shown that instantiating multiple agents and making them debate improves reasoning and factuality, thus reducing the risk of hallucinations which hinder acceptance of deep learning-based models in healthcare [DU *et al.*, 2024 ]. Therefore, as I mentioned earlier in this [Foreword](#), keeping up with AI research on flagship modalities, especially text, is an ever-increasing necessity.

Other Things You Will Find in This Manuscript

This manuscript being an HDR, I want it to be mostly focused on my past research activities and plans for the future. Still, I am very active on the teaching aspect of my work as well. Therefore, you will find some information on these activities and my teaching philosophy in [Appendix A](#). A more detailed list of my teaching responsibilities and in-class activities is also available in [Appendix B](#). You will also find there some details on my background, the people I have supervised, my projects, and other activities related to my research, such as projects and funding. Finally, [Appendix C](#) provides a few figures summarizing some quantifiable elements of my research and teaching activities.

To conclude this [Foreword](#), it seems important to me to mention the importance I attach to open research and the development of public knowledge. I regularly participate in outreach activities aimed at the general public, and a substantial portion of my teaching at IMT Atlantique addresses research and the scientific method. All my articles are freely available online via [[ARXIV](#) (B.P.) , or [[HAL](#) (B.P.) , and the code required to reproduce the experiments is systematically published on [[GITHUB](#) (B.P.) , whenever possible. As mentioned earlier, you will find direct access to these resources through the icons under the references in the [Bibliography](#).

Generative AI Usage

As an AI researcher, it is natural to me to use AI models to assist me in the writing of this manuscript. Still, as a teacher, it is also important to me to be transparent about it, which is what I ask from students. Here is how I used generative AI across this document:

- **Template development** — I started from a blank $\text{\LaTeX}$ template, which I gradually enriched with functionalities I wanted. For this, I initially used mostly [[CHATGPT](#) , (GPT-4.5 to GPT-5.3), [[GEMINI](#) , (Gemini 3), and later switched to [[CLAUDE CODE](#) , (Opus 5). These models served both for code generation and advice on $\text{\LaTeX}$ template coding. All codes were reviewed before integration and extensively tested.
- **References search** — Models have become powerful search engines for finding articles. In addition to general purpose models such as those cited above, I used more specialized tools such as [[GOOGLE SCHOLAR LABS](#) ; [SCHOLAR INBOX](#) ; [FLICKE *et al.*, 2025](#) ; [GEMINI](#) , (Gemini 3 Deep Search). Models were used to find and summarize relevant articles before going through their actual contents.
- **Summarization** — Summaries of my previous contributions in [Chapter II](#) were first produced using [[CLAUDE CODE](#) , (Opus 5), and then reviewed and adapted manually. Full documents are available in the corresponding entries in the [Bibliography](#).
- **Scripting** — The graph in [Figure 21](#) was partially generated using [[CLAUDE CODE](#) , (Opus 4.5 & 4.6). Produced output was checked and manually adapted.
- **Syntax improvement** — I wrote this manuscript in English directly, but I'm still French 🇫🇷 nonetheless, so I used [[CLAUDE CODE](#) , (Opus 5) for syntax improvement and final touches at the end of the process.

Finally, I do not endorse the use of generative AI for artistic content outside private usage (automate the work, keep the fun!), so you will not find such content in this manuscript.

TRANSFERRING ADVANCES FROM MODERN AI TO HEALTHCARE

CHAPTER I

INTRODUCTION



This chapter is a quick introduction about me and my vision of my job. It tells how I came to research ([Section I.1](#)), and what interests me in the two halves of my Associate Professor position ([Section I.2](#)).

It also describes my role as a supervisor ([Section I.3](#)), how I try to give students a taste of research and to share my interest beyond the school and lab, and why I try to publish as openly as possible ([Section I.4](#)).

[Section I.5](#) finally describes how this document is organized.

I.1	How I Got Here	12
I.2	What Interests Me	14
I.3	The People I Work With	14
I.4	Bringing People to Science	15
I.5	Organization of This Manuscript	16

I.1 How I Got Here

I must admit that before learning that computer science was a thing, I was quite lost in my curriculum. After failing my 11th grade in science, I obtained a baccalaureate in economics and social sciences, and then tried a few things that did not lead anywhere, namely a few months at the university to study biology, then English, before quitting and working in a children’s activity center. I then joined a BTS in computer science for management, essentially because a friend told me “*You’ll see, it’s easy, basically the same thing as playing MtG*” – which I’ve been doing since I was 10.

And well, that ended up being very true! I quickly became passionate about everything revolving around computer science: programming of course, but also networks and databases. In my spare time, I made quite a few programming projects – web pages to keep my MtG cards organized, small videogames, etc. – and also got quite interested in computer virology, which I saw as a cool way of conceiving some sort of artificial life that aimed to survive through propagation and self-mutations [LUDWIG, 1991]. I’m quite confident that this early interest in virology is one of the things that drove me to AI when it surfaced again from oblivion thanks to AlexNet [KRIZHEVSKY *et al.*, 2012].

My interests in research and teaching came from my second internship [I₂], in 2010, when I joined JivÉ Création, a one-man company hosted at CERV, in the buildings of ENIB, an engineering school nearby. The objective of Cyrille BAUDOUIN was to develop an MMORPG for teaching mathematics to kids, with problems such as determining the length of the missing chain for a drawbridge, to repair it and be allowed to cross the moat (see Figure 1). This project gave me the basics of pedagogy, while making me discover a research lab from the inside.

After obtaining my BTS, I joined ENIB as a student for a year, during which I consolidated my interest in computer science and learned other elements of engineering such as signal processing or electronics. This is the year when I met Vincent GRIPON, who taught me linear algebra and gave programming courses in OCaml. The latter course took the form of a small maze, in which students would compete by implementing strategies leveraging traversal algorithms. Years later, this small project would serve as a basis for what would become a large part of my PhD: the “*PyRat – Algorithms & Discrete Mathematics*” course (see Section A.1.1).

That year went well, but did not fully meet my expectations regarding research, as most projects were missing the exploratory parts I enjoyed outside classes. I thus decided to change my curriculum once more, and got in touch with Luc BOUGÉ, who enrolled me for the “*Research & Innovation*” track of ENS Rennes (at that time ENS Cachan – Brittany campus), where I served as some sort of prototype of whether it was feasible to open the curriculum to students from backgrounds other than CPGE classes.

From there, the road was finally clear. I did several internships in various labs, exploring both previous topics of interest such as computer virology [I₄] or new ones such as embedded and real-time systems [I₅; I₆], before orienting my research interests toward data analysis and ML. I then pursued a PhD with Vincent GRIPON at Télécom Bretagne (now IMT Atlantique) on GSP, before joining the team of Pascal FROSSARD at EPFL for a year-and-a-half postdoc, where I used GSP and AI for oncology applications.

Finally, I obtained my Associate Professor position in 2019. I started in the DECIDE team, where I had strong collaborations in particular with Patrick MEYER and Romain BILLOT on topics merging operational research and AI. After some time there, I settled in the newly created BRAIN team (formerly BRAIn), whose activities are more centered around modern AI core techniques, with each permanent member having their own favored domain of application.

So here I am, writing my HDR. This not-so-straight trajectory (or maybe straight on a weird manifold) definitely shaped my approach to research and teaching. On the research side, it is only when I got interested in philosophy of science during my PhD that I realized that nobody ever told me what science was: not just a collection of disciplines at school, but an approach to discovery and knowledge establishment. On the teaching side, my variety of experiences gave me strong opinions on what made a good course and what did not. I hope these impacts on my activities are obvious to the reader when going through this manuscript.

«Ce serait tellement bien si on pouvait avoir autant de plaisir à apprendre les maths qu'à jouer aux jeux vidéos...»
Réjouis-toi, c'est enfin arrivé !

Un univers merveilleux... pour apprendre !
Des aventures, des héros, des mystères et des lieux à explorer, tout ça à travers des activités claires et amusantes qui font progresser en mathématiques.

Les maths du collège... mais en jeu.
Toutes les activités sont basées sur le programme scolaire de quatrième, pour que tes progrès te servent en classe.

Tout ce dont tu as besoin :
Grâce à ton profil en ligne, tu vois tes progrès à chaque instant, et le jeu s'adapte pour t'apporter les aides qui conviennent à tes besoins.

C'est toi qui décides !
Sélectionne tes activités en fonction de ce que tu veux travailler. Idéal pour réviser, ou pour avancer en même temps que le cours !

Joue avec tous tes amis !
Le jeu est disponible en ligne, 24h/24, 7j/7, avec des joueurs de partout en France !

Un jeu toujours plus riche !
Chaque mois, de nouveaux contenus accessibles. Activités, accessoires, aventures... des découvertes tout au long de l'année !

«Apprendre ne doit pas nécessairement rimer avec peine, mais plutôt avec plaisir.»

C'est pourquoi JivÉ Création a conçu ce nouvel outil d'apprentissage : un jeu vidéo en ligne avec des activités mathématiques adaptées au soutien scolaire.

JivÉ Math, c'est le jeu vidéo au service des progrès de l'élève. Basé sur le programme officiel, il combine une pédagogie efficace à un vrai plaisir de jeu grâce aux nouvelles technologies.



Vous voulez en voir plus ?
Flashez ce code !



<http://jivemath.com/>

JivÉ
JIVE CREATION

Combine le savoir-faire en conception de jeux vidéo et en sciences de l'éducation, afin d'offrir une expérience d'apprentissage amusante et de qualité.
contact@jivecreation.com

Les Gardiens de l'Oubli

Jouer en ligne pour apprendre les mathématiques

C'est le secret des Héros, celui qui transforme nos apprentis en Grands Maîtres.

On la vit entre amis, en découvrant les Merveilles du monde !

C'est l'Aventure

Des aventures pleines de magie, de personnages mystérieux et de secrets à découvrir

Sur cette île perdue au milieu de l'océan s'est abattue une mystérieuse malédiction.

Élucide toutes les énigmes !

Récolte des indices, explore de nouveaux territoires... découvriras-tu la vérité sur le mystérieux phénomène ?

Une autre façon de faire des maths !

Fabriquer un escalier, des vêtements, des potions... Des activités concrètes où les mathématiques prennent un sens évident.

Seul ou à plusieurs !

Entre amis, partagez vos découvertes et entraidez-vous pour mieux comprendre les maths.

Théorème de Pythagore

Affronte les dangers de l'île !

Vaincs les créatures hostiles et les obstacles qui te séparent des secrets et des trésors perdus.

Crée ton propre héros !

Personnalise ton avatar en obtenant de nouveaux accessoires... et fais admirer le résultat aux autres !

Gagne les plus hautes distinctions !

Remplis des missions et progresse en mathématiques pour monter en grade parmi les Sages et les Chevaliers.

Aide les artisans de la ville !

Résous avec eux leurs soucis du quotidien, participe aux grandes constructions et apporte-leur les ressources dont tout le monde a besoin !

Le jeu s'adapte à tes besoins !

Tous tes progrès sont sauvegardés sur ton profil en ligne.

Pratique pour te guider dans ton apprentissage ! En plus, si tu oublies certaines notions, le jeu t'aide à tout récupérer !

Deviens toujours meilleur(e) !

Des activités de découverte et des aides interactives pour comprendre clairement à ton rythme, jusqu'à tout maîtriser sur le bout des doigts !

Héros, venez secourir ses habitants, et découvrir qui sont les Gardiens de l'Oubli.

Inspirations contractuelles issues de la version en développement

Figure 1 Advertising flyer for the JivÉ Création company.

I.2 What Interests Me

My research revolves around modern AI, both in its fundamental aspects and in its applications to healthcare. The [Foreword](#) already describes this vision in more detail – and this is the main topic of this manuscript – so I will only say here what I look for in a subject: problems where the structure or variability of the data matters more than the amount of it. This is what attracted me to GSP during my PhD, what drove me to multiple modalities throughout my career, and what I now pursue through diversity, sampling and multimodality.

The other half of the job matters to me just as much. One of the arguments that convinced me to join Télécom Bretagne as a PhD student was the opportunity to redesign the first-year computer science course from scratch, and teaching has occupied a large part of my time ever since. I like designing courses that innovate and are useful for students. Above all, one of my objectives is to catch students' interest, to show them that computer science is not just an advanced calculator, but offers a wide set of tools that allow one to learn about the world, simplify daily life problems, and create new things... and that it can be fun! While this manuscript is more about my research activities, I have dedicated [Appendix A](#) to some selected teaching activities, which I believe illustrate this vision well.

I.3 The People I Work With

Almost everything described in this manuscript was done – or will be done – with someone. I like working in teams of complementary people, and I have been lucky enough to be surrounded by great people during my whole career. My work with students is no exception to this principle. This manuscript would not exist without their contributions, and I am deeply grateful to all of them.

I have supervised or co-supervised 10 (soon 11) PhD students, 2 MD theses, 4 postdoctoral researchers and 2 research engineers, along with more than 20 interns and several student projects gathering close to a hundred students; the complete list is given in [Section B.2](#). My role in these works is generally to formulate research directions, to discuss the methodological choices, to challenge the implementations and the conclusions, and to make sure that what is being done fits the students' objectives.

Indeed, be it for a student project or a PhD thesis, one of the first questions I ask students is what they envision for their future. At the project level, for students who are interested or curious about research, I try to orient the project toward exploratory elements, with a publication objective. For those interested in industry, I support them on the technical side, and encourage them to build a proof of concept that showcases their skills.

At the thesis level, this translates into discussions about the student's interests, and what publication strategy to adopt to maximize the chance that they attain their objectives. Overall, the main lines are as follows, depending on the student's target:

- **Professor** — For students who envision such a career, teaching – and in particular designing courses – is a very important part of the job. They need to get started early, to build a CV that will distinguish them from other candidates. As a coordinator of the first-semester [Computer Science](#) course at IMT Atlantique, I am able to accompany them both in classes and in the writing of new materials.
- **Academic researcher** — For students targeting an academic position, I try to make them work on ambitious research questions, with the aim of publishing in high-ranking venues. I also insist that they spend some time in another lab connected to their research interests at some point in their PhD.
- **Industry researcher** — For students more interested in industry, high-ranking publications are also very important, but topics should be more technical, closer to industrial interests. I also strongly encourage them to do an internship in a company at some point, where they can prepare a future collaboration.

These are obviously very general non-exhaustive rules, which need to be adapted to each profile. Overall, I see a thesis as a period during which students get the opportunity to shape their CV, and build a profile that will

distinguish them from other candidates during their future applications, whether through the design of courses, obtaining patents, high-ranking publications, organization of seminars, or anything else of the kind.

In addition, what I like to do for any new PhD student is to quickly aim for a publication, even a minor one. First, this introduces them to the publishing process, with the joys and disappointments that come with it. Second, it is a good pretext to join an existing project and help integrate them into the team. Finally, it serves as a safety net, as they already have a publication to defend. This frees the mind for the rest of the thesis, allowing them to explore more freely and take more risks.

Regarding the work itself, I strongly encourage collaborations within and outside the team. I believe that joining forces toward a common goal is a necessary step for publishing in high-ranking venues such as NeurIPS, ICML, CVPR, or ICLR. It also prepares students for their future careers, where collaboration is the norm.

Overall, I think it's important to be flexible, to adapt to students and their career objectives, but also to the fast-paced evolution of the field. It is not an easy task to write a long-term research plan when working on AI in a small public laboratory in 2026, as large labs or companies may just release a new model that makes months of work obsolete (looking at you, SAM [KIRILLOV *et al.*, 2023 ]). My approach is therefore to favor more niche subtopics of AI and applications that are not obviously on the roadmap of these big actors.

I.4 Bringing People to Science

As mentioned in [Section I.1](#), I was not told what science was when I was a student, which I find quite sad in the current education system, where science is basically perceived as the maths, physics and biology package, and is just where students with good grades go in high school.

In my activities, I try to interest people in science and research, at different levels. In classes, I designed and participated in multiple courses that aim at developing scientific reasoning and at letting students experience research (*e.g.*, [How to Lie with Statistics](#), [Introduction to Research](#)). Also, as already mentioned, I try to orient the projects I propose toward research when students want it, some of which led to publications in conferences [[C₂](#); [C₄](#); [C₇](#); [C₁₂](#); [C₂₆](#); [C₂₈](#); [C₃₀](#)], where students presented their work in front of an audience of researchers.

Still, while this is useful to engineering students, most of whom are already curious about research, I think that the societal impact is limited. I thus try to reach the general public when I can, to let them know that science can be really fun and creative. I regularly participate in dissemination activities (see [Section B.5.4](#)) and give talks in primary school classes. Océanopolis, and especially Fabio PERRUCHET, are extremely valuable partners for this, and I am glad we have them in Brest.

Regarding my publications and codes, I try to make everything I produce freely available. All my articles are on arXiv [[ARXIV \(B.P.\)](#) ] or HAL [[HAL \(B.P.\)](#) ], the code to reproduce the experiments is published on GitHub [[GITHUB \(B.P.\)](#) ] whenever possible, and the software I write – both for work and for personal projects that could serve someone else – is released under free licenses. This is a deliberate position rather than a habit. Gatekeeping an article or a piece of code does not protect much: it mostly forces other people to spend energy – theirs, and that of the machines they run – reproducing something that already exists. In a field whose computational cost is becoming a problem of its own, that waste is hard to justify. I also believe that publicly funded science should serve the world rather than the career of the person who signs it, and that the small advantage gained by keeping a result closed is never worth the collective cost of hiding it.

I.5 Organization of This Manuscript

After these small discussions, let's dive into the core content of this manuscript. [Chapter II](#) presents my past research activities, starting from my early work on GSP to my current activities in AI, followed by their applications in healthcare and by other collaborations. For each article, I tried to situate its context, provide a summary of its contributions, and analyze its impact. [Chapter III](#) then describes what comes next: the projects already under way or about to start, and the longer-term research directions they converge toward.

Three appendices complete the document: [Appendix A](#) on my teaching activities and philosophy, [Appendix B](#) for the exhaustive lists of supervisions, projects, responsibilities and publications, and [Appendix C](#) for a few figures summarizing them.

CHAPTER II

OVERVIEW OF PAST ACTIVITIES



This chapter summarizes the research activities I have been conducting so far. As much as possible, I have tried to group my research interests into meaningful categories.

My research originally revolved around GSP (Section II.1), which I studied as a toolbox to extend signal processing approaches for the purpose of studying signals with complex structures. From there, I naturally shifted toward the study of AI (Section II.2), which offered a more general methodology to achieve a similar objective.

For both domains, my research has been both fundamental and applied. For the latter, I am especially interested in applications related to healthcare (Section II.3). Still, AI tools being prevalent nowadays, I also had the opportunity to help in a few other projects (Section II.4).

II.1	Graph Signal Processing	18
II.1.1	The Domain in a Nutshell	18
II.1.2	Inferring Graphs from Signals	24
II.1.3	The Uncertainty Principle Applied to Signals on Graphs	26
II.1.4	From GSP and CNNs to GCNs	28
II.1.5	Application: GSP for Model Analysis & Improvement	30
II.1.6	Application: GSP for Brain Signals	31
II.2	Artificial Intelligence	33
II.2.1	The Domain in a Coconutshell	33
II.2.2	Few-Shot Learning	34
II.2.3	From FSL to LLMs	40
II.2.4	REVE: A Foundation Model for Brain Signals	45
II.3	Applications in Healthcare	50
II.3.1	Works in Neurology	51
II.3.2	Works in Cardiology	53
II.3.3	Works in Oncology	57
II.4	Other Collaborations	60
II.4.1	Operational Research & Decision Aid	60
II.4.2	Physics & Time Series	64

II.1 Graph Signal Processing

II.1.1 The Domain in a Nutshell

GSP was introduced by [SHUMAN *et al.*, 2013], as a generalization of classical signal processing based on spectral graph theory [CHUNG, 1997]. The article framed several tools leveraging graph theory and spectral theory in a single environment, to process, in a similar way, signals such as time series, images, or more complex ones defined on a domain modeled as a graph. In this framework, vertices of the graph represent elements of interest, where signal values are measured, and edges model relations of similarity between linked vertices.


GSP offers an algebraic vision of signal processing, through an analogy between the eigenvectors of a matrix (generally the Laplacian or the adjacency matrix) representing a line graph – modeling discrete time – and the Fourier modes in classical analysis of signals. Figure 2 illustrates this analogy with an example signal, using a low-pass filter to remove noise.

This example introduces the GFT operator, which allows one to interpret the Fourier domain of a signal as a space defined by the support of the information itself. Building on this central observation, the GSP community has progressively re-developed most of the classical signal processing toolbox for signals on graphs (see [SHUMAN *et al.*, 2013; ORTEGA *et al.*, 2018; LEUS *et al.*, 2023] for major milestones). Tools such as wavelets [HAMMOND *et al.*, 2011], filter banks [NARANG *et al.*, 2012], sampling and reconstruction [CHEN *et al.*, 2015; ANIS *et al.*, 2016] and stationarity of stochastic processes [GIRAULT, 2015; PERRAUDIN *et al.*, 2017] have all found graph counterparts. These developments generally require dedicated theoretical analyses, as many classical results rely on the regularity of the underlying domain. Considerable effort has also been devoted to computational efficiency: since the GFT has no general analog of the fast Fourier transform, approximate transforms based on sparse factorizations of the eigenvector basis [LE MAGOAROU *et al.*, 2017] and polynomial graph filters that avoid Laplacian diagonalization and rely only on local computations [ISUFI *et al.*, 2024] have been proposed. Altogether, these advances establish GSP as a natural generalization of classical 1D signal processing to irregular domains.


Beyond the simple generalization of existing tools, what makes GSP a valuable domain is the observation that developed tools effectively behave as expected for signals evolving on more complex domains, modeled as graphs. For instance, consider the example in Figure 3, illustrating application of a low-pass filter on a random geometric graph. As was the case for the example in Figure 2, we observe a smoothing of the random signal entries over the graph, *i.e.*, vertices linked with edges tend to have more similar signal values after filtering.

Many use cases can be encoded in the GSP framework [ORTEGA *et al.*, 2018]. Social networks are a natural example, where vertices model users, edges encode friendship relations, and a signal can be any quantity measured on profiles (*e.g.*, age, or an opinion on some topic) [SANDRYHAILA *et al.*, 2014]. Similarly, transportation networks provide support graphs for signals such as traffic for congestion analysis [MOHAN *et al.*, 2014]. Another important family of examples comes from brain signals (*e.g.*, EEG or fMRI), where vertices are sensors or voxels, linked according to their physical proximity, or to statistical dependencies between their recordings, and each acquisition provides a new signal on that graph [HUANG *et al.*, 2018]. Other examples include applications to protein analysis [SRIVASTAVA *et al.*, 2023], or computer graphics, where 3D meshes and point clouds are graphs supporting geometric or color information [DINESH *et al.*, 2020], among many more. In all these settings, GSP provides a principled way to exploit the structure of the underlying domain when processing the measured data.


Along the way, the field has also raised issues of its own, some of which remain open challenges today. Most tools assume the support graph to be known, which is not always the case in practice, motivating an entire line of works on inferring plausible graphs from signal observations. Some fundamental results do not directly transpose to irregular domains either, such as the uncertainty principle, which bounds how much a signal can be simultaneously localized in the vertex and spectral domains. Extensions to directed graphs (for which the Laplacian is not symmetric anymore and thus loses its orthonormal eigendecomposition) [SHAFIPOUR *et al.*, 2018];



(a) Consider a periodic time signal $s(t)$ as a support for the example.




(b) This signal can be discretized by regularly sampling its amplitude. The obtained discrete signal can now be seen as a vector $\mathbf{x} \in \mathbb{R}^N$, with N the number of samples.




(c) In the GSP framework, we view $\mathbf{x}$ as a vector associating a value to each node of a support graph. Here, as we consider a periodic time series, the most natural graph to model such a support is a ring graph, where each vertex corresponds to a time instant, and is linked to the previous and next instants in time. We therefore manipulate two distinct objects: the signal $\mathbf{x}$ which contains the information of interest, and the graph $\mathcal{G}$ which models time.

Figure 2 A simple example to illustrate the analogy between classical analysis of time signals and GSP. Plots (a) to (c) illustrate how GSP offers a different point of view from classical analysis, by distinguishing a signal from its support. Plots in (d) make the link between the signal support and the Fourier transform operator, which computes the spectrum of the example signal in (e). Then, (f) to (h) illustrate the analogy through the example of a low-pass filter.




- (d) Now, let us put $\mathbf{x}$ aside to focus only on the support graph $\mathcal{G}$ introduced in (c). Let $\mathbf{W}$ be the adjacency matrix of the graph, and $\mathbf{d}$ the vector of degrees of each vertex. From these, we can compute the graph Laplacian $\mathbf{L} = \text{diag}(\mathbf{d}) - \mathbf{W}$ of the graph, which is a real and symmetric matrix, thus diagonalizable as $\mathbf{L} = \mathbf{V} \cdot \text{diag}(\boldsymbol{\lambda}) \cdot \mathbf{V}^\top$. Assuming that eigenvalues in $\boldsymbol{\lambda}$ are sorted increasingly, we denote by $\mathbf{V} = (\mathbf{v}_1, \dots, \mathbf{v}_N)$ the associated eigenvectors, and plot the first ones. As can be observed, these eigenvectors are analogous to the sines that decompose time signals in classical Fourier analysis, and their frequencies are analogous to the corresponding eigenvalues.




- (e) As a consequence of (d), a spectral decomposition $\hat{\mathbf{x}}$ of $\mathbf{x}$ can be obtained by a change of basis defined by the eigenvectors of the Laplacian, *i.e.*, $\hat{\mathbf{x}} = \mathbf{V}^\top \mathbf{x}$. This operation is called the GFT, and is at the core of GSP. The spectrum of the signal introduced in (c) is depicted here. Each entry $\hat{x}[i]$ measures the correlation between $\mathbf{x}$ and the eigenvector $\mathbf{v}_{i+1}$, *i.e.*, the contribution of the associated graph frequency to the signal. This is the exact analog of the spectrum obtained with the discrete Fourier transform of a time signal.


Figure 2 (continued)



- (f) As in classical Fourier analysis, filters can be conveniently defined in the spectral domain. Here, we define a low-pass filter $\hat{h}[i] = \exp(-0.05i)$, which preserves low graph frequencies and attenuates higher ones.




- (g) Filtering then amounts to an entry-wise multiplication in the spectral domain, *i.e.*, $\hat{y} = \hat{h} \odot \hat{x}$ (followed here by an ℓ_2 renormalization so that the filtered signal keeps the energy of the original one, *i.e.*, $\|\hat{y}\|_2 = \|\hat{x}\|_2$). Compared with the spectrum in (e), entries associated with high graph frequencies are strongly attenuated.




- (h) Going back to the time domain using the GFT⁻¹, *i.e.*, $y = V\hat{y}$, yields a smoothed version of the original signal, in which most of the noise has been removed, as expected from a low-pass filter.

Figure 2 (continued)




(a) A random geometric graph, obtained by drawing $N = 40$ vertices uniformly at random in a rectangle, and linking those closer than a fixed radius. A random signal $\mathbf{x}$ associates a value to each vertex, depicted as a red bar.




(b) First eigenvectors $\mathbf{V} = (\mathbf{v}_1, \dots, \mathbf{v}_N)$ of the Laplacian of the graph introduced in (a), represented on the graph (as the i^{th} entry of an eigenvector corresponds to the i^{th} row/column of the Laplacian matrix). As for the ring graph in Figure 2 (d), eigenvectors associated with low eigenvalues vary smoothly on the graph, thus playing the role of low-frequency Fourier modes on this irregular domain.


Figure 3 Illustration of the application of a low-pass filter on a random signal, evolving on a random geometric graph. Figure (a) introduces the support graph and the studied signal. Plots in (b) depict the first eigenvectors of the graph Laplacian, which generalize Fourier modes to this irregular domain. Then, (c) to (e) reproduce the low-pass filtering procedure of Figure 2 on this support.



(c) The spectrum $\hat{\mathbf{x}} = \mathbf{V}^\top \mathbf{x}$ of the signal, obtained with the GFT associated with the graph.



(d) We filter the spectrum with a low-pass filter $\hat{\mathbf{h}}[i] = \exp(-0.1i)$, i.e., $\hat{\mathbf{y}} = \hat{\mathbf{h}} \odot \hat{\mathbf{x}}$.



(e) The filtered signal $\mathbf{y} = \mathbf{V} \hat{\mathbf{y}}$, obtained with the GFT⁻¹, depicted on the graph. Compared with (a), the signal is smoother: vertices linked by an edge tend to have more similar values.

Figure 3 (continued)

MARQUES *et al.*, 2020 [10]; HAN *et al.*, 2026 [11] and to time-varying topologies [YAN *et al.*, 2024 [12]] remain active research directions, as does the stability of spectral designs under perturbations of the support graph. Finally, defining translations and convolutions – which are trivial operations on regular domains, but not on arbitrary graphs – establishes a direct bridge between GSP and CNNs [BRONSTEIN *et al.*, 2017 [13]], from which GCNs emerged as their generalization to graph-structured data. A large part of the community effort has now shifted toward this interplay between GSP and modern ML [DONG *et al.*, 2020 [14]], epitomized by the recent emergence of graph foundation models [WANG *et al.*, 2025 [15]].

Still, the GSP community remains pretty active, and some works aim at generalizing the approach even more. While GSP has proven remarkably successful for modeling pairwise interactions, many real-world systems exhibit higher-order relationships that cannot be faithfully represented by simple graphs. Examples include social groups involving more than two individuals such as co-authorship networks, biochemical reactions, or higher-order functional interactions in the brain. This observation has motivated the extension of signal processing tools to richer topological domains. Simplicial complexes provide one of the most established frameworks, generalizing graphs by explicitly representing interactions among arbitrary subsets of vertices while preserving their inclusion relationships. Based on the combinatorial Hodge Laplacian, they admit higher-order analogs of the GFT and associated notions of filtering, diffusion, and harmonic analysis [BARBAROSSA *et al.*, 2020 [16]; SCHAUB *et al.*, 2022 [17]]. A closely related formalism is hypergraphs, whose hyperedges naturally encode multi-way interactions and have motivated extensions of spectral signal processing [ZHANG *et al.*, 2019 [18]]. These developments broaden the scope of GSP beyond pairwise interactions, enabling the analysis of signals defined on increasingly expressive relational structures.

II.1.2 Inferring Graphs from Signals

One of the main challenges in GSP is how one can use the tools it provides if the support graph of observed signals is unknown. Indeed, while in some cases the graph is directly available (*e.g.*, a social network where edges model mutual relationships), it is not always the case when vertices are observations of locations in a continuous space (*e.g.*, EEG or fMRI signals), or when the topology is not provided for some reason (*e.g.*, intellectual property).

In such a scenario, a common approach is to rely on priors on the graph or signals, such as sparsity [FRIEDMAN *et al.*, 2008 [19]; LAKE *et al.*, 2010 [20]], smoothness [KALOFOLIAS, 2016 [21]; DONG *et al.*, 2016 [22]], or minimization of edge variation across time [YAMADA *et al.*, 2020 [23]] (among many others). This approach reduces the search space of possible graphs, as the problem is in general ill-defined.

In my previous research, I addressed the question of inferring a graph under a stationarity prior, *i.e.*, assuming the signal is observed after diffusion of white noise on an unknown graph to infer. This led to the following publications, as well as a few workshop presentations [W₂; W₃; W₄]:



[C₁] Bastien PASDELOUP, Michael RABBAT, Vincent GRIPON, Dominique PASTOR, Grégoire MERCIER (2015). “**Graph Reconstruction from the Observation of Diffused Signals**”. *Annual Allerton Conference on Communication, Control, and Computing* (pp. 1386–1390), IEEE

This article introduced a method for reconstructing an unknown graph topology from observations of signals generated by a diffusion process on the graph. It considered the case where observed signals originate from initially *i.i.d.* signals that have been propagated a fixed and known number of times through an unknown graph diffusion operator. Under this assumption, the covariance matrix of the observed signals is a power of that operator, so that its diagonalization directly yields the eigenvectors of the diffusion matrix, but only the absolute values of its eigenvalues. Recovering the signs of these eigenvalues is therefore the core of the problem, and the article formulates it as a set of linear constraints

enforcing the structural properties that any valid diffusion matrix satisfies: symmetry, non-negative off-diagonal entries, and a null diagonal. Graph recovery thus becomes a linear feasibility problem that can be solved efficiently by interior-point methods, and the article further shows that only a fraction of the quadratically many inequality constraints needs to be retained to reach the same solution. Experimental results showed accurate reconstruction of randomly generated graphs from diffused signals when the empirical covariance matrix is assumed to coincide with the true one and the sparsification threshold is chosen by an oracle; from a finite number of signals, the reconstruction error was either close to zero or around 0.5, the noise affecting the empirical covariance being named as the main limitation of the approach. The possibility of learning hidden network topologies without direct access to the underlying graph was therefore validated only under these restrictions, and provided the eigenvalues of the diffusion matrix have pairwise distinct absolute values.



Bastien PASDELOUP, Vincent GRIPON, Grégoire MERCIER, Dominique PASTOR, Michael RABBAT (2017). **“Characterization and Inference of Graph Diffusion Processes from Observations of Stationary Signals”**. *TSIPN (vol. 4, nb. 3, pp. 481–496), IEEE*

This work extended the previous framework by considering a more general setting where signals are stationary and may have undergone an unknown and variable number of diffusion steps. The article characterized the complete set of graph diffusion operators compatible with the observed signals, showed that all admissible operators share the eigenvectors of the covariance matrix, and proved that the set they form is a convex polytope delimited by linear inequality constraints. Instead of directly estimating a single graph, the proposed framework therefore first identifies all admissible diffusion operators, and then selects a point in that polytope according to an additional criterion, such as simplicity or sparsity in the sense of the $\ell_{1,1}$ norm. The article also analyzed the impact of a finite number of observations, and experimentally measured how fast the approximate polytope obtained from an empirical covariance matrix converges toward the limit one as the number of available signals increases. Two further contributions completed the framework: a measure quantifying how well a given graph explains signals under a stationarity assumption, and a method to turn a graph inferred under other priors into a diffusion matrix compliant with that assumption. Experiments on a real dataset of temperatures were finally used to situate state-of-the-art inference methods within this framework.

These works had a significant impact on the GSP community, and served as a reference for numerous subsequent contributions. Among noticeable follow-ups, the spectral templates framework [SEGARRA *et al.*, 2017] generalized this principle to the inference of graph shift operators from spectral information, formulating graph recovery as the selection of an operator satisfying structural constraints within a prescribed spectral basis. The diffusion model was later extended to non-stationary signals in [SHAFIPOUR *et al.*, 2021], which introduced graph-filter identification as an additional step to recover the underlying topology, and to more general graph system-identification settings [EGILMEZ *et al.*, 2018]. Related developments also considered heat-diffusion models [THANOU *et al.*, 2017] and online inference from streaming stationary graph signals [SHAFIPOUR *et al.*, 2019].

These subsequent works illustrate how the spectral characterization introduced in these articles helped establish the relationship between the covariance structure of observed signals and the topology of their underlying graph as a central principle for data-driven graph learning.

II.1.3 The Uncertainty Principle Applied to Signals on Graphs

Another topic of interest I worked on was the study of the uncertainty principle for signals on graphs. Initially introduced in quantum mechanics [HEISENBERG, 1927], and transposed to time signals in [GABOR, 1946], this principle states that a signal cannot be concentrated in the time and spectral domains at the same time, *i.e.*, there is a bounded trade-off between time and spectrum localization. Extending this principle to GSP was initially done in [AGASKAR *et al.*, 2012; AGASKAR *et al.*, 2013], which introduced spread metrics similar to those of [GABOR, 1946], and showed evidence of such a trade-off in the GSP framework as well.

Leveraging these previous works, my first contribution was to extend the proposed metrics to weighted graphs. This work, done in collaboration with Réda ALAMI, led to the following publication:

 [C₂] Bastien PASDELOUP, Réda ALAMI, Vincent GRIPON, Michael RABBAT (2015). “**Toward an Uncertainty Principle for Weighted Graphs**”. *EUSIPCO* (pp. 1496–1500), IEEE

This work extended the graph uncertainty principle from unweighted to weighted graphs. It first showed that a direct extension of the existing definition of graph spread is problematic, as it can lead to discontinuities when graph weights vary continuously: the underlying issue is that the weights matrix encodes a similarity, whereas the definition of the spread uses it as a distance. To overcome this, the article characterized axiomatically the class of admissible distance functions – non-negativity, vanishing only on the diagonal, and continuity together with monotonicity with respect to the weights – and proposed two instances that satisfy it, one based on the geodesic distance computed on an inverse similarity matrix, the other on a diffusion distance. Uncertainty curves were then computed numerically, using the sandwich algorithm, for several families of weighted graphs: randomly weighted complete, star and ring graphs, random geometric graphs built with a Gaussian kernel, and a real graph extracted from an image. The choice of distance was observed to leave the relative ordering of the curves unchanged, affecting mainly their smoothness.

To go beyond a simple numerical bound on the relative spreads in graph and spectral domains, I then proposed a characterization of signals that are Pareto-optimally localized. Results were presented in the following article, as well as in a national conference [C₃]:

 [C₃] Bastien PASDELOUP, Vincent GRIPON, Grégoire MERCIER, Dominique PASTOR (2016). “**Towards a Characterization of the Uncertainty Curve for Graphs**”. *ICASSP* (pp. 4558–4562), IEEE

This work addressed the characterization of the signals that attain the uncertainty curve, rather than approximating the curve through an exhaustive search over the signal space. For a given graph and reference node, the uncertainty curve describes the Pareto-optimal trade-off between localization in the graph domain and localization in the spectral domain, and the signals that reach it are the eigenvectors associated with the lowest eigenvalue of a matrix $M(\alpha) = P - \alpha L$, where P collects the distances to the reference node. The article derived a structural characterization of these signals for a broader class of graphs than the complete and star graphs previously considered: whenever $M(\alpha)$ decomposes into blocks such that one of them is circulant and the block coupling it to the rest is constant along rows, the signals attaining the curve have repeated entries over the corresponding coordinates. This condition reduces the space of candidate signals that needs to be explored, can be applied recursively to refine the characterization further, and was illustrated on star graphs with either the central node or a leaf taken as reference.

Leveraging the uncertainty principle, I later developed an approach to transfer a signal from one graph to another, offering a methodology to evaluate what a phenomenon observed on a particular graph could look like in a different setting. This led to the following communication:

[C₉]

Bastien PASDELOUP, Hermina PETRIC MARETIC, Mireille EL GHECHE, Pascal FROSSARD (2019). **“Une Approche Basée Incertitude pour Transporter un Signal d’un Graphe à un Autre”**. *Colloque du GRETSI* (pp. 37–40)

This article introduced a method for transporting a graph signal from a source graph to a target graph when the two graphs have different topologies or numbers of nodes. Rather than attempting to establish a direct node-to-node correspondence, the approach represents the source signal through its joint localization properties in the graph and spectral domains, relying on the uncertainty region introduced by Agaskar and Lu, whose attaining signals the preceding works had characterized. That region is delimited by four identifiable points, associated with signals that are easy to characterize, and the sandwich algorithm allows an arbitrary number of points to be estimated along its frontier. The method exploits this structure to put the frontiers of the two uncertainty regions in correspondence, and then locates within the target region, by angular triangulation with respect to the barycenters, the uncertainty pair equivalent to that of the source signal. A proxy signal is finally constructed on the target graph by minimizing a criterion that encourages the preservation of chosen statistical properties, under the hard constraints of a unit norm and of the prescribed levels of concentration in both domains. This formulation turns signal transport into a problem of finding a signal on the target graph with prescribed spatial and spectral localization properties, without requiring an explicit alignment between the two graphs, although the solution is not unique. Experiments on random geometric graphs illustrated the interest of the approach by clustering signals representing the same underlying phenomenon observed on the two distinct graphs.

These works were compiled and extended in a book chapter on the uncertainty principle applied to signals on graphs, providing a more complete reading on the topic:

[B₁]

Bastien PASDELOUP, Vincent GRIPON, Réda ALAMI, Michael RABBAT (2019). **“Uncertainty Principle on Graphs”**. *Vertex-Frequency Analysis of Graph Signals* (pp. 317–340), Springer

This chapter provided a comprehensive formulation of the uncertainty principle for graph signals, extending the classical time-frequency uncertainty principle to irregular graph domains. Written as a self-contained tutorial, it first introduced graphs and their matrices, the semantics of edge weights, the properties of the normalized Laplacian, graph signals together with notions of sharpness and of diffusion sequences, and the GFT by analogy with the discrete Fourier transform, before turning to uncertainty. It then defined, for unit-norm signals, a graph spread based on the graph distance and a spectral spread based on the eigenvalues of the normalized Laplacian, both obtained as an infimum over a reference – a vertex for the former, an eigenvalue for the latter – so that neither depends on an arbitrarily chosen reference vertex, unlike the definitions used in the preceding works. The set of achievable pairs of spreads defines the uncertainty domain, shown to be compact, and its lower frontier defines the uncertainty curve, which characterizes the fundamental trade-off between localization in the vertex and spectral domains. The chapter then analyzed this framework for several representative graph families. For complete graphs and star graphs, explicit expressions and bounds show that the uncertainty domain can approach the origin as the number of vertices increases, which suggests that no non-trivial uncertainty principle holds for these degenerate topologies. In contrast, for ring graphs, the spectral structure of the normalized Laplacian is characterized through sinusoidal eigenvectors, and a positive lower bound on the uncertainty domain is

established, proving the existence of a non-trivial uncertainty principle. Overall, the chapter established that the possibility of simultaneously localizing a graph signal in the vertex and spectral domains is fundamentally determined by the graph topology, and identified the characterization of the graph families for which the uncertainty principle does not hold as an important direction, since it is on those graphs that very efficient sampling and reconstruction strategies could be devised.

My work on the uncertainty principle had a moderate impact, as was mostly incremental. Still, it provided theoretical results that contributed to building knowledge on this topic. The book chapter however had a more significant impact, as it provides an easy introduction to the topic, and is therefore often cited as a broad introductory resource (e.g., [STANKOVIĆ *et al.*, 2021 ]).

II.1.4 From GSP and CNNs to GCNs

GSP has played a foundational role in the development of modern graph-based ML, in particular by providing the mathematical framework underlying graph convolution, graph filtering, and spectral representations. Generalizing the convolution operation to graphs inspired architectures such as ChebNet [DEFFERRARD *et al.*, 2016 ] (where graph convolution can be interpreted as the application of a learnable spectral filter to graph signals) and the GCN [KIPF *et al.*, 2017 ] (which acts as a message-passing system, where local aggregation of neighbors acts as a graph filtering operation). The other canonical branch of this lineage replaces the fixed aggregation weights of a graph convolution by attention coefficients learned from the representations of the two linked vertices, as GATs [VELIČKOVIĆ *et al.*, 2018 ] do. GCNs and their variations have since become a classic in the taxonomy of AI models.

In my final year of PhD, I also developed a generalization of convolutions for graph signals. Contrary to the other approaches, I wanted to first approximate a translation operator for graphs, and then use it to define a convolution, as one would classically do in the time domain. These works on translation identification, done in collaboration with Nicolas GRELIER, led to the following publication, also presented in [C₇] and in a more complete – though unpublished – preprint [P₁]:

 [C₄] Nicolas GRELIER, Bastien PASDELOUP, Jean-Charles VIALATTE, Vincent GRIPON (2016). “**Neighborhood-Preserving Translations on Graphs**”. *GlobalSIP* (pp. 410–414), IEEE

This work introduced a definition of translations on graph signals that relies only on the preservation of neighborhoods, in contrast with spectral translations that are built directly from the graph Fourier basis and that do not match the usual translations on regular domains. A perfect graphical translation is defined as an injective map sending each vertex to one of its neighbors while preserving adjacency; because such maps exist only on very particular graphs, a generalized version is also introduced, in which vertices admitting no valid image are sent to an absorbing state, the requirement of preserving as many vertices as possible being imposed vertex by vertex rather than globally – so loosely that, as the article immediately observes, the map sending every vertex to the absorbing state qualifies as well. The central result is that on cyclic grid graphs the only graphical translations are the usual elementary geometrical ones – which are perfect – together with the degenerate map onto the absorbing state; partial results are established for families of non-cyclic grid graphs, together with the conjecture that, whenever every dimension has length at least six, the only graphical translations are those elementary translations and the map collapsing everything onto the absorbing state. The article is purely theoretical and contains no experimental evaluation, its only illustration being a comparison with a spectral shift operator on

an example image. Finding such translations on arbitrary graphs is explicitly left as an open challenge, because of the combinatorial explosion in the number of candidates, with the suggestion that modifying the locality condition – that the image of a vertex be one of its neighbors – should allow other isometries, such as rotations, to be captured in the same way.

Using this approach, the final contribution of my PhD was to develop a pipeline that would leverage my work on graph inference and translation identification to build a graph convolutional model, even in the case where no graph is available. These results, summed up in Figure 4, were presented in multiple workshops and seminars [W₆; S₃; S₄; S₅; S₆].




Figure 4 Proposed methodology to create a CNN from a given set of signals to classify. (a) From a set of signals, we infer a graph; (b) Then, we identify some translation operators on that graph; (c) We instantiate a kernel on the graph, with trainable weights; (d) Using the previously found translations, we translate the kernel so that it gets centered on every possible vertex. Once the kernel has been localized on every vertex, we can define a corresponding convolutional layer; (e) By considering various initializations, we obtain different layer shapes. A CNN can then be constructed by assembling them to form a complex network; (f) Finally, the CNN is trained using the data from which the whole process started.

Overall, this work did not have any real impact, due to the high complexity of identifying translations on an irregular graph – an NP-complete problem – even with the proposed approximations, and performance that did not match other approaches [P₂]. The work on translation identification is sometimes cited as an attempt at generalizing this operator [GHANDEHARI *et al.*, 2021 [43]; SOUP *et al.*, 2026 [44]], but the derived CNN did not compete with existing approaches. Still, I am happy with this work, which was very timely, and marked a shift in my career toward AI.

II.1.5 Application: GSP for Model Analysis & Improvement

A possible application of GSP is improving of performance of AI models, by taking into account a graph that models relationships between some data. In a work with Mounia HAMIDOUCHE, we explored how filtering on a graph of similarities in feature space could help improve classification accuracy:

 [C₁₄] Mounia HAMIDOUCHE, Carlos LASSANCE, Yuqing HU, Lucas DRUMETZ, Bastien PASDELOUP, Vincent GRIPON (2021). **“Improving Classification Accuracy with Graph Filtering”**. *ICIP (pp. 334–338), IEEE*

This work proposed to reduce intra-class noise in the latent space of a trained DNN using graph filtering, in order to improve classification accuracy. For each class of the labeled set, a graph is built whose vertices are the samples of that class and whose edges connect k nearest neighbors according to the cosine similarity of their feature vectors; the features are then seen as graph signals and low-pass filtered, after which only the classifier is retrained on the denoised representations. Under a Gaussian model of the features, and in the ideal case of a complete graph per class and of a perfect step filter, the article showed that this filtering preserves the expectation of the class centroid while reducing its covariance, which vanishes as the number of labeled samples per class grows; extending the result to the k -nearest-neighbor graph and the three-level filter actually used is left as a conjecture. Experiments in the few-shot setting on Mini-ImageNet, CUB, CIFAR-Few-Shot and Tiered-ImageNet, and in the standard setting on CIFAR-10, showed that the procedure slightly improves upon state-of-the-art results in both regimes. The filter itself was hand-tuned and kept fixed across the four few-shot datasets, although the standard regime called for a different graph and different filter parameters, the design of an optimal graph filter being left as an open question.

GSP can also be used for understanding models. As an example, using graph smoothness to measure organization of labels across various layers, [GRIPON *et al.*, 2018 

 [B₂] Carlos LASSANCE, Myriam BONTONOU, Mounia HAMIDOUCHE, Bastien PASDELOUP, Lucas DRUMETZ, Vincent GRIPON (2022). **“Graphs as Tools to Improve Deep Learning Methods”**. *Advances in AI: Reviews (vol. 2), IFSA Publishing*

This chapter reviewed the use of graphs constructed from intermediate representations of DNNs to analyze and improve their behavior. The considered graphs – termed latent geometry graphs – have samples as vertices, and edge weights encoding similarities between their representations at a given network layer. These graphs are then exploited through four families of approaches. First, graph smoothness and spectral quantities are used to interpret intermediate representations: a label smoothness measure is shown to separate correctly optimized architectures from under- and overfitted ones, and to keep evolving after the training accuracy has saturated, which suggests using it as a stopping criterion; related methods identify the training samples or image regions that most influence predictions. Second, graph filtering is used to denoise intermediate representations, exploiting the assumption that semantically similar samples should have similar representations, with a particular interest in the few-data regime. Third, graph-based objective functions are incorporated into the learning process as losses, either to encourage desirable geometric properties of the learned embeddings, or to transfer the geometry of a large network onto a smaller one through graph knowledge distillation, thereby compressing the architecture. Finally, graph regularization is used during training to improve the robustness of the network to input

noise, adversarial perturbations, and perturbations of its own parameters and activations, an idea pushed further by architectures that interleave graph and image filters. The chapter therefore presented graphs as tools for analyzing, controlling, and improving conventional DNNs, rather than as objects to be learned or processed for their own sake.

While the chapter did not have any impact, the work on GSP-based filtering of data for deep learning served as a preprocessing step in some works for performance improvement [PAN *et al.*, 2023 ; KANG *et al.*, 2024 .

Finally, we explored the use of GSP for contrastive learning, by reading a batch of augmented views as a graph. This work was conducted with Ilyass MOUMMAD, PhD student in the BRAIN team, and was presented during a workshop as follows:



Ilyass MOUMMAD, Bastien PASDELOUP, Nicolas FARRUGIA (2023). “**Contrastive Learning Using Random Walk Laplacian Matrix**”. *GSP Workshop*

This work proposed a contrastive SSL objective derived from the random-walk Laplacian of a graph of augmented views. Within a training batch, each augmented view of an image is a vertex, edge weights are given by the exponentiated inner product of the encoder outputs, and the row-normalized adjacency matrix is read as a matrix of transition probabilities between views. The loss then simply pushes the transition probability toward one between two views of the same image, and toward zero between views of different images, which yields a formulation close to the SimCLR objective but without its logarithm and without a temperature parameter. Evaluated by linear probing on CIFAR-10, with a ResNet-18 backbone trained for 100 epochs using standard image augmentations, the method reached 84.03% accuracy against 80.17% for SimCLR under the same budget, while remaining behind a competing Laplacian-eigenmap approach trained eight times longer. This preliminary work therefore suggested that reading a batch of augmentations as a random walk on a graph of views can yield a simpler contrastive objective, and left the study of a temperature parameter, of alternative similarity measures, and of transfer to downstream tasks for future work.

II.1.6 Application: GSP for Brain Signals

My works in GSP gave me opportunities to work on complex signals, evolving on graphs. Such representations are very common in numerous applications, where multiple variables are considered, with their relationships analyzed. This is notably true in neuroscience, where researchers are interested in measuring the activity at some specific locations of the brain, and analyzing co-activity of these areas to classify events, or – more fundamentally – to understand its functioning.

At the time of my PhD, two neuroscientists were present in the team: Nicolas FARRUGIA and Mathilde MÉNORET. I had the opportunity to work with them, to compare graph inference approaches on fMRI datasets. This work led to the following publication:

[C₈]

Mathilde MÉNORET, Nicolas FARRUGIA, Bastien PASDELOUP, Vincent GRIPON (2017). **“Evaluating Graph Signal Processing for Neuroimaging Through Classification and Dimensionality Reduction”**. *GlobalSIP* (pp. 618–622), IEEE

This article evaluated whether GSP can improve the decoding of fMRI recordings, taken volume by volume over the regions of a brain atlas, by using graph-based representations as a dimensionality-reduction step whose quality is measured through downstream classification accuracy. Seven graphs were built over these regions, from their geometric proximity, from resting-state functional connectivity, or from a combination of both, and two reduction strategies were considered: graph sampling, which selects the vertices where the energy of the signal is the most concentrated over a frequency band of interest and classifies in the vertex domain, and a projection onto the graph Fourier basis followed by a statistical selection of the components. Experiments were conducted on a simulated fMRI dataset and on a real open dataset, and compared these representations with conventional dimensionality-reduction techniques. Two tendencies stood out, neither of them without exception: among the mixed graphs combining geometry and functional connectivity, one outperformed the graphs built from either source alone while the other fell behind both, and graph sampling outperformed principal and independent component analysis on the most difficult simulated group, whereas the statistical selection prevailed on the easiest one and no comparison reached significance on the real dataset.

Another use of GSP for brain signals I studied was to try to determine whether an EEG signal would deviate from a stationary process on a graph across time, thus indicating a change in brain activity. This work, which I conducted with Ali DADRAS, led to the following workshop publication:

[W₅]

Ali DADRAS, Bastien PASDELOUP (2018). **“Classifying Non-Stationary Time-Varying Signals on Graphs Through Their Deviation from a Stationary Behavior: Preliminary Results on EEG Signals”**. *GSP Workshop*, EPFL

This article introduced a framework for characterizing non-stationary, time-varying graph signals through their deviation from a stationary behavior. For a given EEG recording, the graph is taken to be the covariance matrix of the electrodes, so that stationarity amounts to the sample covariance sharing its eigenvectors with that matrix. Sliding windows are extracted from the recording, a sample covariance matrix is computed for each of them, and the eigenvalues of these matrices are projected onto the eigenbasis of their average; because the windowed covariances are not aligned with that common basis, the residual left by the best diagonal approximation – measured through its spectral radius – provides a local measure of the deviation from stationarity, while the sequence of projected eigenvalues traces a trajectory expected to be characteristic of the class. These preliminary results, obtained on the motor-imagery dataset of the Berlin BCI group, showed that peaks in the deviation curve coincide with visible patterns in the corresponding EEG signal, suggesting that the measure locates moments characteristic of a particular activity. Extracting discriminant features from these trajectories, and thereby actually classifying the recordings, was identified as the next step of the work.

The work in [C₈] was one of the first successful attempts at demonstrating the interest of GSP for neuroscience, and often serves as an example of applications of GSP in practical situations [ORTEGA *et al.*, 2018 ; MATEOS *et al.*, 2019 ]. It also inspired similar works such as [MORTAHEB *et al.*, 2019 ] on high-density EEG, and served as a basis for co-authors’ subsequent research [LIOI *et al.*, 2021 ].

Additionally, this collaboration with neuroscientists was for me a first step in the world of brain signals and – more generally – healthcare-related applications. During my PhD, I leveraged the experience I gained from this

work to evaluate the methodology proposed in Section II.1.4 on fMRI data [P₂]. Later, fMRI and EEG signals became a recurrent subject for other of my contributions in applied AI.

II.2 Artificial Intelligence

II.2.1 The Domain in a Coconutshell

Before trying to define AI in a short introductory section, let me make it clear that this will be very shallow, and will miss tens of thousands of relevant references. This section will mostly try to situate elements relevant to my work in the constellation of subdomains of AI.

So first, what do I mean by AI? Since the birth of the field [TURING, 1950 ; DARTMOUTH , the term has changed meaning on several occasions over time based on prevalent approaches [RUSSELL *et al.*, 2021 ] and does not carry the same meaning depending on whether you ask a chess player, a science fiction addict, or a student in the LLMs era. In this document, I will focus on technical definitions and – depending on context – will use AI as one or the other of the following definitions:

Definition 1 Artificial Intelligence (AI)

(~ 2010 – 2020)

A broad family of computational methods for performing tasks traditionally associated with intelligent behavior, mostly dominated by ML approaches that learn representations and decision functions from data. Neural network-based approaches – such as CNNs, GCNs or transformers – became particularly prominent during this period.

(~ 2020 – current)

A broad family of computational systems capable of learning, generating, reasoning, and acting toward specified objectives, increasingly built around foundation models, generative models, and autonomous or agentic systems. These systems extend data-driven learning with generation, interaction, tool use, and adaptation across tasks and modalities.

Over the past two decades, the field has undergone several major paradigm shifts. The first was the progressive dominance of DNNs over more traditional approaches based on hand-crafted features and task-specific representations, following breakthroughs in image recognition and subsequently extending to a broad range of domains [KRIZHEVSKY *et al.*, 2012 ; LECUN *et al.*, 2015 ; GOODFELLOW *et al.*, 2016 ]. The success of deep learning was driven not only by advances in architectures and optimization, but also by the increasing availability of computational resources and large-scale datasets [DENG *et al.*, 2009 ; SHOEYBI *et al.*, 2019 ]. This period progressively shifted the emphasis from designing task-specific representations toward learning representations directly from data [BENGIO *et al.*, 2013 ].

A more recent shift has been from training models for individual downstream tasks toward pretraining increasingly large, general-purpose models that can subsequently be adapted to a wide range of tasks. This paradigm, formalized under the notion of “*foundation models*” [BOMMASANI *et al.*, 2021 , has substantially changed the way models are developed and used: rather than learning a new model from scratch for each task, a common pretrained model can serve as a basis for prompting, fine-tuning, few-shot adaptation, or other forms of downstream specialization [BROWN *et al.*, 2020 ; LIU *et al.*, 2023 ]. This evolution has also been accompanied by a rapid increase in the scale of both models and training datasets, with empirical scaling laws demonstrating systematic relationships between model size, dataset size, compute, and performance [KAPLAN *et al.*, 2020

[[1](#)]; [HOFFMANN et al., 2022](#) [[1](#)]. Consequently, the availability and curation of large-scale data have become increasingly central to the development of competitive AI systems.

More recently, these models have increasingly been extended beyond passive prediction or generation toward systems capable of interacting with external environments and software. Agentic systems can combine model-based reasoning with planning, external memory, retrieval, code execution, and tool use [[WANG et al., 2024](#) [[1](#)]]. At the same time, multimodal models increasingly integrate information from several modalities, such as text, images, audio, video, and structured data, within shared or interacting representations [[LI et al., 2024](#) [[1](#)]; [LIU et al., 2024](#) [[1](#)]]. These developments blur the traditional boundaries between subfields such as NLP, computer vision, speech processing, and RL, and contribute to an increasingly unified view of AI systems.

The landscape of contemporary AI is therefore not limited to the development and training of models, but also encompasses numerous non-exclusive research directions concerned with how models are understood, adapted, deployed, and integrated into larger systems. These include, among others, model compression and efficient inference, mechanistic interpretability, world models, data-centric AI, alignment and safety [[OUYANG et al., 2022](#) [[1](#)]], and agentic systems (see *e.g.*, [[SAJADIEH et al., 2026](#) [[1](#)]; [TGRS](#) [\[1\]](#)] for further examples). Rather than forming a strict taxonomy, these directions cut across model architectures, learning paradigms, application domains, and stages of the AI lifecycle.

Despite the rapid progress of AI, several fundamental challenges remain open. Current general-purpose models exhibit highly uneven capabilities: they can solve increasingly complex problems while still failing on seemingly simple tasks, producing hallucinations, making basic reasoning errors, and struggling with long-horizon or physical-world interactions [[AI SAFETY REPORT](#) [\[1\]](#)]. Improving the reliability, robustness, and generalization of these systems therefore remains a major challenge, particularly as models are increasingly deployed in high-stakes settings and as agentic systems gain the ability to autonomously interact with external tools and environments. At the same time, the continued scaling of models raises questions about the availability of high-quality training data, computational resources, and energy, while also motivating research into more efficient architectures, training procedures, inference, and model compression [[KAPLAN et al., 2020](#) [[1](#)]; [HOFFMANN et al., 2022](#) [[1](#)]]. Another increasingly important challenge is evaluation: traditional benchmarks often provide an incomplete picture of real-world capabilities, may suffer from data contamination, and do not necessarily capture the reliability of systems operating over long sequences of actions or in changing environments [[AI SAFETY REPORT](#) [\[1\]](#)]. Finally, the increasing autonomy and societal deployment of AI systems raise broader challenges concerning interpretability, privacy, security, fairness, alignment, and human oversight [[AI SAFETY REPORT](#) [\[1\]](#); [HUGGINGFACE INCIDENT](#) [\[1\]](#)]. Addressing these challenges will likely require progress not only in model architectures and learning algorithms, but also in data curation, evaluation methodologies, system design, and our ability to understand and control increasingly capable models.

My transition toward AI from more traditional areas of computer science and signal processing approximately coincided with my appointment as an Associate Professor. Consequently, many of the contributions presented in this section emerged from research projects conducted with the interns, PhD students, engineers, and post-doctoral researchers whom I supervised. My role in these works has therefore often been to formulate research directions, discuss methodological frameworks, and challenge their implementation, in close collaboration with the researchers involved.

II.2.2 Few-Shot Learning

Within this vast landscape of AI, I have chosen to focus my research on specific topics that I believe are both relevant to my applications of interest and sufficiently fundamental to have broader applicability. I have also sought to connect these topics to flagship domains, particularly vision and language, so as to leverage and contribute to the most recent advances in the field.

FSL was thus a natural candidate for my research. At a time when major advances in AI were increasingly driven by scaling model size, training data, and computational resources [[KAPLAN et al., 2020](#) [[1](#)]; [HOFFMANN et al., 2022](#)

[10]; BOMMASANI *et al.*, 2021 [10]), FSL instead addressed the complementary problem of adapting models to new tasks or domains when only a small number of labeled examples are available [VINYALS *et al.*, 2016 [10]; SNELL *et al.*, 2017 [10]; FINN *et al.*, 2017 [10]; SONG *et al.*, 2023 [10]]. This setting is particularly relevant to real-world applications in which obtaining large annotated datasets may be prohibitively costly, time-consuming, or otherwise impractical. Moreover, FSL has historically been strongly connected to the flagship modalities of AI, with much of its methodological development and benchmarking taking place in computer vision and, increasingly, NLP [WANG *et al.*, 2020 [10]; BROWN *et al.*, 2020 [10]].

The PhD of Yassir BENDOU was central to this topic. His interests in both vision and LLMs, as well as his strong technical and theoretical skills, made several important contributions to FSL possible:



[T₆] Yassir BENDOU (2024). **“A Data Modeling Perspective on Few-Shot Image Classification”**. *École Nationale Supérieure Mines-Télécom Atlantique (PhD)*

In his PhD, Yassir BENDOU proposed to study few-shot classification from the perspective of explicitly modeling the data distribution in the feature space, rather than focusing solely on the design of increasingly complex classifiers. The thesis first investigated how strong representations and simple inference strategies can provide competitive few-shot performance, leading to EASY, an ensemble-based approach combining SSL, data augmentation, multi-crop inference, and simple classifiers. A second contribution revisited the adaptation of large VLMs through a kernel perspective, showing that competitive training-free adaptation can be obtained by choosing an appropriate kernel over the available samples. The thesis then addressed the problem of understanding and predicting generalization in the few-shot regime by introducing a statistical model linking the geometry and class statistics of the feature space to the expected classification accuracy, providing a way to estimate performance without resorting to cross-validation. The fourth and last core contribution introduced zero-shot one-class classification as a problem of its own, together with a dedicated benchmark. Additional contributions explored how information about unseen classes can be inferred from auxiliary modalities – textual descriptions combined with CLIP representations to predict latent class statistics – and the disambiguation of one-shot visual tasks.

A first work we did was to try to identify which elements in an image could create ambiguity in a few-shot image classification task. To do so, we used random crops, and studied the distribution of their embeddings in feature space, with the objective of distinguishing distinct objects at small cost (note that detection/segmentation could also have been used without this constraint, *e.g.*, using ViTs [DOSOVITSKIY *et al.*, 2021 [10]; CARON *et al.*, 2021 [10])). This led to the following contribution:



[C₁₇] Yassir BENDOU, Lucas DRUMETZ, Vincent GRIPON, Giulia LIOI, Bastien PASDELOUP (2022). **“Le Manchot, la Banane et la Bibliothèque... (de la Désambiguïsation d’une Tâche de Classification avec un Exemple)”**. *Colloque du GRETSI (pp. 1313–1316)*

This work identified a source of ambiguity specific to one-shot image classification: a labeled image may contain several objects, while only one of them corresponds to the class label. Rather than treating this ambiguity as noise, the approach exploits it by analyzing the distribution of representations obtained from multiple random crops of the image. Using a pretrained feature extractor, these representations are modeled as lying within a simplex whose vertices correspond to the different objects present in the scene. The vertices are estimated through an optimization problem combining a reconstruction term with a regularization penalizing the perimeter of the simplex, so as to keep its vertices close to actual crops, while the number of vertices is automatically selected from the reconstruction error by an elbow criterion. The estimated vertices are then shrunk toward the average crop representation, and

the resulting object-level representations are used for nearest-neighbor classification between support and query images. Experiments on 10^5 one-shot 5-way problems from Mini-ImageNet showed that this simple disambiguation strategy slightly improves upon the previous state of the art, while also providing interpretable evidence of the different objects contributing to an image’s representation, and a way to identify the ambiguous images that act as confusion factors in a dataset.

As we could achieve good performance with such a simple approach, we asked ourselves what the sufficient elements in the FSL literature were to build an efficient classification pipeline, as it contained numerous tricks, preprocessing operations, normalizations, etc. that were not clearly useful in the general case. This research led to the development of such a simple model, published as follows:



Yassir BENDOU, Yuqing HU, Raphael LAFARGUE, Giulia LIOI, Bastien PASDELOUP, Stéphane PATEUX, Vincent GRIPON (2022). **“EASY: Ensemble Augmented-Shot Y-shaped Learning: State-of-the-Art Few-Shot Classification with Simple Components”**. *Journal of Imaging* (vol. 8, nb. 7, art. 179), MDPI

This article investigated how far few-shot image classification could be pushed through a careful combination of simple and well-established techniques, motivated by the observation that many methods were evaluated with relatively weakly trained backbones. EASY was introduced as a pipeline combining SSL backbone training, data augmentation, multi-crop inference, and ensembling. The backbones are trained on the base classes using a Y-shaped architecture, jointly optimizing the standard classification objective and a rotation-prediction task, together with manifold mixup. Several independently initialized backbones are then frozen and used to extract features from novel images: multiple random crops are processed by each backbone and their representations are averaged, while the resulting representations from different backbones are concatenated, the size of each backbone being reduced so that comparisons are made at an equal total number of parameters. After centering and projection onto the hypersphere, classification is performed using a simple nearest-class-mean classifier in the inductive setting or soft K -means in the transductive setting. Extensive ablation studies showed that the augmented samples obtained by multi-crop inference were the only ingredient to help consistently, the complete pipeline not always being the most efficient combination – the backbones being suspected of carrying too few parameters for the largest datasets. The method nonetheless reached or surpassed state-of-the-art performance on standard few-shot benchmarks, in some cases improving previous results by more than one percentage point, and the advantage persisted in the harder transductive setting where the number of queries per class is unknown. Beyond the resulting performance, the main contribution was to establish a strong, simple, and reproducible baseline, highlighting the importance of evaluating increasingly sophisticated few-shot algorithms against a well-trained representation and a straightforward inference pipeline.

This article – which received the MDPI Journal of Imaging best paper award – was highly influential in the community, establishing a simple and efficient baseline against which subsequent approaches could be compared. For example, [JIANG *et al.*, 2025

the best-performing methods on few-shot image classification benchmarks [[FEW-SHOT LEADERBOARD](#) ; [OPEN CODE PAPERS](#) .

Following these works, we continued investigating properties of representations in the feature space. In the following work, we tried to evaluate whether a pretrained ResNet [[He et al., 2016](#) C₂₃]:

 [[C₂₂](#)] Yassir BENDOU, Vincent GRIPON, Bastien PASDELOUP, Giulia LIOI, Lukas MAUCH, Stefan UHLICH, Fabien CARDINAUX, Ghouthi BOUKLI-HACENE, Javier ALONSO GARCIA (2023). **“A Statistical Model for Predicting Generalization in Few-Shot Classification”**. *EUSIPCO* (pp. 1260–1264), IEEE

This work proposed a statistical model to predict the generalization error of a few-shot classifier directly from the few labeled samples available for the task, thereby avoiding the validation set that few-shot settings cannot afford. Each class is modeled as a Gaussian in the feature space, and for the binary isotropic case an analytical expression of the probability of error of a nearest-class-mean classifier is derived, depending only on the distance between the class means and the within-class standard deviation; beyond two classes, the error is estimated by Monte-Carlo sampling in the $(n - 1)$ -dimensional subspace spanned by the differences of class means, which is the only subspace affecting the decisions of such a classifier. A central observation is that the naive plug-in estimator of the distance between class means is positively biased, which systematically underestimates the probability of error in high dimension with few samples; the article therefore derived the expression of that bias, proposed an unbiased estimator, and realized the corrected distances through multidimensional scaling before sampling. On 10^3 five-class problems drawn from Mini-ImageNet, Tiered-ImageNet and Meta-dataset with a ResNet-18 backbone, the approach predicted accuracy more faithfully – in terms of mean absolute percentage error – than leave-one-out cross-validation in the low-sample regime, and than the Davies-Bouldin index on every dataset but Mini-ImageNet, where that baseline stays competitive, while the variant without bias correction failed, confirming that the correction is the key ingredient. Beyond providing a practical alternative to cross-validation, the work offers a statistical interpretation of few-shot generalization by identifying the separation and compactness of class distributions as the factors governing classification performance.

This work was nominated among the finalists of the best student paper award at EUSIPCO 2023.

Finally, Yassir BENDOU explored the added value of textual information for few-shot image classification. Using a similar idea as in [[C₂₂](#)], we used textual descriptions of classes to estimate the distribution of the feature representations of data. This work was presented in the following workshop:

 [[W₁₁](#)] Yassir BENDOU, Vincent GRIPON, Bastien PASDELOUP, Giulia LIOI, Lukas MAUCH, Fabien CARDINAUX, Ghouthi BOUKLI-HACENE (2023). **“Inferring Latent Class Statistics from Text for Robust Visual Few-Shot Learning”**. *NeurIPS R0-FoMo Workshop*

This work investigated whether semantic information available in text can be used to infer the latent statistics of visual classes and thereby improve few-shot image classification. Using CLIP to obtain a shared representation of images and text, the empirical mean and a diagonal covariance are first estimated for each class of a base dataset, after which two small MLPs are trained to regress these statistics from the textual features alone – either a simple prompt built from the class label, or a more informative visual description

generated by an LLM, the latter proving preferable on datasets with uncommon class names. At test time, the predicted mean is interpolated with the empirical mean of the few available shots, the predicted covariance is shrunk toward the identity, and classification is performed with a Mahalanobis distance, so that the text-derived statistics compensate for the limited information carried by the support set. Experiments on base datasets of natural and fine-grained images, and on nine cross-domain benchmarks, showed that incorporating text-derived class statistics improves robustness, particularly when the visual support set is extremely small or poorly representative of the underlying class distribution, and that the covariance matters most in the one-class setting where a single class has to be distinguished from everything else.

We later extended the idea to the labels only, by prompting an LLM for counter-examples that resemble the target class and help define the decision boundary. This unpublished work is available in [P₄].

Some time after these works, with the emergence of foundation models for image segmentation [KIRILLOV *et al.*, 2023 

 [C₂₈] Jonathan LYS, Frederic LIN, Clément BÉLIVEAU, Jules DECAESTECKER, Yassir BENDOU, Aymane ABDALI, Bastien PASDELOUP (2024). “**FICUS: Few-Shot Image Classification with Unsupervised Segmentation**”. *EUSIPCO (pp. 1791–1795), IEEE*

This article investigated whether unsupervised image segmentation could improve few-shot image classification by explicitly separating the potentially different semantic objects contained in an image. FICUS was introduced as a fully training-free framework that decomposes each query image into a set of regions and extracts a representation for each of them with a pretrained visual encoder. The decomposition chains existing pretrained models: eigenmaps produced by the deep spectral method provide the prompt points for a class-agnostic segmentation model, the best of the three candidate scales it returns for each prompt is selected by IoU with an Otsu thresholding of the eigenmap, redundant masks are removed by non-maximum suppression, and the original image is kept among the resulting crops so that a segmentation failure cannot degrade the prediction. Rather than aggregating the region-level representations into a single descriptor, classification assigns to a query the label of the class prototype closest to any of its crops, which amounts to selecting the most discriminative region of the image. The approach was evaluated in the inductive setting on PascalVOC, CUB and ImageNet, and compared both with an automatic mask generator and with an oracle using ground-truth bounding boxes that bounds the gain achievable by cropping. Experiments showed that explicitly accounting for the spatial structure and multiple visual components of an image improves the performance of state-of-the-art few-shot classifiers, particularly when the object defining the class occupies only part of the image. The improvement over the automatic mask generator is slight, however, and the oracle remains well above both, so that better localization stays an open direction.

This project, as was the case for many of my previous works in FSL, integrated questions regarding ambiguity (here, caused by multiple objects in an image), which is a topic I have been enjoying throughout my research.

This question was also central to a collaboration I did with Raphael LAFARGUE during his PhD in the BRAIN team, though regarding ambiguity caused by classes. The central question of our work was to evaluate whether

some classes could act as confusing factors for the few-shot classification of elements of another class, and to make the model disregard them when producing feature representations, to improve few-shot classification. This work led to the following article and its unpublished extension [P₆]:

 [C_{2I}] Raphaël LAFARGUE, Bastien PASDELOUP, Jean-Philippe DIGUET, Vincent GRIPON (2022). “Classes Adversaires dans l’Apprentissage avec Peu d’Exemples”. *Colloque du GRETSI (pp. 197–200)*

This work questioned the widely held assumption that a pretraining dataset should be as large and as diverse as possible, and showed that in few-shot image classification less can sometimes be more. The notion of “*adversarial class*” was introduced to designate a class of the generic pretraining dataset whose presence during training degrades the ability to discriminate a given pair of novel classes, so that removing it from pretraining improves few-shot accuracy. To explain when this happens, the quality of the feature space is modeled through a signal-to-noise ratio, defined as the distance between the two novel class means divided by their average intra-class standard deviation, and any feature vector is decomposed over the approximately orthogonal directions associated with the generic classes. Three mechanisms follow, according to whether a generic class is orthogonal to the discriminative direction – in which case removing it only affects the noise – or shares features with both novel classes, and is therefore adversarial, or is close to only one of them, and is therefore beneficial. Experiments on Mini-ImageNet trained 65 ResNet-12 extractors, one on all 64 generic classes and one for each class removed, and measured for every triplet the gain in few-shot accuracy over 20,000 episodes together with the gain in signal-to-noise ratio. The two were significantly correlated, and an RF using the signal-to-noise ratio together with the cosine similarities within the triplet predicted the gain with a coefficient of determination of 55%; interpretable cases were exhibited, such as a gain of more than 3.5% on discriminating dalmatians from lions once a class of underwater images is removed. This work therefore shifted the focus from the quantity of pretraining data to its composition, questioning the existence of universal feature extractors and suggesting that domain knowledge used to select pretraining classes can significantly improve few-shot performance.

Later, Raphaël LAFARGUE proposed a very interesting analysis of a widespread benchmarking issue in the literature of FSL, showing a real problem of underestimation of confidence intervals, due to the commonly accepted procedure accounting for the sampler randomness but not the data itself [LAFARGUE *et al.*, 2024 

In addition to this issue, the rapid development of foundation models has progressively changed the role of FSL in ML. Originally motivated by the difficulty of learning new visual concepts from only a handful of labeled examples, FSL increasingly benefits from representations acquired through large-scale pretraining, which can substantially reduce or even eliminate the need for task-specific examples. This transition was particularly evident with VLMs such as CLIP, which demonstrated that large-scale multimodal pretraining can provide strong zero-shot transfer to previously unseen visual classification tasks simply by associating images with textual descriptions of the target classes [RADFORD *et al.*, 2021 et al., 2020 et al., 2022 et al., 2021 et al., 2024 

Overall, my work on FSL has mainly focused on understanding how the structure of the data and its representation influence generalization when only a few labeled examples are available. This and the frequent background question of confusing elements across my works led me to propose the “*Encouraging Diversity in FSL*” project. In a few words, **ENDIVE** aims at exploring what diversity (of data or features) can bring to AI, and especially for FSL, through the use of random sampling procedures. This project is described in more detail in the following sections and in [Section III.I.I.I.](#)

II.2.3 From FSL to LLMs

Following [Yassir BENDOU](#)’s work on the inclusion of text as an important modality for processing images, I decided to focus a larger part of my research on LLMs. This transition was motivated by the increasingly central role of language in multimodal foundation models. The success of large-scale VLMs demonstrated that text can provide not only an additional source of information for visual tasks, but also a semantic interface through which visual concepts can be specified and transferred to previously unseen tasks in a zero-shot manner [[RADFORD et al., 2021](#)]. This role has progressively expanded, with multimodal foundation models using language as an interface to combine visual and textual information and, more broadly, to integrate heterogeneous modalities [[ALAYRAC et al., 2022](#); [LI et al., 2023](#); [GIRDHAR et al., 2023](#)]. In this context, the few-shot setting that motivated much of my earlier work has increasingly become one regime within a broader paradigm in which large pretrained models provide substantial prior knowledge and can be adapted through instructions, prompts, or a small number of examples. This evolution naturally led me to consider LLMs as central research objects, rather than as tools to use as support for other projects.

However, research on LLMs is not an easy domain for a small research team to compete in, given the extremely rapid pace of progress and the substantial computational and data resources required to train and evaluate state-of-the-art models. The strong scaling trends observed for language models [[KAPLAN et al., 2020](#); [HOFFMANN et al., 2022](#)] have contributed to an increasingly competitive environment in which progress is closely tied to access to large-scale computing infrastructures and massive training corpora. This makes it difficult for academic teams without comparable resources to compete directly on model scale. Rather than attempting to reproduce this race toward increasingly large models, it therefore appeared more promising to focus on questions where smaller teams can provide distinctive methodological contributions. For these reasons, I chose to devote greater attention to diffusion models, which have emerged as an interesting alternative to the more common autoregressive models. These two approaches differ as follows:

- **Autoregressive language models** — These models generate text sequentially by factorizing the probability of a sequence into a product of conditional probabilities, with each token predicted from the tokens generated previously. This formulation, initially introduced by [[BENGIO et al., 2003](#)] for neural networks, underlies modern LLMs such as GPT [[RADFORD et al., 2018](#); [BROWN et al., 2020](#)] and naturally supports left-to-right generation and causal modeling [[ZHAO et al., 2026](#)].
- **Diffusion language models** — In contrast, diffusion models formulate generation as an iterative denoising process. As for images [[HO et al., 2020](#)], a forward process progressively corrupts a text sequence, typically by masking or replacing tokens – in the spirit of the masked prediction objective of BERT [[DEVLIN et al., 2019](#)] – while a neural network learns to reverse this corruption and recover the original sequence from an increasingly noisy representation [[HOOGEBOOM et al., 2021](#); [AUSTIN et al., 2021](#)]. Rather than committing to one token at a time, diffusion models can therefore predict and refine multiple positions simultaneously, potentially allowing information from the entire sequence to contribute to each prediction. This difference provides greater flexibility in the generation process, including the possibility of iterative refinement and non-autoregressive generation, although it also introduces additional computational steps during sampling. Recent diffusion language models such as D3PM [[AUSTIN et al., 2021](#)], DiffusionBERT [[HE et al., 2023](#)], and SEDD [[LOU et al., 2024](#)] have demonstrated that this alternative formulation can support competitive language generation while opening different possibilities for controllable generation, infilling, and efficient decoding [[LI et al., 2025](#)].

Three of my current PhD students are focusing a large part of their research on LLMs, and especially on the discrete diffusion ones described in the remainder of this section.

Historically, this axis of my research became very prevalent thanks to [Jonathan Lys](#), whom I recruited for a PhD after his earlier internship, on topics related to the [ENDIVE](#) project. As mentioned earlier, FSL was slowly becoming a subproblem of foundation model adaptation, rather than a paradigm of its own. This change, along with the very strong interest of [Jonathan Lys](#) in diffusion-based approaches, led me to partly reorient the [ENDIVE](#) project to give more space to LLMs, still under the prism of diversity.

Before introducing our work on diversity-related aspects, I want to introduce two contributions we made to try to improve pretrained autoregressive models. First, as data flows through the model, we asked ourselves whether looping on some layers could leverage latent capabilities of the model to improve refinement of the predicted token. This led to the following publication:

 [C₃₅] Jonathan LYS, Vincent GRIPON, Bastien PASDELOUP, Axel MARMORET, Lukas MAUCH, Fabien CARDINAUX, Ghouthi BOUKLI-HACENE (2026). **“Inner Loop Inference for Pretrained Transformers: Unlocking Latent Capabilities Without Training”**. *EUSIPCO, IEEE*

This work introduced a training-free way of allocating additional computation at inference time in a frozen pretrained transformer, by repeatedly re-applying a selected range of its blocks to the propagated hidden state. The starting point is the view of depth as a sequence of refinement steps applied to a latent state carried by the residual stream, which suggests that further refinement steps could be obtained without modifying any parameter. Applied naively, however, such looping degrades accuracy, the interpretation put forward being that re-executing layers pushes activations away from the manifold encountered during standard inference; the contribution is therefore a lightweight regularization that interpolates the looped states with their non-looped counterparts, evaluated in three variants and controlled against an ablation replacing the loop-induced signal with random noise of matched magnitude. On WinoGrande, ARC-Easy and ARC-Challenge, GSM8K, HellaSwag and MMLU, with the loop region selected once on a single benchmark and consistently found at roughly 40 to 60% of the relative depth, inner looping yielded modest but consistent accuracy gains for the two Gemma-2 models, while the pre-norm LLaMA-3-8B responded more variably. Visualizing the latent trajectories suggested that looping induces a small but coherent deviation rather than a qualitatively different computation, which is consistent with the magnitude of the observed gains.

This article already had some impact, and serves as a building block for other approaches aiming at developing new models integrating looping approaches [[CHEN *et al.*, 2026](#) SHIN *et al.*, 2026 

Additionally, we noticed a fundamental misalignment between the causal structure imposed by autoregressive attention and the way information is propagated through the residual connections of transformer architectures [[VASWANI *et al.*, 2017](#) 

[C₃₆]

Jonathan LYS, Vincent GRIPON, Bastien PASDELOUP, Axel MARMORET, Lukas MAUCH, Fabien CARDINAUX, Ghouthi BOUKLI-HACENE (2026). **“Residual Connections and the Causal Shift: Uncovering a Structural Misalignment in Transformers”**. *EUSIPCO, IEEE*

This work investigated a structural tension in autoregressive transformers: the hidden state at a given position is initialized from the embedding of the current token and continuously propagated by the residual connections, whereas the supervision target is the next token. The article first localized where representations stop being anchored to the input token and start being oriented toward the prediction target, by decoding intermediate hidden states through the tied embedding matrix, and by measuring cosine similarities and normalized projections along the input-output axis over 1,000 Wikitext sequences. Three regimes appeared – early layers decoding to the input sequence, intermediate layers decoding to nothing interpretable, and late layers decoding to the one-token-shifted sequence – with the transition occurring deep in the network, around layer 17 out of 26 for Gemma-2-2B, and at qualitatively similar relative depths for LLaMA-3.2-3B and Mistral-7B. Building on this observation, a lightweight mitigation was proposed: attenuating the residual branch by a scalar factor, either as a fixed single-layer intervention or through a mixture-of-depth-inspired gate learned jointly with the model. Evaluated by training 150M-parameter GPT-2-style models from scratch on 10B tokens, the learned gate progressively concentrated on attenuating the final layer, outperformed the fixed single-layer cuts – themselves very sensitive to the layer chosen – on Wikitext, LAMBADA and OpenWebText, and reduced the measured input-output misalignment. The analysis is confined to models whose embedding matrix is tied, which is no longer the rule at the largest scales.

While the works above focused on autoregressive models, I believe that discrete diffusion LLMs have a great positioning in the vast domain of AI, as they are currently overlooked by main actors of the field, who have already put a lot of effort into autoregressive models. Still, at the same time, they provide a basis for working on fundamental results on models that process text which, as introduced earlier, I have identified as an important central modality for multimodal approaches.

Coming back to the [ENDIVE](#) project, an important work we did with [Jonathan Lys](#) was to try to reproduce a mechanism similar to beam search [[LOWERRE, 1976](#) ; [WISEMAN et al., 2016](#)

[R₃]

Jonathan LYS, Vincent GRIPON, Axel MARMORET, Lukas MAUCH, Fabien CARDINAUX, Ghouthi BOUKLI-HACENE, Bastien PASDELOUP (2026). **“D5P4: Partition Determinantal Point Process for Diversity in Parallel Discrete Diffusion Decoding”**. *NeurIPS*

This work introduced D5P4, a decoding framework designed to improve the diversity of candidate generations in discrete diffusion language models. Unlike autoregressive models, where beam search naturally operates over left-to-right prefixes, discrete diffusion models iteratively refine complete sequences through parallel denoising, making standard beam-search strategies difficult to apply and providing limited control over the diversity of retained hypotheses. The proposed framework generalizes beam search to this setting by generating multiple candidate sequences in parallel and performing selection at each diffusion step. D5P4 formulates this selection as a maximum-a-posteriori inference problem over a partitioned DPP, whose kernel combines the quality signal carried by the diffusion logits with a diversity signal computed from the model’s own hidden representations, so that the whole procedure remains model-internal and requires no external verifier or auxiliary scorer. The partition constraint prevents different

hypotheses from collapsing onto the same diffusion trajectory, an effect known as lineage collapse, while a scalable fast greedy solver keeps the selection compatible with distributed multi-GPU inference: the marginal overhead it introduces is more than offset by a reduction of about 40% in peak memory and by a lower post-forward latency, the categorical distributions having to be materialized on a much smaller scale than in standard decoding. Experiments on open-ended text generation, question answering, and mathematical reasoning showed that D5P4 substantially improves the diversity and pass@ k coverage of generated candidates, while matching or surpassing the generation quality and fidelity of standard and diversity-oriented decoding methods, including compute-matched baselines that rely on external models.

More recently, [Joshua Chaz RIVERA](#) started his PhD on LLMs as well. He is particularly interested in more cross-cutting aspects, such as low-resource languages or style understanding. We are currently revising some works on how LLMs handle metaphorical sentences. To do so, [Joshua Chaz RIVERA](#) first assembled a novel dataset of realistic novel metaphors, presented below:

[R₁]

Joshua Chaz RIVERA, Bastien PASDELOUP, Vincent GRIPON (2026). **“Evaluating LLMs on Sentence-Level Metaphors”**. *Findings of the ACL*

This work introduced MP1K, a dataset of 1,000 sentence pairs designed to evaluate LLMs beyond conventional, lexicalized metaphors. Each pair associates an automatically generated, multi-word, and intentionally unconventional metaphorical sentence with a literal equivalent, the generated sentences having been manually filtered for plausibility and clear metaphoricality. On this dataset, autoregressive and discrete diffusion LLMs were evaluated on three complementary tasks: binary metaphor detection, pairwise metaphor discrimination based on the log-probability margin between the metaphorical and the literal sentence of a pair, and metaphor interpretation through literal paraphrasing; a fourth evaluation, on the existing MUNCH dataset, measured context-sensitive lexical interpretation through single-word substitution, allowing comparison with prior token-level work. Experiments showed that autoregressive models, particularly the larger LLaMA and Gemma models, achieved the strongest pairwise metaphor discrimination, and that the wide gap between this margin and raw multiple-choice accuracy indicates that models detect the relative difference between a metaphorical and a literal sentence without being able to turn it into a reliable binary decision. On the lexical substitution task, diffusion models such as LLaDA and TraDo clearly outperformed the comparable autoregressive model, which was near chance, although the best overall accuracy was still obtained by the largest Gemma model; diffusion models also achieved the best semantic preservation in full-sentence paraphrasing, by a margin the article itself describes as modest, all models struggling to produce high-quality literal rewrites. These results demonstrated that metaphor understanding is strongly task- and architecture-dependent, and that conventional token-level benchmarks can fail to reveal important limitations and differences in the semantic capabilities of modern LLMs.

On a similar topic, another work we did – also under review – was to compare the two families of LLMs presented earlier on texts featuring idioms, to find out their respective processing abilities in such a less usual setting:

[R₂]

Joshua Chaz RIVERA, Bastien PASDELOUP, Vincent GRIPON (2026). **“Idiomatic Blindspots in Discrete Diffusion Models”**. *EMNLP*

This work investigated the ability of discrete diffusion language models to process idiomatic language, focusing on whether their bidirectional denoising mechanism provided advantages for semantic under-

standing compared with autoregressive models. The study evaluated diffusion models through four complementary experiments – open-ended idiom generation, cloze completion, multiple-choice semantic understanding, and a probe injection test – designed to disentangle limitations arising from the architecture from those caused by differences in training data. The experiments showed that diffusion models retained substantial idiomatic knowledge, achieving competitive performance on cloze and multiple-choice tasks, including in several non-English languages and occasionally ahead of larger autoregressive models even in English, despite weaker performance in open-ended generation. This discrepancy was traced to the confidence-based demasking strategy: the probe injection experiment showed that a diffusion model is steered by an incorrect early token in 46.3% of cases against 26.5% for a smaller autoregressive model, with the gap widening precisely where the idiomatic prior is weak, so that early commitments interfere with the recovery of multi-token idiomatic expressions and prevent information already encoded in the model from being expressed during generation. Cross-lingual experiments further indicated that, where autoregressive models retain an edge, the difference is driven by training-data coverage rather than by a fundamental limitation of the diffusion architecture. The work highlighted a distinction between the semantic knowledge encoded by diffusion language models and their ability to retrieve this knowledge during generation, identifying the decoding process as an important source of their current linguistic limitations.

Finally, I supervise the PhD of [Nolan Sisouphanthong-Vongkotrattana](#) on diffusion LLMs. His interests lie between reasoning [[Wei et al., 2022](#)] and mechanistic interpretability [[Olah et al., 2020](#); [Turner et al., 2023](#); [Rai et al., 2024](#)], and he explores both for diffusion models.

Overall, my research on LLMs is fairly recent, but has taken a large place in my interests, due to its prevalence in the AI field. In the long run, I want to integrate text encoders in some of my more applied work in collaboration with healthcare researchers and practitioners, as text has become a central modality that would allow them to leverage many kinds of data in a single prediction task.

I mentioned earlier that working on discrete diffusion LLMs was a great opportunity to contribute to text processing from a relatively niche perspective compared with the dominant autoregressive paradigm. At the time I started working on this paradigm, this represented a relatively open research direction, with only a limited number of models and studies exploring the potential of discrete diffusion for language. This situation has changed remarkably quickly. Models such as Dream-7B [[Ye et al., 2025](#)] and LLaDA [[Nie et al., 2025](#)] have demonstrated that discrete diffusion models can achieve competitive performance with autoregressive models of comparable size, while exploiting their ability to generate tokens in arbitrary orders and to iteratively refine complete sequences [[Ye et al., 2025](#)]. More recently, the release of DiffusionGemma provides an even stronger indication that diffusion language models are moving toward the mainstream of LLMs research: this 26B-parameter mixture-of-experts [[Shazeer et al., 2017](#)] model generates blocks of tokens in parallel and is explicitly designed to explore the speed and flexibility advantages of diffusion-based generation [[Taïga et al., 2026](#)]. This rapid development is also accompanied by an increasingly active effort to understand these models: analyses of their internal representations and mechanisms are appearing shortly after their release, extending mechanistic interpretability beyond the established autoregressive models [[Engels et al., 2026](#)]. Consequently, the niche that initially motivated this line of research is becoming considerably more competitive, illustrating the extremely rapid evolution of the LLM field and the short time scale over which new research directions can move from emerging to mainstream.

Looking ahead, I believe that the future of diffusion LLMs lies on two main axes:

- **Multimodality** — A particularly promising extension of this paradigm is the development of multimodal diffusion models, in which the same iterative generation and refinement framework can operate across text, images, audio, video, and other modalities. While diffusion has already become a dominant paradigm for image and video generation, recent unified models such as Show-o [[Xie et al., 2025](#)] and Transfusion [[Zhou et al., 2025](#)] illustrate complementary approaches, combining autoregressive and diffusion

objectives to handle both discrete text and continuous visual data within a single model. This perspective is especially interesting for language models because continuous representations naturally provide a common interface between heterogeneous modalities: rather than requiring each modality to be generated by a separate architecture, a model could progressively denoise a joint multimodal representation and resolve the dependencies between modalities during inference.

- **Continuous diffusion** — A particularly compelling perspective for the future of continuous diffusion language models is provided by ELF [HU *et al.*, 2026]. Rather than repeatedly projecting continuous representations back onto the discrete vocabulary during denoising, ELF performs almost the entire generation process in a continuous embedding space in continuous time, and postpones discretization to the final step through a shared-weight decoder. This approach to reasoning in continuous space opens several avenues for future research: in particular, understanding the geometry of the embedding space (*e.g.*, for diversity-aware decoding), or investigating whether continuous trajectories can be used for multimodal generation.

The first axis will most likely have a mid-term impact on my applications in healthcare, as data handled in such settings is highly multimodal (*e.g.*, scanners, tabular data, text reports, EEG, ECG). The second axis definitely raises questions that are of interest for Nolan SISOUPHANTHONG-VONGKOTRATTANA, that we will explore together in the near future.

II.2.4 REVE: A Foundation Model for Brain Signals

A final major axis of my research is related to brain signals, and the study of AI for neuroscience. This line of research, initiated during my PhD (see Section II.1.6), took on another dimension with the start of Yassine EL OUAHIDI's PhD, which I co-supervised along with Giulia LIOI, neuroscientist in the BRAIN team.

Neuroscience brings several interesting challenges to AI, because of the specificities of the data handled (generally EEG, MEG or fMRI). These modalities provide complementary views of brain activity, but differ substantially in their spatial and temporal resolutions, acquisition mechanisms, and associated sources of noise. EEG measures electrical activity at the scalp through electrodes, offering millisecond-scale temporal resolution but relatively poor spatial resolution and strong sensitivity to artifacts such as eye movements, muscle activity, electrode impedance, and environmental noise [CHADDAD *et al.*, 2023]. MEG similarly captures neural activity through the magnetic fields generated by neuronal currents, combining high temporal resolution with better spatial localization than EEG, but requiring highly sensitive sensors and controlled acquisition environments because of its vulnerability to external magnetic interference [BAILLET, 2017]. In contrast, fMRI measures the BOLD response indirectly through changes in blood oxygenation, providing high spatial resolution but much lower temporal resolution, and introducing additional variability related to head motion, physiological fluctuations, scanner characteristics, acquisition protocols, and preprocessing choices [LOGOTHETIS, 2008].

These differences have important methodological consequences for AI: EEG and MEG naturally motivate temporal signal processing, source localization, spectral analysis, and spatio-temporal representations, whereas fMRI is commonly approached through voxel- or region-level activation analysis, functional connectivity, and graph-based representations (which led to a strong interest in the GSP community [ORTEGA *et al.*, 2018]). Across all three modalities, the number of measured variables can be very large compared with the number of independent subjects, making dimensionality reduction, regularization, representation learning, and careful cross-validation essential to avoid overfitting. Moreover, neural measurements are not *i.i.d.*: observations are structured by subjects, sessions, trials, acquisition sites, devices, and experimental conditions, resulting in substantial within- and between-subject variability and potentially strong domain shifts. This creates a particularly important challenge for AI, as models may inadvertently exploit subject- or acquisition-specific signatures rather than neural phenomena of interest. Consequently, robust neuroscience AI requires methodologies that explicitly exploit the spatial, temporal, spectral, and relational structure of brain signals while accounting for the specificities of such data – *i.e.*, noise, limited sample sizes, heterogeneous acquisition conditions – and the reproducibility of learned representations across subjects and sessions.

In his PhD, [Yassine EL OUAHIDI](#) was interested in transferring recent advances in flagship domains of AI to brain data, especially EEG signals for BCI tasks. Throughout his PhD, he made numerous contributions at the frontier of AI and neuroscience, described hereafter:

[T₇]

Yassine EL OUAHIDI (2025). **“Deep Learning for BCI Applications: Leveraging Recent Advances in Computer Vision”**. *École Nationale Supérieure Mines-Télécom Atlantique (PhD)*

The thesis of [Yassine EL OUAHIDI](#) investigated how recent advances in deep learning for computer vision can be transferred to BCI applications, with a particular focus on EEG decoding. The work addressed the limited amount of labeled data and strong variability across subjects and recording conditions that make BCI particularly challenging for deep learning. Observing that many BCI architectures had drifted far from the standards of computer vision by accumulating domain-specific modules, the first contribution took the opposite direction with EEG-SimpleConv, a deliberately simple 1D convolutional network inspired by classical vision models, designed so that its latent space would share the properties that make vision techniques transferable, and shown to reach state-of-the-art motor-imagery decoding once paired with an adapted training routine; it was then reused as a backbone for merging datasets with heterogeneous electrode configurations through graph signal interpolation, for selecting meaningful EEG channels, and for aligning EEG with speech in a CLIP-like framework. Building on that latent space, the second contribution transferred FSL techniques from vision to BCI in order to reduce the calibration time required from a new user, and comprised an unsupervised adaptive decoding method that closes the gap between offline and online operation, a few-shot method reaching state-of-the-art results on the BEETL challenge, and a new realistic cross-dataset benchmark limiting the amount of data available from an unseen subject. The third contribution moved beyond BCI to EEG at large, with a review of existing EEG foundation models and with REVE, a pretrained model designed to learn general-purpose representations from large-scale and heterogeneous EEG data, providing a foundation for transferring knowledge across subjects and downstream tasks.

This PhD manuscript builds upon several contributions. First, we took a GSP-based approach, and tried to find meaningful graph frequencies for fMRI signals, using a brain atlas as a support graph. The idea was to see whether known results in the neuroscience literature could be found using signal processing tools alone, to interpret a more systemic approach. This work was published as follows, and also presented in [C₁₈]:

[C₁₉]

Yassine EL OUAHIDI, Hugo TESSIER, Giulia LIOI, Nicolas FARRUGIA, Bastien PASDELOUP, Vincent GRIPON (2022). **“Pruning Graph Convolutional Networks to Select Meaningful Graph Frequencies for fMRI Decoding”**. *EUSIPCO (pp. 937–941), IEEE*

This work exploited GSP to investigate which graph frequencies were most informative for decoding fMRI activation maps, and whether explicitly selecting these frequencies could improve both performance and interpretability. It introduced a spectral ResNet, whose graph convolutions operate in the graph Fourier domain, together with an MLP taking the graph spectrum as input, so as to check that the conclusions did not depend on a single architecture, and progressively removed frequencies with low importance during training using a selective weight decay pruning scheme. Two large datasets of activation maps defined over the regions of a brain atlas were considered, along with a structural and a functional graph. Across the datasets, architectures, and graph types, the experiments consistently showed that low graph frequencies were the most informative: retaining the first 60 frequencies was generally sufficient, while including more frequencies provided little additional benefit. Unlike manually selecting a contiguous

low-frequency band, the pruning procedure identified a sparse and non-contiguous subset of frequencies and consistently improved decoding performance, whereas band selection degraded it in most scenarios. The selected frequencies were also reproducible across datasets and architectures, with IoU scores typically above 50%, and the functional graph outperformed the structural one overall, though not on every dataset and without any universal ranking being claimed, indicating that functional connectivity is generally better aligned with the decoded tasks. Examination of the most frequently selected functional eigenvectors further linked the learned frequencies to interpretable brain systems, including visual, auditory, motor, and default-mode regions.

Continuing on the graph-related approaches, we addressed one of the main limitations for EEG signals: the variability of settings. Indeed, acquisition is performed with devices of various numbers of electrodes, depending on what is available in acquisition centers. Consequently, the variability in data dimension makes it hard to train large models, making it an important bottleneck in AI for neuroscience. In a second work, we thus proposed a method based on interpolation on a learned brain graph to unify data shape. This led to the following publication, also presented in [C₂₅]:

[C₂₄]

Yassine EL OUAHIDI, Lucas DRUMETZ, Giulia LIOI, Nicolas FARRUGIA, Bastien PASDELOUP, Vincent GRIPON (2023). **“Spatial Graph Signal Interpolation with an Application for Merging BCI Datasets with Various Dimensionalities”**. *ICASSP (pp. 1–5), IEEE*

This work addressed the lack of large and homogeneous datasets in EEG-based BCIs, where datasets often differed in the number and spatial arrangement of electrodes, sampling frequencies, and experimental protocols. It proposed a spatial graph signal interpolation method that represented EEG electrodes as nodes of a graph and considered each dataset as a partial observation of a common, larger electrode space. A graph was learned from a dataset of EEG signals containing a nearly complete electrode configuration, and the missing electrodes in the other datasets were reconstructed by minimizing the graph variation of the signal, which admits a closed-form solution. The approach therefore allowed datasets with different electrode configurations to be mapped onto a common virtual electrode layout without discarding electrodes, unlike approaches based on retaining only the intersection of electrode sets. Experiments on five MI EEG datasets compared the proposed approach with spherical-spline interpolation and with graphs built from sensor positions or weighted by a phase-lag index. The learned graph provided the most accurate graph-based reconstruction and remained effective when a substantial number of electrodes were missing, but only after being fine-tuned on the target dataset: transferred directly it performed poorly, and the two other graph constructions both stayed behind spherical splines. Using the proposed unification approach, a multi-dataset training setting inspired by the BEETL challenge was then set up, jointly training a 1D CNN on multiple homogenized datasets. The results were instructive rather than uniformly favorable: interpolating up to the electrode set of the target dataset clearly outperformed restricting all datasets to their common electrodes on one target network, whereas on the other the intersection remained the best option, and interpolating up to the union of all electrodes was never worthwhile. Graph-based interpolation therefore offers an interpretable mechanism for transferring information between BCI datasets of different dimensionalities while preserving the spatial structure of the signals, but its benefit depends on how far apart the electrode configurations actually are.

Extending the idea initiated in this previous work, we then designed a more advanced model to process MI EEG data. Similar to the approach we had in Yassir BENDOU’s work [J₂], we tried to keep the model as simple as possible, and extensively validated all chosen components. This work led to the development of EEG-SimpleConv, published as follows, and also presented in [W₁₂; W₁₃]:



Yassine EL OUAHIDI, Vincent GRIPON, Bastien PASDELOUP, Ghaith BOUALLEGUE, Nicolas FARRUGIA, Giulia LIOI (2024). **“A Strong and Simple Deep Learning Baseline for BCI Motor Imagery Decoding”**. *Transactions on Neural Systems and Rehabilitation Engineering* (vol. 32, pp. 3338–3347), IEEE

This work introduced EEG-SimpleConv, a deliberately simple 1D CNN for MI decoding from EEG. The architecture relied exclusively on standard components – stacked 1D convolutions, batch normalization, ReLU activations, and pooling – and a training procedure was specifically designed to address the strong inter-subject and inter-session variability encountered in BCI data. In particular, the training strategy combined Euclidean alignment of the input trials, normalization statistics computed per session rather than per subject, batch-normalization statistics recomputed on the data of the target subject or session at test time, data augmentation through mixup, and an auxiliary subject-classification head acting as a regularizer. An extensive ablation study showed that normalization dominates: recomputing the batch-normalization statistics was the single most valuable ingredient, followed by mixup, Euclidean alignment, and session-level statistics, whereas the auxiliary head contributed almost nothing. Removing every normalization ingredient at once – the configuration compatible with a genuinely online setting – cost more than six accuracy points on average, and twelve points on the datasets with few subjects against two and a half on the largest ones, which suggests that pretraining on more subjects would help in online scenarios. The model was evaluated on four public MI EEG datasets and compared with established ML and deep learning approaches under within-subject, cross-subject, cross-subject with fine-tuning, and simulated online settings. EEG-SimpleConv achieved competitive or state-of-the-art performance across these evaluation settings while maintaining low inference latency, and the experiments also demonstrated strong cross-subject transfer, the gain brought by fine-tuning on a short calibration being largest for the subjects that are hardest to classify.

This model had a strong impact on the EEG community, as it provided a simple and practical tool to study brain signals. As such, it was often used as a baseline for comparison in subsequent work [SEDI NZAKUNA *et al.*, 2025 ; WANG *et al.*, 2025 ].

Using the EEG-SimpleConv model – which had state-of-the-art performance when released – we tackled another specificity of brain signals: the cross-subject variability. This bottleneck seriously limits the use of models in real-life applications, as pretraining a model does not guarantee its performance on a new subject. The adaptation procedure we developed was published as follows:



Yassine EL OUAHIDI, Giulia LIOI, Nicolas FARRUGIA, Bastien PASDELOUP, Vincent GRIPON (2024). **“Unsupervised Adaptive Deep Learning Method for BCI Motor Imagery Decoding”**. *EUSIPCO* (pp. 1626–1630), IEEE

This work addressed the distribution shift between subjects that limits the practical deployment of EEG-based BCIs. It proposed an unsupervised adaptive decoding method that combined a frozen EEG-SimpleConv backbone with continuous alignment of incoming EEG data, avoiding both retraining and the need for labels from the target subject. The method adapted the input representation using Euclidean alignment, whose reference covariance matrix is recomputed each time a new trial arrives, while also dynamically updating the batch-normalization statistics within the network from the current trial and all previous ones of the same session, thereby progressively adapting the decoder to the statistics of the target subject. Because both quantities are unreliable when estimated from the very first trials, a warm-up buffer of data was introduced to initialize them. Experiments focused on challenging cross-subject MI decoding scenarios and compared the proposed approach with conventional offline and online adaptation strategies, including a recent deep learning method based on normalization. The results showed

that continuously updating the normalization and alignment statistics allowed the model to approach the performance obtained in offline settings, while remaining suitable for online use and requiring no supervision during deployment. The study therefore demonstrated that substantial improvements in cross-subject BCI decoding could be obtained without modifying the learned model itself, by instead adapting the representation to the statistics of the incoming data stream.

Finally, with the emergence of foundation models in flagship domains of AI, and the experience we gained from our previous works, we decided to develop our own foundation model for EEG data. [Yassine EL OUAHIDI](#) – assisted by [Jonathan Lys](#) – did a massive amount of work to gather more than 90 publicly available EEG datasets, resulting in over 25,000 subjects and 60,000 hours of measurements. Using this largest assembled dataset in the community, and proposing a novel approach for integrating heterogeneous spatio-temporal data in a normalized positional encoding, we developed REVE, a foundation model for EEG. This work was published as follows:



[\[C_{3i}\]](#) Yassine EL OUAHIDI, Jonathan Lys, Philipp THÖLKE, Nicolas FARRUGIA, Bastien PASDELOUP, Vincent GRIPON, Karim JERBI, Giulia LIOI (2025). **“REVE: A Foundation Model for EEG: Adapting to Any Setup with Large-Scale Pretraining on 25,000 Subjects”**. *NeurIPS*

This work introduced REVE, a foundation model designed to address the strong heterogeneity of EEG recordings across electrode configurations, recording devices, protocols, and subjects. The model used a transformer-based architecture with a novel 4D positional encoding that jointly represented the 3D spatial location of electrodes and the temporal position of signal patches, allowing the same model to process recordings with different numbers and arrangements of electrodes, and of variable duration. REVE was pretrained using a masked autoencoding objective on more than 60,000 hours of EEG from 92 datasets and approximately 25,000 subjects, making it the largest EEG pretraining effort to date. The resulting representations were evaluated on 10 downstream tasks, covering MI decoding, seizure detection, sleep staging, cognitive-load estimation, emotion recognition, and other EEG analysis problems. REVE achieved state-of-the-art results across these benchmarks, including under linear probing and with little or no task-specific fine-tuning, demonstrating that large-scale pretraining could produce representations that transferred across substantially different EEG setups. The work therefore addressed a central obstacle to EEG foundation models – the lack of a common representation across heterogeneous acquisition settings – by making the electrode layout itself part of the model’s input rather than requiring recordings to be projected onto a fixed configuration. The release of pretrained weights, code, and tutorials further positioned REVE as a reusable foundation for EEG representation learning and subsequent neuroscience and BCI applications.

REVE was a resounding success in the neuroscience community. At the time of writing, it has already received more than 70 citations [[GOOGLE SCHOLAR \(B.P.\)](#)], and was integrated into Neural-Bench [[BANVILLE et al., 2026](#)], a large-scale initiative aimed at unifying the benchmarking of models in neuroscience, where it ranks first overall and on several downstream tasks. This model established a new scale and a new direction for EEG foundation models, by demonstrating that large-scale pretraining could be performed across highly heterogeneous acquisition setups rather than requiring a standardized electrode montage.

Since then, numerous other foundation models have emerged, making REVE a very timely contribution. Giulia LIOI recently proposed a review of these models, published as follows:

[C₃₄]

Giulia LIOI, Yassine EL OUAHIDI, Jonathan LYS, Philipp THÖLKE, Bastien PASDELOUP, Vincent GRIPON, Nicolas FARRUGIA, Karim JERBI (2026). **“Key Ingredients for EEG Foundation Models: What Works - and What Doesn’t”**. *EUSIPCO, IEEE*

This work presented a systematic review of recent EEG foundation models, with the objective of identifying the architectural and methodological choices that enable transferable EEG representations. It first analyzed the main ingredients shared by these models, including EEG tokenization, transformer-based architectures, positional encodings, and SSL objectives, and then examined how these choices interacted with the scale and diversity of pretraining data. The review covered a rapidly expanding set of models, including BIOT, EEGzREP, NeuroGPT, LaBraM, NeuroLM, ELASTIQ, CBraMod, EEG-Mamba, LUNA, NeurIPT, CSBrain, and REVE, and organized them according to their input representations, architectures, pretraining strategies, and ability to handle heterogeneous electrode configurations. It specifically highlighted the transition from models trained on a fixed electrode montage and a single dataset toward architectures such as LaBraM, LUNA, and REVE that could accommodate varying numbers and arrangements of electrodes, thereby improving their potential for cross-dataset and cross-subject transfer. The review also compared models on common downstream EEG tasks and showed substantial variability in reported performance, with the sobering conclusion that increasing model scale does not automatically yield better generalization than a smaller model trained from scratch for the task at hand. It consequently identified several unresolved issues, including the lack of standardized evaluation protocols, the limited availability of fair comparisons across models, and the need for data-efficient and lightweight adaptation methods in low-data settings. Overall, the work provided a structured analysis of the emerging EEG foundation-model landscape and identified data diversity, flexible input handling, and appropriate SSL objectives as central design considerations for developing transferable EEG representations.

In addition to its impact on neuroscience, REVE also had a structuring impact on the BRAIN team. It was the first time that we managed to have a very strong collaboration between PhD students on a common topic, and the first NeurIPS publication in the team after several attempts across years. This work motivated new collaborations between PhD students, who are now working on an extension of REVE to MEG data.

Overall, while not a neuroscientist myself, studying brain signals gave me invaluable support for evaluating approaches both in GSP and model development. Collaborations with neuroscientists – especially Giulia LIOI and Nicolas FARRUGIA – also gave me the basic knowledge on brain signals, which allowed me to collaborate with neurologists in more applied contexts.

II.3 Applications in Healthcare

As mentioned earlier, the way I like to conduct my work is by keeping up to date with modern AI research through fundamental projects on subfields of AI, and then adapting this knowledge to healthcare applications, which I believe deserve more attention from the AI community due to their societal importance.

The choice of healthcare subdomains I studied was not really made on purpose, but rather happened based on opportunities. My earlier research in collaboration with neuroscientists gave me the background I needed to talk to neurologists at the CHU of Brest; my postdoc at EPFL allowed me to work on an oncology project with CHUV; and meeting Yassine EL OUAHIDI’s brother Amine EL OUAHIDI led to new collaborations in cardiology. I also studied ophthalmology for a moment with the thesis of Arsalan NAJIB, but the project was abandoned for various reasons.

Throughout all these collaborations, I found two general rules that seem to apply broadly in healthcare:

- **The necessity of data and anticipation** — Healthcare is obviously a critical domain, and several regulations [CNIL ; GDPR ] guarantee the safety and protection of the persons whose data is acquired through medical acts. This important protection however comes with heavy procedures to share data, which can sometimes delay the start of a project. It is thus necessary to anticipate data collection and project preparation before its official start. Also, there are several datasets online that already underwent anonymization and open publication. These are very valuable resources that allow working on model development while data access procedures run their course.
- **The complementarity of modalities** — When working on the various healthcare projects I was involved in, I generally focused on one particular modality that had the potential to provide highly informative insights for the target task. However, the data often came with a lot of extra material (*e.g.*, annotations, reports, questionnaires, blood analyses), which brought complementary information. My experience as a researcher in AI, where multimodality has become a prevalent research direction, makes me strongly believe that complementary – if not all – available data should be used for a subject’s diagnosis, to improve the understanding of their medical state, and allow better support.

II.3.1 Works in Neurology

My collaborations in neurology started when I met **Victor QUÉRÉ**, a Master’s student in neurology at the CHU of Brest, who was doing an internship at IMT Atlantique under the supervision of Serge TIMSIT (CHU of Brest) and Philippe LENCA & Sorin MOGA (IMT Atlantique). In his internship, **Victor QUÉRÉ** was interested in strokes, and wanted to evaluate whether ML tools could be used to predict the risk of spontaneous ICH in presence of missing data.

Strokes are a major cause of death and disability in the world. It is estimated that in the event of a stroke, a patient loses about 1,900,000 neurons per minute left untreated [SAVER, 2006 ]. Among all strokes, ICH is particularly devastating, and is fatal in 40% of cases, 50% of which happen in 2 days or less. Survivors of such an event only regain autonomy in 12 to 39% of cases [VAN ASCH *et al.*, 2010 ].

To fulfill his project, **Victor QUÉRÉ** worked on the BSR, a large population-based collection of data monitoring all strokes in the area of Brest, France [RETHO *et al.*, 2024 ]. This dataset is one of the largest collections of stroke-related information in Europe, and consists of both ischemic and ICH strokes (see [Figure 5](#)). For each event, clinical data, analysis results and imaging data are archived, making the BSR a highly valuable resource for research on strokes.


Following his internship work, **Victor QUÉRÉ** decided to pursue under my supervision for his MD thesis. We extended his prior work beyond a simple comparison of model performance, and explored variable selection and data imputation techniques to improve interpretability of results:



[T₂]

Victor QUÉRÉ (2021). “**Comparison of Logistic Regression Analysis with Random Forest for Prediction of ICH Early Mortality in the Population-Based Brest Stroke Registry**”. *UBO, Faculté de Médecine et des Sciences de la Santé de Brest (MD)*

In his thesis, **Victor QUÉRÉ** investigated the use of AI and statistical learning to predict early mortality following spontaneous ICH, with particular emphasis on the impact of missing clinical data. Using the population-based BSR, the study developed and compared four prognostic scores for 30-day mortality: two regression-based models and two RF models, each trained either on complete-case data or on a dataset in which missing predictors were imputed using predictive mean matching. The models incorporated demographic and medical history, treatments, clinical and biological variables, and imaging-derived characteristics, including hematoma volume and localization. The results showed that handling missing



(a) CT scan of a brain showing an ischemic stroke.

(b) CT scan of a brain showing an ICH stroke.

Figure 5 Illustration of the two main categories of strokes available in the BSR. Ischemic strokes (a) are caused by an interruption of the blood supply in parts of the brain, generally due to a clot. Hemorrhagic strokes (b) are caused by the rupture of a blood vessel, and can be categorized into two subcategories: ICH, and subarachnoid hemorrhage. Images taken from [WIKIPEDIA (STROKE) 

data through imputation substantially improved predictive performance: the best regression model achieved an AUC-ROC of 0.83 on the imputed dataset, compared with 0.78 after complete-case analysis, while the RF model reached an AUC-ROC of 0.81 versus 0.78. The analysis also showed that excluding patients with missing observations could reduce the sample size by 65% and introduce selection bias, as patients with more severe strokes were more likely to have missing information. Beyond prediction, the study analyzed the contribution of individual variables and confirmed the importance of established prognostic factors such as GCS, age, and hematoma volume, while highlighting additional predictors including admission blood glucose, NIHSS, anticoagulant treatment, and, more tentatively, prior fibrate treatment. The thesis therefore treated missing data as a first-class methodological question rather than a preprocessing detail, and argued that future studies on ICH prognosis should record why a variable is missing instead of discarding the corresponding patients.

This collaboration allowed me to meet Serge TMSIT, head of the department of neurology at the CHU of Brest and of the BSR, who later invited me to join the STRATIF-AI project. This large-scale European project aims at developing better models for patient stratification, using modern approaches such as digital twins [ARMENI *et al.*, 2022 et al., 2017 et al., 2022 

Within this project, I supervise the postdoc of Amine SADIKINE on the technical side, who works on multimodal data integration for the prediction of mortality at 30 days following an ICH. We are in particular interested in models that can provide embeddings for tabular data [HOLLMANN *et al.*, 2022 et al., 2026 




Figure 6 Illustration of the workflow of our work with [Amine SADIKINE](#). Tabular and clinical data is embedded through a tabular foundation model, while clinical images go through a specialized model (here a CNN). Separate representations are then fused into a single decision process. Figure taken from [\[S₁₃\]](#).

II.3.2 Works in Cardiology

My work in cardiology was conducted in partnership with multiple practitioners from the CHU of Brest. The cardiology department features researchers with various interests, reflecting the diversity in our collaborations. My work as a researcher in AI was to orient the choices of ML solutions, and provide technical support.

The first problem in cardiology I studied was in the context of [Amine EL OUAHIDI](#)'s Master internship, where we wanted to determine whether it was possible or not to predict the risk of pacemaker implantation after a TAVI procedure (see [Figure 7](#)). To do so, we proposed multiple realistic scenarios – before implantation, during implantation, and after – that restricted availability of data. We compared multiple normalization and variable selection techniques, and showed that ML models could improve prediction of the risk of pacemaker implantation over previous baselines, with kept variables on par with existing medical knowledge.




Figure 7 Illustration of the implantation of a TAVI (a), and its principal complication (b).

After his internship, [Amine EL OUAHIDI](#) pursued an MD thesis under my supervision. We extended his previous work by integrating imaging data in our analysis, through handcrafted features measuring specific elements of interest. In particular, we showed that the size of the membranous septum was especially of interest in the prediction, and that imaging alone was nearly as informative as integrating all relevant clinical modalities found during [Amine EL OUAHIDI](#)'s internship:

[T₄]

Amine EL OUAHIDI (2023). **“Utilisation de Méthodes d’Apprentissage Automatique pour Prédire l’Implantation d’un Pacemaker dans les Suites d’une Procédure de TAVI”**. UBO, Faculté de Médecine et des Sciences de la Santé de Brest (MD)

The thesis investigated the use of ML and deep learning to predict permanent pacemaker implantation within 28 days following TAVI, with the objective of identifying patients at high risk of post-procedural conduction disturbances. The study retrospectively included 520 patients who underwent TAVI at the CHU of Brest between January 2019 and December 2021, and compared two complementary approaches based on heterogeneous clinical data and cardiac CT images. The first approach used 67 numerical and categorical variables covering clinical characteristics, electrocardiography, echocardiography, procedural information, and CT measurements. Eight ML models were systematically compared using nested cross-validation, leading to the selection of an SVM, whose performance was subsequently improved through feature selection combining SHAP and RFE. This analysis identified CT-derived anatomical variables as particularly informative, including the diastolic length of the membranous septum and its variation throughout the cardiac cycle, which emerged as highly predictive of pacemaker implantation; training the same model on each modality separately confirmed the central role of imaging, CT variables alone almost matching the full selection. The second approach investigated whether the prediction could be performed directly from raw 3D CT data, using a 3D ResNet trained independently on systolic, diastolic, and non-contrast (calcium) CT sequences. While the image-based deep learning models remained less accurate than the feature-based SVM (AUC-ROC of 70.2% for the calcium sequence, 62.3% for systole, and 60.1% for diastole), they demonstrated the feasibility of predicting the outcome directly from medical images without explicit anatomical feature extraction. The final SVM achieved an AUC-ROC of $92.1\% \pm 4.7$, an AUC-PR of $77.6\% \pm 9.9$ and an F1-score of $71.8\% \pm 9.9$, outperforming an established multivariate risk score evaluated on the same data as well as the four published ML studies on this question – a comparison the thesis itself qualifies, since those studies were not re-evaluated on this cohort and report no confidence interval – and was made available online for potential clinical use.

Amine EL OUAHIDI was invited to present his thesis work during the Research Days of the CHU of Brest [W₈], for which he received the clinical research award. His work was later published in [J₅], which got some impact on the clinical cardiology research community, and is generally cited as a work marking transition from clinical scores to multimodal ML prediction [JAMHOUR-CHELH *et al.*, 2025 [E5]; ABRAS *et al.*, 2026 [E6]].

Our work also led to a subsequent collaboration with researchers from AP-HP, on the use of ML tools for prediction of adverse events in CICUs. This led to the following publication, as well as communication [W₁₆]:

[J₁₀]

Orianne WEIZMAN, Kenza HAMZI, Patrick HENRY, Guillaume SCHURTZ, Marie HAUGUEL-MOREAU, Antonin TRIMAILLE, Marc BEDOSSA, Jean-Claude DIB, Sabir ATTOU, Tanissia BOUKERTOUTA, Franck BOCCARA, Thibaut POMMIER, Pascal LIM, Thomas BOCHATON, Damien MILLISCHER, Benoit MERAT, Fabien PICARD, Nissim GRINBERG, David SULMAN, Bastien PASDELOUP, Yassine EL OUAHIDI, Treçy GONÇALVES, Eric VICAUT, Jean-Guillaume DILLINGER, Solenn TOUPIN, Théo PEZEL (2025). **“Machine Learning Score to Predict In-Hospital Outcomes in Patients Hospitalized in Cardiac Intensive Care Unit”**. *European Heart Journal, Digital Health* (vol. 6, nb. 2, pp. 218–227), Oxford University Press

This work investigated the use of ML to predict major adverse events in patients admitted to CICUs. The study used a prospective multicenter cohort of 1,499 consecutive patients admitted to 39 French CICUs during a two-week nationwide inclusion period, and defined the primary outcome as an in-hospital

major adverse event, comprising death, resuscitated cardiac arrest, or cardiogenic shock. 28 variables were collected at admission, including clinical, biological, electrocardiographic, and echocardiographic measurements, and the study compared several ML approaches. Feature selection identified 7 variables associated with the outcome – illicit drug use, mean arterial pressure, Killip class, exhaled carbon monoxide level, left ventricular ejection fraction, tricuspid annular plane systolic excursion, and peak E/e' ratio – four clinical and three echocardiographic, all of which the final model – a boosted cost-sensitive decision-tree ensemble, the best of the four ML approaches compared – retained. The resulting score achieved an AUC-ROC of 0.90 in the testing cohort, with a sensitivity of 0.42 and a specificity of 0.99, and retained good discrimination in an independent external cohort of patients from 8 centers that had not contributed to model development. The model also outperformed several established risk scores. These results demonstrated that a relatively small set of heterogeneous clinical and echocardiographic variables could be combined through ML to provide accurate early risk stratification in CICU patients, although the small number of major adverse events, a proportion of missing echocardiographic measurements reaching a quarter of the cohort for the E/e' ratio, the confinement of the validation to French centers, and the fact that variables such as exhaled carbon monoxide are not routinely collected everywhere, all limited the conclusions regarding its immediate clinical applicability.

On a different topic, collaborating with the cardiology department of the CHU of Brest gave me the opportunity to meet [Yahia BELLOUCHE](#), whose PhD I supervised. The interests of [Yahia BELLOUCHE](#) revolved around cardiovascular modeling and CFD, bringing a lot of physics elements into his work, with the objective of integrating a mechanistic understanding into predictions as a complement to imaging data.

In a first work, we developed a novel index for the assessment of coronary lesions, using CCTA for its non-invasiveness. This was published as follows:



Yahia BELLOUCHE, Clément BENIC, Sinda HANNACHI, Pierre-Philippe NICOL, Christopher JOUSSE, Florent LE VEN, Jacques MANSOURATI, Bastien PASDELOUP, Romain DIDIER (2025). **“Performance of a Novel Computational Hyperemic Resistance Index Derived from Cardiac CT in Coronary Chronic Syndromes”**. *Journal of Clinical Medicine* (vol. 14, nb. 20, art. 7270), MDPI

This work investigated a novel non-invasive computational equivalent of the HSR index, derived from CCTA, for the functional assessment of coronary lesions in patients with suspected chronic coronary syndrome. The study retrospectively included 64 patients with 106 coronary lesions who underwent both CCTA and invasive coronary angiography at the CHU of Brest. Patient-specific 3D coronary models were reconstructed from CCTA using open-source tools, and CFD simulations were performed under basal and hyperemic conditions. Importantly, the simulations incorporated patient-specific boundary conditions derived from CCTA, left-ventricular and aortic flow measurements, and invasive pressure data, which served both to initialize the peripheral resistances and as targets when tuning the lumped-parameter models, allowing the resulting computational hemodynamics to account for individual anatomical and physiological characteristics. The validity of the pipeline was supported by the strong correlation between the computational fractional flow reserve it produces and its invasive measurement ($r = 0.87$, $p < 0.0001$). The proposed cHSR index was evaluated against computational fractional flow reserve and quantitative flow ratio, using the clinical decision to perform coronary revascularization as the reference outcome. At an optimal cut-off of $0.75 \text{ mmHg/cm} \cdot \text{s}^{-1}$, the cHSR reached 96.2% sensitivity and 96.2% specificity and an overall diagnostic accuracy of 96.2%, providing a large accuracy improvement over the two other indices, although the corresponding gap in AUC-ROC was narrower (0.96 against 0.93 and 0.81). The study additionally investigated whether cHSR was associated with local hemodynamic

parameters related to atherosclerosis, but found no significant correlation with TAWSS or OSI. The retrospective design, the small cohort, the cost of the simulations and their assumption of rigid arterial walls were all acknowledged as limitations.

In a second work, we used patient-specific CFD to derive a non-invasive hemodynamic biomarker for the detection and characterization of CTEPH and CTEPD. This work – published as follows – highlighted the potential of CFD to extract clinically relevant information beyond that provided by anatomical imaging alone:



[9] Yahia BELLOUCHE, Cécile TROMEUR, Romain DIDIER, Philippe ROBIN, Christophe GUT-GOBERT, Pierre-Yves LE ROUX, Irène FRACHON, Michel NONENT, Alexandre FAUCHÉ, Jamal ELHASNAOUI, Aurélia LE HIR, Marie GUEGAN, Christophe LEROYER, Pierre-Yves SALAÜN, Bastien PASDELOUP, Francis COUTURAUD (2025). **“Proximal Pulmonary Artery Wall Shear Stress Derived from Computational Fluid Dynamics: A Non-Invasive Biomarker for CTEPH and Perfusion Mismatch”**. *Physiological Reports* (vol. 13, nb. 22, pp. e70664), Wiley Online Library

This work investigated whether CFD could provide a non-invasive hemodynamic biomarker for the detection of CTEPH and CTEPD. The study retrospectively analyzed 90 patients with a previous pulmonary embolism in a case-control design, with three matched groups of 30: patients with CTEPH, patients with perfusion mismatches but without pulmonary hypertension, and patients without perfusion mismatch. Patient-specific 3D models of the pulmonary arteries were reconstructed from CT pulmonary angiography, and CFD simulations were performed using individualized resistance-compliance boundary conditions tuned on right-heart catheterization and cardiac output measurements. The simulations were used to characterize proximal pulmonary artery hemodynamics, with particular emphasis on TAWSS and OSI. The study found that median TAWSS was substantially reduced in both CTEPH and CTEPD patients compared with patients without perfusion mismatch (16.5, 27.5 and 42.0 dyn/cm² respectively), with the lowest values observed in CTEPH. TAWSS was also strongly associated with the presence of ventilation/perfusion mismatch ($\rho = -0.673$) and predicted it with an AUC-ROC of 0.918, an optimal cut-off of 27.0 dyn/cm² yielding full specificity for a sensitivity of 70%, whereas OSI did not significantly differ between the groups – in contrast with earlier reports of elevated OSI in CTEPH. These findings suggested that reduced wall shear stress in the proximal pulmonary arteries reflected the chronic thromboembolic burden and could also identify perfusion abnormalities in patients who had not yet developed pulmonary hypertension. The retrospective design and limited cohort size nevertheless warranted further validation in larger prospective populations before clinical implementation.

Throughout these collaborations, my experience in cardiology allowed me to extend my expertise in ML toward patient-specific medical applications, while contributing to the development of imaging-based and computational biomarkers for cardiovascular diagnosis and risk assessment.

These collaborations have highlighted the importance of imaging, but have also shown that integrating other modalities (clinical data, CFD-based indices) could lead to an increase in model performance. This reinforces my initial statement in [Section II.3](#), *i.e.*, that multimodality is an important direction for healthcare applications involving AI techniques.

II.3.3 Works in Oncology

The final healthcare application I worked on was oncology, in the context of personalized medicine. During my postdoc, I worked in collaboration with doctors from CHUV in Lausanne, who were interested in immunotherapy, and particularly in finding biomarkers that would indicate a higher risk of death, based on the distribution of cells within cancerous tissues.

Leveraging my experience in GSP, a natural approach for me was to model the tissue as a graph (see Figure 8), and to study the properties of this graph.




Figure 8 Example of the construction of a cell graph from a tissue. The adjacent slide is colored using H&E for determining tumor cells and stromal environment. Multiplex immunofluorescence provides more information on cell phenotyping. In [C₁₂], this information was estimated using a CNN as only H&E was available on TCGA [GDC]. The graph is then constructed using Delaunay triangulation [DELAUNAY, 1934]. Figure taken from [R₄].

Due to data acquisition delays (see my point in Section II.3), this work was completed after I moved to IMT Atlantique. There, I continued this research, assisted by Yassine EL OUAHIDI *et al.*. Using the graphs of cells, we measured distance-based indicators of T-cell infiltration in the tumor. We then showed that the distribution of such distances obtained for a subject correlates with their vital prognosis, and used optimal transport [VILLANI, 2008] to compute distances between subjects for clustering. This work led to the following publication:



Yassine EL OUAHIDI, Matis FELLER, Matthieu TALAGAS, Bastien PASDELOUP (2021). **“Une Approche pour Regrouper des Sujets Atteints de Cancers, sur la Base des Similarités des Répartitions Cellulaires au Sein de Leurs Biopsies”**. *EGC (pp. 333–340), RNTI*

This work investigated whether the spatial organization of cells in histopathological biopsies could be used to identify patient subgroups with different survival outcomes. It proposed an interpretable, unsupervised framework that represented each patient through histograms describing several aspects of the spatial distribution of cancer cells, stromal cells, and lymphocytes, rather than relying only on conventional measures such as cell counts or densities. Five complementary features were considered: the distance between lymphocytes and cancer cells, the distance between lymphocytes and the cancer-stroma interface, the distance between lymphocyte aggregates, the size of these aggregates, and the density of lymphocytes at different distances from the cancer-stroma interface. A Delaunay triangulation of the detected cells was used to identify spatial relationships between neighboring cells and to characterize interfaces and lymphocyte aggregates. The resulting distributions were compared between patients using the Wasserstein distance, which accounted for differences in both the location and shape of the distributions, and patients were subsequently grouped using agglomerative hierarchical clustering. The methodology was evaluated on H&E-stained biopsies from patients with stage I lung adenocarcinoma from TCGA. After filtering the initial cohort, 94 patients were retained, each associated with several regions of interest containing both tumor and stromal tissue. When considered individually, four of the five proposed spatial features produced clusters with significantly different survival profiles; in particular, the spatial organization and density of lymphocytes relative to the tumor-stroma interface were strongly associated with prognosis. Combining the informative features into a single weighted similarity measure further improved the separation between patient subgroups, yielding a highly significant difference in survival. Importantly, the framework remained interpretable, as representative histograms could be examined to identify the spatial characteristics distinguishing the groups. Unlike existing approaches, which assign subjects to predefined classes, the framework was fully unsupervised, which both allows new candidate features to be explored for correlation with survival and lets a newly seen subject be matched against the patients most similar to them.

While I am particularly proud of this work, which concluded my postdoctoral experience, and initiated collaborations with [Yassine EL OUAHIDI](#) – who would later do a PhD under my supervision – it did not have any impact, due to its publication in a national conference. Still, I was recently contacted by a former collaborator from CHUV, who extended this work to more advanced phenotyping techniques (multiplex immunofluorescence), paving the way for personalized and effective treatment strategies in melanoma. This article, currently under review, is summarized below:



[R₄] Alexandre WICKY, Jonatan BONJOUR, Beril BESBINAR, Petros LIAKOPOULOS, Bastien PASDELOUP, Mireille EL GHECHE, Gojko CUTURA, Asmi SOUHAIL, Carl HATOUM, Charilaos ZISOUS, Sevda OGUT, Robin LIECHTI, Julien DORIER, Ekaterina FORTIS, Christel VOIZE, Rita DE MICHELI, Clarisse DROMAIN, Julien DAGHER, Igor LETOVANEC, Laurence DE LEVAL, Paweł TETERYCZ, Anna SZUMERA-CIEĆKIEWICZ, Piotr RUTKOWSKI, Sofiya LATIFYAN, Andrew JANOWCZYK, Sylvain PRADERVAND, Mikaël J. PITTET, Stephanie TISSOT, Sylvie RUSAKIEWICZ, George COUKOS, Dorina THANOU, Pascal FROSSARD, Michel A. CUENDET, Olivier MICHIELIN (2026). **“Spatial Organization of the Immune Microenvironment Predicts Checkpoint Inhibitor Response in Advanced Melanoma”**. *Nature Communications*, Nature Publishing Group

This work investigated whether the spatial organization of immune cells within the tumor microenvironment could predict response to immune checkpoint inhibitor therapy in patients with advanced melanoma. The study analyzed pretreatment tumor samples from 77 patients, comprising a discovery cohort of 47 patients and an independent validation cohort of 30 patients, using two panels of deliberately low-plex multiplex immunofluorescence – an affordability constraint chosen for clinical viability – to identify several immune cell populations, including CD4⁺ and CD8⁺ T cells, B cells, natural killer cells, dendritic cells, and macrophages, together with markers of activation and proliferation. Rather than relying on the abundance of individual cell types, the tissue was represented as a cell graph obtained by Delaunay triangulation, partitioned into tumor, border, and stroma compartments, from which a few hundred interpretable spatial features describing cell proximity, clustering, and interactions were extracted. Features were ranked using MRMR, and an ℓ_2 -regularized logistic regression model was trained to distinguish patients who remained progression-free for two years from those who did not. Cross-validated performance on the discovery cohort peaked with only two features, and decreased as further ones were added, so that the final model was deliberately kept minimal: the number of contacts between dendritic cells and T cells per dendritic cell within the tumor, and the number of B-cell clusters in the stroma, the latter being strongly correlated with markers of early tertiary lymphoid structures. The resulting model achieved an AUC-ROC of 0.74 on the independent validation cohort, whereas the amount of CD8⁺ T-cell infiltration alone was not predictive, and it was also associated with progression-free and overall survival in that cohort. Replacing each of the two features by its most correlated alternative left the performance essentially unchanged, indicating that the model captures biological characteristics rather than idiosyncrasies of the selection procedure. Where immune cells were located and how they interacted therefore carried more predictive information than how many of them were present. Several limitations were acknowledged, however: the cohorts remain small, single markers cannot resolve the functional diversity of macrophages or CD4⁺ subsets whose opposing effects may cancel out, and above all the absence of an untreated control group leaves open whether the biomarker is genuinely predictive of response or partly prognostic of a better course regardless of treatment.

This concludes my past work on healthcare topics. Overall, these applications have brought me opportunities to work on various data types (*e.g.*, graphs, images, tabular), providing me with interesting and useful support for exploring the literature on more fundamental aspects of AI.

I strongly believe that multimodality revolving around text, and agentic AI, are the next directions that would enable important breakthroughs in healthcare. Datasets exist, but access to them remains too difficult for training general healthcare foundation models. Federated learning, as used in the [STRATIF-AI](#) project, offers some solutions, but the system is very heavy to put in place. I think the future of AI-assisted healthcare will leverage agents developed in independent hospitals, and then made available (with anonymization guarantees). Such collaborations of specialized units would probably allow the detection of early markers of some diseases, coming from a mix of healthcare subdomains.

II.4 Other Collaborations

To conclude this chapter on my past research and collaborations, I want to mention some other works I have done, in which I generally had the role of a supervisor, bringing AI expertise to a different domain.

These other collaborations can mainly be grouped into two categories. After my arrival at IMT Atlantique, I worked in the DECIDE team, which featured researchers with a strong interest in operational research. In this team, I worked on this topic, but also on applications, often related to time series or physics. The subsequent sections detail these works.

II.4.1 Operational Research & Decision Aid

As an expert in AI in a team of operational research experts, I had the opportunity to supervise the PhD thesis of [Maryam KARIMI-MAMAGHAN](#), who wanted to develop hybrid approaches mixing metaheuristics and ML for combinatorial optimization. Her PhD manuscript is summed up below, and ideas that led to her work were initially presented in [C₁₀]:



[T₃]

Maryam KARIMI-MAMAGHAN (2022). **“Hybridizing Metaheuristics with Machine Learning for Combinatorial Optimization: A Taxonomy and Learning to Select Operators”**. *École Nationale Supérieure Mines-Télécom Atlantique (PhD)*

The thesis of [Maryam KARIMI-MAMAGHAN](#) investigated how ML could be integrated into metaheuristic algorithms to improve combinatorial optimization. First, it provided a comprehensive review of existing approaches and proposed a unified taxonomy describing the different ways in which ML could interact with metaheuristics: algorithm selection, fitness evaluation, initialization, evolution – itself split into operator selection, learnable evolution models, and neighbor generation – parameter setting, and cooperation between algorithms. The thesis then focused on adaptive operator selection, addressing the problem of choosing, during the search, which perturbation or local-search operator should be applied based on the current state of the optimization process and the operators’ past performance. To this end, [Maryam KARIMI-MAMAGHAN](#) proposed a general framework integrating Q-learning into an ILS, allowing the algorithm to learn online which operators were most effective rather than relying on a fixed or random selection strategy. The approach was evaluated on two combinatorial optimization problems, the TSP and the PFSP, where it consistently improved solution quality and convergence compared with random operator selection, with particularly strong results on large instances. In the PFSP, the resulting algorithm also showed competitive performance with state-of-the-art methods.

This manuscript builds on several publications. First, we explored the integration of ML into a classical ILS algorithm, with the goal of estimating solutions to the TSP problem. This problem is NP-hard, and even the best exact solvers struggle to find optimal solutions for large instances. Finding good approximations that are fast to compute is still a challenge. This work led to the following publication:



[C₁₅]

Maryam KARIMI-MAMAGHAN, Bastien PASDELOUP, Mehrdad MOHAMMADI, Patrick MEYER (2021). **“A Learning-Based Iterated Local Search Algorithm for Solving the Traveling Salesman Problem”**. *OLA (pp. 45–61), Springer*

This work proposed a learning-based ILS algorithm for the TSP, with the objective of making the search process more adaptive than traditional ILS approaches. Whereas a standard ILS relies on a single local-search and a single perturbation operator, the method incorporated several of each and used Q-learning

to select the most appropriate one at each step of the search, according to the state of the search and the reward accumulated by every operator. The two levels of integration were studied separately: one variant learns to choose among three local-search operators – the basic 2-opt and two operators proposed in the article, an extended 2-opt and an extended insertion operator – while the other chooses among three perturbation operators – double-bridge, shuffle-sequence, and reversion-sequence. Experiments on 24 symmetric TSP instances from the TSPLIB library showed that both variants recover the optimal solution on small and medium instances and stay within small optimality gaps on the largest ones, and that they converge faster and toward better solutions than the corresponding single-operator and random-selection variants against which they were compared. The comparison also suggested that selecting perturbation operators is the more promising of the two levels, since it is the one that addresses the weakness of ILS: intrinsically strong at exploitation, but prone to being trapped in local optima.

We later extended similar ideas to new metaheuristics and combinatorial problems, such as the IG for the PFSP problem, which is also NP-hard. This work led to the following publication, with intermediate developments presented in [C₁₆; W₇; W₉]:



[U₃] Maryam KARIMI-MAMAGHAN, Mehrdad MOHAMMADI, Bastien PASDELOUP, Patrick MEYER (2023). **“Learning to Select Operators in Meta-Heuristics: An Integration of Q-Learning into the Iterated Greedy Algorithm for the Permutation Flowshop Scheduling Problem”**. *EJOR* (vol. 304, nb. 3, pp. 1296–1330), Elsevier

This work developed a learning-based IG algorithm for the PFSP, with the objective of making the selection of perturbation operators adaptive. The PFSP required finding an ordering of jobs to be processed on a set of machines so as to minimize the makespan, and represented a challenging NP-hard combinatorial optimization problem. The approach extended the standard IG framework by incorporating multiple perturbation operators with different exploration strengths, rather than relying on a single fixed operator. The main contribution was the integration of Q-learning into the perturbation phase. At each iteration, the RL mechanism selected an operator according to the current state of the search and the performance history of the available operators. Successful operators were therefore increasingly favored in relevant search situations, while the learning process retained exploration of alternative operators. This turned operator selection into an online decision-making problem and allowed the metaheuristic to adapt its exploration behavior throughout the optimization process. The resulting Q-learning-based IG algorithm was evaluated on three well-known collections of PFSP benchmark instances and compared with non-learning IG variants, including approaches using individual or randomly selected perturbation operators, as well as 7 state-of-the-art algorithms. The experiments showed that adaptive operator selection improved the exploration ability of IG and led to faster convergence toward higher-quality solutions. The method obtained better optimality gaps while introducing no significant computational overhead, and even required less computational time on large instances due to its improved convergence behavior. Two complementary analyses supported the mechanism: a sensitivity study showed the algorithm to be almost insensitive to the number of operators provided, automatically ignoring the redundant ones, and a decomposition of the improvement per operator showed that each of them contributes significantly at some stage of the search. Overall, this work demonstrated that RL could be integrated directly into a metaheuristic to learn how to search, rather than learning the solution itself.

The works of Maryam KARIMI-MAMAGHAN had a substantial impact on the operational research community. They were published at a time when the interest of the community for the combination of RL and combinatorial

optimization was resurging [JOHNN *et al.*, 2024], and paved the way for several follow-up works aiming at using RL for metaheuristics improvement [LU *et al.*, 2024; ZHAO *et al.*, 2024].

Interestingly, newer approaches have emerged since then that aim at exploiting LLMs and other large models for solving complex optimization problems, which I believe is a strong direction for future work [MASOUD *et al.*, 2024]. In particular, [PEREIRA *et al.*, 2025] proposes an interesting way of approaching the problem, by using a specific positional encoding method inspired by ALiBi [PRESS *et al.*, 2021] and RoPE [SU *et al.*, 2024].

In addition to hybridization of metaheuristics with ML for a controlled exploration of the search space, another direction of research we explored was to use ML models for generating candidate solutions in population-based metaheuristics, leading to novel crossover operators. This work led to the following publication:



[C_{II}]

Bastien PASDELOUP, Maryam KARIMI-MAMAGHAN, Mehرداد MOHAMMADI, Patrick MEYER (2020). “**Autoencoder-Based Generation of Individuals in Population-Based Metaheuristics**”. *Congrès Annuel de la ROADEF*

This work explored a novel way of generating new individuals in population-based metaheuristics by using autoencoders to learn a continuous latent representation of candidate solutions. Instead of relying exclusively on traditional crossover and mutation operators, the approach encodes existing solutions into a lower-dimensional latent space, combines their representations, and decodes the resulting point to generate a new candidate solution. The underlying idea was that the learned latent space could capture characteristics shared by good solutions and therefore provide a meaningful space in which to perform recombination, while removing the need for expert-designed breeding rules. The approach was investigated in the context of the TSP problem. Randomly generated tours were used to train the autoencoder, after which each new individual was obtained by drawing several parents with probability proportional to their fitness, averaging their latent representations with the same weights, and decoding the result; since the decoded matrix does not in general describe a Hamiltonian cycle, a greedy regularization step reconstructs a valid tour, and the child is inserted in the population only if it outperforms all of its parents. The article also observed that the imperfect reconstruction of the autoencoder plays a role analogous to mutation in a classical genetic algorithm. Being a short conference abstract, it presented the idea and the pipeline rather than a quantitative evaluation, the experiments and comparisons with existing methods being deferred to the oral presentation.

Related to my interests in diversity within algorithms – which led to the [ENDIVE](#) project – I extended my work to population-based metaheuristics with the exploration of the impact of diversity in candidate solutions on convergence of genetic algorithms. It was also the opportunity for me to work on decision aid, a subfield of operational research that aims at producing tools for supporting informed decision making in complex situations. This work was conducted in close collaboration with Patrick MEYER, Arwa KHANNOUSSI and Alexandru-Liviu OLTEANU from the DECIDE team, and was pursued in a student project by [Kadriye Nur BAKIRCI *et al.*](#). This line of research led to the following publication:

[C₂₆]

Miguel GUZMAN, Kadriye Nur BAKIRCI, Martin ROUESNÉ, Hélène SAVATIER-DUPRÉ, Bastien PASDELOUP, Patrick MEYER (2023). **“Utilisation de Processus Ponctuels Déterminantaux pour la Sélection de Parents dans un Algorithme Génétique : Application à l’Aide Multi-Critère à la Décision”**. *Colloque du GRETSI* (pp. 141–1144)

This work investigated the use of DPPs for parent selection in genetic algorithms, with an application to multi-criteria decision aid. The study considered the RMP model, for which a genetic algorithm was used to optimize the model parameters. The work identified a limitation of conventional selection mechanisms: by repeatedly favoring the fittest individuals, they could reduce population diversity and lead to premature convergence. To address this issue, k -DPP sampling was introduced as a diversity-aware parent selection mechanism, which draws a subset of individuals according to their pairwise dissimilarity and therefore deliberately ignores their fitness; a hybrid variant was also proposed, in which a diverse mating pool is first sampled by k -DPP and the parents are then drawn from it by conventional roulette selection, so as to recover a quality criterion. Two similarity kernels were investigated, one based on the ℓ_1 distance between the vectorized model parameters and one based on the Kendall-Tau correlation between the rankings that the corresponding models produce. Experiments showed that DPP-based selection increased population diversity across generations – significantly so with the Kendall-Tau kernel – and slightly improved the generalization of the resulting decision models to previously unseen alternatives, the ℓ_1 kernel performing best; combining the two selection mechanisms proved less effective than using either of them alone. These modest gains came at a considerable computational cost, from around 50 seconds for roulette selection to more than an hour for the ℓ_1 kernel and several hours for the Kendall-Tau one, the more expensive kernel being no more accurate.

Following this research on genetic algorithms for inferring RMP models, we extended the approach by developing faster crossover and mutation operators in the population-based algorithm. This work was published as follows, with preliminary results presented in [C₂₀; W₁₀]:

[J₆]

Arwa KHANNOUSSI, Alexandru-Liviu OLTEANU, Patrick MEYER, Bastien PASDELOUP (2024). **“A Metaheuristic for Inferring a Ranking Model Based on Multiple Reference Profiles”**. *Annals of Mathematics and Artificial Intelligence* (vol. 92, nb. 6, pp. 1467–1488), Springer

This work developed an evolutionary metaheuristic for inferring the parameters of a ranking model based on RMP in the context of multi-criteria decision aiding, focusing on the simple variant of the model in which the importance of criteria is represented by additive weights. The objective was to learn a preference model from holistic judgments provided by a decision maker, here pairwise comparisons between alternatives. The work addressed the difficulty of learning the parameters of such models, whose number increases rapidly with the number of criteria and reference profiles, which prevents existing approaches based on MILP or constraint programming from handling realistic instances. A metaheuristic was therefore designed to search directly for model parameters that reproduced the observed preferences. The approach was evaluated on extensive synthetically generated datasets, where a reference model was used to generate training preferences and the inferred model was then assessed both on the training alternatives and on previously unseen alternatives, including in the presence of noisy comparisons. The experiments showed that the method attained a generalization power comparable to that reported for exact formulations, while requiring significantly lower computation times and remaining applicable to instances that those formulations cannot address; on the small instances they do handle, and without a time budget, those exact methods remain the more accurate. In particular, the results demonstrated that

the inferred models could generalize the decision maker’s preferences beyond the alternatives used during learning, and that this ability degrades gracefully when the training comparisons are noisy.

This work served as a reference for subsequent work in the DECIDE team [MALLÉGOL, 2023 ; JOURDIN *et al.*, 2024 ], where operational research for decision aiding is still an active line of research.

II.4.2 Physics & Time Series

I also had the opportunity to collaborate with researchers interested in applications of ML approaches to physics. In such domains, time series are often an object of interest, so I have chosen to group some non-physics-specific related works in this section too.

Although no publication came out of these collaborations, I supervised the works of Mervyn GUILLOU and Behnam VAHDANI, who both worked on projects in partnership with Elliptika, a company located in Brest, France, which wanted to use AI to learn the parameters of signal filters from time series – instead of performing a full costly simulation – with the objective of 3D printing such filters afterward. This first experience with physicists made me discover a very frequent concern in applied physics: the need for fast models to replace already existing simulations. This problem arose several times in my research.

Regarding physics-related applications, I had the opportunity to supervise the PhD theses of two students. The first one was Fred Michael GONSALVES, who was doing a CIFRE PhD with Chantiers de l’Atlantique, a large French-based ship manufacturer. In this collaboration, the goal was to reduce the energy consumption of ships, through the use of operational research and AI. His manuscript is presented below:



[T₅]

Fred Michael GONSALVES (2023). **“Leveraging Machine Learning and Operations Research to Enhance Cruise Ship Energy Efficiency”**. *École Nationale Supérieure Mines-Télécom Atlantique (PhD)*

The thesis of Fred Michael GONSALVES investigated the use of ML and operational research to improve the energy efficiency of cruise ships, considering both operational optimization and the design of onboard energy systems. The work addressed the growing need to reduce the environmental impact of cruise shipping by developing data-driven models to characterize ship energy consumption and optimization methods to determine more efficient operating strategies. On the operational side, which concentrated on propulsion as the largest energy consumer aboard, ML models were developed to predict propulsion power from historical measurements and environmental conditions, providing a data-driven representation of the relationship between ship operation, weather conditions, and energy consumption. Particular attention was given to the methodology behind these models: the integration of hindcast rather than forecast weather data, and a temporally consistent splitting of training and validation data, both aimed at obtaining parsimonious models whose measured performance reflects their real ability to predict future operating conditions. Two applications were then built on these models: a fluid deadweight management tool, which optimally distributes the ship’s liquid loads so as to improve its hydrodynamic condition in operation, and a methodology to quantify the effect of biofouling on propulsion efficiency over time. On the design side, the thesis developed a generic MILP formulation of the cruise-ship multi-energy system, catering to the heterogeneous propulsion, electrical, and thermal demands of the vessel, and optimizing technological choices and operating decisions jointly rather than treating them as separate problems. This formulation was addressed as a bi-objective problem over lifetime cost and lifetime greenhouse-gas emissions, resolved both through an epsilon-constraint Pareto front giving a spectrum of intermediate solutions, and by internalizing carbon taxes into the cost objective.

His manuscript builds on two main contributions. First, we worked on numerous time series corresponding to the propulsion power of ships across several months. This allowed us to produce a fast ML-based model for predicting power consumption, thus replacing a full simulation. This work led to the following publication:

[C₁₃]

Fred Michael GONSALVES, Bastien PASDELOUP, Romain BILLOT, Patrick MEYER, Arnaud JACQUES, Matthieu LORANG (2021). **“New Insights into the Propulsion Power Prediction of Cruise Ships”**. *ICTAI (pp. 846–850), IEEE*

This work investigated the use of ML for predicting the propulsion power of cruise ships, with a particular focus on the methodological challenges associated with evaluating predictive models on high-frequency operational data. The study used 27 months of sensor and weather data from a large cruise ship and considered MLPs trained from variables related to ship speed, propulsion, draft, acceleration, wind, and waves. Rather than focusing primarily on proposing a new neural network architecture, the work examined how the way in which data was divided into training, validation, and test sets affected the apparent performance and generalization of the models. In particular, it built on the cross-talk problem, which arises when highly correlated observations from the same time periods are randomly distributed between training and validation sets, and showed that it is further aggravated by the overlapping windows commonly used when aggregating high-frequency measurements. This could lead to artificially low validation errors and an overly optimistic assessment of model performance. To address this issue, the work compared conventional shuffle-based splitting with a temporal splitting strategy, in which earlier observations were used for training and later observations for validation. The experiments showed that temporal splitting provided a more reliable estimate of performance on future unseen data and consistently generalized better than shuffle-based validation, the advantage persisting even at a six-month prediction horizon where the error inevitably grows. The study also investigated the effect of model complexity and showed that, for the considered MLP architecture, relatively small and parsimonious models performed as well as or better than more complex architectures, while exhibiting lower variability in their test errors.

One of the main impacts of this work was to highlight a fundamental flaw in existing literature on time-series forecasting using ML in the naval industry. Existing work achieved high performance on the validation set, which did not reflect the true ability to predict future operating conditions [GKEREKOS *et al.*, 2019]. Indeed, these works performed random (shuffle-based) validation, which introduced leakage between the training set and the test set. This gave the illusion of a model that was able to predict the future, while it was only interpolating within known operating conditions [NGUYEN *et al.*, 2025].

Fred Michael GONSALVES also explored more operational research-based approaches, with the development of an MILP formulation to design energy systems in large cruise ships. This led to the following unpublished work:

[P₃]

Fred Michael GONSALVES, Patrick MEYER, Romain BILLOT, Bastien PASDELOUP, Matthieu LORANG (2023). **“Optimal Design and Operation of Cruise Ship Multi-Energy Systems: An MILP Formulation”**. *HAL preprint hal-04310676*

This work developed an MILP formulation for the integrated design and operation of cruise-ship multi-energy systems. The objective was to jointly determine the configuration and sizing of the onboard energy system and its operating strategy while satisfying the ship’s heterogeneous energy demands, including propulsion, electrical, and thermal loads. The formulation represented the main energy flows and operational constraints of the ship while accounting for discrete decisions concerning the selection and use of energy-system components. It therefore allowed design and operational decisions to be optimized simultaneously, rather than treating them as separate problems, and was deliberately written in a generic way so that different technologies can be substituted without modifying the underlying model. The

formulation was instantiated with internal combustion engines and solid oxide fuel cells, and solved under three objective functions – lifecycle greenhouse-gas emissions, lifecycle cost, and lifecycle cost including a carbon tax – with the resulting design and operation results validated by domain experts. Consistently with the literature, solid oxide fuel cells yielded significant reductions in greenhouse-gas emissions. More unexpectedly, the carbon-tax analysis over a 15-year horizon showed a much weaker effect of such measures than anticipated, which raises questions about their ability to drive the adoption of cleaner but capital-intensive technologies in the cruise-ship industry.

Leaving the sea for the land, I later had the opportunity to supervise the PhD of [Dubon RODRIGUE](#), who wanted to accelerate the simulation of DHNs for the development of centralized heating systems in modern cities. His PhD manuscript is summed up below:

 [T₈] Dubon RODRIGUE (2025). **“Topology Simplification of District Heating Networks via Artificial Intelligence-Based Node Aggregation”**. *École Nationale Supérieure Mines-Télécom Atlantique (PhD)*

The thesis of [Dubon RODRIGUE](#) investigated the use of ML to simplify the topology of DHNs and reduce the computational cost of their dynamic simulation. Detailed DHN models required the simulation of heat transport through pipes, energy consumption at substations, and heat generation from multiple sources, leading to increasingly large and computationally expensive systems as the network size grew. The work addressed this limitation by replacing clusters of substations with supervised ML models, primarily based on artificial neural networks, trained to reproduce the physical dynamics of the aggregated parts of the network. This created a hybrid representation in which the original physical model was retained for the main network while selected groups of substations were represented by learned surrogate models, thereby reducing the number of physical components and equations that had to be simulated, and achieving up to an 82% reduction in simulation time on a large-scale network. The thesis also proposed a graph-based clustering methodology to automatically identify suitable groups of substations for aggregation, exploiting the network structure and the similarity of their dynamic behavior, and grounded in a statistical analysis of 2,696 distinct clusters drawn from 17 different networks. The resulting approaches were evaluated on several DHNs and showed that substantial reductions in model complexity and simulation cost could be obtained while preserving comparable physical accuracy. Along the way, the work also contributed two software tools: HeatGrid, a pseudo-dynamic physical simulation model whose predictions were experimentally validated, and SyntheticDHN, a generator of realistic third-generation networks used to build the case studies.

A first work consisted in replacing substations with surrogate models, to evaluate the feasibility of the approach. To do so, we first worked on hand-defined clusters, and trained models to represent them as a single node. This work led to the following article:

 [J₇] Dubon RODRIGUE, Mohamed Tahar MABROUK, Bastien PASDELOUP, Patrick MEYER, Bruno LACARRIÈRE (2024). **“Topology Reduction Through Machine Learning to Accelerate Dynamic Simulation of District Heating”**. *Energy and AI (vol. 17, art. 100393), Elsevier*

This work proposed an ML-based methodology for reducing the computational cost of dynamic simulations of DHNs. The approach addressed the increasing complexity of DHN simulations, particularly for

large networks integrating renewable heat sources and thermal storage, by replacing selected clusters of consumer substations with surrogate models. The DHN was represented as a graph, with the selected clusters treated as aggregated nodes whose thermal dynamics were learned from physical simulations. Three neural network architectures were compared on 220 diverse clusters extracted from four synthetic DHNs with different topologies and heating-demand levels: MLPs proved clearly inadequate, with a mean absolute error of 0.797°C against 0.121 for convolutional models and 0.172 for recurrent ones, confirming that temporal dependencies must be modeled explicitly; since the latter two were statistically indistinguishable, the GRU-based recurrent architecture was retained for its lower internal complexity. The results also showed that the surrogate models were able to reproduce the dynamics of the replaced clusters across a wide range of configurations, including clusters of more than 100 substations, with the external connectivity of a cluster having a stronger influence on prediction performance than its size or heating demand, since it determines how many independent temperature signals have to be learned. The learned models were then integrated with the remaining physical simulation to form a hybrid simulation framework, combining data-driven predictions for aggregated clusters with physics-based modeling for the rest of the network. This approach achieved a substantial reduction in simulation time while maintaining acceptable accuracy in the thermal power generated by the heat sources: replacing 39% of the substations of a medium-sized network cut the simulation time by 27% for an error below 1%, while replacing 83% of a network of 700 substations cut it by 82% for an error of 1.42%.

Details on the integration of the approach in an MPC optimization framework were published as follows:



[C₂₉]

Dubon RODRIGUE, Mohamed Tahar MABROUK, Bastien PASDELOUP, Patrick MEYER, Bruno LACARRIÈRE (2024). **“Leveraging Neural Networks in a Hybrid Model Predictive Control Framework for District Heating Networks”**. *ECOS* (pp. 287–298)

This work proposed a hybrid MPC framework for the optimal operation of large DHNs, combining physics-based dynamic simulation with neural network surrogate models. The approach addressed the high computational cost of conventional MPC, which repeatedly simulated the complete network during the optimization of the mass flow rates and supply temperatures injected by the heat sources. Rather than replacing the entire physical model with a data-driven model, the work selectively replaced clusters of consumer substations with RNNs trained to reproduce their aggregated heat consumption and outlet-water temperature dynamics. This localized strategy reduced the amount of data required for learning while preserving the physical representation of the rest of the network. The neural networks accounted for the thermal inertia of the clusters by processing sequences of past network states, and the resulting hybrid model was integrated directly into the rolling-horizon MPC optimization. Experiments on a synthetic 150-node DHN showed that the neural networks reproduced the cluster dynamics with mean absolute percentage errors below 0.6% on the outlet temperatures and below 2.7% on the consumed thermal power, the error on the delivered energy remaining lower than on the instantaneous power since over- and under-estimations partly cancel out. Replacing approximately one-third of the network with learned models reduced the computational time required by the MPC by 9% while maintaining comparable control performance to the fully physics-based approach. The results also showed that the achievable computational reduction grew with the size of the replaced clusters – substituting the smallest one brought no gain at all, its inference time canceling the equations spared – whereas prediction accuracy varied between clusters without following their size, which the article attributes to differences in physical complexity rather than to scale.

After this controlled work on hand-defined clusters, we wanted to extend our work to automatically defined sets of nodes to aggregate, offering a flexible approach to balance between accuracy and acceleration. This work led to the following article, with early results presented in [W₁₅]:

[J₁₁]

Dubon RODRIGUE, Mohamed Tahar MABROUK, Bastien PASDELOUP, Patrick MEYER, Bruno LACARRIÈRE (2026). **“Flexible Clustering of Substations for Accurate and Rapid Hybrid Simulation of District Heating”**. *Energy Informatics, Springer*

This work proposed a flexible, task-driven graph clustering methodology for improving the accuracy and computational efficiency of ML-based spatial reduction of DHNs. The approach addressed a key limitation of previous hybrid simulation methods: the quality of the reduced model depended strongly on how substations were grouped into clusters, while clusters of similar size could lead to substantially different surrogate-model accuracy. To address this issue, the work introduced physics-informed distance metrics designed to capture the three factors that drive the difficulty of learning the dynamics of a cluster, namely its thermal losses, its transport time, and its external connectivity. These metrics were integrated into a hierarchical agglomerative clustering procedure, allowing the methodology to automatically identify clusters that were suitable for replacement by surrogate models. The approach was evaluated on 16 DHNs with diverse topologies and thermal characteristics and was further validated on an independent network. The experiments revealed a clear trade-off between simulation accuracy and computational reduction, which could be adjusted according to user preferences. On the validation network, configurations prioritizing physical accuracy reduced the simulation time by 7% while limiting the deviation in total heat production to 6.45 MWh, or 3.76% of the network thermal losses, whereas configurations prioritizing computational efficiency achieved a 91% reduction in simulation time, at the cost of a 45.24 MWh deviation, or 26.37% relative to thermal losses.

Finally, we validated our approach using Dubon RODRIGUE’s tools, on two real-world DHNs. This showed the practical relevance of the approach, and the importance of validating DHN simulation tools in actual operational conditions rather than on the usual synthetic benchmarks. This work led to the following publication:

[C₃₃]

Dubon RODRIGUE, Roberto BOGHETTI, Jérôme Henri KÄMPF, Bastien PASDELOUP, Mohamed Tahar MABROUK, Patrick MEYER, Bruno LACARRIÈRE (2025). **“Validation of Two Distinct Simulation Models of District Heating Networks: Application to Efficient Looping Analysis”**. *Journal of Physics: Conference Series (vol. 3140, nb. 4, art. 042021), IOP Publishing*

This work investigated the validation and practical use of two distinct quasi-dynamic simulation models for DHNs, namely HeatGrid and PyDHN, which differ in how they digitalize the network, in the hydraulic model used to compute mass flow rates and pressure losses, and in the thermal equations governing water temperature. The study addressed a key limitation of many DHN simulation approaches, namely the lack of validation against measurements from real-world networks. Both models were evaluated using on-site measurements from a real meshed DHN and, despite their methodological differences, produced closely consistent results, with mean absolute differences of only 0.0268 kg/s for the mass flow rate and 0.75 °C for the return temperature at the heating station. The validated models were subsequently used to assess the potential of adding new looping pipes to the network, with the objective of reducing pumping energy and improving hydraulic performance. The simulations showed that although additional loops reduced pumping requirements, they simultaneously increased heat transport losses. Consequently, the resulting energy savings were insufficient to compensate for the additional thermal losses, making all investigated

looping scenarios economically unattractive – a conclusion the article underlines as counter-intuitive, and which illustrates the value of pre-validated simulation models for the rapid assessment of network extension scenarios.

The works of [Dubon RODRIGUE](#) are often cited in review articles as representative examples of clustering-based simulation approaches [[BOGHETTI et al., 2026](#) ; [RINCHI et al., 2026](#) , which are expected to grow in importance in years to come in physics-informed simulation [[THUMMERER, 2026](#) .

Leaving physics for the not-so-distant study of time series for industrial cases, I had the opportunity to supervise the work of [Manon RENAULT](#) in a collaboration with Goodflow, a company that develops systems to integrate into reusable packages. Such systems integrate several sensors, and monitor events through regular measurements. In the work we did, we explored the potential of ML tools for classification of an event of interest from sensor measurements, in the following publication:



[Manon RENAULT](#), [Hamoud YOUNES](#), [Hugo TESSIER](#), [Ronan LE ROY](#), [Bastien PASDELOUP](#), [Mathieu LÉONARDON](#) (2025). **“Event Classification of Accelerometer Data for Industrial Package Monitoring with Embedded Deep Learning”**. *COMPSAC* (pp. 1671–1676), *IEEE*

This work investigated the use of embedded deep learning for monitoring the state and location of reusable industrial packages from accelerometer data. The objective was to develop a low-power sensing system capable of operating for several years while distinguishing different package situations, such as being transported by a forklift or a truck. Since maximizing the lifespan of the device amounts to minimizing its wake time, the classifier must combine very high precision, so that no unnecessary operation is triggered by a false alarm, with high recall, so that no event requires a second inference. The study addressed the challenges associated with classifying imbalanced, multiclass time-series data collected from an embedded IoT sensor, while also considering the severe computational and memory constraints of the target device. A 1D CNN was developed to classify the accelerometer signals, and two synthetic oversampling techniques, SMOTE and ADASYN, were investigated to compensate for the imbalance between classes; predictions whose confidence fell below a threshold were routed to an additional class rather than assigned to an event, so that uncertain windows could be rejected instead of misclassified. The resulting model was subsequently compressed to reduce its memory footprint and make embedded deployment possible. On the considered classification problem, the approach achieved precisions of about 95% for each class, while model compression reduced the model size by a factor of four. The compressed model was finally deployed directly on the microcontroller-based IoT device, where inference required 316 mW of power.

Finally, totally leaving industry for a more academic project, I joined Axel MARMORET for a project on audio. We were interested in studying the functioning of ASR models [[BAEVSKI et al., 2020](#) ; [RADFORD et al., 2023](#) , and took the opportunity to evaluate how such models would behave on out-of-distribution data. To do so, we assembled a corpus of subgenres of Heavy Metal – from the softest to the most extreme ones – with the goal of assessing the ability of the main models to transcribe the lyrics. This work led to the following publication:

[C₃₂]

Bastien PASDELOUP, Axel MARMORET (2025). **“Raining Words : Les Modèles d’ASR Peuvent-Ils Retranscrire les Sous-Genres de Metal ?”**. *Colloque du GRETSI* (pp. 357–360)

This work investigated the ability of recent ASR models to transcribe extreme vocal styles found in Metal music, with a particular focus on their robustness in an out-of-distribution setting. The study evaluated five state-of-the-art ASR models on three complementary corpora: isolated recordings of extreme vocals, a purpose-built collection of complete Metal songs covering 14 subgenres, and the vocal tracks extracted from those songs by source separation, the latter being intended to distinguish difficulties caused by the vocal characteristics from those introduced by the musical accompaniment. Transcription quality was assessed with both syntactic and semantic metrics, and the study was restricted to English lyrics. No model transcribed the whole range of styles reliably, and three conclusions emerged. First, among the models considered, the two most recent Whisper versions stood out, if only slightly, and did so whatever the subgenre. Second, performance varied enormously with the vocal style, and did so consistently with human perception: four groups appeared, ordered by extremeness, gathering clear vocals (Heavy, Symphonic, Punk, Doom), moderately saturated ones (Thrash, Melodic Death), saturated ones (Black, Death, Deathcore), and extreme ones (War, Brutal Death, DSBM, Grindcore, Goregrind). Third, and more surprisingly, isolating the vocals through source separation did not significantly improve transcription, which the article relates to Whisper having been trained on audio recorded under widely varying noise conditions, and therefore being comparatively robust to acoustic perturbations.

Unexpectedly, this work was of interest to some people on social media, as it addressed similar problems to research on audio models for understanding deaf people, whose speech is also considered out-of-distribution. We plan to extend this work at some point, using LoRA [HU *et al.*, 2022 

I also took the opportunity to participate in the reproducible science initiative organized by the GRETSI conference, which improved my code publication habits. For this work, we received the silver badge for reproducibility.

We extended our collaboration with Axel MARMORET through the supervision of a student project, aiming at understanding what makes a good musical mashup, and at creating a system for automatically doing so from two source songs. The project, conducted by Marine DELABAERE *et al.*, was presented as follows:

[C₃₀]

Marine DELABAERE, Léa MIQUEU, Michael MORENO, Gautier BIGOIS, Gautier DUONG, Ella FERNANDEZ, Flavie MANENT, Maria SALGADO-HERRERA, Bastien PASDELOUP, Nicolas FARRUGIA, Axel MARMORET (2025). **“AutoMashup: Automatic Music Mashups Creation”**. *Colloque du GRETSI* (pp. 701–704)

This work introduced AutoMashup, an open-source system for automatically creating music mashups by combining vocal and instrumental stems from different songs. The proposed pipeline combined source separation, musical analysis, and compatibility estimation to automate the main stages of mashup creation: a source separation model splits each song into its constituent stems, while musical information such as beat positions, tonality, and song structure is extracted to align and combine them harmonically and rhythmically. A central contribution was the use of COCOLA, a representation model specifically designed to assess the coherence of musical accompaniments, to estimate the compatibility between candidate song pairs. Because COCOLA requires a pairwise comparison of the audio itself, and therefore scales quadratically with the size of a song library, the work also investigated whether general-purpose pretrained audio representations, namely CLAP and MERT, could perform this compatibility estimation in a zero-shot setting, which would reduce the search to a linear scan of precomputed embeddings.

The experiments showed that mashup compatibility was asymmetric: the suitability of combining two songs depended on which song provided the vocals and which provided the accompaniment, so that compatibility cannot be derived from symmetric musical features such as tonality or tempo alone. COCOLA was taken as the reference for compatibility – an assumption supported only by the authors’ own listening, and whose formal validation is left as future work – and the distances between CLAP or MERT embeddings showed essentially no correlation with it. These results highlighted a limitation of general-purpose audio representations for music recombination tasks, despite their strong performance on broader audio-understanding problems, on a corpus deliberately kept narrow: 21 popular-music tracks in major keys close to C and of comparable duration.

CHAPTER III

FUTURE WORK & DIRECTIONS



This chapter outlines my plans for the coming years, both in terms of research projects and future directions. It is organized in two main parts, starting with short-term projects ([Section III.1](#)), followed by the longer-term directions they converge toward ([Section III.2](#)).

III.1	Short-Term Perspectives	74
III.1.1	Projects on Fundamental AI	74
III.1.1.1	Next Steps in ENDIVE	74
III.1.1.2	Agentic AI & Trust Between Agents	77
III.1.1.3	Development of REVE 2	79
III.1.2	Applications of AI	79
III.1.2.1	The STRATUM Project	79
III.1.2.2	The SIGMAS Project	80
III.1.2.3	Anticipating Cardiology Admissions from Weather Conditions	80
III.1.2.4	Applications to Audio Processing	81
III.2	Future Research Directions	82
III.2.1	Text as the Pivot of Multimodality	82
III.2.2	Agentic AI & Trust	83
III.2.3	Genomics as the Next Modality	84
III.2.4	A Cross-Service Agentic Environment	85

III.1 Short-Term Perspectives

In this section, I describe the near future of ongoing projects, as well as the ones that are already on the table for the coming years. As in [Chapter II](#), I have organized them in two main categories: projects on fundamental AI ([Section III.1.1](#)), and projects on applications of AI ([Section III.1.2](#)).

Once again, this organization reflects my approach to research. The former aim to keep up with recent advances in AI, while the latter aim to bring these advances to real-world problems, in collaboration with specialized researchers and practitioners from the application domains concerned.

III.1.1 Projects on Fundamental AI

Regarding fundamental aspects of my research, the next steps will be both direct continuations of existing projects on diversity and model development, and the exploration of new research directions related to agentic AI and quantification of trust between agents.

III.1.1.1 Next Steps in ENDIVE

The [ENDIVE](#) project – which I already briefly introduced in [Section II.2.2](#), and fully described in [[ENDIVE](#) KULESZA et al., 2011 ; [KULESZA et al., 2012](#) ]. DPPs make it possible to account both for the relevance of the points sampled and for their similarity, yielding a diverse form of importance sampling that covers the distribution better and estimates its mean with lower variance. [Figure 9](#) illustrates this sampling approach, compared with uniform sampling, on a random point cloud.

The project was initially designed for FSL, but the evolution of research in AI led me to reorient it. The rapid development of LLMs has progressively caused FSL to shift from an isolated paradigm to a subproblem of foundation model adaptation. While this does not change the main objectives of the project, integrating LLMs more prominently into it followed naturally.

The revised project explores the benefits of diversity along two main axes:

Diversity in the Data

Models are currently trained on massive amounts of data, with scaling laws showing that more data generally leads to better models [[KAPLAN et al., 2020](#) ; [HOFFMANN et al., 2022](#) ]. However, data curation is increasingly recognized as a crucial component of the scaling process, as the benefits of additional data depend not only on its quantity, but also on its quality, relevance, and diversity [[LI et al., 2023](#) ; [UDANDARAO et al., 2025](#) ]. In particular, naively selecting data according to individual quality or importance can result in highly redundant subsets, limiting the effective diversity of the training signal [[ZHANG et al., 2025](#) ]. Recent approaches therefore explicitly combine data quality with diversity, showing that carefully curated subsets can achieve competitive or even superior performance while using substantially less data [[WU et al., 2024](#) ; [ZHANG et al., 2025](#) ; [YANG et al., 2025](#) ], with recent applications to fields such as oncology [[BOSCH et al., 2026](#) ].

In this context, the [ENDIVE](#) project aims to explore the following research questions:

- **How can diversity-aware data curation make training of diffusion LLMs more data-efficient?** — It is not obvious that previously mentioned approaches extend to discrete diffusion LLMs: a curation criterion designed for next-token prediction scores an example by what it teaches about the continuation of a prefix, whereas the masked-denoising objective supervises every position of a sequence at once and across corruption levels, so the quantity that makes a training example informative is not the same one. Additionally, none of




Figure 9 Illustration of the differences between DPP sampling (a) and uniform *i.i.d.* sampling (b). DPP-sampled points exhibit interesting properties in terms of diversity (c) and estimation of the mean of the distribution (d). Here, a Gaussian kernel as in [TREMBLAY *et al.*, 2023] is used to model similarities between points as a decreasing function of the Euclidean distance. In practical cases, this should depend on the desired diversity criteria. Figure taken from [ENDIVE].

these works make use of DPPs, which are natural candidates for such applications. Comparing and combining these two lines of work is one of the next goals of the project.

- **How does data diversity shape the behavior of diffusion LLMs?** — Recent work on continuous diffusion language models such as ELF [HU *et al.*, 2026] highlights that language generation can be understood as a trajectory through a continuous representation space, progressively transforming noisy representations into linguistically meaningful ones. Data diversity in training sets may be a key element for altering these decoding trajectories. In addition, mechanistic interpretability offers approaches to reveal the internal mechanisms of these models, and investigating the link between the diversity of the training data and the diversity of the mechanisms it gives rise to is a new and exciting research direction. Making this second notion measurable is part of the question: the overlap between the features recovered by SAEs trained on models with different data diets, and the consistency of those features across training runs [SONG *et al.*, 2025], are the two proxies I would start from.

Diversity in the Representations

The emergence of foundation models has substantially changed the role of learned representations in ML. Rather than learning task-specific representations from scratch, downstream tasks can exploit representations acquired from large and diverse pretraining corpora through frozen embeddings, prompting, retrieval, or lightweight adaptation. In language, the contextual nature of transformer representations further means that the representation of an input depends on its surrounding context and on the computation performed by the model. At the same time, recent work suggests that diversity is an important property of both the representations and the data used to construct them: diverse prompts can elicit complementary reasoning paths, while diverse models can provide complementary strengths in ensemble settings [HU *et al.*, 2025]. Recent advances in mechanistic interpretability also seek to decompose these representations into meaningful features and computational circuits [ELHAGE *et al.*, 2021; BRICKEN *et al.*, 2023], opening the possibility of studying not only whether representations are useful, but also how their diversity relates to the diversity of the computations they support.

These developments raise the following research questions:

- **Which features are relevant, and how can diversity improve representations?** — In the transfer-based FSL formulation of the project, I considered pretrained representations as high-dimensional feature spaces and formulated this question as the study of whether restricting these representations to relevant dimensions could provide a more concise and discriminative description of the data. Foundation models offer a richer perspective on this problem: their representations encode a large variety of concepts and features acquired from heterogeneous pretraining corpora, many of which may be irrelevant to a particular downstream task. Mechanistic interpretability approaches decompose representations into meaningful features and functional components – most often through SAEs [BRICKEN *et al.*, 2023; HUBEN *et al.*, 2024; GAO *et al.*, 2025] – while also highlighting the importance of understanding how reliably such features are represented across models and training runs [SONG *et al.*, 2025]. Diversity may provide a useful criterion for this decomposition: retaining highly correlated features may provide little additional information, whereas diverse task-relevant features could offer a more compact coverage of the information encoded by the model.
- **How can diversity of context be exploited for downstream tasks?** — Transformer representations are inherently contextual: the representation associated with an input depends on the surrounding sequence and on the transformations performed by the model. This is particularly important for LLMs, where prompting and retrieval (*e.g.*, with RAG [LEWIS *et al.*, 2020]) can modify the context in which an input is processed without modifying model parameters. In the few-shot setting, this provides a mechanism for generating multiple representations of the same example by varying demonstrations, retrieved documents, or instructions. Rather than seeking a single optimal context, diversity can therefore be used to explore complementary interpretations or computational trajectories induced by the model. Recent work illustrates this principle at inference time: diverse prompts can produce complementary reasoning paths and improve ensemble performance [HU *et al.*, 2025]. The diversity of contextual representations can similarly be exploited for downstream tasks, where selecting complementary rather than merely similar contexts can improve prediction,

robustness, and adaptation. This could lead to demonstration and retrieval strategies that optimize both relevance to the input and diversity of the information and representations elicited by the selected contexts.

- **How can diversity guide the selection and combination of foundation models?** — The rapid proliferation of foundation models creates a new model-selection problem: given a collection of pretrained models, which model – or combination of models – provides the most appropriate representation for a target task? Models may differ in their pretraining data, architectures, training objectives, scales, modalities, and learned representations. Diversity is particularly relevant here because models with similar predictions or representations are unlikely to provide much additional information when combined, whereas complementary models may compensate for each other’s weaknesses. Recent LLM ensemble studies explicitly exploit this principle, using diversity in model behavior or response fingerprints to construct more robust ensembles [COHEN *et al.*, 2026]. I therefore aim to investigate whether diversity between foundation models and their representations can serve as a principled criterion for model selection and combination, identifying complementary models rather than simply selecting the individually strongest ones.

These questions are ambitious, and some will very likely be studied after the end of the project. In the current state of this project, both Jonathan LYS and Joshua Chaz RIVERA are working on diffusion LLM-related aspects, and Nolan SISOUPHANTHONG-VONGKOTRATTANA will probably contribute to elements related to mechanistic interpretability. I also plan to hire a postdoctoral researcher at the beginning of 2027 for 18 months, as well as a Master 2 student for 6 months. They will work primarily on questions related to mechanistic interpretability, both because I feel this is currently the most promising axis of research, and because it will encourage teamwork between Nolan SISOUPHANTHONG-VONGKOTRATTANA and the new recruits.

III.1.1.2 Agentic AI & Trust Between Agents

The next short-term project on fundamental AI I want to describe concerns agentic systems, *i.e.*, AI systems able to plan sequences of actions, call external tools, and make decisions with limited human supervision [YAO *et al.*, 2023]. This project, entitled “*Supervision of Trust in Agentic AI Systems Applied to the Medical Domain*”, is conducted in collaboration with Jean-Charles VIALATTE, head of the startup Cerevial, which develops the Alive Engine agent platform described below. It is a maturation project, funded for 6 months and staffed by Rayan DAKHIA, whose purpose is to establish feasibility ahead of the CIFRE thesis discussed at the end of this section.

The motivation for this project comes from the speed at which agentic AI is being adopted in healthcare, compared with the maturity of the tools available to supervise it. An industry survey of healthcare executives reports that many healthcare organizations already develop or deploy agentic AI initiatives, and that most of them plan to increase their investments in the coming years [MICROSOFT AI]. At the same time, analyses insist that in the agentic era, organizations can no longer restrict themselves to controlling what models say, but also have to guard against what they do (*e.g.*, unintended actions, inappropriate tool use, or violations of established guardrails [AGENTIC AI TRUST]). This is particularly critical in the medical domain, where many clinicians have already been confronted with medical hallucinations produced by foundation models, and consider that such errors can directly harm patients [KIM *et al.*, 2025]. Interestingly, clinical audits reported in the same work suggest that most residual hallucinations do not come from missing factual knowledge, but from failures in causal or temporal reasoning, which makes them particularly hard to detect from the output alone.

The project builds on Alive Engine [ALIVE ENGINE], which orchestrates persistent agents: it manages their background tasks, their sleep and wake cycles, and the persistence of their internal state. One of its design choices is especially interesting from a research point of view: the personality of an agent is encoded as a structured context block, mixing static parts with dynamic parts injected during interactions, and revised during regular introspection phases, where the agent consolidates its memory and adapts its own representations. Extending this idea, one could adapt the system so that each agent maintains memory blocks profiling the other actors it interacts with, be they humans or other agents, which opens the possibility of encoding an explicit notion of inter-agent trust that evolves with the history of exchanges.

Existing work only partially addresses the supervision of such systems, and the HuggingFace incident has recently shown how critical that can be [HUGGINGFACE INCIDENT]. Chain-of-thought prompting [WEI *et al.*, 2022] is known to reduce hallucinations in most settings, but has also been shown to paradoxically increase major hallucinations and omissions in some clinical tasks, in particular when generating medical notes from transcriptions [ASGARI *et al.*, 2025]. Self-improving agents, able to read and modify their own source code, additionally exhibit a risk of behavioral drift, where small successive modifications accumulate until the system silently departs from its initial objective [REN *et al.*, 2026]. Available tooling mostly focuses on detecting anomalies in outputs, and provides no integrated mechanism to supervise the introspection cycles themselves, *i.e.*, to verify that self-evaluation and self-improvement phases do not degrade previously acquired capabilities [REBEDEA *et al.*, 2023].

The project therefore aims to develop a methodological and experimental framework for the supervision of trust in medical agentic systems, organized around three complementary contributions:

- **A regression detection protocol** — This will be based on test suites executed periodically, to verify that introspection and self-improvement cycles do not degrade performance on tasks that were previously mastered.
- **A reasoning audit trail mechanism** — This will attach to each conclusion produced by an agent a machine-checkable provenance record – sources consulted, tool calls issued, intermediate claims and decision points – using knowledge graphs as a provenance support. Such a record establishes what the agent really did; the faithfulness of self-reported reasoning is a distinct and largely open problem.
- **A monitoring system for behavioral drift** — This will aim to detect the evolution – intended or not – of biases, tone, inter-agent trust and decision orientations of the modeled personalities of agents.

These contributions will be validated on a concrete use case, namely a medical agent deployed on Alive Engine and evaluated on its ability to assist a practitioner in the analysis of intensive care records (differential diagnosis support, detection of drug interactions, literature synthesis), in the spirit of multi-agent medical reasoning systems such as [TANG *et al.*, 2024]. Numerous medical benchmarks and datasets are available – MedQA [JIN *et al.*, 2021], PubMedQA [JIN *et al.*, 2019], MedMCQA [PAL *et al.*, 2022], BIOASQ [TSATSARONIS *et al.*, 2015], MedS datasets [WU *et al.*, 2025], or MIMIC-IV for intensive care records [JOHNSON *et al.*, 2023] – to support our experiments. These supply the fixed task suites of the first contribution, replayed after each introspection cycle so that a drop on a previously mastered task can be detected. They cannot, however, measure provenance or drift: being static and single-turn, they carry neither the time axis nor the execution trace that the two other contributions target. Part of the work will therefore consist in instrumenting the deployed agent itself – logging its introspection cycles, its tool calls and its inter-agent exchanges – so as to obtain the longitudinal measurements these contributions require.

Beyond its own objectives, this project is at the intersection of several directions I want to develop in the coming years. It naturally extends the reflections on agentic AI in healthcare I described in Section II.3, and the supervision mechanisms it targets are, in my opinion, a prerequisite to the deployment of the collaborative networks of hospital-specific agents I envision in Section III.2. It also provides a first opportunity to study trust as a quantity that circulates between agents, rather than as a property of a single model. In the shorter term, this maturation project is intended to lead to a CIFRE PhD thesis with Cereval, widening the scope to other critical domains and to formal certification approaches for autonomous agents. If the proof of concept is successful, it could hopefully lead to a tripartite collaboration with the CHU of Brest, which would provide evaluation scenarios in real clinical settings [S₁₄].

III.1.1.3 Development of REVE 2

Finally, following the reception of the REVE model [C₃₁] described in Section II.2.4, we have started to work on extensions of the model. This work is currently carried out by a team of current or ex-PhD students from the BRAIN team – Jonathan LYS, Nolan SISOUPHANTHONG-VONGKOTRATTANA, Yassine EL OUAHIDI, Lucas BENOÎT and Venkatesh SUBRAMANI – under the close guidance of Giulia LIOI and other permanent members of the team.

We are currently exploring how MEG can be integrated into the model, and are also working on architectural adaptations of the previous versions (alternative positional encoding strategies, new loss functions, etc.). Additionally, new datasets are being gathered, and new downstream tasks are being explored. Data curation may become a line of research as well, which could feed into the work presented in Section III.1.1.1.

While this was already a dynamic we had wanted to establish in the team, REVE has shown PhD students that collaborating – rather than working individually on their own subject – is the key to reaching top-tier conferences in AI. The current collaboration on REVE 2 is the largest collective effort I have seen in any team I have been part of since I started my career as a researcher, which I believe is one of the most beneficial consequences of REVE.

III.1.2 Applications of AI

Continuing the collaborations introduced in Section II.3 and Section II.4, I am currently involved in several applied projects, either already under way or about to start, described below.

III.1.2.1 The STRATUM Project

As a complement to the STRATIF-AI project, Serge TIMSIT (CHU of Brest) proposed the STRATUM project – currently under phase 2 review of ANR PRCE – which aims to identify blood-based markers of stroke measurable in the very first hours after the onset of symptoms, and to turn them into models able to support triage and prognosis at the bedside.

The clinical gap it addresses is easy to state. Distinguishing an ischemic stroke from an ICH, or from a stroke mimic – a seizure, a complicated migraine, a conversion disorder – currently relies entirely on neuroimaging [KARLIŃSKI *et al.*, 2015]. No blood test is used in routine care for diagnosis, subtype discrimination or prognostic stratification – candidate proteins have been reviewed for years, but their incremental value over clinical models remains modest [MONTELLANO *et al.*, 2021] – and since recanalization therapies are extremely time-sensitive, every hour spent on an uncertain diagnosis is an hour lost for treatment, with a strong impact on brain health [SAVER, 2006]. Beyond the acute phase, prediction of the functional outcome at three months still rests almost exclusively on demographic and clinical variables [NTAIOS *et al.*, 2012].

The project's bet is that whole-blood transcriptomics carries part of this information early enough to be useful. A signature of 9 genes, common to ischemic brain tissue and to peripheral blood, was first identified in a primate model of focal ischemia [RAMSAY *et al.*, 2019], then measured in the IBIS-CT1 study [IBIS-CT1] on 20 ischemic strokes, 20 ICH and 20 controls, with repeated blood samples collected from less than six hours after onset up to seven days. Combined, these genes discriminate stroke from controls with an AUC-ROC of 0.914, and a model trained on the corresponding RNA sequencing data reaches an F1-score of 0.94 for the prediction of the functional outcome at three months. These encouraging results – though coming from a small sample size of 60 subjects – suggest strong promise for the project. They also leave the harder question open: IBIS-CT1 included no stroke mimic, so the reported discrimination is against healthy controls rather than against the differential diagnosis that actually confronts a clinician at admission. Establishing that the signature separates strokes from mimics is precisely what the first work package below is for.

The project would run for 48 months, structured in three work packages:

- **A prospective multicenter trial** — The new IBIS-CT₂ study will include 720 patients admitted with a suspected stroke within six hours of onset, and will reduce the signature to the 3 to 5 genes that best separate strokes from mimics – a constraint dictated by the point-of-care assays that would carry the test to the bedside [MOEHLING *et al.*, 2021 - **A longitudinal multi-omic analysis** — This analysis will exploit the longitudinal nature of IBIS-CT₁ by adding repeated proteomic measurements to the transcriptomic ones, and will integrate them to look for early biological trajectories associated with long-term outcome.
- **Translational AI** — Harmonizing four heterogeneous cohorts, with clinical, biological, radiomic and transcriptomic data, quantifying how much each modality actually adds over the clinical baseline, and finally turning the retained models into a decision-support prototype.

This last work package is where I would contribute, and it studies exactly the kind of question I believe is important in healthcare: whether multimodality pays off on cohorts where some modalities are missing for many patients. It would also be my first practical contact with omics data – here transcriptomics and proteomics – a family of modalities I have identified as the next one to add to my applied work, in continuity with the collaborations of Section II.3.1.

III.1.2.2 The SIGMAS Project

Along similar lines, I will participate in the SIGMAS project, led by Quentin LE GRAND (CHU of Brest), whose purpose is to quantify imaging markers of cSVD directly on clinical acquisitions, and then to relate them to the clinical course of stroke patients and to their genomic determinants. cSVD affects the small penetrating vessels of the brain and accounts for up to one third of all strokes, counting both small-vessel ischemic strokes and ICH cases [LAM *et al.*, 2026 et al., 2013 

The project – currently in its development phase – is organized around three aims:

- **Segmentation** — Adaptation of existing tools [TSUCHIDA *et al.*, 2024 et al., 2024 - **Marker importance** — Association of the resulting markers with recurrence, mortality and functional outcome, and quantification of how much they add to existing predictive models.
- **Genomic determinants** — Association of the markers with genomic data in large clinical samples, comparing them with what is observed in the general population. This is the aim where the project meets the longer-term direction of Section III.2.3, even though my own contribution to it would be limited.

In this project, I will bring my expertise in AI to assist the clinicians involved. My contribution will most likely concern the first two aims, and may leverage elements from my past experience in image segmentation [C₂₈; P₅]. Still, the third aim will give me the opportunity to familiarize with genomic data, and may pave the way for future works in Section III.2.3.

III.1.2.3 Anticipating Cardiology Admissions from Weather Conditions

Moving from neurology to cardiology, we are discussing the development of a new project with Amine EL OUAHIDI (CHU of Brest), Romain DIDIER (CHU of Brest) and Bouabid EL OUAHIDI (Mohammed V University of Rabat, Morocco). This project aims to explore the links between myocardial infarction and weather conditions, and to exploit these links for practical purposes. Indeed, several epidemiological studies report that variations of atmospheric pressure and temperature are associated with the occurrence of myocardial infarction and coronary deaths [DANET *et al.*, 1999 et al., 2020 

The objective of this project is to predict the daily number of myocardial infarction admissions at the hospital's cardiology department well enough to be useful in practice: knowing, a few days in advance and from the weather forecast alone, how many doctors should be on duty at the CHU of Brest to account for the weather-induced risks. The data for this project will come from the ORBI registry [ORBI ; AUFFRET *et al.*, 2018 ], which records STEMI cases across Brittany from 2009 to 2016. Two properties of this source will have to be handled explicitly: it is region-wide, whereas the quantity of interest is the admission load of a single center; and it stops in 2016, so any model fitted on it will have to be confronted with recent admission counts.

This project may lead to a joint PhD thesis between the CHU of Brest, IMT Atlantique and the Mohammed V University of Rabat. The objectives above would then most likely form the first year of that PhD.

III.1.2.4 Applications to Audio Processing

Finally, the last short-term project I am involved in is the soon-to-start PhD of Quentin BANET, which I will co-supervise with Axel MARMORET. This PhD takes place in the context of a CIFRE agreement with Orange Labs, and will revolve around neural audio codecs, for mono and stereo signals.

Audio coding is an old field, born from source coding, whose successive codecs (*e.g.*, MP3, AAC) have long structured the way music is stored and transmitted. Like most fields, it was recently transformed by the rise of deep learning: neural networks reach compression rates that classical codecs cannot approach, and a new generation of methods has emerged from this observation [KIM *et al.*, 2025 ]. Autoencoder architectures trained with a GAN objective give the best results so far, with codecs such as SoundStream [ZEGHIDOUR *et al.*, 2021 ], EnCodec [DÉFOSSEZ *et al.*, 2022 ] or DAC [KUMAR *et al.*, 2023 ], all of which compress an audio signal into a sequence of discrete tokens and reconstruct it from them, at a perceptual quality that classical codecs only attain at much higher rates [MULLER *et al.*, 2024 ].

Two limitations of this generation motivate the thesis. Current neural codecs are essentially mono, whereas the signals to be transmitted are not. And they are considerably more complex than the codecs they aim to replace, both in computation and in storage: models in the range of 10M to 80M parameters are the norm, which is a real obstacle to deployment on the devices that would actually run them.

The objective of the thesis is therefore to design neural coding methods that address these points. To reach this goal, we will work on three challenges:

- **One representation for mono and stereo** — Obtaining a neural representation that covers both cases, rather than treating stereo as two mono streams or as a hack of a mono architecture.
- **Reduced complexity** — Bringing the size of the representation models down to something compatible with real deployment constraints, without giving up the perceptual quality that motivated the neural approach in the first place.
- **An interpretable latent space** — Obtaining a latent space that separates along interpretable directions, be they frequency-wise or content-wise, rather than an opaque sequence of tokens.

Recent architectures will be studied along the way, in particular transformers [VASWANI *et al.*, 2017 ] – which have been shown to scale well in this setting, at very low bitrates [PARKER *et al.*, 2025 ; XU *et al.*, 2024 ] – and diffusion models; new designs will be explored as well.

This project is, for me, an opportunity of a different nature from the others described in this chapter. The constraints are those of a company that has to ship a codec, so the trade-off between quality, size and interpretability is not negotiable in the way it is in an academic setting. The interpretability challenge echoes the questions on representations of Section III.1.1.1, on a modality I have so far approached only through a student project and a study of ASR models (see Section II.4.2).

It is also an opportunity to strengthen collaborations with other members of the BRAIN team: Axel MARMORET – whose interests lie in audio processing using AI – and to some extent Mathieu LÉONARDON, who specializes

in model compression and may help with the second challenge. On a personal note, music is a very big part of my life, so working on it directly – on topics that do not contribute to the degradation of the status of artists due to generative AI – is a pleasure.

To conclude with less formal audio-related projects, Axel MARMORET and I will continue proposing student projects and will pursue our exploration of ASR models for out-of-distribution data [C₃₂]. These student projects, while being smaller in scope than the projects described previously, are important to me. They allow me to stay in close contact with the students to give them a glimpse of what research looks like, while also being a good pretext to work on various modalities and models.

III.2 Future Research Directions

Most of the projects in Section III.1 are already bounded, or close to being so. I would like to close this chapter with longer-term directions I believe are important for the coming years. Still, given how fast AI evolves, it is possible that some of them will already be obsolete by the time I defend this HDR. The reorientation of ENDIVE described in Section III.1.1.1 is a good reminder that a research program written today has to survive that pace.

These directions nevertheless all converge toward the same objective, which I see as the natural continuation of the two halves of my research: a collaboration of models for multimodal healthcare – with text as a pivot modality – that a hospital can actually train or adapt, deploy, and trust.

My initial thought regarding this objective was to envision extremely large models with easy integration of various modalities – and robust to missing ones [LIU *et al.*, 2023 [10]; JIANG *et al.*, 2025 [11]] – to provide a useful representation of a person’s health, along the lines of the recent Gemma-4-12B LLM [EL ABD *et al.*, 2026 [12]], which has demonstrated the possibility of designing encoder-free mechanisms for raw integration of modalities such as audio and images. However, training such models is realistically out of reach for an academic team, and it runs counter to the idea of democratizing AI in healthcare.

I therefore think that a pragmatic way forward is to move toward collaborative, specialized agents, which are deployed at the hospital scale, and would share their individual expertise to jointly build a global representation of a patient’s health, which could be queried in natural language.

Overall, my project for the coming years will therefore revolve around the axes described in the following sections.

III.2.1 Text as the Pivot of Multimodality

Medical data is heterogeneous by nature – images, tabular records, physiological signals, free-text reports – and the collaborations of Section II.3 have confronted me with almost all of these modalities, each time with an ad-hoc model. I now want to exploit the ability of text to describe, and sometimes replace, other modalities [PAZ-ARGAMAN *et al.*, 2020 [13]; SUSHIL *et al.*, 2024 [14]], together with the emergent capabilities of LLMs [WEI *et al.*, 2022 [15]], so that text becomes the pivot around which the other modalities are aligned rather than one modality among others. Two properties matter more to me here than raw performance: the model must degrade gracefully when modalities are missing, which is the rule rather than the exception in clinical routine [LIU *et al.*, 2023 [16]] and for which solutions exist outside the medical domain [LIAQAT *et al.*, 2025 [17]]; and its architecture must accept a variable set of inputs, in the spirit of [JAEGLER *et al.*, 2021 [18]].

If the objective is a collaboration of specialized models rather than a single large one, the interface between them matters as much as their individual quality. Text is not the only interface that is at once machine-consumable and human-auditable – structured clinical terminologies such as SNOMED CT or FHIR resources are both – but it is the only one expressive enough to carry a hedged, qualified or partial finding, which is most often what a specialized model has to communicate: an orchestrator can query it, and a practitioner can read it and disagree

with it. This is the choice made by recent orchestrated medical agents [KLANG *et al.*, 2026 ; ZHOU *et al.*, 2026 ], and it is, in my opinion, what separates a black-box multimodal encoder from a system whose intermediate steps a clinician can inspect.

Two experiments would tell me whether the idea survives contact with data. The first one is a controlled comparison on a problem I already know: in Section II.3.1 tabular and imaging data was fused for the prognosis of ICH using ad-hoc encoders and a tabular foundation model [HOLLMANN *et al.*, 2022 ; S₁₃]; replacing the tabular branch with a textual description produced by a specialized model, and measuring what is lost, gives a direct estimate of the cost of the text bottleneck. The second one would require some development, to extend REVE [C₃₁] with text to make it produce a description of an EEG signal rather than an embedding. This second experiment is not in the plans yet, and the obstacle is the absence of an evaluation target rather than a modeling difficulty: there is no corpus pairing EEG with free-text descriptions at anything like the scale REVE was pretrained on, so neither training nor scoring a generated description is currently well defined. The clinical reports attached to diagnostic recordings are the natural material from which to build one – they do describe the signal, though in the vocabulary of a diagnosis rather than of a waveform – and assembling such a corpus would be a data effort of its own, in the spirit of the collection work that made REVE possible.

The obvious risk of the ideas introduced in this section is that a fluent description is also an unfaithful one, and that errors introduced by one expert propagate to every model that reads it. This is why this axis cannot be separated from the supervision one described in Section III.2.2.

III.2.2 Agentic AI & Trust

The short-term project presented in Section III.1.1.2 is a first step toward solutions to an important problem in applied AI: agents are being deployed in hospitals faster than the tools able to supervise them are being built. Beyond the regression detection, audit trail and drift monitoring mechanisms of that project, I also want to study trust as a quantity that circulates between agents and is revised over the history of their exchanges, rather than as a static property of a single model.

A possible formulation of this problem relies on GSP. A set of agents that exchange, evaluate and rely on one another can be modeled as a graph, and trust updates could be derived from exchanges with neighboring nodes, along the lines of message-passing systems defining GCNs [KIPF *et al.*, 2017 ]. In such an analogy, the aggregation weights should however not be dictated by the support graph, as they are in a graph convolution: what an agent grants to a neighbor depends on the history of their exchanges rather than on the mere existence of a link between them. This is precisely what GATs [VELIČKOVIĆ *et al.*, 2018 ] bring to graph-structured data, by computing the aggregation coefficients from the representations of the two nodes involved – which makes the attention coefficients of a GAT a natural candidate to carry trust, and its learning procedure a natural candidate to revise it over interactions. Moreover, if trust is modeled as a signal on a graph, smoothness could measure whether trust is locally coherent, and a change in the spectrum could describe a drift of the population rather than of an individual model. Whether these quantities are informative in practice is an open question, and would make a great playground for a PhD student.

Trust also has to be grounded in something more precise than the confidence an agent reports about itself. In the clinical setting of Section III.2.1, the question a practitioner asks is not only what the answer is, but whether it can be relied upon for this particular patient, which calls for a quantification of the uncertainty attached to it. Conformal prediction is, in my opinion, the neighboring literature that fits this need best [ANGELOPOULOS *et al.*, 2023 ]: given an off-the-shelf predictor and a small calibration set of labeled samples, it returns a set of candidate answers containing the correct one with a prescribed probability, under the sole assumption that calibration and test samples are exchangeable. Two properties make it particularly attractive here. First, it is model-agnostic and requires no retraining, which is exactly the regime of a hospital deploying frozen foundation models it cannot afford to train. Second, it produces an object an orchestrator can act upon: the size of a prediction set is

a comparable measure of how informative each specialized agent is on a given case, and abstention becomes a decision the system takes explicitly rather than a failure someone has to detect afterward.

The difficulty is that the guarantee is marginal, *i.e.*, it holds on average over the joint draw of the calibration set and of the test sample, whereas what a practitioner experiences is the coverage conditioned on the one calibration set that was actually used. This conditional coverage fluctuates around the nominal level according to a Beta distribution whose concentration is governed by the number of calibration samples [VOVK, 2012], and the usual rule of thumb of about a thousand of them [ANGELOPOULOS *et al.*, 2023] is out of reach when every label has to be produced by a clinician. This is where this axis meets the **ENDIVE** project: under a labeling budget, the calibration set becomes something one chooses rather than receives, and drawing it jointly with a fixed-size DPP [KULESZA *et al.*, 2011] would spend that budget on cases that repel one another. Such a draw keeps those cases exchangeable among themselves, but not exchangeable with the test sample, so it breaks the very symmetry the guarantee rests upon – exchangeability being weaker than independence is of no help here. The question is thus whether validity can be recovered by weighting the calibration scores by the sampling law, as is done under covariate shift [TIBSHIRANI *et al.*, 2019], and how many labels that would save. Pseudo-labeling would enlarge the set for free, but errors would corrupt a calibration score and run into the same obstacle.

Prior discussions with Rémi BARDENET (CNRS, CRISAL) have convinced me that this direction is worth investigating. Conformal prediction in the few-shot regime is an active subject [FISCH *et al.*, 2021; PARK *et al.*, 2023], but the interaction between diverse sampling and the validity of the resulting predictors seems to me largely unexplored, and would combine several of my topics of interest naturally.

III.2.3 Genomics as the Next Modality

Among the modalities I have not worked with yet, genomic data is the one I am interested in exploring first. Given the large size of a genome and the number of available datasets, foundation models for this modality are emerging [FENG *et al.*, 2025], thus producing new complementary representations that could enrich a patient’s record in a multimodal analysis.

This modality is also appealing for reasons that connect directly to my past and current work. A genome is a sequence over a four-letter alphabet, formally close to text and directly amenable to the kind of language models discussed throughout this manuscript [JI *et al.*, 2021; NGUYEN *et al.*, 2023]. But the way the field actually manipulates these sequences is through their factorization into k -mers, and through the assembly and indexing structures built on them – de Bruijn graphs foremost [COMPEAU *et al.*, 2011; MARCHET *et al.*, 2021] – so that a graph is introduced by the representation itself rather than imposed on the data afterward. A modality that is at once a sequence and a graph is an unusually good fit for the two halves of my methodological background, Section II.1 on one side and Section II.2.3 on the other.

The difficulty is of the same nature as the opportunity. A genome carries orders of magnitude more variables than a clinical cohort has patients, which is the regime where selection, regularization and representation learning stop being refinements and become the problem itself – the same tension I described for brain signals in Section II.2.4, one scale further. This is also what makes the modality worth the effort: a compact and task-relevant representation of a genome would add to a patient’s record information that no other exam provides.

I have had discussions with Anthony BOUREUX and Jérôme REBOUL (IHU of Montpellier), Camille MARCHET and Antoine LIMASSET (CNRS, University of Lille), and Daniel GAUTHERET (University of Paris-Saclay). These researchers – whose primary interest is in genomics – are increasingly interested in LLMs and their adaptation to DNA sequences. When the time comes to focus more on this modality, they would be great partners to collaborate with.

III.2.4 A Cross-Service Agentic Environment

Finally, I would like the opportunity to put the ideas developed above into practice. This could be done in partnership with the CHU of Brest, but also on open medical datasets for a proof of concept.

Concretely, such an environment would host one (or more) agent per service, each deployed where its data lives. An orchestrator would query them in natural language, assemble their answers, and assist the patient's general practitioner. The representation of a patient's health would not be a vector produced by a single model, but the trace of a conversation between specialized ones, which has the pleasant property of being readable: a practitioner can follow which expert was consulted, what it answered, and where the suggestion comes from.

This is a deliberate alternative to the federated approach used in STRATIF-AI [McMAHAN *et al.*, 2017 ; NGUYEN *et al.*, 2022 ]. Federated learning is a working answer to the impossibility of centralizing medical data, but it is heavy to deploy and it supposes that all partners agree on a common model. In the environment I describe, what circulates between hospitals is text rather than weights: it can be logged, inspected and audited after the fact, it does not require any architectural agreement, and each service remains free to update its own expert. This is an argument about auditability, and not about privacy: free clinical text is in fact harder to de-identify than a model update, re-identification from narrative content being a well-documented risk, and this is one of the reasons federated learning exists in the first place. The legal constraints therefore remain in full – everything here is subject to the GDPR, and a version deployed to assist actual clinical decisions would fall under the high-risk regime [AI ACT ] – but the object that leaves the hospital is at least one whose content can be examined, which is not true of a gradient.

A proof of concept does not require a hospital deployment to begin with. Question answering benchmarks and open medical datasets [JOHNSON *et al.*, 2023 ] can support a first version with a few artificial services, on which the questions that matter can already be studied: does the orchestrator degrade gracefully when an expert is unavailable (or trust issues prevent it from using its answers), and does its answer improve on a decision that the clinical baseline already makes [KRONES *et al.*, 2025 ]?

While I have no illusions that this environment will exist in its complete form within the next five years, I do believe that each of its pieces is a legitimate research subject, and that working on these pieces has a chance of contributing to what I believe are important societal challenges.

CHAPTER IV

CONCLUSION



I am not sure that a chapter entitled “*Conclusion*” is the best way to end a manuscript whose main goal is to open the discussion on the future of a research program, but “*Opening*” seemed even weirder. First of all, thank you for reading this manuscript, which I enjoyed putting together, and which brought back many memories.

Looking back at [Chapter II](#), the trajectory summarized in this manuscript may look like a succession of unrelated topics: GSP, FSL, diffusion LLMs, foundation models for brain signals, and a series of clinical collaborations in neurology, cardiology and oncology. I hope to have convinced the reader that there is in fact a cross-cutting question that underpins all of them: how to take advantage of the structure or variability of the data, rather than of its amount. GSP was one answer, built on the geometry of the domain on which signals live. FSL was another, built on the geometry of representations. Diversity, as explored in [ENDIVE](#), is a third one, built on what distinguishes the samples one selects. And the applied projects were, each time, great use cases.

This is also what makes me confident that the directions of [Chapter III](#) are worth the next few years. They contribute to a system I find both desirable and reachable: not a single very large multimodal model producing an opaque vector for a patient, but a collaboration of specialized models, each trained and deployed where its data lives, exchanging natural language descriptions that an orchestrator can assemble and a practitioner can read. Text as a pivot modality ([Section III.2.1](#)), trust and uncertainty as quantities circulating between agents ([Section III.2.2](#)), genomics as the next modality to bring in ([Section III.2.3](#)), and a cross-service environment to put all of this to the test ([Section III.2.4](#)). Each of these axes has its own open questions, and none is solved.

I am aware that writing a research program in AI in 2026 is a non-trivial exercise. The reorientation of [ENDIVE](#) described in [Section III.1.1.1](#) is a good reminder of it: a paradigm that was a research field of its own when the project was submitted had become, three years later, one regime of foundation model adaptation. I expect this to happen again, probably more than once, and possibly before this manuscript is defended. My answer to this is not to try to predict where the field goes or to pretend we can compete with big actors of AI – which nobody in a small public laboratory does anyway – but rather to identify a broad objective, which serves a valuable purpose, to which I will contribute as much as I can.

I will therefore conclude simply. The two halves of my research – staying close to the core of modern AI, and bringing what happens there to healthcare – have so far reinforced each other, and I intend to keep them coupled. If, in a few years, a general practitioner at the CHU of Brest can ask a question in plain language, obtain an answer assembled from several specialized models, see which one said what, and disagree with it, then the program described in [Chapter III](#) will have done its job. If not, the questions it raises along the way – on diversity, on calibration under a labeling budget, on trust between models – will have been worth asking anyway.

Thank you for reading.

APPENDICES

APPENDIX A

TEACHING ACTIVITIES



This manuscript being mostly about my research and projects, I have chosen to avoid mentioning my teaching activities too much in the main chapters. Still, I am very active on this side as well, and I like experimenting with approaches to engage students in their learning process. Therefore, this appendix exists to give an idea of my views and achievements on these activities.

I have chosen to focus on some courses that I particularly value, and specific pedagogy topics I find interesting to discuss. For an exhaustive list of my teaching activities and responsibilities, please refer to [Appendix B \(Section B.4\)](#). The complete list of internships and student projects I supervised is also given in [Section B.2.5](#) and subsequent sections.

These examples are complemented by a Q&A on specific topics that I find interesting to address. Overall, these two parts together should draw a good sketch of my teaching philosophy.

A.1	Focus on a Few Courses	92
A.1.1	PyRat	92
A.1.1.1	Context	92
A.1.1.2	Main Challenges	92
A.1.1.3	Beyond the First-Year Class Course	100
A.1.1.4	Perspectives	100
A.1.2	How to Lie with Statistics	101
A.1.2.1	Context	101
A.1.2.2	Main Challenges	101
A.1.3	DS Track	104
A.1.3.1	Context	104
A.1.3.2	Main Challenge	104
A.2	Q&A on Specific Topics	105

A.1 Focus on a Few Courses

A.1.1 PyRat

A.1.1.1 Context

When I was looking for a PhD position after completing my Master's degree in 2014, one of the main arguments that convinced me to join Télécom Bretagne (besides research topics, obviously) was the possibility of completely refactoring the Computer Science teaching unit for first-year engineering students.

At that time, there was a national reform of CPGE programs [[CPGE 2013 REFORM](#) , which introduced Python as the main programming language, among other changes. To prepare the arrival in 2015 of students who followed this new curriculum, Télécom Bretagne gave me the responsibility of designing a new course that would replace the existing first-year programming course in C language.

A.1.1.2 Main Challenges

Challenge 1 Engaging students

In the first year of general engineering school, many students enter with little to no interest in computer science. This was particularly true at that time, as most computer science courses in CPGE were given by maths or physics teachers, and programming was essentially reduced to using MATLAB (or equivalents) as an advanced calculator.

To address [Challenge 1](#), I relied on two main elements:

The PyRat Game

I chose to build the course around a video game that takes the form of a customizable maze, where characters have to grab pieces of cheese faster than an opponent. Cells in the maze may be separated by mud, requiring more than one turn to cross. To control the characters, students have to write programs (in Python to capitalize on previous skills) that describe their behavior on a turn-by-turn basis.

The customization options offered by PyRat allow users to easily create diverse learning situations, where students can study many standard algorithms (*e.g.*, traversals, exhaustive searches, heuristics), related notions (*e.g.*, complexity, approximation bounds, data structures), as well as programming paradigms (*e.g.*, recursion, OOP), using graph theory as a central concept.

With this approach, I wanted students to discover a more fun aspect of computer science than the use of a few lines of code for one-time problem solving, and therefore to spark more interest in the domain. Since its inception, the course has always been very well received by the students in general, and former students mostly remember me as the PyRat guy when I meet them later after they graduate (see [Figure 10](#)).

The software has gone through many iterations since its creation, to adapt to students' feedback (*e.g.*, usage simplification, improved graphics) or evolution of the context (*e.g.*, cross-campus matches, integration of OOP). [Figure 11](#) to [Figure 14](#) show the various versions of the GUI, with a short description of the main changes.



(a) Organization of a PyRat party in 2016. Students built a maze in the student common room and used the algorithms seen in class to reach the bar.



(b) Announcement of the end of semester PyRat tournament, with my face photoshopped on it (with authorization).

Figure 10 Some references to the PyRat course in the local CDV journal at IMT Atlantique.





Figure 11 The first version of PyRat was pretty rudimentary: a color per player, coins to grab, a few walls and fences, and that's it!



- (a) The main game interface. Coins became pieces of cheese, and red/green circles nice animated characters. Tiles were designed by Christina ROBERTS and used with authorization [NEYJOUR .




- (b) The PyRat logo and characters were made for the course by Lauren LEFUMEUR-PASDELOUP [NERUAL .



- (c) Each player could monitor their own output in a dedicated shell, for easier debugging when working locally.

Figure 12 The second version brought many changes. First of all, it introduced many graphical elements to improve the user experience. The core was also recoded, and the whole architecture changed to a client-server design to prepare cross-campus matches between Télécom Bretagne and Mines Nantes, which would merge as IMT Atlantique in 2017.




- (a) The third version of PyRat was programmed by Vincent GRIPON in 2018, with new graphics by Mooncat Studio [MOONCAT STUDIO 🌐].




- (b) The PyRat MOOC consisted of 19 videos covering all notions seen in class.


Figure 13 During my postdoc, a MOOC based on PyRat was developed, to open the course to a broader public [PYRAT MOOC 🌐]. It was discontinued in 2023 due to a lack of time to maintain it.



- (a) Later versions of PyRat came back to previous graphics. It is also now possible to add holes in the maze for more customization, and the interface now shows cell indices and mud value for an easier analysis of trajectories. PyRat now supports more than 2 players, and can be played in teams.



- (b) PyRat now integrates fallback modes in ASCII/ANSI to improve portability.



- (c) I recently added a maze editor, which allows students to design their own maze for PyRat, and import it into the game. It makes it easy to design specific configurations (e.g., highlighting the difference between a BFS and Dijkstra's algorithm).

Figure 14 Current version of the PyRat software, now registered on PyPI as `pyrat-game` [PyRat PyPI .

Emphasis on Hands-On Practical Activities

As I believe programming is best learned with a computer, I wanted to focus mostly on practical activities. I initially designed the lectures to be as short as possible (~ 20 minutes), with a comprehensive website to dig further into concepts between sessions. The rest of the time would be dedicated to hands-on practical sessions.

Later, I experimented with fully flipped classrooms – and even a MOOC [PyRat MOOC 🌐] – where students had to study articles before the class, which would start with a short quiz to check understanding of notions. I provided Mattermost or Discord servers for out-of-class support from the teachers, and in-class activities focused on leveraging algorithms to solve problems in Python.

Challenge 2 Handling heterogeneity

In addition to various appetites for computer science, first-year students also come from diverse backgrounds, with some already very familiar with a few programming languages (often OCaml or Python at that time), and others not at all. This challenge has become increasingly important over the years, as the population of students has only been growing.

To address [Challenge 2](#), I relied on three elements:

Increasingly Difficult Objectives

As the flexibility of the PyRat software allows easy customization of mazes, it can be used to produce multiple learning situations, of gradually increasing complexity. It also provides visual feedback to students, allowing them to see their progression.

To reach the final objective of the course – designing a strategy to face an opponent in a PyRat game – the course introduces intermediate sub-objectives:



Improve a random program

Starting from a program that moves uniformly at random among cardinal directions, students learn to take the maze map into account to avoid walls.

This preliminary objective allows them to familiarize themselves with the software, learn to manipulate objects, and identify parts of a graph.



Catch a single piece of cheese in an unweighted graph

This objective can be attained using simple traversal algorithms such as DFS or BFS [MOORE, 1959 📖].


This first objective allows students to discover algorithms on graphs and some of their properties (shortest path, complexity, termination), as well as code factorization.



Catch a single piece of cheese in a weighted graph

Introducing mud in the maze forces the use of more advanced algorithms – and associated data structures – like Dijkstra’s algorithm [DIJKSTRA, 1959 

This objective highlights that algorithms have pre-requisites, and shows the importance of the choice of data structures.



Catch multiple pieces of cheese in a minimum number of turns

Solving this problem requires testing all possible permutations of pieces of cheese using exhaustive algorithms and small optimizations of these such as backtracking and branch-and-bound.

Introducing the TSP allows students to study classes of complexity, and the existence of hard-to-solve problems.



Catch multiple pieces of cheese, but faster

As the previous objective is intractable for large quantities of cheese – as will be the case in the final objective – we introduce heuristics and greedy algorithms.

This objective shows how one can use approximate algorithms to estimate solutions to a complex problem, leading to a trade-off between tractability and quality.

The final objective, as mentioned earlier, is to face an opponent in a PyRat game using a strategy of the students’ choice. The semester ends with a friendly tournament, where volunteering students register their best program to try to win a trophy (see [Figure 15](#)) and become the final boss for the next year. It is generally a fun moment and a milestone in the first year (see videos referenced in [M₁; M₃]).

I have seen all kinds of submissions to the tournament since it was created, including deep learning approaches, ant colonies, genetic algorithms... and less orthodox ones such as opponent brainwashing and teleportation!



Figure 15 The PyRat trophy, made locally at the school's FabLab.

Optional Advanced Content

The objectives presented above constitute the core content that any student should understand. Still, for more advanced students, I have organized the course website so that, for each concept, there are two optional extensions:

- **Advanced concepts** — These can be for instance alternative algorithms for a specific problem, advanced mathematical results, or extra exercises. These additional materials are fully guided.
- **Links** — I end each course material with a list of references that students can study on their own. These references can be online tutorials or research papers, and generally serve as an entry point to new domains not seen in class (*e.g.*, spectral graph theory).

Extra content is easily identifiable so that interested students can choose to explore it. They can make the most of them in the various evaluations, but not doing so and mastering the basic content is sufficient to pass. To this extent, each class would start with a short quiz, with questions covering both necessary and advanced concepts. Questions are organized in pairs (Q_{basic} , Q_{advanced}) to reward the level of involvement of students. The system works as follows:

- All answers to basic questions need to be correct to pass.
- For a given pair, Q_{advanced} is counted in the grade only if both Q_{basic} and Q_{advanced} are correct.
- Any student who fails a question Q_{basic} has the possibility of coming back to the teacher and explaining their mistake. If the explanation is convincing, the teacher retro-validates Q_{basic} . In such a case, the associated Q_{advanced} is not counted though, even if it was correct.

This approach has the advantage of rewarding students who go beyond basic expectations, while ensuring fundamental notions are mastered. It also helps detect students who struggle more than others and assist them more during sessions.

Out-of-Class Support & Solidarity

Finally, for all courses I manage, I create a Mattermost or Discord server, to provide the students with a place to ask questions between classes. I believe this to be essential, especially in the context of partially or fully flipped classes, where students have a significant preparation to make at home.

Besides easier access to the teachers, these platforms let students help each other. Every year, a few students – generally with a prior computer science background – are very active, and play the role of teaching assistants.

Along these lines, I ask students to work in groups in the PyRat classes. They are required to discuss together and to form a triplet with approximately the same level in computer science. The goal is threefold:

- **Homogeneous teams** — This ensures – or at least improves the probability – that no student is left behind, or suffers from a larger quantity of work.
- **Heterogeneous classes** — Having homogeneous teams in a more diverse class encourages advanced students to help teams that struggle more when they are done with their own work.
- **Triplets** — Having three students per group gives them responsibility for specific aspects of the project. One student would be in charge of the code, another of the documentation, and the other of the unit tests. While the grade for a deliverable would mostly be common, it can be adapted based on individual responsibilities. Having three deliverables to produce (catch one cheese, catch multiple pieces of cheese, face an opponent) makes it possible to rotate responsibilities, so that each student is responsible for each aspect once.

A.1.1.3 Beyond the First-Year Class Course

The flexibility of the PyRat software also allowed its use in the [Introduction to AI](#) course for Master 1 & 2 students at IMT Atlantique, where it was used to teach RL. It also served as a support for the internship of [Guillaume BURET](#) (hacking the platform), the TIPE of [Romane CASTEL](#) (traversal algorithms) and the student project of [Quentin AVARE et al.](#) (prototyping a maze editor).

I was also invited to [ECAM Louis de Broglie](#) to present the course and train teachers who wanted to integrate the software in their teaching activities. Also, when Télécom Bretagne and Mines Nantes merged as IMT Atlantique in 2017, I actively participated in porting the course to the new program.

Finally, I made several communications on the course [[C₆](#); [S₁](#)] and the associated game [[W₁](#)], and open sourced all materials and versions [[M₁](#); [M₃](#)] for public usage.

PyRat was a standalone course from 2015 to 2024, and is now integrated as a project in a more global [Computer Science](#) teaching unit. This change occurred following the 2021 CPGE reform [[CPGE 2021 REFORM](#) , when students studying the new programs started at IMT Atlantique.

A.1.1.4 Perspectives

Although not currently used in the classes I give, I have written a few example programs that illustrate the use of solvers for finding solutions to the TSP. This could serve as a basis for teaching more advanced operational research concepts. Also, the PyRat software itself being designed following OOP paradigms, it can be used as an example to introduce software engineering and UML class diagrams.

I also want to work on a platform where PyRat matches would be made constantly, and produce a leaderboard that would show on the screens at IMT Atlantique. This would allow the most motivated students to continue having fun after the course is completed.

A.1.2 How to Lie with Statistics

A.1.2.1 Context

My teaching experience with the PyRat course during my PhD revealed that many students – even with a scientific background – are prone to biases. I observed that they often assumed a correlation between the amount of work invested in a project and the quality of the produced solution. In addition, when asking whether their program stood a chance of beating a greedy opponent in a PyRat game, the answer was often “*We played 2 or 3 games and won each time*”. I addressed the problem by introducing statistical tests as a tool for making confident and statistically significant statements. Still, with these observations I developed an interest in cognitive biases, and started to think about how to integrate a more scientific approach to problems in my courses.

Later came 2020 and all the misinformation spreading about COVID-19. I realized that the impact of biases on sound reasoning was a lot more prevalent than just a few assertions in a PyRat class, and thus I took the opportunity to build a teaching unit around biases at IMT Atlantique, named “*How to Lie with Statistics*” in reference to the book of the same name [HUFF, 1954 

Obviously, the course is not about lying with statistics, but aims at raising students’ awareness of the scientific method and how engineering skills and scientific research can help trust (or not) an assertion.

At the end of the module, students are awarded a diploma, which reminds them that anyone can say random things, which may look official or scientific (see Figure 16).

A.1.2.2 Main Challenges

Challenge 3 Making it interesting

I am obviously not the only one who worked on courses about biases and the scientific method (e.g., [HAUTEURS UGA 

I addressed Challenge 3 with two main axes:

Focus on Personalized Projects

For a course whose objective is to develop critical thinking, it seems natural to focus as much as possible on practical situations. Besides a short introduction to science, research and the scientific method, and a full day on biases (explained below), the course mostly consists of two projects:

- **έποχή (withholding of assent)** — In this first project, students have to choose a falsifiable question [POPPER, 1989 Is Earth flat?” are always very popular, but students often take the opportunity to explore more personal questions like “*Is coffee (or red wine) good for health?*”.



1. Ne le perdez pas !

2. Le présent diplôme* atteste d'une expertise strictement non-nulle en (liste non-exhaustive) :

- Méthodologie scientifique et pseudo-scientifique ;
- Simpsons ;
- Affaires judiciaires en environnement froid ;
- Synergologie®** ;
- (Enologie thérapeutique ;
- Néo-évhémérisme ;
- Platisme ;
- Pyramidisme ;
- Et. cæterisme...

Le présent diplôme fait également de vous un formateur pyramidal OFFICIEL en coaching Trans-Méthodologique®, lié par le sang à l'ITM® pour 18 générations.

Enfin, le présent diplôme fait également de vous un (une) Super-Héros™ (Super-Héroïne™), en totale liberté de pensée cosmique vers un nouvel âge réminiscent. Le grade de Maître du Monde™ (Maîtresse du Monde™) est toutefois débloable en suivant une formation secrète au tarif préférentiel de \$7259, payable uniquement en cryptomonnaies.

*. Au cas où ça ne serait pas assez évident***, ce diplôme est satirique et vise à parodier les "diplômes" attestant de "compétences" de certains pseudo-praticiens, délivrés par des organismes sans caution de validité scientifique.

** . Oui, il y a un vrai ® dans le nom de cette "science".

***. Et pour couvrir le Gourou Trans-Méthodologique®.

Figure 16 Fake ITM® certification delivered to all students who attend the course. This is made as a parody of expensive and useless pseudo-scientific training courses, and to remind students that the value of any diploma depends on the quality and recognition of the associated training.

This first project is for many students a first contact with a research article. Also, it illustrates the necessity to check (and cross) sources and to evaluate their seriousness.

- **Lie with statistics** — In the second project, students have to choose a question and a dataset, and manage to get an answer of their choice. They also have to make the proper analysis, without cheating. At the end of the project, students present the false analysis, and then exchange with other students, who try to debunk the presentation. Then the presenting students show the actual analysis and results.

This project aims at training the students' critical thinking, and experimentally illustrates how they can be manipulated in real-life situations. Also, when performing the actual analysis of the question, students can also make mistakes. This is a nice opportunity to show that producing sound results is not straightforward, and to mention the importance of involving domain experts.

Involve Actual Experts

As mentioned earlier, the course includes a full day on statistical and cognitive biases. For the former, I introduce statistical curiosities such as Simpson's paradox [SIMPSON, 1951] and application of Bayesian statistics to the Monty Hall problem [BARBEAU, 1993]. I am assisted by Romain BILLOT, who introduces statistical tests and p -values. The goal is to present tools that can help determine the soundness of an assumption, and highlight the necessity to involve actual statisticians in studies.

For the latter though, I cannot do it myself as it requires a very distinct background. To provide interesting content, I thus invite an actual psychologist, Marie-Claude VEILLEUX, to introduce how the human brain can deviate from rational thinking and lead to biases. Her talks are always praised by the students, who really value her expertise.

Challenge 4 Evaluating that course

In a context where students may choose to explore projects as diverse and personally engaging as *"Is there a gene that causes homosexuality?"* or *"Did we actually land on the Moon?"*, how can we give a grade that has no risk of being biased?

Short answer: I don't think we can... so I didn't. One of the first things I tell students at the beginning of the course is that they just passed the course. With this approach, I wanted to solve the following problems:

- **Risk of biases** — Even though as a teacher I keep my opinions to myself, I may unconsciously be biased in favor of presentations that address interests of mine (*i.e.*, coffee), or confirm my opinions (*i.e.*, that coffee is good). Indeed, it has been shown that even those with a knowledge of cognitive biases are still subject to them [PRONIN *et al.*, 2002].
- **Self-restriction** — In projects that are graded, choosing a topic is generally a trade-off between a personal interest in the topic and ambition, as a project too risky may lead to bad implications for the students' curriculum. In the context of this course, I want to remove this apprehension, and give the opportunity to students to go all in.

Overall, I believe that it makes very little sense to give grades to students at a Master's level⁽¹⁾. They should be treated as adults who are given time to learn new things and diversify their profile. In practice, since I created this course, it never happened that a student decided not to come because the course was essentially free. On the contrary, students have really appreciated being seen as volunteers, and many have confirmed that disabling the fear of a bad grade allowed them to explore topics they would have otherwise put aside.

⁽¹⁾ Note that I believe that grades are the problem, because of the pressure they create, not evaluations. However, these should be formative, with the objective of giving feedback and helping students get better at what they do.

A.1.3 DS Track

A.1.3.1 Context

It is no secret that AI is a very attractive term for students in engineering schools, given its prevalence in modern society, as well as the career perspectives it may open. IMT Atlantique offers 2 TAFs at Master's level (MCE & DS) on the topic, which are oversubscribed every year. In addition, some elective units exist to give students who could not follow those tracks the opportunity to study ML, or who just wanted a small exposure to AI in their curriculum to check the box.

One of these elective units is the [Introduction to AI](#) course, in which I regularly teach. Initially, the course was given once a year during the autumn, but demand increased to the point where it was run 3 times in 2024 – 2025. To face this increase in demand, I was asked to design a full [DS Track](#), which I decomposed into three units:

- **Tools for DS** — This module covers the essential tools that are necessary for a good approach to AI techniques. Its content is both technical (*e.g.*, setting up an IDE with good agentic practices) and fundamental (*e.g.*, algebra, statistics) to focus both on efficiency and on developing intuition.
- **AI** — This unit covers fundamentals of AI techniques, *i.e.*, unsupervised and supervised learning, representation learning and model adaptation. Emphasis is put on how to choose the appropriate approach, given a problem to solve and data availability.
- **Personalized project** — A track-wide project, which puts into practice all elements seen in the other units (more details below).

The track welcomes 10 to 20 Master 1 & 2 students, coming from diverse backgrounds, such as engineering schools or foreign universities. Students are assumed to have basic programming skills, and should have found an AI-related internship, starting right after the track ends.

A.1.3.2 Main Challenge

The track offers similar challenges to those described earlier, especially [Challenge 2](#) due to students' backgrounds. In addition to solutions similar to those already put in place in the PyRat course (*e.g.*, a Discord server for out-of-class support), I have chosen to focus here on a particular attempt at engaging students by taking into account their diversity.

Challenge 5 Matching everyone's interests

Students attending the track have very different objectives, although they all revolve around AI technologies. Some are interested in finance, others in medical applications, and others in 3D point clouds or weather forecasting. This diversity raises the challenge of whether we can propose a project that suits everyone's interests, and offer practical opportunities to start an internship efficiently.

To address [Challenge 5](#), I chose to leverage the future internships of the students, and to design a project made of general guidelines, instantiated per student, to work on data resembling what they would later manipulate. Following the same order as the courses of the AI module, the project consists of the following steps:

1. **Find a proxy dataset** — Because the data handled during the internship is not yet available, and probably overly complex, students are first asked to determine what they will work on (*e.g.*, modalities, missing data, quantities). With that in mind, they find a more polished open dataset online with similar characteristics. They work on this dataset during the module, as a proxy of the actual dataset they will later manipulate.
2. **Find clusters in data** — To illustrate the interest of unsupervised learning approaches, students are encouraged to explore clustering algorithms as well as dimension reduction techniques on their datasets.

3. **Predict a decision variable** — Similarly, students have to explore a classification or regression task on their dataset, using basic ML algorithms.
4. **From data to embeddings** — After that, we transition to neural networks and models based on them. The goal is to show that models trained to provide a good representation of the chosen modality can be used off-the-shelf jointly with a classifier/regressor to improve performance thanks to relevant embeddings.
5. **Adapt pretrained models** — Finally, the last objective of the project is to take one of these pretrained models and to adapt it to their dataset, either with fine-tuning or PEFT [HAN *et al.*, 2024 

For the project to succeed, a major difficulty was to have a basic knowledge of multiple modalities. While I could assist students with the main modalities such as text, vision or tabular data, involving multiple teachers in the pedagogical team helped mitigate this difficulty, as some conduct their research in medical imaging, others in EEG signals, or in audio time series.

Overall, students have appreciated the personalization of the project, and felt more confident about starting their internship. Their feedback after the internship will be invaluable for improving the module next year.

A.2 Q&A on Specific Topics

While I feel the previous section should draw a good sketch of my teaching philosophy, there are still a few pedagogy-related questions that do not fit there and that I find interesting to address below.

Question 1 AI & teaching

How to teach efficiently in a world where all students can simply use generative AI to solve problems in less time than it takes to acquire the necessary knowledge?

It is quite obvious that we cannot just pretend generative AI is not a thing students will use in their curriculum, especially in a context where their grades determine parts of their scholarship, such as the possibility of doing a semester exchange at a foreign university. Conversely, fully embracing AI as a standard tool such as a computer or the Internet erodes the fundamentals. As a result, students may fail to master basic knowledge required to use those tools efficiently, let alone build intuition and understand the core concepts.

I think the correct spot is somewhere in between those two extremes. On the one hand, it makes no sense to be an engineer who cannot use AI tools after graduating, but on the other hand, missing a systemic vision of how to address and decompose problems makes usage of such tools sub-optimal.

In a context without grades, my view would be to give students responsibility, and to tell them the argument above, insisting on the fact that they need to master a high-level understanding of a topic to be able to efficiently delegate code production to automatic tools. They would then have three years available to learn what teachers consider fundamental, and apply those concepts using AI tools when it makes sense to do so. However, as mentioned above, grades put pressure on students, and because of that I do not think this is a realistic view. My experience as a teacher in first-year [Computer Science](#) courses has shown that for numerous students, passing the course prevails over understanding the concepts.

More realistically – and while I still think that raising students’ awareness of the need for fundamental skills for a proper use of AI needs to be done – I believe that we should refocus on the development of high-level skills, and accompany students in mastering agentic tools. In more detail, here are some actions I have put in place:

A Course on Agentic AI Integration in the IDE

Recent experience has shown me that many students use AI badly. In the PyRat project of the [Computer Science](#) course, some have been producing the documentation, code and tests independently, hoping that all three models would come up with something coherent, without checking the produced outputs. Others have just tried to prompt ChatGPT [[CHATGPT](#) Give me a Dijkstra implementation for PyRat”. I have also observed that students were not aware of the models they were using, assuming that GitHub Copilot [[GITHUB COPILOT](#) 

I believe that these misuses of AI tools are easy to solve by integrating a small hands-on practical session early in the curriculum. Disambiguating the terms, insisting on the importance of context, and experimenting to discover the power and limits of these tools in a controlled environment should mostly be sufficient to orient students toward a better use of AI. Studies have also been conducted – in particular by some major AI companies [[SHEN et al., 2026](#) ; [DEEPMIND BLOG](#) 

For first-year students in [Computer Science](#), I have set up a practical session for this purpose. A similar course is part of the Tools for DS module of the [DS Track](#). I also help students set up a free [[GITHUB EDUCATION](#) 

Emphasis on High-Level Skills

The course on agentic AI described above aims at correcting mistakes in the basic usage of AI for development. However, it is not sufficient by itself. To be able to use models properly, it is important to understand what an agent can automate efficiently, and what it cannot. This is obviously subject to change as models improve, but at the current time, I think a correct delimitation is code production versus software design. While the former can easily be delegated to an AI, the produced output is generally messy or inefficient without a proper well-thought-out project organization.

To put an emphasis on high-level skills, I have chosen to transform my teaching activities in three ways:

- **Focus on code structure** — Structuring a program is a good way to learn to decompose a problem into sub-problems and to keep things organized. Putting an emphasis on this rather than on code production develops the fundamental skills that allow students to keep their projects under control.
- **Insist on unit testing** — Code generated by an agent can be seen in some ways as a library one would use without checking the code details. In both cases, errors can happen, and it is the responsibility of the project developer to make sure that this external code does what it is supposed to in the context of the project where it is used. I have thus chosen to put more emphasis on unit tests, and on the necessity to validate solutions to sub-problems identified previously.
- **Deliverable production** — PyRat is for most students the first computer science project they are confronted with. Therefore, it provides a good context to make first experiments with AI tools in a production-like context. However, as mentioned earlier, improper use of AI can lead to a deliverable that is not coherent, both in form and functionality. I have thus integrated elements into the first-year courses to teach students what properties a deliverable should respect, by taking the role of a client with specific requirements.

A Two-Step Approach to Master both Fundamentals and Tools

Finally, to insist on the message that both knowledge and efficiency are important, I have chosen to conduct all programming practical classes in two consecutive phases:

1. **Understand concepts** — After indicating to the students that we have specifically designed problems to develop their understanding of what we consider an important problem, we ask them to perform the exercises with their agents turned off, in order to develop a good fundamental understanding.
2. **Be efficient** — Then, after this series of exercises, we ask students to redo them, but putting the focus on efficiency and code generation. The goal is to learn to master agent usage, and to exploit the knowledge

obtained in the first step to do it more efficiently. In addition, this second step can also serve as a way to discover better solutions compared to those manually developed initially.

Question 2 Flipped classes & MOOCs

What is my view on alternative teaching approaches after several years of practice?

Interestingly, I have observed that alternative teaching methods like flipped classes or MOOCs seem to attract less interest over time. While the approach was pretty original 10+ years ago, it has become very common nowadays, and informal discussions with the students reveal that having too many flipped classes develops a perception of laziness on the part of the teachers. This feeling is exacerbated by the rise of very powerful AI agents that take the role of the teacher outside classes, even in the presence of platforms allowing support at any time.

With that in mind, students seem to value having short in-class lectures a lot, as high-level understanding of a domain is what a teacher still provides in a world where AI agents can take care of small implementation details. I therefore believe that my initial approach when I built the course – with short high-level courses and an emphasis on practical activities – is what we should lean toward nowadays.

In addition, I have been disappointed by the lack of flexibility in the MOOC model. Indeed, if an error or a dated choice is captured in a video, it is very hard to update it. This requires a lot of time to process, and in some cases it is not even feasible. A concrete example is the PyRat project, which received funding to develop the videos used in the course. These videos integrate complex animations, which I am not able to reproduce. Therefore, redoing a video would clash with the rest of the course and be perceived as a patch. On the contrary, the websites I developed for the [Computer Science](#) course and [DS Track](#) provide a very flexible way of keeping content up to date.

I do not want to paint a bad portrait of alternative teaching approaches though, as I feel they play an important role in engaging students. In addition to flipped classes and MOOCs, I regularly use interactive quizzes, which I find very valuable to check global understanding of a concept, and to start conversations on misunderstood aspects. Still, I think it is important to question these choices regularly, as ideas that were good in a context may be detrimental in a new one several years later.

Question 3 Making the most of internships & projects

Internships and student projects are a key part of any curriculum. How can we help students make the most of their experience?

When I supervise a project, one of my first discussions with students is what they want to do after their studies, especially whether they lean toward an industrial career, or if they want to take the opportunity to have a more research-oriented experience. Based on their choices, I try to orient the project toward a deliverable that they can showcase in their CV.

For the former category, I generally insist on the need to produce a tool to demonstrate concrete results in addition to standard project requirements as asked by the school. For the latter, when possible, I ask students to publish the results of their project in a research article, and to submit it to an actual conference. If the article is accepted, I offer the authors the chance to attend the conference to present it, covering all costs.

This approach has allowed some students vaguely interested in research to have a first experience in that domain, and sometimes a publication before the end of their studies. Among the research-oriented projects I supervised, I would like to highlight the following three, of which I am very proud:

- “*Study of the distribution of cells in a tumor*” (Yassine EL OUAHIDI *et al.*) — This Master 2 project led to a publication at the EGC conference [C₁₂]. After this project, Matis FELLER started a PhD at ENS Paris; Yassine EL OUAHIDI did a PhD under my supervision and is now Lead Applied Scientist at Mistral AI.
- “*Adaptation of segmentation models*” (Clément BÉLIVEAU *et al.*) — This Master 1 & 2 project led to a publication at the EUSIPCO conference [C₂₈]. After this project, Jules DECAESTECKER entered the MVA Master with the goal of later pursuing a career in research; both Clément BÉLIVEAU and Frederic LIN started a PhD, respectively with Naval Group and with CEA-LIST; finally, Jonathan LYS started a PhD under my supervision, and will join Yassine EL OUAHIDI’s team at Mistral AI for a 6-month internship from September of 2026.
- “*ML for cuneiform writing recognition*” (Quentin ANDRÉ *et al.*) — While this Master 1 & 2 project did not lead to any publication, it gave the opportunity to engineering students passionate about archaeology – who did not pursue a career in this domain due to a lack of job prospects – to explore their topic of interest. After the project, Maëlle GABENS pursued a Research Master at Keio University; Quentin ANDRÉ is going to start a PhD this year with Sorbonne University on a topic at the crossroads of AI and archaeology.

APPENDIX B

CURRICULUM VITAE



B.1	General Information	110
B.1.1	Personal Details	110
B.1.2	Higher Education	110
B.1.3	Academic Positions	110
B.1.4	Details on My PhD	111
B.2	Supervisions	111
B.2.1	PhD Students	111
B.2.2	MD Students	113
B.2.3	Postdoctoral Fellows	114
B.2.4	Research Engineers	115
B.2.5	Interns	115
B.2.6	Academic Student Projects	120
B.2.7	Industrial Student Projects	127
B.2.8	Tutored Students	128
B.3	Funded Projects	130
B.3.1	Research Projects I Wrote or Co-Wrote	130
B.3.2	Research Projects I Contribute(d) To	133
B.3.3	Other Funding	135
B.4	Teaching Experience	136
B.4.1	Responsibilities	136
B.4.2	In-Class Teaching Activities	137
B.4.3	Certifications	141
B.4.4	Organization of Programming Challenges	142
B.5	Other Activities	143
B.5.1	Research Stays	143
B.5.2	Invited Talks	143
B.5.3	Organization of Scientific Conferences & Symposiums	145
B.5.4	Scientific Dissemination	146
B.5.5	Reviewing Activities	148
B.5.5.1	Research Articles	148
B.5.5.2	Research Projects	150

B.1 General Information

B.1.1 Personal Details

Name	Bastien PASDELOUP
Born	October 3 rd , 1988
Nationality	French
Links	    
Email	bastien.pasdeloup@imt-atlantique.fr
Phone	02.29.00.15.65 / 06.32.59.01.18

B.1.2 Higher Education

2007	Baccalaureate “ <i>Economics & Social Sciences</i> ” Charles de Foucauld (Brest, France)
2008 – 2010	BTS 1 & 2 “ <i>Computer Science for Management</i> ” Charles de Foucauld (Brest, France) Internships: • ISEN (Brest, France) [I ₁] • JiVé Création, CERV (Brest, France) [I ₂]
2010 – 2011	Bachelor 3 “ <i>Engineering Track for Computer Scientists</i> ” ENIB (Brest, France) Internship: Diateam (Brest, France) [I ₃]
2011 – 2014	Bachelor 3, Master 1 & 2 “ <i>Research & Innovation</i> ” ENS Rennes & University of Rennes 1 (Rennes, France) Internships: • DGA/MI (Bruz, France) [I ₄] • CISTER (Porto, Portugal) [I ₅] • Irisa (Rennes, France) [I ₆]

B.1.3 Academic Positions

2014 – 2017	PhD Laboratory: Lab-STICC, UMR CNRS 6285 School: Télécom Bretagne / IMT Atlantique (Brest, France)
2018 – 2019	Postdoctorate Laboratory: LTS4 School: EPFL (Lausanne, Switzerland)
Since 2019	Associate Professor Laboratory: Lab-STICC, UMR CNRS 6285 School: IMT Atlantique (Brest, France)

B.1.4 Details on My PhD

Title	“Extending Convolutional Neural Networks to Irregular Domains Through Graph Inference”
Defended	December 12 th , 2017
Direction	Dominique PASTOR, IMT Atlantique (director) Vincent GRIPON, IMT Atlantique (supervisor)
Jury	Christine GUILLEMOT, Inria Rennes (president) Pierre BORGNAT, ENS Lyon (reviewer) Pascal FROSSARD, EPFL (reviewer) Rémi GRIBONVAL, Inria Rennes (examiner) Michael RABBAT, McGill University (invited)
Manuscript	[T ₁]

B.2 Supervisions

B.2.1 PhD Students

As Associate Professor	Maryam KARIMI-MAMAGHAN   	
	Period	2019 – 2022 (36 months)
	Defended	July 20 th , 2022
	Funding	Lab-STICC, DECIDE (personal funds)
	Subject	Hybridization of metaheuristics and ML for combinatorial optimization
	Manuscript	[T ₃]
	Publications	[J ₃ ; C ₁₀ ; C ₁₅ ; C ₁₆ ; W ₇ ; W ₉]
⋮	Fred Michael GONSALVES   	
	Period	2020 – 2023 (36 months)
	Defended	November 29 th , 2023
	Funding	Chantiers de l’Atlantique (CIFRE)
	Subject	ML and operational research for cruise ship energy efficiency
	Manuscript	[T ₅]
	Publications	[C ₁₃ ; P ₃]
⋮	Yassir BENDOU   	
	Period	2021 – 2024 (36 months)
	Defended	December 2 nd , 2024
	Funding	Fondation Mines-Télécom & Institut Carnot (“Futur & Ruptures” grant)
	Subject	Few-shot image classification

Manuscript	[T ₆]
Publications	[J ₂ ; C ₁₇ ; C ₂₂ ; C ₂₃ ; W ₁₁ ; P ₄]

Yassine EL OUAHIDI

Period	2021 – 2024 (36 months)
Defended	March 20 th , 2025
Funding	Région Bretagne (“ARED” grant) IMT & ANR (“AI@IMT” grant)
Subject	Vision-inspired deep learning approaches for BCI applications
Manuscript	[T ₇]
Publications	[J ₄ ; J ₅ ; J ₁₀ ; C ₁₈ ; C ₁₉ ; C ₂₄ ; C ₂₅ ; C ₂₇ ; C ₃₁ ; W ₈ ; W ₁₂ ; W ₁₃ ; W ₁₆]

Dubon RODRIGUE

Period	2022 – 2025 (36 months)
Defended	October 23 rd , 2025
Subject	Simplification of energy network models through AI
Funding	IMT Atlantique (CDE) IMT & ANR (“AI@IMT” grant)
Manuscript	[T ₈]
Publications	[J ₇ ; C ₂₉ ; C ₃₃ ; W ₁₅ ; J ₁₁]

Arsalan NAJIB

Period	2023 – 2026 (36 months, abandoned after 11 months)
Funding	CHU Rouen (salary)
Subject	Classification of ocular diseases using deep learning approaches
Abandon	Arsalan did not manage to obtain a reduction in working hours at CHU. This and family reasons led him to choose to abandon his PhD to focus on his work at the hospital

Yahia BELLOUCHE

Period	2024 – 2026 (24 months, ongoing)
Funding	CHU Brest (salary)
Subject	Computational fluid dynamics in cardiovascular medicine: from application to a clinical perspective
Publications	[J ₈ ; J ₉]




Jonathan Lys   

Period	2024 – 2027 (36 months, ongoing)
Funding	ANR (“JCJC” grant)
Subject	Exploring diversity in FSL and diffusion LLMs
Publications	[C ₃₁ ; C ₃₄ ; C ₃₅ ; C ₃₆ ; R ₃]




Joshua Chaz RIVERA   

Period	2025 – 2028 (36 months, ongoing)
Funding	IMT Atlantique (CDE) SequoIA (“ <i>Doctoral Projects</i> ” grant)
Subject	Diffusion LLMs for unusual data
Publications	[R ₁ ; R ₂]



Nolan SISOUPHANTHONG-VONGKOTRATTANA   

Period	2025 – 2028 (36 months, ongoing)
Funding	IMT Atlantique (CDE) PRACom
Subject	Diffusion LLMs & Mechanistic interpretability



Quentin BANET   

Period	2026 – 2029 (36 months, starting soon)
Funding	Orange (CIFRE)
Subject	Neural coding for audio signals

B.2.2 MD Students



Victor QUÉRÉ   

Period	2020 – 2021 (12 months)
Defended	October 13 th , 2021
Funding	CHU Brest (salary)
Subject	Comparison of ML approaches for prediction of ICH early mortality
Manuscript	[T ₂]

Amine EL OUAHIDI   

Period	2022 – 2023 (12 months)
Defended	September 9 th , 2023
Subject	Use of ML to predict pacemaker implantation following a TAVI procedure
Funding	CHU Brest (salary)
Manuscript	[T ₄]
Publications	[J ₅ ; W ₈]

B.2.3 Postdoctoral Fellows

As Associate Professor

Mounia HAMIDOUCHE   

Period	2020 (9 months)
Funding	Carnot TSN (“ <i>Numérique 2020</i> ” grant)
Subject	Using GSP to improve few-shot classifiers
Publications	[B ₂ ; C ₁₄]

Behnam VAHDANI   

Period	2020 (10 months)
Funding	Carnot TSN (“ <i>Numérique 2020</i> ” grant)
Subject	Radio-frequency design via 3D printing

Hafsa EL HAFYANI   

Period	2022 – 2023 (18 months)
Funding	ANR & InPixal (“ <i>ASTRID</i> ” grant)
Subject	Multimodal approaches for image segmentation
Publications	[P ₅]

Amine SADIKINE   

Period	2025 – 2027 (18 months, ongoing)
Funding	Horizon Europe (“ <i>RLA</i> ” grant)
Subject	AI for stroke recovery analysis
Publications	[S ₁₃]

B.2.4 Research Engineers

As Associate Professor

Mervyn GUILLOU   

Period	2021 – 2023 (24 months)
Funding	ANR & Elliptika (“ <i>Plan Relance</i> ” grant)
Subject	Codesign of hyperfrequency components using AI

...

Manon RENAULT   

Period	2023 – 2025 (18 months)
Funding	ADEME & Goodflood (“ <i>Logistique 4.0</i> ” grant)
Subject	Prediction from sensor data using embedded deep learning models
Publications	[W ₁₇]

B.2.5 Interns

As PhD Student

Chiraz NAFOUKI   

Period	2014 – 2015 (1 month)
Curriculum	Télécom Bretagne (Bachelor 3)
Subject	Visualization of the diffusion of a signal on a graph

...

Khalid AITELAMINE   

Period	2014 – 2015 (2 months)
Curriculum	Télécom Bretagne (Master 1)
Subject	Graph signal processing for classification
Report	[I ₇]

...

Maëlys ROBERT   

Period	2014 – 2015 (2 months)
Curriculum	Télécom Bretagne (Master 1)
Subject	Graph signal processing

Axel KERINEC

Period	2015 – 2016 (3 months)
Curriculum	ENS Lyon (Bachelor 3)
Subject	Spatio-temporal detection of patterns in multivariate signals
Report	[I ₁₅]

Loïc CARR

Period	2015 – 2016 (1 month)
Curriculum	Télécom Bretagne (Bachelor 3)
Subject	ML for object recognition in images

Guillaume BURET

Period	2015 – 2016 (2 months)
Curriculum	Télécom Bretagne (Bachelor 3)
Subject	Hacking the PyRat platform

Martin DORNIER

Period	2015 – 2016 (1 month)
Curriculum	Télécom Bretagne (Bachelor 3)
Subject	BCI systems for hand movement classification

Ladislav DUCERF

Period	2015 – 2016 (1 month)
Curriculum	Télécom Bretagne (Bachelor 3)
Subject	ML and GSP

Houda ABICHO

Period	2016 – 2017 (1 month)
Curriculum	Télécom Bretagne (Bachelor 3)
Subject	Spectral graph theory

As Postdoc

Wenqi SHU-QUARTIER-DIT-MAIRE   

Period	2017 – 2018 (3 months)
Curriculum	ENS Paris-Saclay (Master 1)
Funding	ENS Paris-Saclay (salary)
Subject	Classification of cancer types from medical stained images
Report	[L ₁₈]

..

Ali DADRAS   

Period	2018 – 2019 (6 months)
Curriculum	Tabriz University (Master 2)
Funding	LTS4, EPFL (personal funds)
Subject	Classification of patterns in time series using GSP
Report	[L ₁₇]
Publications	[W ₅]

As Associate Professor

Yassine EL OUAHIDI   

Period	2019 – 2020 (1 month)
Curriculum	IMT Atlantique (Master 2)
Subject	Study of the distribution of cells in a tumor

..

Matis FELLER   

Period	2019 – 2020 (5 months)
Curriculum	IMT Atlantique (Master 2)
Funding	Lab-STICC, DECIDE (personal funds)
Subject	Study of the distribution of cells in a tumor

..

Clément RIBES   

Period	2019 – 2020 (2 months)
Curriculum	IMT Atlantique (Master 1)
Subject	Analysis of spatio-temporal signals with GSP

Victor QUÉRÉ   

Period	2019 – 2020 (6 months)
Curriculum	UBO (Master 2)
Funding	ARS & CHU Brest (“ <i>Année Recherche</i> ” grant)
Subject	Comparison of ML approaches for prediction of ICH early mortality
Report	[I ₂₆]

Gaspard CANEVET   

Period	2020 – 2021 (2 months)
Curriculum	IMT Atlantique (Master 1)
Subject	Integrating RL in metaheuristics

Marc SERRE   

Period	2020 – 2021 (2 months)
Curriculum	IMT Atlantique (Bachelor 3)
Subject	Discovery of adversarial classes in few-shot problems
Report	[I ₂₉]

Amine EL OUAHIDI   

Period	2021 – 2022 (6 months)
Curriculum	UBO (Master 2)
Funding	ARS & CHU Brest (“ <i>Année Recherche</i> ” grant)
Subject	Use of ML to predict post-TAVI risk
Report	[I ₃₀]

Maëlle GABENS   

Period	2021 – 2022 (2 months)
Curriculum	IMT Atlantique (Master 1)
Subject	Application of pruning to graph signal sampling
Report	[I ₃₁]

Nathan DARJANA   

Period	2021 – 2022 (5 months)
Curriculum	IMT Atlantique (Master 1)
Funding	Lab-STICC, DECIDE (personal funds)
Subject	3D segmentation of CT scanner heart imaging using CNNs

Zakaria EL OUAHIDI   

Period	2022 – 2023 (1 month)
Curriculum	UBO (Bachelor 2)
Subject	Real-time decoding of EEG signals for BCI using ML

Frederic LIN   

Period	2022 – 2023 (2 months)
Curriculum	IMT Atlantique (Master 1)
Subject	Extension of a segmentation algorithm using SAM

Jonathan LYS   

Period	2022 – 2023 (2 months)
Curriculum	IMT Atlantique (Master 2)
Funding	ANR & InPixal (“ASTRID” grant)
Subject	Adaptation of segmentation models

Anatole LEHÉ   

Period	2023 – 2024 (5 months)
Curriculum	IMT Atlantique (Master 1)
Funding	BRAIN, Lab-STICC (personal funds)
Subject	Integration of a DPP sampling procedure in a genetic algorithm

Rayan DAKHIA   

Period	2026 – 2027 (6 months, starting soon)
Curriculum	IMT Atlantique (Master 2)
Funding	France 2030 & Cerevial (“In the Box” grant)
Subject	Supervision of trust in agentic AI systems applied to the medical domain

B.2.6 Academic Student Projects

As PhD Student

Karim BOUALILA   

Alexandre DELAROQUE   

Natacha DESERT   

Julien MATTEI   

Period	2014 – 2015 (5 months)
Curriculum	Télécom Bretagne (Bachelor 3)
Type	Development project
Subject	Improvement of the Enigmatic platform
Report	[I ₈]

Jinglei SHI   

Yihong XU   

Period	2015 – 2016 (3 months)
Curriculum	Télécom Bretagne (Master 1)
Type	Research track project
Subject	Construction of the graph structure from graph signals
Report	[I ₉]

Nicolas GRELIER   

Period	2015 – 2016 (5 months)
Curriculum	Télécom Bretagne (Bachelor 3)
Type	Semester project
Subject	Identification of translation in irregular graphs
Report	[I ₁₃]
Publications	[C ₄ ; C ₇ ; P ₁]

Réda ALAMI   

Period	2015 – 2016 (5 months)
Curriculum	Télécom Bretagne (Master 2)
Type	Semester project
Subject	Uncertainty principle on graphs
Report	[I ₁₀]
Publications	[B ₁ ; C ₂]

Amine ECHRAIBI   

Lotfi MOUBAKIR   

Louis Gildas NDZAMBA   

Period	2015 – 2016 (5 months)
Curriculum	Télécom Bretagne (Bachelor 3)
Type	Development project
Subject	Software execution in controlled environment
Report	[I ₁₂]

Ladislav DUCERF   

Sili Zuo   

Period	2016 – 2017 (3 months)
Curriculum	Télécom Bretagne (Master 1)
Type	Research track project
Subject	Visualization and classification of signals through dimension reduction
Report	[I ₁₁]

Nicolas GRELIER   

Period	2016 – 2017 (3 months)
Curriculum	Télécom Bretagne (Master 1)
Type	Research track project
Subject	Quantum random walks
Report	[I ₁₄]

Romane CASTEL   

Period	2016 – 2017 (5 months)
Curriculum	CPGE Kerichen (Bachelor 2)
Type	TIPE
Subject	Traversal algorithms in the PyRat game

Maven ZANOLI   

Period	2017 – 2018 (5 months)
Curriculum	EPFL (Master 1)
Type	Semester project

Subject | Discovery of meaningful features for classification of tumor cells
Report | [I₁₉]

Murtaza Eren AKBIYIK   

Period | 2018 – 2019 (5 months)
Curriculum | EPFL (Bachelor 3)
Type | Semester project
Subject | Interpretable ML in personalized medicine
Report | [I₁₆]

Firas BEN OTHMAN   

Period | 2018 – 2019 (5 months)
Curriculum | EPFL (Master 1)
Type | Semester project
Subject | Improving classification performance of CNNs by changing the kernel shape
Report | [I₂₀]

Bassel BELHADJ   

Period | 2018 – 2019 (5 months)
Curriculum | EPFL (Master 1)
Type | Semester project
Subject | Graph convolutional networks

Frédéric BOUDJEMAA   

Albane COLIN DE VERDIÈRE   

Joséphine SCHMUTZ   

Period | 2019 – 2020 (4 months)
Curriculum | IMT Atlantique (Master 1 & 2)
Type | Research track project
Subject | Analysis of weather data using graph spectrograms
Report | [I₂₁]

Quentin AVARE   

Théo LE GALL   

Aurélien PAGEAU   

Period	2019 – 2020 (4 months)
Curriculum	IMT Atlantique (Bachelor 3)
Type	CODEVSI project
Subject	Design of a maze editor for the PyRat game
Report	[I ₂₂]

Thibaut ELOY   

Anthony FRION   

Period	2020 – 2021 (4 months)
Curriculum	IMT Atlantique (Master 2)
Type	Research track project
Subject	Graph spectrograms and the minimum bandwidth sum problem
Report	[I ₂₅]

Yassine EL OUAHIDI   

Matis FELLER   

Period	2020 – 2021 (6 months)
Curriculum	IMT Atlantique (Master 2)
Type	PROCOM project
Subject	Study of the distribution of cells in a tumor
Report	[I ₂₄]
Publications	[C ₁₂]

Ilias AMAL   

Diego Andres CARREÑO   

Jordan MEJIA   

Tatiana MORENO PEREA   

Period	2020 – 2021 (6 months)
Curriculum	IMT Atlantique (Master 2)
Type	PROCOM project
Subject	Generalization of the joint Fourier transform with sensor data
Report	[I ₂₇]

Jack BOURDIN WHITE  

Gabriel GROS   

Period	2020 – 2021 (4 months)
Curriculum	IMT Atlantique (Bachelor 3)
Type	CODEVSI research project
Subject	Semi-supervised classification
Report	[I ₂₃]

Quentin ANDRÉ   

Maëlle GABENS   

Period	2021 – 2022 (4 months)
Curriculum	IMT Atlantique (Master 1 & 2)
Type	Research track project
Subject	ML for cuneiform writing recognition
Report	[I ₂₈]

Kadriye Nur BAKIRCI   

Miguel GUZMAN   

Martin ROUESNÉ   

Hélène SAVATIER-DUPRÉ   

Period	2022 – 2023 (6 months)
Curriculum	IMT Atlantique (Master 1 & 2)
Type	PROCOM project
Subject	DPPs for MCDA
Publications	[C ₂₆]

Mouhamed LEYE   

Guillaume LODÉ   

Thomas MONTAGNE   

Period	2022 – 2023 (4 months)
Curriculum	IMT Atlantique (Bachelor 3)
Type	CODEVSI project
Subject	Design of a portable guitar amp
Report	[I ₃₂]

⋮	Paul DEVOUGE-BOYER   	
	Astrid FLICHY   	
	Lucas POIRIER   	
	Pauline THIBAUT   	
	Period	2023 – 2024 (4 months)
⋮	Curriculum	IMT Atlantique (Bachelor 3)
	Type	CODEVSI project
	Subject	A website for borrowing games from the Trolls club at IMT Atlantique
	Report	[I ₃₄]
	Yanis ADJA   	
⋮	Florent BERGE   	
	Aeddan WILLIAMS   	
	Period	2023 – 2024 (4 months)
	Curriculum	IMT Atlantique (Bachelor 3)
	Type	CODEVSI project
⋮	Subject	Design of a portable guitar amp (follow-up)
	Report	[I ₃₃]
	Clément BÉLIVEAU   	
	Jules DECAESTECKER   	
	Frederic LIN   	
⋮	Jonathan LYS   	
	Period	2023 – 2024 (6 months)
	Curriculum	IMT Atlantique (Master 1 & 2)
	Type	PROCOM project
	Subject	Adaptation of segmentation models
⋮	Publications	[C ₂₈]
	Marine DELABAERE   	
	Léa MIQUEU   	
	Michael MORENO   	
	Period	2024 – 2025 (6 months)
⋮	Curriculum	IMT Atlantique (Master 1 & 2)
	Type	PROCOM project

Subject

Publications

Automatic generation of music mashups
[C₃₀]

Mohamed-Ali BOULLOUZ



Axel JONARD



Brahim ONASSER



Aymane OUNZAR



Period

Curriculum

Type

Subject

Report

2024 – 2025 (4 months)
IMT Atlantique (Bachelor 3)
PRONTO project
Conditioning an LLM with RAG for the space observatory
[I₃₅]

Quentin ASSIÉ



Fouad LBAKALI



Loïck MARCHAL



Mouad SOUISSI



Period

Curriculum

Type

Subject

Report

2025 – 2026 (6 months)
IMT Atlantique (Master 1 & 2)
PROCOM project
Generative AI for music mashups
[I₃₆]

Raphaël HIERSO-IGLESIAS



Period

Curriculum

Type

Subject

2025 – 2026 (4 months)
IMT Atlantique (Master 2)
Research track project
Graph foundation models for De Bruijn graphs

Mohamed-Walid CHAABANE



Ayoub DAOUDI



Salomé GOULIN



Sami GUESSOUM



Elias HASNAOU



...

Christian TRANDAFIR



Period

2025 – 2026 (4 months)

Curriculum

IMT Atlantique (Bachelor 3)

Type

PRONTO project

Subject

Automatic recognition of Pokémon Cards using AI

Report

[I₃₇]

...

Loïc DERRIEN



Théodore DIERS



Jihad OUARD



Period

2025 – 2026 (2 months)

Curriculum

IMT Atlantique (Bachelor 3)

Type

PRAMeL project

Subject

Fine-tuning LLMs for minority languages: an application to Breton

Report

[I₃₈]

B.2.7 Industrial Student Projects

As Associate Professor

Victor BENNINI



Damien CARIOU



Camille FERRERO



Mathieu FRÉMAUX



Benjamin SCHYNS



Xavier SELLAM



Period

2019 – 2020 (4 months)

Curriculum

IMT Atlantique (Master 1)

Type

Industry-partnered tutored project

Employer

Thales

Subject

Design of a chatbot for naval mine warfare

...

Ghita BENKIRANE



Runyuan CAI



Hanchen GUO



Luther OLLIER



Lei WANG   

Period	2019 – 2020 (4 months)
Curriculum	IMT Atlantique (Master 1)
Type	Industry-partnered tutored project
Employer	Zahnrad Fabrik Friedrichshafen
Subject	Classification of defective parts

Fatima-Ezzahra ALAOUI-MRANI   

Emna JMAL   

Soufiane LACHGUER   

Erwann LALE NY   

⋮

Period	2025 – 2026 (6 months)
Curriculum	IMT Atlantique (Master 1 & 2)
Type	PROCOM project
Employer	Arkéa
Subject	Evaluation of prediction uncertainty of LLMs

B.2.8 Tutored Students

As Associate Professor

Fanfol Ismael Raoul COULIBALY   

Period	2019 – 2020 (6 months)
Curriculum	IMT Atlantique (Master 2)
Type	End-of-studies internship
Employer	Ifremer
Subject	Study of upwelling variability from satellite data

Quentin ANDRÉ   

⋮

Period	2021 – 2022 (6 months)
Curriculum	IMT Atlantique (Master 2)
Type	End-of-studies internship
Employer	Hytech-Imaging
Subject	ML for archaeology

Adrien PELLISSIER

Period	2021 – 2022 (6 months)
Curriculum	IMT Atlantique (Master 2)
Type	End-of-studies internship
Employer	Implicity
Subject	Creation of markers for detecting heart failure

Simon Roux

Period	2021 – 2024 (36 months)
Curriculum	IMT Atlantique (FIP)
Type	Apprenticeship program
Employer	Schneider Electric
Subject	Continuous integration

Tom MOALIC

Period	2023 – 2024 (6 months)
Curriculum	IMT Atlantique (Master 2)
Type	End-of-studies internship
Employer	Claude Bernard
Subject	ML for accelerated diagnosis in dermatology

Nicolas RIOU

Period	2024 – 2027 (36 months, ongoing)
Curriculum	IMT Atlantique (FIP)
Type	Apprenticeship program
Employer	EDF
Subject	Networks and cybersecurity

Khaoula AROUSSI

Period	2025 – 2026 (6 months, ongoing)
Curriculum	IMT Atlantique (Master 2)
Type	End-of-studies internship
Employer	Criteo
Subject	Product embeddings search

Agnes Carmela NTANKEU   

Period	2025 – 2026 (6 months, ongoing)
Curriculum	IMT Atlantique (Master 2)
Type	End-of-studies internship
Employer	IMT Atlantique
Subject	Learning analytics with GCNs

B.3 Funded Projects

B.3.1 Research Projects I Wrote or Co-Wrote

How Can One Compare Signals on Different Graphs?

Period	2019 (6 months)
Funding	Hasler Foundation (“ <i>Small Project</i> ” grant)
Type	Private
Amount	~ 50 kCHF
Role	PI

Using Graphs to Constrain Class Dynamics in Deep Learning

Period	2020 (10 months)
Funding	Carnot TSN (“ <i>Numérique 2020</i> ” grant)
Type	Public
Amount	~ 50 k€
Role	Co-PI
Hired	Mounia HAMIDOUCHE

Theoretical Foundations of Deep Learning

Period	2021 (6 months)
Funding	Carnot TSN (“ <i>Numérique 2021</i> ” grant)
Type	Public
Amount	~ 28 k€
Role	Co-PI
Hired	Mounia HAMIDOUCHE

FSL: Toward a Cheap Labeling of Datasets

Period	2021 – 2024 (36 months)
Funding	Fondation Mines-Télécom & Institut Carnot (“ <i>Futur & Ruptures</i> ” grant)
Type	Public
Amount	~ 108 k€
Role	PhD advisor
Hired	Yassir BENDOU

FSL for BCI Systems

Period	2021 – 2024 (36 months)
Funding	Région Bretagne (“ <i>ARED</i> ” grant) IMT & ANR (“ <i>AI@IMT</i> ” grant)
Type	Public
Amount	~ 106 k€
Role	PhD advisor
Hired	Yassine EL OUAHIDI

FUSCHIA

Period	2022 – 2023 (24 months)
Funding	ANR & InPixal (“ <i>ASTRID</i> ” grant)
Type	Public
Amount	~ 142 k€ (total project funding: ~ 295 k€)
Role	Academic partner
Hired	Hafsa EL HAFYANI Jonathan LYS

Research Partnership on AI

Period	2022 – 2026 (1 year, renewed 3 times, ongoing)
Funding	Sony Research
Type	Private
Amount	~ 240 k€
Role	Co-PI

Autonomous Pack

Period	2023 – 2026 (36 months)
Funding	ADEME & Goodflow (“ <i>Logistique 4.0</i> ” grant)
Type	Public
Amount	~ 260 k€ (total project funding: ~ 1,849 k€)
Role	Academic partner
Hired	Manon RENAULT

ENDIVE

Period	2024 – 2028 (48 months, ongoing)
Funding	ANR (“ <i>JCJC</i> ” grant)
Type	Public
Amount	~ 294 k€
Role	PI
Hired	Jonathan Lys + 1 Postdoctoral fellow + 1 Master 2 student

MOVIOLIA

Period	2024 – 2026 (24 months, abandoned after 3 months)
Funding	Région Bretagne & Pôle Images et Réseaux (“ <i>PME 2024</i> ” grant)
Type	Public
Amount	~ 191 k€ (total project funding: ~ 963 k€)
Role	Academic partner
Abandon	The project was abandoned shortly after its acceptance as one of the partners (Tydeo) was placed in liquidation

Diffusion LLMs for Unusual Data

Period	2025 – 2028 (36 months, ongoing)
Funding	IMT Atlantique (CDE) SequoIA (“ <i>Doctoral Projects</i> ” grant)
Type	Public
Amount	~ 139 k€
Role	PhD advisor
Hired	Joshua Chaz RIVERA

Diffusion LLMs & Mechanistic Interpretability

Period	2025 – 2028 (36 months, ongoing)
Funding	IMT Atlantique (CDE) PRACom
Type	Public
Amount	~ 139 k€
Role	PhD advisor
Hired	Nolan SISOUPHANTHONG-VONGKOTRATTANA

Supervision of Trust in Agentic AI Systems Applied to the Medical Domain

Period	2026 – 2027 (6 months, starting soon)
Funding	France 2030 & Cerevial (“ <i>In the Box</i> ” grant)
Type	Public
Amount	~ 10 k€
Role	Academic partner
Hired	Rayan DAKHIA

Advanced Neural Coding for Audio and Stereo Signals

Period	2026 – 2029 (36 months, starting soon)
Funding	Orange (CIFRE)
Type	Private
Amount	PhD salary + 48 k€
Role	PhD director (from HDR)
Hired	Quentin BANET

B.3.2 Research Projects I Contribute(d) To**ML for the Prognosis of Cerebral Hemorrhages**

Period	2019 – 2020 (12 months)
Funding	ARS & CHU Brest (“ <i>Année Recherche</i> ” grant)
Type	Public
Amount	~ 20 k€
Role	Internship advisor
Hired	Victor QUÉRÉ

DS for Energy Efficiency of Cruise Ships

Period	2020 – 2023 (36 months)
Funding	Chantiers de l'Atlantique (CIFRE)
Type	Private
Amount	PhD salary + ~ 75 k€
Role	PhD advisor
Hired	Fred Michael GONSALVES

DARFI3D

Period	2020 (10 months)
Funding	Carnot TSN (“ <i>Numérique 2020</i> ” grant)
Type	Public
Amount	~ 50 k€
Role	AI expert
Hired	Behnam VAHDANI

CODHYPIA

Period	2021 – 2023 (24 months)
Funding	ANR & Elliptika (“ <i>Plan Relance</i> ” grant)
Type	Public
Amount	~ 132 k€
Role	AI expert
Hired	Mervyn GUILLOU

ML for Post-TAVI Risk Assessment

Period	2021 – 2022 (12 months)
Funding	ARS & CHU Brest (“ <i>Année Recherche</i> ” grant)
Type	Public
Amount	~ 20 k€
Role	Internship advisor
Hired	Amine EL OUAHIDI

STRATIF-AI

Period	2022 – 2026 (48 months, ongoing)
Funding	Horizon Europe (“ <i>RIA</i> ” grant)
Type	Public
Amount	~ 401 k€ to CHU Brest (total project funding: ~ 5,969 k€)
Role	AI expert
Hired	Amine SADIKINE

AI for Accelerating the Simulation of District Heating Networks

Period	2022 – 2025 (36 months)
Funding	IMT Atlantique (CDE) IMT & ANR (“ <i>AI@IMT</i> ” grant)
Type	Public
Amount	~ 115 k€
Role	PhD advisor
Hired	Dubon RODRIGUE

B.3.3 Other Funding**Research Stay with Michael RABBAT at McGill University**

Period	2015 (2 months)
Funding	UEB & Région Bretagne (“ <i>International Mobility</i> ” grant) GdR ISIS
Type	Public
Amount	~ 3 k€
Role	Invited researcher

Implementation of the PyRat Course at ECAM Louis de Broglie

Period	2017 (1 day)
Funding	ECAM Louis de Broglie
Type	Private
Amount	~ 1 k€
Role	Instructor

As Associate Professor	Organization of an FSL Workshop	
	Period	2024 (5 days)
	Funding	Brest Métropole
		Département Finistère
	Type	Public
	Amount	~ 5 k€
	Role	Co-organizer

B.4 Teaching Experience

B.4.1 Responsibilities

As PhD Student	PyRat – Algorithms & Discrete Mathematics	
	Period	2015 – 2017
	Curriculum	Télécom Bretagne (Bachelor 3, FISE)
	Role	Course creator
		Development of the video game central to the course
As Associate Professor		Management of a team of ~ 8 teachers
	Publications	[C ₆ ; W ₁ ; S ₁ ; S ₂]
	PyRat – Algorithms & Discrete Mathematics	
	Period	2019 – 2024
	Curriculum	IMT Atlantique (Bachelor 3, FISE)
...	Role	Course leader
		Development of the video game central to the course
		Management of a team of ~ 16 teachers across 2 campuses
	Introduction to Research	
	Period	2019 – 2025
	Curriculum	IMT Atlantique (Master I & 2)
	Role	Course co-leader



Fundamentals of ML

Period	2020 – 2024
Curriculum	IMT Atlantique (Master 1 & 2, TAF DS & TAF DigIC)
Role	Course co-leader

How to Lie with Statistics

Period	2021 – 2026 (ongoing)
Curriculum	IMT Atlantique (Master 1 & 2)
Role	Course creator

Computer Science

Period	2024 – 2026 (ongoing)
Curriculum	IMT Atlantique (Bachelor 3, FISE)
Role	Course co-leader Management of a team of ~ 16 teachers across 2 campuses

DS Track

Period	2024 – 2026 (ongoing)
Curriculum	IMT Atlantique (Master 1 & 2)
Role	Track co-leader

B.4.2 In-Class Teaching Activities



PyRat – Algorithms & Discrete Mathematics

Period	2015 – 2018 (~ 3 × 42 hours)
Curriculum	Télécom Bretagne (Bachelor 3, FISE)
Role	Supervision of 2 classes per year Flipped classes, practical sessions, project oversight
Contents	Graph theory, programming, problem solving
Resources	[M ₁]

Computer Science Major

Period	2015 – 2018 ($\sim 3 \times 22$ hours)
Curriculum	Télécom Bretagne (Master 1 & 2, FISE)
Role	Occasional talks Tutorials, practical sessions, project oversight
Contents	Object-oriented programming, software development

Elementary Signal Processing

Period	2018 – 2019 (~ 5 hours)
Curriculum	EPFL (Bachelor 1)
Role	Occasional talks Practical sessions
Contents	Fourier analysis, filtering

Project Management

Period	2019 – 2020 (~ 63 hours)
Curriculum	IMT Atlantique (Master 1)
Role	Supervision of 2 project teams Tutorials, project oversight
Contents	Planning, management

Business Intelligence

Period	2019 – 2020 (~ 41 hours)
Curriculum	IMT Atlantique (Master 1 & 2, TAF DS)
Role	Supervision of 1 class Tutorials, practical sessions
Contents	Data visualization, databases

Introduction to ML

Period	2019 – 2020 (~ 10 hours)
Curriculum	IMT Atlantique (Master 1 & 2, TAF DS & TAF DigIC)
Role	Occasional talks Lectures, practical sessions, project oversight
Contents	General ML, neural networks

Advanced Data Mining

Period	2019 – 2023 ($\sim 4 \times 6$ hours)
Curriculum	IMT Atlantique (Master 1 & 2, TAF DS)
Role	Occasional talks Lectures, practical sessions
Contents	Time series, LSTM and GRU models
Resources	[M ₂]

PyRat – Algorithms & Discrete Mathematics

Period	2019 – 2024 ($\sim 5 \times 24$ hours)
Curriculum	IMT Atlantique (Bachelor 3, FISE)
Role	Supervision of 1 class per year Flipped classes, practical sessions, project oversight
Contents	Graph theory, programming, problem solving
Resources	[M ₃]

Networked Systems

Period	2019 – 2024 ($\sim 5 \times 23$ hours)
Curriculum	IMT Atlantique (Bachelor 3, FISE)
Role	Supervision of 1 class per year Lectures, tutorials, practical sessions, project oversight
Contents	Optimization, decision aiding, graphs

Advanced ML

Period	2019 – 2025 ($\sim 6 \times 5$ hours)
Curriculum	IMT Atlantique (Master 1 & 2, TAF DS)
Role	Occasional talks Lectures, practical sessions, project oversight
Contents	Deep learning, FSL
Resources	[M ₄]

Introduction to Research

Period	2019 – 2025 ($\sim 6 \times 6$ hours)
Curriculum	IMT Atlantique (Master 1 & 2)

Role	Supervision of 1 class per year Lectures, project oversight
Contents	Research, science

Statistical Learning and Sparse Representations

Period	2019 – 2025 ($\sim 6 \times 3$ hours)
Curriculum	IMT Atlantique (Master 1 & 2, TAF MCE)
Role	Occasional talks Lectures, practical sessions
Contents	Graph signal processing
Resources	[M ₅]

Fundamentals of ML


Period	2020 – 2024 ($\sim 4 \times 36$ hours)
Curriculum	IMT Atlantique (Master 1 & 2, TAF DS & TAF DigIC)
Role	Occasional talks Lectures, practical sessions, project oversight
Contents	ML, deep learning
Resources	[M ₆]

Introduction to AI

Period	2020 – 2026 ($\sim 5 \times 40$ hours, ongoing)
Curriculum	IMT Atlantique (Master 1 & 2)
Role	Supervision of 2 classes per year Lectures, practical sessions, project oversight
Contents	ML, deep learning, foundation models
Resources	[M ₇]


How to Lie with Statistics

Period	2021 – 2026 ($\sim 4 \times 40$ hours, ongoing)
Curriculum	IMT Atlantique (Master 1 & 2)
Role	Supervision of 1 class per year Lectures, tutorials, project oversight
Contents	Science, critical thinking, scientific methodology
Resources	[M ₈]




Advanced Deep Learning

Period	2022 – 2023 (~ 3 hours)
Curriculum	IMT Atlantique (Master 1 & 2, TAF MCE)
Role	Occasional talks Lectures, practical sessions
Contents	RL



Computer Science

Period	2024 – 2026 (~ 60 hours, ongoing)
Curriculum	IMT Atlantique (Bachelor 3, FISE)
Role	Supervision of 1 class Lectures, tutorials, practical sessions, project oversight
Contents	Graph theory, algorithms, programming, problem solving, computer environment
Resources	[M ₉]



DS Track

Period	2024 – 2026 (~ 40 hours, ongoing)
Curriculum	IMT Atlantique (Master 1 & 2)
Role	Supervision of 1 class Lectures, practical sessions, project oversight
Contents	Deep learning, foundation models, unsupervised and supervised models
Resources	[M ₁₀]

B.4.3 Certifications



Computer Science

Period	2018
Type	Qualification
Delivered By	CNU (Section 27)



Computer Engineering, Control Systems, and Signal Processing

Period	2018
Type	Qualification
Delivered By	CNU (Section 61)

Flipping a Class

Period	2019
Type	Training
Delivered By	EPFL

Course Design

Period	2019
Type	Training
Delivered By	RCFE & EPFL

B.4.4 Organization of Programming Challenges

As PhD Student

Enigmatic

Period	2014
Organization	Télécom Bretagne
Role	Co-organizer

TaupIC

Period	2015
Organization	Télécom Bretagne & ENS Rennes
Role	Co-organizer

TaupIC

Period	2016
Organization	Télécom Bretagne & ENS Rennes
Role	Co-organizer

B.5 Other Activities

B.5.1 Research Stays

As PhD Student	LTS₂, EPFL (Lausanne, Switzerland)	
	Period	2014 (2 weeks)
	Invited By	Pierre VANDERGHEYNST
...	ECE Laboratory, McGill University (Montréal, Canada)	
	Period	2015 (2 months)
	Invited By	Michael RABBAT
As Associate Prof.	Sony Research (Stuttgart, Germany)	
	Period	2022 (1 week)
	Invited By	Fabien CARDINAUX
...	CRIS_tAL (Lille, France)	
	Period	2025 (1 week)
	Invited By	Rémi BARDENET

B.5.2 Invited Talks

As PhD Student	Workshop on Serious Games & Pedagogy in Higher Education (UBL)	
	Year	2016
	Subject	Teaching Python with the PyRat serious game
	Presentation	[S ₁]
...	ECAM Louis de Broglie	
	Year	2017
	Subject	Integration of the PyRat game into the training
	Presentation	[S ₂]

GIPSA-Lab

Year	2017
Subject	Extension of CNNs to data supported by graphs
Presentation	[S ₃]

University of Nantes

Year	2017
Subject	Extension of CNNs to data supported by graphs
Presentation	[S ₄]

Inria Rennes

Year	2017
Subject	Extension of CNNs to data supported by graphs
Presentation	[S ₅]

LTS₄, EPFL

Year	2017
Subject	Extension of CNNs to data supported by graphs
Presentation	[S ₆]

ENS Lyon

Year	2018
Subject	Applications of GSP to ML and healthcare
Presentation	[S ₇]

CRIS_tAL

Year	2018
Subject	Applications of GSP to ML and healthcare
Presentation	[S ₈]

As Postdoc

As Associate Professor	Tutorial Day on Foundation Models (IMT Atlantique) <table><tr><td>Year</td><td>2024</td></tr><tr><td>Subject</td><td>Foundation models for text</td></tr><tr><td>Presentation</td><td>[S₉]</td></tr></table>	Year	2024	Subject	Foundation models for text	Presentation	[S ₉]
Year	2024						
Subject	Foundation models for text						
Presentation	[S ₉]						
...	Lab-STICC Presentation <table><tr><td>Year</td><td>2024</td></tr><tr><td>Subject</td><td>Foundation models for text</td></tr><tr><td>Presentation</td><td>[S₁₀]</td></tr></table>	Year	2024	Subject	Foundation models for text	Presentation	[S ₁₀]
Year	2024						
Subject	Foundation models for text						
Presentation	[S ₁₀]						
...	Research Committee JCJC Presentation (IMT Atlantique) <table><tr><td>Year</td><td>2024</td></tr><tr><td>Subject</td><td>Diversity in FSL</td></tr><tr><td>Presentation</td><td>[S₁₁]</td></tr></table>	Year	2024	Subject	Diversity in FSL	Presentation	[S ₁₁]
Year	2024						
Subject	Diversity in FSL						
Presentation	[S ₁₁]						
...	CRIS_tAL <table><tr><td>Year</td><td>2025</td></tr><tr><td>Subject</td><td>Diversity in FSL</td></tr><tr><td>Presentation</td><td>[S₁₂]</td></tr></table>	Year	2025	Subject	Diversity in FSL	Presentation	[S ₁₂]
Year	2025						
Subject	Diversity in FSL						
Presentation	[S ₁₂]						
...	STRATIF-AI Consortium <table><tr><td>Year</td><td>2026</td></tr><tr><td>Subject</td><td>Multimodal ML and agentic AI for health</td></tr><tr><td>Presentation</td><td>[S₁₄]</td></tr></table>	Year	2026	Subject	Multimodal ML and agentic AI for health	Presentation	[S ₁₄]
Year	2026						
Subject	Multimodal ML and agentic AI for health						
Presentation	[S ₁₄]						

B.5.3 Organization of Scientific Conferences & Symposiums

As Postdoc	GSP Workshop 						
	<table><tr><td>Year</td><td>2018</td></tr><tr><td>Type</td><td>Workshop</td></tr><tr><td>Location</td><td>EPFL (Lausanne, Switzerland)</td></tr></table>	Year	2018	Type	Workshop	Location	EPFL (Lausanne, Switzerland)
Year	2018						
Type	Workshop						
Location	EPFL (Lausanne, Switzerland)						

...

Signal Processing over Graphs, with a Focus on Neuroscience Data

Year	2019
Type	GdR ISIS seminar
Location	CNRS (Paris, France)

As Associate Professor

Tutorial Day on Foundation Models

Year	2024
Type	Tutorial
Location	IMT Atlantique (Brest, France)

...

Workshop on FSL

Year	2024
Type	Workshop
Location	IMT Atlantique (Brest, France)

B.5.4 Scientific Dissemination

A Long Long Time Ago

Scientific Dissemination Semester Break

Year	2010
Type	In-class intervention
Organization	ENIB & Les Petits Débrouillards
Public	CM 2 class (École de l'Argoat, Lesneven, France)
Subject	Sea shores
Resources	[A _I]

As Associate Professor

Nuit Européenne des Chercheures

Year	2019
Type	Booth
Organization	Océanopolis
Public	General
Subject	GSP & AI

“Brest is AI” Meetup

Year	2021
Type	Round table
Organization	Brest is AI
Public	General
Subject	AI explainability
Resources	[A ₂]

Nuit Européenne des Chercheurs

Year	2022
Type	Booth
Organization	Océanopolis
Public	General
Subject	Few-shot classification of images
Resources	[A ₃]

Nuit Européenne des Chercheurs

Year	2023
Type	Booth
Organization	Océanopolis
Public	General
Subject	Introduction to AI
Resources	[A ₄]

Fête de la Science

Year	2024
Type	Booth
Organization	Océanopolis
Public	Primary school students
Subject	Building musical instruments

Une Chercheuse à l'École

Year	2025
Type	In-class interventions
Organization	Océanopolis

Public	CM 2 class (Groupe Scolaire Kerbernard, Brest, France) CM 1 & 2 class (École Jean Rostand, Brest, France) CM 1 & 2 class (École Jacques Prévert, Brest, France) CM 2 class (École Keribin, Ploudalmézeau, France) CM 1 & 2 class (École Roz Avel, Plougonvelin, France) CM 1 & 2 class (École Kerinou, Brest, France) CM 1 & 2 class (École Anita Conti, Plouzané, France)
Subject	Understanding AI with Legos
Resources	[A ₅]

Nuit des Chercheurs

Year	2025
Type	Speed searching
Organization	Océanopolis
Public	General
Subject	Research in AI

Fête de la Science

Year	2025
Type	Booth
Organization	Océanopolis
Public	Primary school students
Subject	AI for sound & music

B.5.5 Reviewing Activities

B.5.5.1 Research Articles

IEEE TSP

Period	2016 – 2026
Type	Journal
Nb. Reviews	10

Elsevier Signal Processing

Period	2018
Type	Journal
Nb. Reviews	1

As PhD Student

As Postdoc




IEEE SPL

Period	2018
Type	Journal
Nb. Reviews	1




IEEE TSIPN

Period	2019 – 2026
Type	Journal
Nb. Reviews	4




EURASIP Advances in Signal Processing

Period	2020
Type	Journal
Nb. Reviews	1




GRETSI

Period	2022 – 2025
Type	Conference
Nb. Reviews	7



ACM SIGMETRICS

Period	2022
Type	Conference
Nb. Reviews	1



Oxford Academic - Bioinformatics

Period	2025
Type	Journal
Nb. Reviews	1



ACL ARR

Period	2026
Type	Conference
Nb. Reviews	3

B.5.5.2 ... Research Projects



ANR JCJC

Period	2025 – 2026
Type	Public funding agency
Nb. Reviews	2

APPENDIX C

DATA ANALYTICS



C.1	Teaching	152
C.1.1	In-Class Teaching Hours Per Year	152
C.1.2	Supervised Internships & Student Projects Per Year	152
C.2	Research	153
C.2.1	Publications Per Year	153
C.2.2	Citations Profile	153
C.2.3	Graph of Co-Authors	154

C.1 Teaching

C.1.1 In-Class Teaching Hours Per Year




Figure 17 In-class teaching hours (*i.e.*, projects not included) per school year. Label color indicates when I was a PhD student, a Postdoc or an Associate Professor. The breakdown by course is provided in Section B.4.2.

C.1.2 Supervised Internships & Student Projects Per Year




Figure 18 Number of internships and student projects supervised per school year. Label color indicates when I was a PhD student, a Postdoc or an Associate Professor. See Section B.2.5 to Section B.2.7 for details on all internships and projects.

C.2 Research

C.2.1 Publications Per Year




Figure 19 Number of publications (which underwent a reviewing process) per year. Label color indicates when I was a PhD student, a Postdoc or an Associate Professor. See Bibliography for all details.

C.2.2 Citations Profile




Figure 20 Number of citations per year (as of August 31st, 2026) according to [GOOGLE SCHOLAR (B.P.) Total: 1,048. **H-index:** 12.

RESOURCES

BIBLIOGRAPHY



References I Contributed To

PhD & MD Theses

[T₁ | PASDELOUP, 2017]



III

Bastien PASDELOUP (2017). “**Extending Convolutional Neural Networks to Irregular Domains Through Graph Inference**”. *École Nationale Supérieure Mines-Télécom Atlantique (PhD)*.

[T₂ | QUÉRÉ, 2021]



51, II3

Victor QUÉRÉ (2021). “**Comparison of Logistic Regression Analysis with Random Forest for Prediction of ICH Early Mortality in the Population-Based Brest Stroke Registry**”. *UBO, Faculté de Médecine et des Sciences de la Santé de Brest (MD)*.

[T₃ | KARIMI-MAMAGHAN, 2022]



60, III

Maryam KARIMI-MAMAGHAN (2022). “**Hybridizing Metaheuristics with Machine Learning for Combinatorial Optimization: A Taxonomy and Learning to Select Operators**”. *École Nationale Supérieure Mines-Télécom Atlantique (PhD)*.

[T₄ | EL OUAHIDI, 2023]



54, II4

Amine EL OUAHIDI (2023). “**Utilisation de Méthodes d’Apprentissage Automatique pour Prédire l’Implantation d’un Pacemaker dans les Suites d’une Procédure de TAVI**”. *UBO, Faculté de Médecine et des Sciences de la Santé de Brest (MD)*.

[T₅ | GONSALVES, 2023]



64, III

Fred Michael GONSALVES (2023). “**Leveraging Machine Learning and Operations Research to Enhance Cruise Ship Energy Efficiency**”. *École Nationale Supérieure Mines-Télécom Atlantique (PhD)*.

[T₆ | BENDOU, 2024]



35, II2

Yassir BENDOU (2024). “**A Data Modeling Perspective on Few-Shot Image Classification**”. *École Nationale Supérieure Mines-Télécom Atlantique (PhD)*.

[T₇ | EL OUAHIDI, 2025]



46, II2

Yassine EL OUAHIDI (2025). “**Deep Learning for BCI Applications: Leveraging Recent Advances in Computer Vision**”. *École Nationale Supérieure Mines-Télécom Atlantique (PhD)*.

[T₈ | RODRIGUE, 2025]



66, II2

Dubon RODRIGUE (2025). “**Topology Simplification of District Heating Networks via Artificial Intelligence-Based Node Aggregation**”. *École Nationale Supérieure Mines-Télécom Atlantique (PhD)*.

Book Chapters

[B₁ | PASDELOUP *et al.*, 2019]



5, 27, 120

Bastien PASDELOUP, Vincent GRIPON, Réda ALAMI, Michael RABBAT (2019). “**Uncertainty Principle on Graphs**”. *Vertex-Frequency Analysis of Graph Signals* (pp. 317–340), Springer.

[B₂ | LASSANCE *et al.*, 2022]



30, 114

Carlos LASSANCE, Myriam BONTONOU, Mounia HAMIDOUCHE, Bastien PASDELOUP, Lucas DRUMETZ, Vincent GRIPON (2022). “**Graphs as Tools to Improve Deep Learning Methods**”. *Advances in AI: Reviews* (vol. 2), IFSA Publishing.

Journals

[J₁ | PASDELOUP *et al.*, 2017]



5, 25

Bastien PASDELOUP, Vincent GRIPON, Grégoire MERCIER, Dominique PASTOR, Michael RABBAT (2017). “**Characterization and Inference of Graph Diffusion Processes from Observations of Stationary Signals**”. *TSIPN* (vol. 4, nb. 3, pp. 481–496), IEEE.

[J₂ | BENDOU *et al.*, 2022]



36, 47, 112

Yassir BENDOU, Yuqing HU, Raphael LAFARGUE, Giulia LIOI, Bastien PASDELOUP, Stéphane PATEUX, Vincent GRIPON (2022). “**EASY: Ensemble Augmented-Shot Y-shaped Learning: State-of-the-Art Few-Shot Classification with Simple Components**”. *Journal of Imaging* (vol. 8, nb. 7, art. 179), MDPI.

[J₃ | KARIMI-MAMAGHAN *et al.*, 2023]



61, 111

Maryam KARIMI-MAMAGHAN, Mehrdad MOHAMMADI, Bastien PASDELOUP, Patrick MEYER (2023). “**Learning to Select Operators in Meta-Heuristics: An Integration of Q-Learning into the Iterated Greedy Algorithm for the Permutation Flowshop Scheduling Problem**”. *EJOR* (vol. 304, nb. 3, pp. 1296–1330), Elsevier.

[J₄ | EL OUAHIDI *et al.*, 2024]



48, 112

Yassine EL OUAHIDI, Vincent GRIPON, Bastien PASDELOUP, Ghaith BOUALLEGUE, Nicolas FARRUGIA, Giulia LIOI (2024). “**A Strong and Simple Deep Learning Baseline for BCI Motor Imagery Decoding**”. *Transactions on Neural Systems and Rehabilitation Engineering* (vol. 32, pp. 3338–3347), IEEE.

[J₅ | EL OUAHIDI *et al.*, 2024]



54, 112, 114

Amine EL OUAHIDI, Yassine EL OUAHIDI, Pierre-Philippe NICOL, Sinda HANNACHI, Clément BENIC, Jacques MANSOURATI, Bastien PASDELOUP, Romain DIDIER (2024). “**Machine Learning for Pacemaker Implantation Prediction After TAVI Using Multimodal Imaging Data**”. *Scientific Reports* (vol. 14, nb. 1, art. 25008), Nature Publishing Group.

[J₆ | KHANNOUSSI *et al.*, 2024]



63

Arwa KHANNOUSSI, Alexandru-Liviu OLTEANU, Patrick MEYER, Bastien PASDELOUP (2024). “**A Metaheuristic for Inferring a Ranking Model Based on Multiple Reference Profiles**”. *Annals of Mathematics and Artificial Intelligence* (vol. 92, nb. 6, pp. 1467–1488), Springer.

[J₇ | RODRIGUE *et al.*, 2024]



66, 112

Dubon RODRIGUE, Mohamed Tahar MABROUK, Bastien PASDELOUP, Patrick MEYER, Bruno LACARRIÈRE (2024). “**Topology Reduction Through Machine Learning to Accelerate Dynamic Simulation of District Heating**”. *Energy and AI* (vol. 17, art. 100393), Elsevier.

[J₈ | BELLOUCHE *et al.*, 2025]



55, 112

Yahia BELLOUCHE, Clément BENIC, Sinda HANNACHI, Pierre-Philippe NICOL, Christopher JOUSSE, Florent LE VEN, Jacques MANSOURATI, Bastien PASDELOUP, Romain DIDIER (2025). **“Performance of a Novel Computational Hyperemic Resistance Index Derived from Cardiac CT in Coronary Chronic Syndromes”**. *Journal of Clinical Medicine* (vol. 14, nb. 20, art. 7270), MDPI.

[J₉ | BELLOUCHE *et al.*, 2025]



56, 112

Yahia BELLOUCHE, Cécile TROMEUR, Romain DIDIER, Philippe ROBIN, Christophe GUT-GOBER, Pierre-Yves LE ROUX, Irène FRACHON, Michel NONENT, Alexandre FAUCHÉ, Jamal ELHASNAOUI, Aurélia LE HIR, Marie GUEGAN, Christophe LEROYER, Pierre-Yves SALAÜN, Bastien PASDELOUP, Francis COUTURAUD (2025). **“Proximal Pulmonary Artery Wall Shear Stress Derived from Computational Fluid Dynamics: A Non-Invasive Biomarker for CTEPH and Perfusion Mismatch”**. *Physiological Reports* (vol. 13, nb. 22, pp. e70664), Wiley Online Library.

[J₁₀ | WEIZMAN *et al.*, 2025]



54, 112

Orianne WEIZMAN, Kenza HAMZI, Patrick HENRY, Guillaume SCHURTZ, Marie HAUGUEL-MOREAU, Antonin TRIMAILLE, Marc BEDOSSA, Jean-Claude DIB, Sabir ATTOU, Tanissia BOUKERTOUTA, Franck BOCCARA, Thibaut POMMIER, Pascal LIM, Thomas BOCHATON, Damien MILLISCHER, Benoit MERAT, Fabien PICARD, Nissim GRINBERG, David SULMAN, Bastien PASDELOUP, Yassine EL OUAHIDI, Tręcy GONÇALVES, Éric VICAUT, Jean-Guillaume DILLINGER, Solenn TOUPIN, Théo PEZEL (2025). **“Machine Learning Score to Predict In-Hospital Outcomes in Patients Hospitalized in Cardiac Intensive Care Unit”**. *European Heart Journal, Digital Health* (vol. 6, nb. 2, pp. 218–227), Oxford University Press.

[J₁₁ | RODRIGUE *et al.*, 2026]



68, 112

Dubon RODRIGUE, Mohamed Tahar MABROUK, Bastien PASDELOUP, Patrick MEYER, Bruno LACARRIÈRE (2026). **“Flexible Clustering of Substations for Accurate and Rapid Hybrid Simulation of District Heating”**. *Energy Informatics, Springer*.

Conferences

[C₁ | PASDELOUP *et al.*, 2015]



24

Bastien PASDELOUP, Michael RABBAT, Vincent GRIPON, Dominique PASTOR, Grégoire MERCIER (2015). **“Graph Reconstruction from the Observation of Diffused Signals”**. *Annual Allerton Conference on Communication, Control, and Computing* (pp. 1386–1390), IEEE.

[C₂ | PASDELOUP *et al.*, 2015]



15, 26, 120

Bastien PASDELOUP, Réda ALAMI, Vincent GRIPON, Michael RABBAT (2015). **“Toward an Uncertainty Principle for Weighted Graphs”**. *EUSIPCO* (pp. 1496–1500), IEEE.

[C₃ | PASDELOUP *et al.*, 2015]



26

Bastien PASDELOUP, Vincent GRIPON, Grégoire MERCIER, Dominique PASTOR (2015). **“Vers une Caractérisation de la Courbe d’Incertitude pour des Graphes Portant des Signaux”**. *Colloque du GRETSI* (pp. 989–992).

[C₄ | GRELIER *et al.*, 2016]



15, 28, 120

Nicolas GRELIER, Bastien PASDELOUP, Jean-Charles VIALATTE, Vincent GRIPON (2016). **“Neighborhood-Preserving Translations on Graphs”**. *GlobalSIP* (pp. 410–414), IEEE.

[C₅ | PASDELOUP *et al.*, 2016]



26

Bastien PASDELOUP, Vincent GRIPON, Grégoire MERCIER, Dominique PASTOR (2016). **“Towards a Characterization of the Uncertainty Curve for Graphs”**. *ICASSP* (pp. 4558–4562), IEEE.

- [C₆ | PASDELOUP *et al.*, 2016]
 
 100, 136
- [C₇ | GRELIER *et al.*, 2017]
  
 15, 28, 120
- [C₈ | MÉNORET *et al.*, 2017]
 
 32
- [C₉ | PASDELOUP *et al.*, 2019]
  
 27
- [C₁₀ | KARIMI-MAMAGHAN *et al.*, 2020]
 
 60, III
- [C₁₁ | PASDELOUP *et al.*, 2020]
 
 62
- [C₁₂ | EL OUAHIDI *et al.*, 2021]
   
 15, 57, 58, 108, 123
- [C₁₃ | GONSALVES *et al.*, 2021]
 
 65, III
- [C₁₄ | HAMIDOUCHE *et al.*, 2021]
 
 30, II4
- [C₁₅ | KARIMI-MAMAGHAN *et al.*, 2021]
 
 60, III
- [C₁₆ | KARIMI-MAMAGHAN *et al.*, 2021]
 
 61, III
- [C₁₇ | BENDOU *et al.*, 2022]
   
 35, 38, II2
- Bastien PASDELOUP, Marine KARMANN (2016). **“PyRat : Un Jeu Sérieux pour l’Enseignement de l’Informatique”**. *Colloque QPES* (pp. 713–722).
- Nicolas GRELIER, Bastien PASDELOUP, Jean-Charles VIALATTE, Vincent GRIPON (2017). **“Translation sur Graphe avec Préservation du Voisinage”**. *Colloque du GRETSI* (pp. 821–824).
- Mathilde MÉNORET, Nicolas FARRUGIA, Bastien PASDELOUP, Vincent GRIPON (2017). **“Evaluating Graph Signal Processing for Neuroimaging Through Classification and Dimensionality Reduction”**. *GlobalSIP* (pp. 618–622), IEEE.
- Bastien PASDELOUP, Hermina PETRIC MARETIC, Mireille EL GHECHE, Pascal FROSSARD (2019). **“Une Approche Basée Incertitude pour Transporter un Signal d’un Graphe à un Autre”**. *Colloque du GRETSI* (pp. 37–40).
- Maryam KARIMI-MAMAGHAN, Mehrdad MOHAMMADI, Bastien PASDELOUP, Romain BILLOT, Patrick MEYER (2020). **“An Online Learning-Based Metaheuristic for Solving Combinatorial Optimization Problems”**. *Congrès Annuel de la ROADEF*.
- Bastien PASDELOUP, Maryam KARIMI-MAMAGHAN, Mehrdad MOHAMMADI, Patrick MEYER (2020). **“Autoencoder-Based Generation of Individuals in Population-Based Metaheuristics”**. *Congrès Annuel de la ROADEF*.
- Yassine EL OUAHIDI, Matis FELLER, Matthieu TALAGAS, Bastien PASDELOUP (2021). **“Une Approche pour Regrouper des Sujets Atteints de Cancers, sur la Base des Similarités des Répartitions Cellulaires au Sein de Leurs Biopsies”**. *EGC* (pp. 333–340), RNTI.
- Fred Michael GONSALVES, Bastien PASDELOUP, Romain BILLOT, Patrick MEYER, Arnaud JACQUES, Matthieu LORANG (2021). **“New Insights into the Propulsion Power Prediction of Cruise Ships”**. *ICTAI* (pp. 846–850), IEEE.
- Mounia HAMIDOUCHE, Carlos LASSANCE, Yuqing HU, Lucas DRUMETZ, Bastien PASDELOUP, Vincent GRIPON (2021). **“Improving Classification Accuracy with Graph Filtering”**. *ICIP* (pp. 334–338), IEEE.
- Maryam KARIMI-MAMAGHAN, Bastien PASDELOUP, Mehrdad MOHAMMADI, Patrick MEYER (2021). **“A Learning-Based Iterated Local Search Algorithm for Solving the Traveling Salesman Problem”**. *OLA* (pp. 45–61), Springer.
- Maryam KARIMI-MAMAGHAN, Patrick MEYER, Mehrdad MOHAMMADI, Bastien PASDELOUP (2021). **“A Reinforcement Learning-Based Operator Selection in Iterated Local Search for Solving Scheduling Problems”**. *EURO*.
- Yassir BENDOU, Lucas DRUMETZ, Vincent GRIPON, Giulia LIOI, Bastien PASDELOUP (2022). **“Le Manchot, la Banane et la Bibliothèque... (de la Désambiguïsation d’une Tâche de Classification avec un Exemple)”**. *Colloque du GRETSI* (pp. 1313–1316).

[C₁₈ | EL OUAHIDI *et al.*, 2022]



46, 112

[C₁₉ | EL OUAHIDI *et al.*, 2022]



46, 112

[C₂₀ | KHANNOUSSI *et al.*, 2022]



63

[C₂₁ | LAFARGUE *et al.*, 2022]



39

[C₂₂ | BENDOU *et al.*, 2023]



37, 112

[C₂₃ | BENDOU *et al.*, 2023]



37, 112

[C₂₄ | EL OUAHIDI *et al.*, 2023]



47, 112

[C₂₅ | EL OUAHIDI *et al.*, 2023]



47, 112

[C₂₆ | GUZMAN *et al.*, 2023]



15, 63, 124

[C₂₇ | EL OUAHIDI *et al.*, 2024]



48, 112

[C₂₈ | LYS *et al.*, 2024]



15, 38, 80, 108, 125

[C₂₉ | RODRIGUE *et al.*, 2024]



67, 112

Yassine EL OUAHIDI, Hugo TESSIER, Giulia LIOI, Nicolas FARRUGIA, Bastien PASDELOUP, Vincent GRIPON (2022). “**Élagage de Réseaux de Neurones Convolutifs sur Graphes pour la Sélection de Fréquences Significatives pour le Décodage d’IRMf**”. *Colloque du GRETSI* (pp. 661–664).

Yassine EL OUAHIDI, Hugo TESSIER, Giulia LIOI, Nicolas FARRUGIA, Bastien PASDELOUP, Vincent GRIPON (2022). “**Pruning Graph Convolutional Networks to Select Meaningful Graph Frequencies for fMRI Decoding**”. *EUSIPCO* (pp. 937–941), IEEE.

Arwa KHANNOUSSI, Patrick MEYER, Alexandru-Liviu OLTEANU, Bastien PASDELOUP (2022). “**Un Algorithme Génétique pour l’Apprentissage d’un Modèle de Rangement Multi-Critère à Base de Profils de Références**”. *Congrès Annuel de la ROADEF*.

Raphael LAFARGUE, Bastien PASDELOUP, Jean-Philippe DIGUET, Vincent GRIPON (2022). “**Classes Adversaires dans l’Apprentissage avec Peu d’Exemples**”. *Colloque du GRETSI* (pp. 197–200).

Yassir BENDOU, Vincent GRIPON, Bastien PASDELOUP, Giulia LIOI, Lukas MAUCH, Stefan UHLICH, Fabien CARDINAUX, Ghouthi BOUKLI-HACENE, Javier Alonso GARCIA (2023). “**A Statistical Model for Predicting Generalization in Few-Shot Classification**”. *EUSIPCO* (pp. 1260–1264), IEEE.

Yassir BENDOU, Giulia LIOI, Bastien PASDELOUP, Vincent GRIPON (2023). “**Un Modèle Statistique pour Prédire la Généralisation dans des Tâches de Classification avec Peu d’Exemples**”. *Colloque du GRETSI* (pp. 1005–1008).

Yassine EL OUAHIDI, Lucas DRUMETZ, Giulia LIOI, Nicolas FARRUGIA, Bastien PASDELOUP, Vincent GRIPON (2023). “**Spatial Graph Signal Interpolation with an Application for Merging BCI Datasets with Various Dimensionalities**”. *ICASSP* (pp. 1–5), IEEE.

Yassine EL OUAHIDI, Lucas DRUMETZ, Giulia LIOI, Nicolas FARRUGIA, Bastien PASDELOUP, Vincent GRIPON (2023). “**Interpolation Spatiale de Signaux sur Graphes avec une Application pour la Fusion d’Ensembles de Données EEGs de Diverses Dimensions**”. *Colloque du GRETSI* (pp. 1077–1080).

Miguel GUZMAN, Kadriye Nur BAKIRCI, Martin ROUESNÉ, Hélène SAVATIER-DUPRÉ, Bastien PASDELOUP, Patrick MEYER (2023). “**Utilisation de Processus Ponctuels Déterminantaux pour la Sélection de Parents dans un Algorithme Génétique : Application à l’Aide Multi-Critère à la Décision**”. *Colloque du GRETSI* (pp. 141–144).

Yassine EL OUAHIDI, Giulia LIOI, Nicolas FARRUGIA, Bastien PASDELOUP, Vincent GRIPON (2024). “**Unsupervised Adaptive Deep Learning Method for BCI Motor Imagery Decoding**”. *EUSIPCO* (pp. 1626–1630), IEEE.

Jonathan LYS, Frederic LIN, Clément BÉLIVEAU, Jules DECAESTECKER, Yassir BENDOU, Aymane ABDALI, Bastien PASDELOUP (2024). “**FICUS: Few-Shot Image Classification with Unsupervised Segmentation**”. *EUSIPCO* (pp. 1791–1795), IEEE.

Dubon RODRIGUE, Mohamed Tahar MABROUK, Bastien PASDELOUP, Patrick MEYER, Bruno LACARRIÈRE (2024). “**Leveraging Neural Networks in a Hybrid Model Predictive Control Framework for District Heating Networks**”. *ECOS* (pp. 287–298).

[C₃₀ | DELABAERE *et al.*, 2025]



15, 70, 126

[C₃₁ | EL OUAHIDI *et al.*, 2025]



6, 49, 79, 83, 112, 113

[C₃₂ | PASDELOUP *et al.*, 2025]



70, 82

[C₃₃ | RODRIGUE *et al.*, 2025]



68, 112

[C₃₄ | LIOI *et al.*, 2026]



50, 113

[C₃₅ | LYS *et al.*, 2026]



41, 113

[C₃₆ | LYS *et al.*, 2026]



42, 113

Marine DELABAERE, Léa MIQUEU, Michael MORENO, Gautier BIGOIS, Gautier DUONG, Ella FERNANDEZ, Flavie MANENT, Maria SALGADO-HERRERA, Bastien PASDELOUP, Nicolas FARRUGIA, Axel MARMORET (2025). **“AutoMashup: Automatic Music Mashups Creation”**. *Colloque du GRETSI* (pp. 701–704).

Yassine EL OUAHIDI, Jonathan LYS, Philipp THÖLKE, Nicolas FARRUGIA, Bastien PASDELOUP, Vincent GRIPON, Karim JERBI, Giulia LIOI (2025). **“REVE: A Foundation Model for EEG: Adapting to Any Setup with Large-Scale Pretraining on 25,000 Subjects”**. *NeurIPS*.

Bastien PASDELOUP, Axel MARMORET (2025). **“Raining Words : Les Modèles d’ASR Peuvent-Ils Retranscrire les Sous-Genres de Metal ?”**. *Colloque du GRETSI* (pp. 357–360).

Dubon RODRIGUE, Roberto BOGHETTI, Jérôme Henri KÄMPF, Bastien PASDELOUP, Mohamed Tahar MABROUK, Patrick MEYER, Bruno LACARRIÈRE (2025). **“Validation of Two Distinct Simulation Models of District Heating Networks: Application to Efficient Looping Analysis”**. *Journal of Physics: Conference Series* (vol. 3140, nb. 4, art. 042021), IOP Publishing.

Giulia LIOI, Yassine EL OUAHIDI, Jonathan LYS, Philipp THÖLKE, Bastien PASDELOUP, Vincent GRIPON, Nicolas FARRUGIA, Karim JERBI (2026). **“Key Ingredients for EEG Foundation Models: What Works - and What Doesn’t”**. *EUSIPCO, IEEE*.

Jonathan LYS, Vincent GRIPON, Bastien PASDELOUP, Axel MARMORET, Lukas MAUCH, Fabien CARDINAUX, Ghouthi BOUKLI-HACENE (2026). **“Inner Loop Inference for Pretrained Transformers: Unlocking Latent Capabilities Without Training”**. *EUSIPCO, IEEE*.

Jonathan LYS, Vincent GRIPON, Bastien PASDELOUP, Axel MARMORET, Lukas MAUCH, Fabien CARDINAUX, Ghouthi BOUKLI-HACENE (2026). **“Residual Connections and the Causal Shift: Uncovering a Structural Misalignment in Transformers”**. *EUSIPCO, IEEE*.

Workshops

[W₁ | PASDELOUP, 2016]



5, 100, 136

[W₂ | PASDELOUP *et al.*, 2016]



24

[W₃ | PASDELOUP *et al.*, 2016]



24

[W₄ | PASDELOUP *et al.*, 2016]



24

Bastien PASDELOUP (2016). **“PyRat : Un Jeu pour l’Apprentissage de l’Informatique en Python”**. *PyConFr, AFPy*.

Bastien PASDELOUP, Vincent GRIPON, Dominique PASTOR, Grégoire MERCIER, Michael RABBAT (2016). **“Characterizing Graphs from Diffused Signals”**. *GSP Workshop, University of Pennsylvania*.

Bastien PASDELOUP, Vincent GRIPON, Grégoire MERCIER, Dominique PASTOR, Michael RABBAT (2016). **“Characterization and Inference of Graph Diffusion Processes from Observations of Stationary Processes”**. *Bellairs Workshop on GSP, McGill University*.

Bastien PASDELOUP, Vincent GRIPON, Dominique PASTOR, Grégoire MERCIER, Michael RABBAT (2016). **“Characterizing Graphs from Diffused Signals”**. *Data Driven Approach to Networks and Language, ENS Lyon*.

[W₅ | DADRAS *et al.*, 2018]



32, 117

Ali DADRAS, Bastien PASDELOUP (2018). “**Classifying Non-Stationary Time-Varying Signals on Graphs Through Their Deviation from a Stationary Behavior: Preliminary Results on EEG Signals**”. *GSP Workshop, EPFL*.

[W₆ | PASDELOUP *et al.*, 2018]



29

Bastien PASDELOUP, Jean-Charles VIALATTE, Vincent GRIPON (2018). “**An Extension of Convolutional Neural Networks to Irregular Domains Through Graph Inference and Translations Identification**”. *GSP Workshop, EPFL*.

[W₇ | KARIMI-MAMAGHAN *et al.*, 2021]



61, 111

Maryam KARIMI-MAMAGHAN, Mehrdad MOHAMMADI, Bastien PASDELOUP, Romain BILLOT, Patrick MEYER (2021). “**An Adaptive Iterated Local Search Using Reinforcement Learning for Solving Combinatorial Optimization Problems**”. *CORS Annual Conference (art. OPT₅)*.

[W₈ | EL OUAHIDI *et al.*, 2022]



54, 112, 114

Amine EL OUAHIDI, Yassine EL OUAHIDI, Jacques MANSOURATI, Romain DIDIER, Bastien PASDELOUP (2022). “**Utilisation de Méthodes d'Apprentissage Automatique pour Prédire l'Implantation d'un Pacemaker dans les Suites d'une Procédure de TAVI**”. *Journées Recherche, CHU Brest*.

[W₉ | KARIMI-MAMAGHAN *et al.*, 2022]



61, 111

Maryam KARIMI-MAMAGHAN, Mehrdad MOHAMMADI, Bastien PASDELOUP, Patrick MEYER (2022). “**Learning to Select Operators in Meta-Heuristics: Application to Flowshop Scheduling Problem**”. *CORS Annual Conference (art. TA16)*.

[W₁₀ | PASDELOUP *et al.*, 2022]



63

Bastien PASDELOUP, Arwa KHANNOUSSI, Alexandru-Liviu OLTEANU, Patrick MEYER (2022). “**A Genetic Algorithm for Learning the Parameters of an SRMP Preference Model**”. *DA2PL*.

[W₁₁ | BENDOU *et al.*, 2023]



37, 112

Yassir BENDOU, Vincent GRIPON, Bastien PASDELOUP, Giulia LIOI, Lukas MAUCH, Fabien CARDINAUX, Ghouthi BOUKLI-HACENE (2023). “**Inferring Latent Class Statistics from Text for Robust Visual Few-Shot Learning**”. *NeurIPS R0-FoMo Workshop*.

[W₁₂ | EL OUAHIDI *et al.*, 2023]



47, 112

Yassine EL OUAHIDI, Nicolas FARRUGIA, Bastien PASDELOUP, Vincent GRIPON, Giulia LIOI (2023). “**Cross-Dataset Few-Shot Learning for Motor Imagery BCI Classification**”. *International BCI Meeting, BCI Society*.

[W₁₃ | EL OUAHIDI *et al.*, 2023]



47, 112

Yassine EL OUAHIDI, Nicolas FARRUGIA, Bastien PASDELOUP, Vincent GRIPON, Giulia LIOI (2023). “**Cross-Dataset Few-Shot Learning for Motor Imagery BCI Classification**”. *CORTICO Days, CORTICO*.

[W₁₄ | MOUMMAD *et al.*, 2023]



31

Ilyass MOUMMAD, Bastien PASDELOUP, Nicolas FARRUGIA (2023). “**Contrastive Learning Using Random Walk Laplacian Matrix**”. *GSP Workshop*.

[W₁₅ | RODRIGUE *et al.*, 2023]



68, 112

Dubon RODRIGUE, Mohamed Tahar MABROUK, Bastien PASDELOUP, Patrick MEYER, Bruno LACARRIÈRE (2023). “**Learning the Dynamic of Clusters of Vertices in a Graph for District Heating Network Simplification**”. *GSP Workshop*.

[W₁₆ | WEIZMAN *et al.*, 2023]



54, 112

Orianne WEIZMAN, Théo PEZEL, Kenza HAMZI, Guillaume SCHURTZ, Marie HAUGUEL-MOREAU, Antonin TRIMAILLE, Solenn TOUPIN, Thomas BOCHATON, Sabir ATTOU, Christophe MEUNE, Franck BOCCARA, Bastien PASDELOUP, Yassine EL OUAHIDI, Jean-Guillaume DILLINGER, Patrick HENRY (2023). “**Machine-Learning Score to Predict In-Hospital Outcomes in Patients Hospitalized in Intensive Cardiac Care Unit**”. *ESC Congress*.

[W₁₇ | RENAULT *et al.*, 2025]



69, 115

Manon RENAULT, Hamoud YOUNES, Hugo TESSIER, Ronan LE ROY, Bastien PASDELOUP, Mathieu LÉONARDON (2025). “**Event Classification of Accelerometer Data for Industrial Package Monitoring with Embedded Deep Learning**”. *COMPSAC (pp. 1671–1676), IEEE*.

Seminars & Invited Talks

[S₁ | PASDELOUP, 2016]



100, 136, 143

Bastien PASDELOUP (2016). “**PyRat : Un Jeu pour l’Apprentissage de l’Informatique avec Python**”. *Jeux Sérieux et Pédagogie de l’Enseignement Supérieur, Université Bretagne Loire*.

[S₂ | PASDELOUP, 2017]



136, 143

Bastien PASDELOUP (2017). “**PyRat : Un Jeu pour l’Apprentissage de l’Informatique avec Python**”. *ECAM Louis de Broglie*.

[S₃ | PASDELOUP, 2017]



29, 144

Bastien PASDELOUP (2017). “**Extending Convolutional Neural Networks to Irregular Domains Through Graph Inference**”. *GIPSA-Lab*.

[S₄ | PASDELOUP, 2017]



29, 144

Bastien PASDELOUP (2017). “**Extending Convolutional Neural Networks to Irregular Domains Through Graph Inference**”. *Université de Nantes*.

[S₅ | PASDELOUP, 2017]



29, 144

Bastien PASDELOUP (2017). “**Extending Convolutional Neural Networks to Irregular Domains Through Graph Inference**”. *Inria*.

[S₆ | PASDELOUP, 2017]



29, 144

Bastien PASDELOUP (2017). “**Extending Convolutional Neural Networks to Irregular Domains Through Graph Inference**”. *EPFL*.

[S₇ | PASDELOUP, 2018]



144

Bastien PASDELOUP (2018). “**A Few Applications of GSP to Machine Learning and Medicine**”. *ENS Lyon*.

[S₈ | PASDELOUP, 2018]



144

Bastien PASDELOUP (2018). “**A Few Applications of GSP to Machine Learning and Medicine**”. *CRISTAL*.

[S₉ | LIOI *et al.*, 2024]



145

Giulia LIOI, Bastien PASDELOUP (2024). “**Natural Language Processing with LLaMA-2**”. *Tutorial Day on Foundation Models, IMT Atlantique*.

[S₁₀ | LIOI *et al.*, 2024]



145

Giulia LIOI, Bastien PASDELOUP (2024). “**Natural Language Processing with LLaMA-2**”. *Lab-STICC*.

[S₁₁ | PASDELOUP, 2024]



145

Bastien PASDELOUP (2024). “**JCJC ENDIVE**”. *IMT Atlantique*.

[S₁₂ | PASDELOUP, 2025]



145

Bastien PASDELOUP (2025). “**JCJC ENDIVE**”. *CRISTAL*.

[S₁₃ | SADIKINE *et al.*, 2026]



53, 83, 114

[S₁₄ | VIALATTE *et al.*, 2026]



78, 145

Preprints

[P₁ | PASDELOUP *et al.*, 2017]



28, 120

[P₂ | PASDELOUP *et al.*, 2017]



29, 33

[P₃ | GONSALVES *et al.*, 2023]



65, 111

[P₄ | BENDOU *et al.*, 2024]



38, 112

[P₅ | EL HAFYANI *et al.*, 2024]



38, 80, 114

[P₆ | LAFARGUE *et al.*, 2024]



39

Under Review

[R₁ | RIVERA *et al.*, 2026]



43, 113

[R₂ | RIVERA *et al.*, 2026]



43, 113

[R₃ | LYS *et al.*, 2026]



42, 113

Amine SADIKINE, Serge TIMSIT, Bastien PASDELOUP (2026). “**Multimodal Data Integration for 30-Day Mortality Prediction in Hemorrhagic Stroke**”. *STRATIF-AI Consortium*.

Jean-Charles VIALATTE, Bastien PASDELOUP (2026). “**Multimodal Machine Learning and Agentic AI for Health**”. *STRATIF-AI Consortium*.

Bastien PASDELOUP, Vincent GRIPON, Jean-Charles VIALATTE, Nicolas GRELLIER, Dominique PASTOR (2017). “**A Neighborhood-Preserving Translation Operator on Graphs**”. *arXiv preprint arXiv:1709.03859*.

Bastien PASDELOUP, Vincent GRIPON, Jean-Charles VIALATTE, Dominique PASTOR, Pascal FROSSARD (2017). “**Convolutional Neural Networks on Irregular Domains Based on Approximate Vertex-Domain Translations**”. *arXiv preprint arXiv:1710.10035*.

Fred Michael GONSALVES, Patrick MEYER, Romain BILLOT, Bastien PASDELOUP, Matthieu LORANG (2023). “**Optimal Design and Operation of Cruise Ship Multi-Energy Systems: An MILP Formulation**”. *HAL preprint hal-04310676*.

Yassir BENDOU, Giulia LIOI, Bastien PASDELOUP, Lukas MAUCH, Ghouthi BOUKLI-HACENE, Fabien CARDINAUX, Vincent GRIPON (2024). “**LLM Meets Vision-Language Models for Zero-Shot One-Class Classification**”. *arXiv preprint arXiv:2404.00675*.

Hafsa EL HAFYANI, Bastien PASDELOUP, Camille YVER, Pierre ROMENTEAU (2024). “**Multimodal Object Detection via Probabilistic a priori Information Integration**”. *arXiv preprint arXiv:2405.15596*.

Raphael LAFARGUE, Yassir BENDOU, Bastien PASDELOUP, Jean-Philippe DUGUET, Ian REID, Vincent GRIPON, Jack VALMADRE (2024). “**Few and Fewer: Learning Better from Few Examples Using Fewer Base Classes**”. *arXiv preprint arXiv:2401.15834*.

Joshua Chaz RIVERA, Bastien PASDELOUP, Vincent GRIPON (2026). “**Evaluating LLMs on Sentence-Level Metaphors**”. *Findings of the ACL*.

Joshua Chaz RIVERA, Bastien PASDELOUP, Vincent GRIPON (2026). “**Idiomatic Blindspots in Discrete Diffusion Models**”. *EMNLP*.

Jonathan LYS, Vincent GRIPON, Axel MARMORET, Lukas MAUCH, Fabien CARDINAUX, Ghouthi BOUKLI-HACENE, Bastien PASDELOUP (2026). “**D5P4: Partition Determinantal Point Process for Diversity in Parallel Discrete Diffusion Decoding**”. *NeurIPS*.

[R₄ | WICKY *et al.*, 2026]

■ 57, 59

Alexandre WICKY, Jonatan BONJOUR, Beril BESBINAR, Petros LIAKOPOULOS, Bastien PASDELOUP, Mireille EL GHECHE, Gojko CUTURA, Asmi SOUHAIL, Carl HATOUM, Charilaos ZISOU, Sevda OGUT, Robin LIECHTI, Julien DORIER, Ekaterina FORTIS, Christel VOIZE, Rita DE MICHELI, Clarisse DROMAIN, Julien DAGHER, Igor LETOVANEC, Laurence DE LEVAL, Paweł TETERYCZ, Anna SZUMERA-CIEĆKIEWICZ, Piotr RUTKOWSKI, Sofiya LATIFYAN, Andrew JANOWCZYK, Sylvain PRADERVAND, Mikaël J. PIT-TET, Stephanie TISSOT, Sylvie RUSAKIEWICZ, George COUKOS, Dorina THANOU, Pascal FROSSARD, Michel A. CUENDET, Olivier MICHIELIN (2026). **“Spatial Organization of the Immune Microenvironment Predicts Checkpoint Inhibitor Response in Advanced Melanoma”**. *Nature Communications*, Nature Publishing Group.

Internship & Project Reports

[I₁ | PASDELOUP, 2009]



■ IIO

Bastien PASDELOUP (2009). **“Développement d’un Site Intranet de Gestion des Formations”**. *ISEN (BTS 1, Charles de Foucauld)*.

[I₂ | PASDELOUP, 2010]



■ I2, IIO

Bastien PASDELOUP (2010). **“Étude et Développement de Solutions 3D pour la Réalisation d’un Prototype de Jeu Éducatif Basé sur les Mathématiques”**. *JiVé Création, CERV (BTS 2, Charles de Foucauld)*.

[I₃ | PASDELOUP, 2011]



■ IIO

Bastien PASDELOUP (2011). **“Développement C++/Python d’un Logiciel de Contrôle d’Actions sur une Machine Distante”**. *Diateam (Bachelor 3, ENIB)*.

[I₄ | PASDELOUP, 2012]



■ I2, IIO

Bastien PASDELOUP (2012). **“Détection de Perte de Précision dans le Cadre de l’Analyse Statique d’un Fichier Binaire Exécutable”**. *DGA/MI (Magistère 1, ENS Rennes & Bachelor 3, Université de Rennes 1)*.

[I₅ | PASDELOUP, 2013]



■ I2, IIO

Bastien PASDELOUP (2013). **“Reduction of the Worst-Case Execution Time of Real-Time Tasks Targeting Many-Core Platforms”**. *CISTER (Magistère 2, ENS Rennes & Master 1, Université de Rennes 1)*.

[I₆ | PASDELOUP, 2014]



■ I2, IIO

Bastien PASDELOUP (2014). **“Static Probabilistic Timing Analysis of Worst-Case Execution Time for Random Replacement Caches”**. *IRISA, ALF (Magistère 3, ENS Rennes & Master 2, Université de Rennes 1)*.

[I₇ | AITELAMINE, 2015]



■ IIS

Khalid AITELAMINE (2015). **“Application du Traitement de Signal sur Graphe à la Classification”**. *Télécom Bretagne (Master 1, Télécom Bretagne)*.

[I₈ | BOUALILA *et al.*, 2015]



■ I2O

Karim BOUALILA, Alexandre DELAROCHE, Natacha DESERT, Julien MATTEI (2015). **“Amélioration de la Plateforme Enigmatic”**. *Télécom Bretagne (Bachelor 3, Télécom Bretagne)*.

[I₉ | SHI *et al.*, 2015]



■ I2O

Jinglei SHI, Yihong XU (2015). **“Comparison and Improvement of Two Methods for Reconstruction of a Graph from Signals”**. *Télécom Bretagne (Master 2, Télécom Bretagne)*.

[I₁₀ | ALAMI, 2016]



■ I2O

Réda ALAMI (2016). **“Graph Signal Processing: Towards the Diffused Spectral Clustering”**. *Télécom Bretagne (Master 2, Télécom Bretagne)*.

[I₁₁ | DUCERF *et al.*, 2016]



121

[I₁₂ | ECHRAIBI *et al.*, 2016]



121

[I₁₃ | GRELIER, 2016]



120

[I₁₄ | GRELIER, 2016]



121

[I₁₅ | KERINEC, 2016]



116

[I₁₆ | AKBIYIK, 2018]



122

[I₁₇ | DADRAS, 2018]



117

[I₁₈ | SHU-QUARTIER-DIT-MAIRE, 2018]



117

[I₁₉ | ZANOLI, 2018]



122

[I₂₀ | BEN OTHMAN, 2019]



122

[I₂₁ | BOUDJEMAA *et al.*, 2019]



122

[I₂₂ | AVARE *et al.*, 2020]



123

[I₂₃ | BOURDIN WHITE *et al.*, 2020]



124

[I₂₄ | EL OUAHIDI *et al.*, 2020]



123

Ladislav DUCERF, Sili ZUO (2016). **“MNIST and EEG Signal Visualization and Classification Through Dimension Reduction”**. *Télécom Bretagne (Master 1, Télécom Bretagne)*.

Amine ECHRAIBI, Louis Gildas NDZAMBA, Lotfi MOUBAKIR (2016). **“Déploiement de Logiciel de Machine Virtuelle et Exécution de Programme en Environnement Contrôlé”**. *Télécom Bretagne (Bachelor 3, Télécom Bretagne)*.

Nicolas GRELIER (2016). **“Identification de Translations dans des Graphes”**. *Télécom Bretagne (Bachelor 3, Télécom Bretagne)*.

Nicolas GRELIER (2016). **“Quantum Random Walk”**. *Télécom Bretagne (Master 1, Télécom Bretagne)*.

Emma KERINEC (2016). **“Détection de Motifs Spatio-Temporels dans des Signaux Multivariés”**. *Télécom Bretagne (Bachelor 3, ENS Lyon)*.

Murtaza Eren AKBIYIK (2018). **“Cell Extraction from Staining Images”**. *EPFL (Bachelor 3, EPFL)*.

Ali DADRAS (2018). **“Classifying Non-Stationary Time-Varying Signals on Graphs Through Their Deviation from a Stationary Behavior: Preliminary Results on EEG Signals”**. *EPFL (Master 2, Tabriz University)*.

Wenqi SHU-QUARTIER-DIT-MAIRE (2018). **“Classification of Cancer Types from Medical Stained Images”**. *EPFL (Master 1, ENS Paris-Saclay)*.

Silvio ZANOLI (2018). **“Extraction of Features from Signals on Cell-Graphs to Model Biological Image Information”**. *EPFL (Master 1, EPFL)*.

Firas BEN OTHMAN (2019). **“Improving Classification Performance of Convolutional Neural Networks by Changing the Convolutional Kernel Shape”**. *EPFL (Master 1, EPFL)*.

Frédéric BOUDJEMAA, Albane COLIN DE VERTIÈRE, Joséphine SCHMUTZ (2019). **“Analyse de Données Météorologique avec un Spectrogramme”**. *IMT Atlantique (Master 1, IMT Atlantique)*.

Quentin AVARE, Théo LE GALL, Aurélien PAGEAU (2020). **“Éditeur de Labyrinthe PyRat”**. *IMT Atlantique (Bachelor 3, IMT Atlantique)*.

Jack BOURDIN WHITE, Gabriel GROS (2020). **“Classification de Données pour de l'Apprentissage Semi-Supervisé”**. *IMT Atlantique (Bachelor 3, IMT Atlantique)*.

Yassine EL OUAHIDI, Matis FELLER (2020). **“Segmentation et Analyse de la Répartition des Cellules dans des Tissus Cancéreux”**. *IMT Atlantique (Master 2, IMT Atlantique)*.

- [I₂₅ | ELOY *et al.*, 2020]
  I23
 Thibaut ELOY, Anthony FRION (2020). “**Analysis of the Relation Between the MinLA Problem and the Visualisation of a Graph Spectrogram**”. *IMT Atlantique (Master 2, IMT Atlantique)*.
- [I₂₆ | QUÉRÉ, 2020]
  I18
 Victor QUÉRÉ (2020). “**Hybrid Machine Learning Approach to Evaluate the Prognosis of Spontaneous Cerebral Hemorrhages**”. *IMT Atlantique (Master 2, UBO)*.
- [I₂₇ | AMAL *et al.*, 2021]
  I23
 Ilias AMAL, Diego Andres CARREÑO, Johan MEJIA, Tatiana MORENO PEREA (2021). “**Spatial-Temporal Analysis of Signals to Detect Drift of Air Quality Sensors**”. *IMT Atlantique (Master 2, IMT Atlantique)*.
- [I₂₈ | ANDRÉ *et al.*, 2021]
  I24
 Quentin ANDRÉ, Maëlle GABENS (2021). “**Language and Dialect Recognition for Cuneiform Texts**”. *IMT Atlantique (Master 1 & 2, IMT Atlantique)*.
- [I₂₉ | SERRE, 2021]
  I18
 Marc SERRE (2021). “**Identification de Classes Adversaires pour l’Apprentissage Few-Shot**”. *IMT Atlantique (Bachelor 3, IMT Atlantique)*.
- [I₃₀ | EL OUAHIDI, 2022]
  53, I18
 Amine EL OUAHIDI (2022). “**Utilisation de Méthodes d’Apprentissage Automatique pour Prédire l’Implantation d’un Pacemaker dans les Suites d’une Procédure de TAVI**”. *IMT Atlantique (Master 2, UBO)*.
- [I₃₁ | GABENS, 2022]
  I18
 Maëlle GABENS (2022). “**Pruning Appliqué à l’Échantillonnage de Signal sur Graphe**”. *IMT Atlantique (Master 1, IMT Atlantique)*.
- [I₃₂ | LEYE *et al.*, 2023]
  I24
 Mouhamed LEYE, Guillaume LODÉ, Thomas MONTAGNE (2023). “**Amplificateur de Guitare Portatif**”. *IMT Atlantique (Bachelor 3, IMT Atlantique)*.
- [I₃₃ | ADJA *et al.*, 2024]
  I25
 Yanis ADJA, Florent BERGÉ, Aeddan WILLIAMS (2024). “**Amplificateur de Guitare Portatif**”. *IMT Atlantique (Bachelor 3, IMT Atlantique)*.
- [I₃₄ | DEVOUGE-BOYER *et al.*, 2024]
  I25
 Paul DEVOUGE-BOYER, Astrid FLICHY, Lucas POIRIER, Pauline THIBAUT (2024). “**Création d’un Site Web Permettant l’Emprunt de Jeux de Société du Club Jeux d’IMT Atlantique**”. *IMT Atlantique (Bachelor 3, IMT Atlantique)*.
- [I₃₅ | BOULLOUZ *et al.*, 2025]
  I26
 Mohamed-Ali BOULLOUZ, Axel JONARD, Brahim ONASSER, Aymane OUNZAR (2025). “**Conception d’un Agent IA pour l’Observatoire Astronomique**”. *IMT Atlantique (Bachelor 3, IMT Atlantique)*.
- [I₃₆ | ASSIÉ *et al.*, 2026]
  I26
 Quentin ASSIÉ, Fouad LBAKALI, Loïck MARCHAL, Mouad SOUSSI (2026). “**AutoMashup – Mashups Musicaux Automatiques avec l’IA**”. *IMT Atlantique (Master 1 & 2, IMT Atlantique)*.
- [I₃₇ | CHAABANE *et al.*, 2026]
  I27
 Mohamed-Walid CHAABANE, Ayoub DAOUDI, Salomé GOULIN, Sami GUESOUM, Elias HASNAOU, Christian TRANDAFIR (2026). “**Reconnaissance Automatique de Cartes Pokémon – Identification Visuelle par IA**”. *IMT Atlantique (Bachelor 3, IMT Atlantique)*.
- [I₃₈ | DERRIEN *et al.*, 2026]
  I27
 Loïc DERRIEN, Théodore DIERS, Jihad OUARD (2026). “**Fine-Tuning de LLMs pour le Breton**”. *IMT Atlantique (Master 1 & 2, IMT Atlantique)*.

Teaching Materials

[M₁ | PASDELOUP *et al.*, 2015 – 2018]



98, 100, 137

Bastien PASDELOUP, Vincent GRIPON (2015 – 2018). “**PyRat – Algorithms & Discrete Mathematics**”. *Télécom Bretagne*.

[M₂ | PASDELOUP, 2019 – 2023]



139

Bastien PASDELOUP (2019 – 2023). “**Advanced Data Mining – Deep Learning for Time Series**”. *IMT Atlantique*.

[M₃ | PASDELOUP, 2019 – 2024]



98, 100, 139

Bastien PASDELOUP (2019 – 2024). “**PyRat – Algorithms & Discrete Mathematics**”. *IMT Atlantique*.

[M₄ | PASDELOUP *et al.*, 2019 – 2025]



139

Bastien PASDELOUP, Yassir BENDOU (2019 – 2025). “**Advanced ML – Deep Learning & FSL**”. *IMT Atlantique*.

[M₅ | PASDELOUP, 2019 – 2025]



140

Bastien PASDELOUP (2019 – 2025). “**Statistical Learning & Sparse Representations – GSP**”. *IMT Atlantique*.

[M₆ | PASDELOUP, 2020 – 2024]



140

Bastien PASDELOUP (2020 – 2024). “**Fundamentals of ML – (Deep) Neural Networks**”. *IMT Atlantique*.

[M₇ | MARMORET *et al.*, 2020 – 2026]



140

Axel MARMORET, Nicolas FARRUGIA, Bastien PASDELOUP, Giulia LIOI (2020 – 2026). “**Introduction to AI**”. *IMT Atlantique*.

[M₈ | PASDELOUP *et al.*, 2021 – 2026]



140

Bastien PASDELOUP, Marie-Claude VEILLEUX, Romain BILLOT (2021 – 2026). “**How to Lie with Statistics**”. *IMT Atlantique*.

[M₉ | PASDELOUP *et al.*, 2024 – 2026]



141

Bastien PASDELOUP, Grégory SMITS, Charles PRUD’HOMME, Remous Aris KOUTSIAMANIS (2024 – 2026). “**Computer Science**”. *IMT Atlantique*.

[M₁₀ | MATTA *et al.*, 2024 – 2026]



141

Sarah MATTA, Bastien PASDELOUP (2024 – 2026). “**DS Track**”. *IMT Atlantique*.

Animations

[A₁ | HANCE *et al.*, 2010]



146

Mickaël HANCE, Mathieu LETURCQ, Bastien PASDELOUP, Maxime RAFFRAY (2010). “**Le Littoral**”. *ENIB & Les Petits Débrouillards*.

[A₂ | CHRAIBI KAADOUH *et al.*, 2021]



147

Ikram CHRAIBI KAADOUH, Aymeric POULAIN MAUBANT, Lina FAHED, Romain BILLOT, Bastien PASDELOUP (2021). “**Explicabilité des Systèmes d’IA**”. *Brest is AI*.

[A₃ | BENDOU *et al.*, 2022]



36, 147

Yassir BENDOU, Raphael LAFARGUE, Bastien PASDELOUP (2022). “**Low-Latency Few-Shot Classification**”. *Nuit Européenne des Chercheurs*.

[A₄ | PASDELOUP *et al.*, 2023]



147

[A₅ | PASDELOUP, 2025]



148

Bastien PASDELOUP, Reda BENSAD, Yassir BENDOU (2023). “L’IA, c’est Quoi ?”. *Nuit Européenne des Chercheures*.

Bastien PASDELOUP (2025). “La Grande Découverte des Legoalgues”. *Une Chercheure à l’École*.

Other References

[ABRAS *et al.*, 2026]



— 54

Marcel ABRAS, Daniela BURSACOVSKI, Ecaterina PASAT, Maria-Magdalena VICOL, Tatiana ABRAS, Lucia MAZUR-NICORICI, Oleg ARNAUT (2026). “Machine Learning-Driven Probability of Permanent Pacemaker Implantation After Transcatheter Aortic Valve Replacement”. *Diagnostics* (vol. 16, nb. 11, art. 1720), MDPI.

[AGASKAR *et al.*, 2012]



— 26

Ameya AGASKAR, Yue M. LU (2012). “Uncertainty Principles for Signals Defined on Graphs: Bounds and Characterizations”. *ICASSP* (pp. 3493–3496), IEEE.

[AGASKAR *et al.*, 2013]



— 26

Ameya AGASKAR, Yue M. LU (2013). “A Spectral Graph Uncertainty Principle”. *Transactions on Information Theory* (vol. 59, nb. 7, pp. 4338–4356), IEEE.

[ALAYRAC *et al.*, 2022]



— 40

Jean-Baptiste ALAYRAC, Jeff DONAHUE, Pauline LUC, Antoine MIECH, Iain BARR, Yana HASSON, Karel LENC, Arthur MENSCH, Katherine MILLICAN, Malcolm REYNOLDS, Roman RING, Eliza RUTHERFORD, Serkan CABI, Tengda HAN, Zhitao GONG, Sina SAMANGOOEI, Marianne MONTEIRO, Jacob L. MENICK, Sebastian BORGEAUD, Andy BROCK, Aida NEMATZADEH, Sahand SHARIFZADEH, Mikołaj BIŃKOWSKI, Ricardo BARREIRA, Oriol VINYALS, Andrew ZISSERMAN, Karén SIMONYAN (2022). “Flemingo: A Visual Language Model for Few-Shot Learning”. *NeurIPS* (vol. 35, pp. 23716–23736).

[ANGELOPOULOS *et al.*, 2023]



— 83, 84

Anastasios N. ANGELOPOULOS, Stephen BATES (2023). “Conformal Prediction: A Gentle Introduction”. *Foundations and Trends in Machine Learning* (vol. 16, nb. 4, pp. 494–591), Emerald Publishing Limited.

[ANIS *et al.*, 2016]



— 18

Aamir ANIS, Akshay GADDE, Antonio ORTEGA (2016). “Efficient Sampling Set Selection for Bandlimited Graph Signals Using Graph Spectral Proxies”. *TSP* (vol. 64, nb. 14, pp. 3775–3789), IEEE.

[ARMENI *et al.*, 2022]



— 52

Patrizio ARMENI, Irem POLAT, Leonardo Maria DE ROSSI, Lorenzo DIAFERIA, Severino MEREGALLI, Anna GATTI (2022). “Digital Twins in Healthcare: Is It the Beginning of a New Era of Evidence-Based Medicine? A Critical Review”. *Journal of Personalized Medicine* (vol. 12, nb. 8, art. 1255), MDPI.

[ASGARI *et al.*, 2025]



— 78

Elham ASGARI, Nina MONTAÑA-BROWN, Magda DUBOIS, Saleh KHALIL, Jasmine BALLOCH, Joshua Au YEUNG, Dominic PIMENTA (2025). “A Framework to Assess Clinical Safety and Hallucination Rates of LLMs for Medical Text Summarisation”. *NPJ Digital Medicine* (vol. 8, nb. 1, art. 274), Nature Publishing Group.

[AUFFRET *et al.*, 2018]

 —  81

[AUSTIN *et al.*, 2021]

 —  40

[BAEVSKI *et al.*, 2020]

 —  69

[BAILLET, 2017]

 —  45

[BANVILLE *et al.*, 2026]

 —  49

[BARBAROSSA *et al.*, 2020]

 —  24

[BARBEAU, 1993]

 —  103

[BENGIO *et al.*, 2013]

 —  33

[BENGIO *et al.*, 2003]

 —  40

[BOGHETTI *et al.*, 2026]

 —  69

[BOMMASANI *et al.*, 2021]

 —  33, 35, 39

Vincent AUFFRET, Yves COTTIN, Guillaume LEURENT, Martine GILARD, Jean-Claude BEER, Amer ZABALAWI, Frederic CHAGUE, Emanuelle FILIPPI, Damien BRUNET, Jean-Philippe HACOT, Philippe BRUNEL, Mourad MEJRI, Luc LORGIS, Gilles ROUAULT, Philippes DRUELLES, Jean-Christophe CORNILY, Romain DIDIER, Emilie BOT, Bertrand BOULANGER, Isabelle COUDERT, Aurélie LOIRAT, Marc BEDOSSA, Dominique BOULMIER, Maud MAZA, Marielle LE GUELLEC, Rishi PURI, Marianne ZELLER, Hervé LE BRETON (2018). “**Predicting the Development of In-Hospital Cardiogenic Shock in Patients with ST-Segment Elevation Myocardial Infarction Treated by Primary Percutaneous Coronary Intervention: the ORBI Risk Score**”. *European Heart Journal* (vol. 39, nb. 22, pp. 2090–2102), Oxford University Press.

Jacob AUSTIN, Daniel D. JOHNSON, Jonathan HO, Daniel TARLOW, Ri-
anne VAN DEN BERG (2021). “**Structured Denoising Diffusion Models in Discrete State-Spaces**”. *NeurIPS* (vol. 34, pp. 17981–17993).

Alexei BAEVSKI, Yuhao ZHOU, Abdelrahman MOHAMED, Michael AULI (2020). “**wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations**”. *NeurIPS* (vol. 33, pp. 12449–12460).

Sylvain BAILLET (2017). “**Magnetoencephalography for Brain Electro-physiology and Imaging**”. *Nature neuroscience* (vol. 20, nb. 3, pp. 327–339), Nature Publishing Group.

Hubert BANVILLE, Stéphane D’ASCOLI, Simon DAHAN, Jérémy RAPIN, Mar-
lène CAREIL, Yohann BENCHETRIT, Jarod LÉVY, Saarang PANCHAVATI, Antoine RATOUCHNIAK, Mingfang Lucy ZHANG, Elisa CASCARDI, Kate-
lyn BEGANY, Teon BROOKS, Jean-Rémi KING (2026). “**NeuralBench: A Unifying Framework to Benchmark NeuroAI Models**”. *arXiv preprint arXiv:2605.08495*.

Sergio BARBAROSSA, Stefania SARDELLITTI (2020). “**Topological Signal Processing over Simplicial Complexes**”. *TSP* (vol. 68, pp. 2992–3007), IEEE.

Edward BARBEAU (1993). “**Mathematical Fallacies, Flaws, and Flimflam**”. *The College Mathematics Journal* (vol. 24, nb. 2, pp. 149–154), MAA.

Yoshua BENGIO, Aaron COURVILLE, Pascal VINCENT (2013). “**Representation Learning: A Review and New Perspectives**”. *TPAMI* (vol. 35, nb. 8, pp. 1798–1828), IEEE.

Yoshua BENGIO, Réjean DUCHARME, Pascal VINCENT, Christian JAUVIN (2003). “**A Neural Probabilistic Language Model**”. *JMLR* (vol. 3, nb. Feb, pp. 1137–1155).

Roberto BOGHETTI, Jean-Marc ODOBEZ, Jérôme H. KÄMPF (2026). “**Graph Neural Network-Based Surrogate Modeling for Fast and Scalable Simulations of Meshed District Heating Networks**”. *Energy and AI* (art. 100721), Elsevier.

Rishi BOMMASANI, (2021). “**On the Opportunities and Risks of Founda-
tion Models**”. *arXiv preprint arXiv:2108.07258*.

[BOSCH *et al.*, 2026]

 —  74

Christoph BOSCH, John KL WONG, Martin PAULIKAT, Myroslav ZAPUKHLYAK, Bharti ARORA, Manasi AICHMÜLLER-RATNAPARKHE, Jens BAUMANN, Shivani KARN, Rutuja KAMBLE, Swapnil KARNIK, Bhushan KHEDKAR, Chhut SEREY VATHANA, Witali ASWOLINSKIY, Christian AICHMÜLLER (2026). **“Diversity over Scale: Whole-Slide Image Variety Enables H&E Foundation Model Training with Fewer Patches”**. *Journal of Pathology Informatics* (art. 100648), Elsevier.

[BOYLE, 1665]

 —  5

Robert BOYLE (1665). **“An Account of a Very Odd Monstrous Calf”**. *Philosophical Transactions of the Royal Society of London* (vol. 1, nb. 1, pp. 10–10), The Royal Society London.

[BRICKEN *et al.*, 2023]

 —  76

Trenton BRICKEN, Adly TEMPLETON, Joshua BATSON, Brian CHEN, Adam JERMYN, Tom CONERLY, Nick TURNER, Cem ANIL, Carson DENISON, Amanda ASKELL, Robert LASENBY, Yifan WU, Shauna KRAVEC, Nicholas SCHIEFER, Tim MAXWELL, Nicholas JOSEPH, Alex TAMKIN, Karina NGUYEN, Brayden MCLEAN, Josiah E. BURKE, Tristan HUME, Shan CARTER, Tom HENIGHAN, Chris OLAH (2023). **“Towards Monosemanticity: Decomposing Language Models with Dictionary Learning”**. *Transformer Circuits Thread*.

[BRONSTEIN *et al.*, 2017]

 —  24

Michael M. BRONSTEIN, Joan BRUNA, Yann LECUN, Arthur SZLAM, Pierre VANDERGHEYNST (2017). **“Geometric Deep Learning: Going Beyond Euclidean Data”**. *Signal Processing Magazine* (vol. 34, nb. 4, pp. 18–42), IEEE.

[BROWN *et al.*, 2020]

 —  33, 35, 39, 40

Tom BROWN, Benjamin MANN, Nick RYDER, Melanie SUBBIAH, Jared D. KAPLAN, Sandhini VOSS, Gretchen KRUEGER, Tom HENIGHAN, Rewon CHILD, Aditya RAMESH, Daniel ZIEGLER, Jeffrey WU, Clemens WINTER, Chris HESSE, Mark CHEN, Eric SIGLER, Mateusz LITWIN, Scott GRAY, Benjamin CHES, Jack CLARK, Christopher BERNER, Sam MCCANDLISH, Alec RADFORD, Ilya SUTSKEVER, Dario AMODEI (2020). **“Language Models Are Few-Shot Learners”**. *NeurIPS* (vol. 33, pp. 1877–1901).

[CARON *et al.*, 2021]

 —  35

Mathilde CARON, Hugo TOUVRON, Ishan MISRA, Hervé JÉGOU, Julien MAIRAL, Piotr BOJANOWSKI, Armand JOULIN (2021). **“Emerging Properties in Self-Supervised Vision Transformers”**. *ICCV* (pp. 9630–9640), IEEE.

[CHADDAD *et al.*, 2023]

 —  45

Ahmad CHADDAD, Yihang WU, Reem KATEB, Ahmed BOURIDANE (2023). **“Electroencephalography Signal Processing: A Comprehensive Review and Analysis of Methods and Techniques”**. *Sensors* (vol. 23, nb. 14, art. 6434), MDPI.

[CHEN *et al.*, 2026]

 —  41

Lizhang CHEN, Jonathan LI, Chen LIANG, Ni LAO, Qiang LIU (2026). **“Training-Free Looped Transformers”**. *arXiv preprint arXiv:2605.23872*.

[CHEN *et al.*, 2015]

 —  18

Siheng CHEN, Aliaksei SANDRYHAILA, José M. F. MOURA, Jelena KOVAČEVIĆ (2015). **“Discrete Signal Processing on Graphs: Sampling Theory”**. *TSP* (vol. 63, nb. 24, pp. 6510–6523), IEEE.

[CHUNG, 1997]

 —  18

Fan R. K. CHUNG (1997). **“Spectral Graph Theory”**. (vol. 92) *American Mathematical Society*.

[COHEN *et al.*, 2026]

 —  77

Seffi COHEN, Nurit Cohen INGER, Niv GOLDSHLAGER, Bracha SHAPIRA, Lior ROKACH (2026). **“DFPE: A Diverse Fingerprint Ensemble for Enhancing LLM Performance”**. *Findings of the ACL* (pp. 5326–5336).

[COMPEAU *et al.*, 2011]

 —  84

Phillip E. C. COMPEAU, Pavel A. PEVZNER, Glenn TESLER (2011). **“How to Apply de Bruijn Graphs to Genome Assembly”**. *Nature Biotechnology* (vol. 29, nb. 11, pp. 987–991), Nature Publishing Group.

[DANET *et al.*, 1999]

 —  80

[DEFFERRARD *et al.*, 2016]

 —  28

[DÉFOSSEZ *et al.*, 2022]

 —  81

[DELAUNAY, 1934]

 —  57

[DENG *et al.*, 2009]

 —  33

[DEVLIN *et al.*, 2019]

 —  40

[DIDIER *et al.*, 2020]

 —  80

[DIJKSTRA, 1959]

 —  98

[DINESH *et al.*, 2020]

 —  18

[DONG *et al.*, 2016]

 —  24

[DONG *et al.*, 2020]

 —  24

[DOSOVITSKIY *et al.*, 2021]

 —  35

Sandrine DANET, Florence RICHARD, Michèle MONTAYE, Stephanette BEAUCHANT, Brigitte LEMAIRE, Catherine GRAUX, Dominique COTTEL, Nadine MARÉCAUX, Philippe AMOUYEL (1999). “**Unhealthy Effects of Atmospheric Temperature and Pressure on the Occurrence of Myocardial Infarction and Coronary Deaths: A 10-Year Survey: The Lille-World Health Organization MONICA Project (Monitoring Trends and Determinants in Cardiovascular Disease)**”. *Circulation* (vol. 100, nb. 1, art. e1–e7), Lippincott Williams & Wilkins.

Michaël DEFFERRARD, Xavier BRESSON, Pierre VANDERGHEYNST (2016). “**Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering**”. *NeurIPS* (vol. 29).

Alexandre DÉFOSSEZ, Jade COPET, Gabriel SYNNAEVE, Yossi ADI (2022). “**High Fidelity Neural Audio Compression**”. *TMLR*.

Boris DELAUNAY (1934). “**Sur la sphère vide. À la mémoire de Georges Voronoï**”. *Известия Российской академии наук. Серия математическая* (nb. 6, pp. 793–800), Российская академия наук, Математический институт им. В.А. Стеклова Российской.

Jia DENG, Wei DONG, Richard SOCHER, Li-Jia LI, Kai LI, Li FEI-FEI (2009). “**ImageNet: A Large-Scale Hierarchical Image Database**”. *CVPR* (pp. 248–255), IEEE.

Jacob DEVLIN, Ming-Wei CHANG, Kenton LEE, Kristina TOUTANOVA (2019). “**BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding**”. *Conference of the North American Chapter of the ACL: Human Language Technologies* (vol. 1, pp. 4171–4186).

Romain DIDIER, Florent LE VEN, Mehdi OUCHIHA, Pierre Philippe NICOL, Vincent AUFFRET, Chaker OUESLATI, Bahaa NASR, Yannick JOBIC, Antoine NOEL, Michel AIDONIDIS, Edward KOIFMAN, Jacques MANSOURATI, Martine GILARD (2020). “**Analysis of Weather Exposure 7 Days Before Occurrence of ST-Segment Elevation Myocardial Infarction**”. *Archives of Cardiovascular Diseases* (vol. 113, nb. 1, pp. 22–30), Elsevier.

Edsger W. DIJKSTRA (1959). “**A Note on Two Problems in Connexion with Graphs**”. *Numerische Mathematik* (vol. 1, pp. 269–271).

Chinthaka DINESH, Gene CHEUNG, Ivan V. BAJIĆ (2020). “**Point Cloud Denoising via Feature Graph Laplacian Regularization**”. *Transactions on Image Processing* (vol. 29, pp. 4143–4158), IEEE.

Xiaowen DONG, Dorina THANOU, Pascal FROSSARD, Pierre VANDERGHEYNST (2016). “**Learning Laplacian Matrix in Smooth Graph Signal Representations**”. *TSP* (vol. 64, nb. 23, pp. 6160–6173), IEEE.

Xiaowen DONG, Dorina THANOU, Laura TONI, Michael BRONSTEIN, Pascal FROSSARD (2020). “**Graph Signal Processing for Machine Learning: A Review and New Perspectives**”. *Signal Processing Magazine* (vol. 37, nb. 6, pp. 117–127), IEEE.

Alexey DOSOVITSKIY, Lucas BEYER, Alexander KOLESNIKOV, Dirk WEISENBORN, Xiaohua ZHAI, Thomas UNTERTHINER, Mostafa DEHGHANI, Matthias MINDERER, Georg HEIGOLD, Sylvain GELLY, Jakob USZKOREIT, Neil HOULSBY (2021). “**An Image Is Worth 16x16 Words: Transformers for Image Recognition at Scale**”. *ICLR*.

- [DU *et al.*, 2024]
 —  6, 7
- [EGILMEZ *et al.*, 2018]
 —  25
- [EL ABD *et al.*, 2026]
 —  52, 82
- [ELHAGE *et al.*, 2021]
 —  76
- [ENGELS *et al.*, 2026]
 —  44
- [FENG *et al.*, 2025]
 —  84
- [FINN *et al.*, 2017]
 —  35
- [FISCH *et al.*, 2021]
 —  84
- [FLICKE *et al.*, 2025]
 —  7
- [FRIEDMAN *et al.*, 2008]
 —  24
- [GABOR, 1946]
 —  26
- [GANSNER *et al.*, 2004]
 —  154
- [GAO *et al.*, 2025]
 —  76
- Yilun DU, Shuang LI, Antonio TORRALBA, Joshua B. TENENBAUM, Igor MORDATCH (2024). **“Improving Factuality and Reasoning in Language Models Through Multiagent Debate”**. *ICML*.
- Hilmi E. EGILMEZ, Eduardo PAVEZ, Antonio ORTEGA (2018). **“Graph Learning from Filtered Signals: Graph System and Diffusion Kernel Identification”**. *TSIPN (vol. 5, nb. 2, pp. 360–374)*, *IEEE*.
- Sherif EL ABD *et al.* (> 200 co-authors) (2026). **“Gemma 4 Technical Report”**. *arXiv preprint arXiv:2607.02770*.
- Nelson ELHAGE, Neel NANDA, Catherine OLSSON, Tom HENIGHAN, Nicholas JOSEPH, Ben MANN, Amanda ASKELL, Yuntao BAI, Anna CHEN, Tom CONERLY, Nova DASSARMA, Dawn DRAIN, Deep GANGULI, Zac HATFIELD-DODDS, Danny HERNANDEZ, Andy JONES, Jackson KERNION, Liane LOVITT, Kamal NDOUSSE, Dario AMODEI, Tom BROWN, Jack CLARK, Jared KAPLAN, Sam MCCANDLISH, Chris OLAH (2021). **“A Mathematical Framework for Transformer Circuits”**. *Transformer Circuits Thread (vol. 1, nb. 1, art. 12)*, *Anthropic*.
- Joshua ENGELS, Callum MCDUGALL, Bilal CHUGHTAI, Janos KRAMAR, Senthoran RAJAMANOHRAN, Cindy WU, Arthur CONMY, Asic Q. CHEN, Jean TARBOURIECH, Min MA, Brendan O'DONOGHUE, João GABRIEL LOPES DE OLIVEIRA, Rohin SHAH, Neel NANDA (2026). **“How Transparent Is DiffusionGemma?”**. *arXiv preprint arXiv:2606.20560*.
- Haonan FENG, Lang WU, Bingxin ZHAO, Chad HUFF, Jianjun ZHANG, Jia WU, Lifeng LIN, Peng WEI, Chong WU (2025). **“Benchmarking DNA Foundation Models for Genomic and Genetic Tasks”**. *Nature communications (vol. 16, nb. 1, art. 10780)*, *Nature Publishing Group*.
- Chelsea FINN, Pieter ABBEEL, Sergey LEVINE (2017). **“Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks”**. *ICML (pp. 1126–1135)*, *PMLR*.
- Adam FISCH, Tal SCHUSTER, Tommi JAAKKOLA, Regina BARZILAY (2021). **“Few-Shot Conformal Prediction with Auxiliary Tasks”**. *ICML (pp. 3329–3339)*, *PMLR*.
- Markus FLICKE, Glenn ANGRABEIT, Madhav IYENGAR, Vitalii PROTSENKO, Illia SHAKUN, Jovan CICVARIC, Bora KARGI, Haoyu HE, Lukas SCHULER, Lewin SCHOLZ, Kavyanjali AGNIHOTRI, Yong CAO, Andreas GEIGER (2025). **“Scholar Inbox: Personalized Paper Recommendations for Scientists”**. *Annual Meeting of the ACL (vol. 3, pp. 307–317)*.
- Jerome FRIEDMAN, Trevor HASTIE, Robert TIBSHIRANI (2008). **“Sparse Inverse Covariance Estimation with the Graphical Lasso”**. *Biostatistics (vol. 9, nb. 3, pp. 432–441)*, *Oxford University Press*.
- Dennis GABOR (1946). **“Theory of Communication. Part I: The Analysis of Information”**. *Journal of the Institution of Electrical Engineers – Part III: Radio and Communication Engineering (vol. 93, nb. 26, pp. 429–441)*, *IET*.
- Emden R. GANSNER, Yehuda KOREN, Stephen NORTH (2004). **“Graph Drawing by Stress Majorization”**. *International Symposium on Graph Drawing (pp. 239–250)*, *Springer*.
- Leo GAO, Tom DUPRÉ LA TOUR, Henk TILLMAN, Gabriel GOH, Rajan TROLL, Alec RADFORD, Ilya SUTSKEVER, Jan LEIKE, Jeffrey WU (2025). **“Scaling and Evaluating Sparse Autoencoders”**. *ICLR*.

[GHANDEHARI *et al.*, 2021]

 —  29

[GIRAULT, 2015]

 —  18

[GIRDHAR *et al.*, 2023]

 —  6, 40

[GKEREKOS *et al.*, 2019]

 —  65

[GOODFELLOW *et al.*, 2016]

 —  33

[GRIPON *et al.*, 2018]

 —  30

[HAMMOND *et al.*, 2011]

 —  18

[HAN *et al.*, 2026]

 —  24

[HAN *et al.*, 2024]

 —  105

[HE *et al.*, 2016]

 —  37

[HE *et al.*, 2023]

 —  40

[HEISENBERG, 1927]

 —  26

[HO *et al.*, 2020]

 —  40

[HOFFMANN *et al.*, 2022]

 —  34, 40, 74

Mahya GHANDEHARI, Dominique GUILLOT, Kris HOLLINGSWORTH (2021). **“Gabor-Type Frames for Signal Processing on Graphs”**. *Journal of Fourier Analysis and Applications* (vol. 27, nb. 2, art. 25), Springer.

Benjamin GIRAULT (2015). **“Stationary Graph Signals Using an Isometric Graph Translation”**. *EUSIPCO* (pp. 1516–1520), IEEE.

Rohit GIRDHAR, Alaaeldin EL-NOUBY, Zhuang LIU, Mannat SINGH, Kalyan Vasudev ALWALA, Armand JOULIN, Ishan MISRA (2023). **“ImageBind: One Embedding Space to Bind Them All”**. *CVPR* (pp. 15180–15190).

Christos GKEREKOS, Iraklis LAZAKIS, Gerasimos THEOTOKATOS (2019). **“Machine Learning Models for Predicting Ship Main Engine Fuel Oil Consumption: A Comparative Study”**. *Ocean Engineering* (vol. 188, art. 106282), Elsevier.

Ian GOODFELLOW, Yoshua BENGIO, Aaron COURVILLE, Yoshua BENGIO (2016). **“Deep Learning”**. (vol. 1, nb. 2) MIT Press Cambridge.

Vincent GRIPON, Antonio ORTEGA, Benjamin GIRAULT (2018). **“An Inside Look at Deep Neural Networks Using Graph Signal Processing”**. *ITA* (pp. 1–9), IEEE.

David K. HAMMOND, Pierre VANDERGHEYNST, Rémi GRIBONVAL (2011). **“Wavelets on Graphs via Spectral Graph Theory”**. *Applied and Computational Harmonic Analysis* (vol. 30, nb. 2, pp. 129–150), Elsevier.

Yongsheng HAN, Raj Thilak RAJAN, Geert LEUS (2026). **“A Covariance Matching Approach to Graph Topology Identification”**. *arXiv preprint arXiv:2601.15999*.

Zeyu HAN, Chao GAO, Jinyang LIU, Jeff ZHANG, Sai Qian ZHANG (2024). **“Parameter-Efficient Fine-Tuning for Large Models: A Comprehensive Survey”**. *TMLR*.

Kaiming HE, Xiangyu ZHANG, Shaoqing REN, Jian SUN (2016). **“Deep Residual Learning for Image Recognition”**. *CVPR* (pp. 770–778).

Zhengfu HE, Tianxiang SUN, Qiong TANG, Kuanning WANG, Xuan-Jing HUANG, Xipeng QIU (2023). **“DiffusionBERT: Improving Generative Masked Language Models with Diffusion Models”**. *Annual Meeting of the ACL* (vol. 1, pp. 4521–4534).

Werner HEISENBERG (1927). **“Über den Anschaulichen Inhalt der Quantentheoretischen Kinematik und Mechanik”**. *Zeitschrift für Physik* (vol. 43, nb. 3, pp. 172–198), Springer.

Jonathan HO, Ajay JAIN, Pieter ABBEEL (2020). **“Denoising Diffusion Probabilistic Models”**. *NeurIPS* (vol. 33, pp. 6840–6851).

Jordan HOFFMANN, Sebastian BORGEAUD, Arthur MENSCH, Elena BUCHATSKAYA, Trevor CAI, Eliza RUTHERFORD, Diego de Las CASAS, Lisa Anne HENDRICKS, Johannes WELBL, Aidan CLARK, Tom HENNIGAN, Eric NOLAND, Katie MILLICAN, George VAN DEN DRIESSCHE, Bogdan DAMOC, Aurelia GUY, Simon OSINDERO, Karen SIMONYAN, Erich ELSSEN, Jack W. RAE, Oriol VINYALS, Laurent SIFRE (2022). **“Training Compute-Optimal Large Language Models”**. *arXiv preprint arXiv:2203.15556*.

- [HOLLMANN *et al.*, 2022]
 —  52, 83
- [HOOGEBOOM *et al.*, 2021]
 —  40
- [HU *et al.*, 2022]
 —  70, 105
- [HU *et al.*, 2026]
 —  45, 76
- [HU *et al.*, 2025]
 —  76
- [HUANG *et al.*, 2018]
 —  18
- [HUANG *et al.*, 2022]
 —  7
- [HUBEN *et al.*, 2024]
 —  76
- [HUFF, 1954]
 —  101
- [ISUFI *et al.*, 2024]
 —  18
- [JAEGLE *et al.*, 2021]
 —  6, 82
- [JAMHOUR-CHELH *et al.*, 2025]
 —  54
- [JI *et al.*, 2021]
 —  84
- Noah HOLLMANN, Samuel MÜLLER, Katharina EGGENSBERGER, Frank HUTTER (2022). “**TabPFN: A Transformer That Solves Small Tabular Classification Problems in a Second**”. *NeurIPS Table Representation Learning Workshop*.
- Emiel HOOGEBOOM, Didrik NIELSEN, Priyank JAINI, Patrick FORRÉ, Max WELLING (2021). “**Argmax Flows and Multinomial Diffusion: Learning Categorical Distributions**”. *NeurIPS* (vol. 34, pp. 12454–12465).
- Edward J. HU, Yelong SHEN, Phillip WALLIS, Zeyuan ALLEN-ZHU, Yuanzhi LI, Shean WANG, Lu WANG, Weizhu CHEN (2022). “**LoRA: Low-Rank Adaptation of Large Language Models**”. *ICLR*.
- Keya HU, Linlu QIU, Yiyang LU, Hanhong ZHAO, Tianhong LI, Yoon KIM, Jacob ANDREAS, Kaiming HE (2026). “**ELF: Embedded Language Flows**”. *arXiv preprint arXiv:2605.10938*.
- Wenyang HU, Gregory Kang Ruey LAU, Diwen LIU, Jizhuo CHEN, See Kiong NG, Bryan Kian Hsiang LOW (2025). “**Dipper: Diversity in Prompts for Producing Large Language Model Ensembles in Reasoning Tasks**”. *EMNLP* (pp. 35546–35560), *ACL*.
- Weiyu HUANG, Thomas A. W. BOLTON, John D. MEDAGLIA, Danielle S. BASSETT, Alejandro RIBEIRO, Dimitri VAN DE VILLE (2018). “**A Graph Signal Processing Perspective on Functional Brain Imaging**”. *Proceedings of the IEEE* (vol. 106, nb. 5, pp. 868–885).
- Wenlong HUANG, Pieter ABBEEL, Deepak PATHAK, Igor MORDATCH (2022). “**Language Models as Zero-Shot Planners: Extracting Actionable Knowledge for Embodied Agents**”. *ICML* (pp. 9118–9147), *PMLR*.
- Robert HUBEN, Hoagy CUNNINGHAM, Logan SMITH, Aidan EWART, Lee SHARKEY (2024). “**Sparse Autoencoders Find Highly Interpretable Features in Language Models**”. *ICLR* (vol. 2024, pp. 7827–7845).
- Darrell HUFF (1954). “**How to Lie with Statistics**”. *Victor GOLLANCZ*.
- Elvin ISUFI, Fernando GAMA, David I. SHUMAN, Santiago SEGARRA (2024). “**Graph Filters for Signal Processing and Machine Learning on Graphs**”. *TSP* (vol. 72, pp. 4745–4781), *IEEE*.
- Andrew JAEGLE, Sebastian BORGEAUD, Jean-Baptiste ALAYRAC, Carl DOERSCH, Catalin IONESCU, David DING, Skanda KOPPULA, Daniel ZORAN, Andrew BROCK, Evan SHELHAMER, Olivier HÉNAFF, Matthew M. BOTVINICK, Andrew ZISSEMAN, Oriol VINYALS, João CARREIRA (2021). “**Perceiver IO: A General Architecture for Structured Inputs & Outputs**”. *arXiv preprint arXiv:2107.14795*.
- Karim JAMHOUR-CHELH, Dabit ARZAMENDI, Lluís ASMARATS (2025). “**Artificial Intelligence for Risk Prediction in Transcatheter Heart Valve Interventions**”. *Expert Review of Cardiovascular Therapy* (vol. 23, nb. 7, pp. 329–331), *Taylor & Francis*.
- Yanrong JI, Zhihan ZHOU, Han LIU, Ramana V. DAVULURI (2021). “**DNABERT: Pre-Trained Bidirectional Encoder Representations from Transformers Model for DNA-Language in Genome**”. *Bioinformatics* (vol. 37, nb. 15, pp. 2112–2120), *Oxford University Press*.

[JIANG *et al.*, 2025]

 —  36, 82

[JIN *et al.*, 2021]

 —  78

[JIN *et al.*, 2019]

 —  78

[JOHNN *et al.*, 2024]

 —  62

[JOHNSON *et al.*, 2023]

 —  78, 85

[JOURDIN *et al.*, 2024]

 —  64

[KALOFOLIAS, 2016]

 —  24

[KANG *et al.*, 2024]

 —  31

[KAPLAN *et al.*, 2020]

 —  33, 34, 40, 74

[KARLIŃSKI *et al.*, 2015]

 —  79

[KIM *et al.*, 2025]

 —  81

[KIM *et al.*, 2025]

 —  77

Zhen JIANG, Na TANG, Jianlong SUN, Yongzhao ZHAN (2025). “Combining Various Training and Adaptation Algorithms for Ensemble Few-Shot Classification”. *Neural Networks* (vol. 185, art. 107211), Elsevier.

Di JIN, Eileen PAN, Nassim OUFATTOLE, Wei-Hung WENG, Hanyi FANG, Peter SZOLOVITS (2021). “What Disease Does This Patient Have? A Large-Scale Open Domain Question Answering Dataset from Medical Exams”. *Applied Sciences* (vol. 11, nb. 14, art. 6421), MDPI.

Qiao JIN, Bhuwan DHINGRA, Zhengping LIU, William COHEN, Xinghua LU (2019). “PubMedQA: A Dataset for Biomedical Research Question Answering”. *EMNLP* (pp. 2567–2577).

Syu-Ning JOHNN, Victor-Alexandru DARVARIU, Julia HANDL, Jörg KALCSICS (2024). “A Graph Reinforcement Learning Framework for Neural Adaptive Large Neighbourhood Search”. *Computers & Operations Research* (vol. 172, art. 106791), Elsevier.

Alistair E. W. JOHNSON, Lucas BULGARELLI, Lu SHEN, Alvin GAYLES, Ayad SHAMMOUT, Steven HORNG, Tom J. POLLARD, Sicheng HAO, Benjamin MOODY, Brian GOW, Li-Wei H. LEHMAN, Leo A. CELI, Roger G. MARK (2023). “MIMIC-IV, a Freely Accessible Electronic Health Record Dataset”. *Scientific Data* (vol. 10, nb. 1, art. 1), Nature Publishing Group.

Yann JOURDIN, Arwa KHANNOUSSI, Alexandru-Liviu OLTEANU, Patrick MEYER (2024). “A Simulated Annealing Algorithm to Learn an RMP Preference Model”. *Metaheuristics International Conference* (pp. 113–119), Springer.

Vassilis KALOFOLIAS (2016). “How to Learn a Graph from Smooth Signals”. *AISTATS* (pp. 920–929), PMLR.

Zhao KANG, Xuanning XIE, Bingheng LI, Erlin PAN (2024). “CDC: A Simple Framework for Complex Data Clustering”. *TNNLS* (vol. 36, nb. 7, pp. 13177–13188), IEEE.

Jared KAPLAN, Sam McCANDLISH, Tom HENIGHAN, Tom B. BROWN, Benjamin CHESSE, Rewon CHILD, Scott GRAY, Alec RADFORD, Jeffrey WU, Dario AMODEI (2020). “Scaling Laws for Neural Language Models”. *arXiv preprint arXiv:2001.08361*.

Michał KARLIŃSKI, Marcin GLUSZKIEWICZ, Anna CZŁONKOWSKA (2015). “The Accuracy of Prehospital Diagnosis of Acute Cerebrovascular Accidents: An Observational Study”. *Archives of Medical Science* (vol. 11, nb. 3, pp. 530–535), Termedia.

Minje KIM, Jan SKOGLUND (2025). “Neural Speech and Audio Coding: Modern AI Technology Meets Traditional Codecs”. *Signal Processing Magazine* (vol. 41, nb. 6, pp. 85–93), IEEE.

Yubin KIM, Hyewon JEONG, Shan CHEN, Shuyue Stella LI, Chanwoo PARK, Mingyu LU, Kumail ALHAMOUD, Jimin MUN, Cristina GRAU, Minseok JUNG, Rodrigo GAMEIRO, Lizhou FAN, Eugene PARK, Tristan LIN, Joonsik YOON, Wonjin YOON, Maarten SAP, Yulia TSVETKOV, Paul LIANG, Xuhai XU, Xin LIU, Chunjong PARK, Hyeonhoon LEE, Hae Won PARK, Daniel McDUFF, Samir TULEBAEV, Cynthia BREAZEL (2025). “Medical Hallucination in Foundation Models and Their Impact on Healthcare”. *medRxiv*.

- [KIPF *et al.*, 2017]
 —  28, 83
- [KIRILLOV *et al.*, 2023]
 —  15, 38
- [KLANG *et al.*, 2026]
 —  6, 83
- [KLINE *et al.*, 2022]
 —  6
- [KRIZHEVSKY *et al.*, 2012]
 —  12, 33
- [KRONES *et al.*, 2025]
 —  6, 85
- [KULESZA *et al.*, 2011]
 —  74, 84
- [KULESZA *et al.*, 2012]
 —  74
- [KUMAR *et al.*, 2023]
 —  81
- [LAFARGUE *et al.*, 2024]
 —  39
- [LAKE *et al.*, 2010]
 —  24
- [LAM *et al.*, 2026]
 —  80
- Thomas N. KIPF, Max WELLING (2017). “**Semi-Supervised Classification with Graph Convolutional Networks**”. *ICLR*.
- Alexander KIRILLOV, Eric MINTUN, Nikhila RAVI, Hanzi MAO, Chloe ROLLAND, Laura GUSTAFSON, Tete XIAO, Spencer WHITEHEAD, Alexander C. BERG, Wan-Yen LO, Piotr DOLLÁR, Ross GIRSHICK (2023). “**Segment Anything**”. *ICCV* (pp. 3992–4003), *IEEE*.
- Eyal KLANG, Mahmud OMAR, Ganesh RAUT, Reem AGBAREIA, Prem TIMSINA, Robert FREEMAN, Nicholas GAVIN, Lisa STUMP, Alexander W. CHARNEY, Benjamin S. GLICKSBERG, Girish N. NADKARNI (2026). “**Orchestrated Multi Agents Sustain Accuracy Under Clinical-Scale Workloads Compared to a Single Agent**”. *NPJ Health Systems* (vol. 3, nb. 1, art. 23), *Nature Publishing Group*.
- Adrienne KLINE, Hanyin WANG, Yikuan LI, Saya DENNIS, Meghan HUTCH, Zhenxing XU, Fei WANG, Feixiong CHENG, Yuan LUO (2022). “**Multimodal Machine Learning in Precision Health: A Scoping Review**”. *Digital Medicine* (vol. 5, nb. 1, art. 171), *Nature Publishing Group*.
- Alex KRIZHEVSKY, Ilya SUTSKEVER, Geoffrey E. HINTON (2012). “**ImageNet Classification with Deep Convolutional Neural Networks**”. *NeurIPS* (vol. 25).
- Felix KRONES, Umar MARIKKAR, Guy PARSONS, Adam SZMUL, Adam MAHDI (2025). “**Review of Multimodal Machine Learning Approaches in Healthcare**”. *Information Fusion* (vol. 114, art. 102690), *Elsevier*.
- Alex KULESZA, Ben TASKAR (2011). “***k*-DPPs: Fixed-Size Determinantal Point Processes**”. *ICML* (pp. 1193–1200).
- Alex KULESZA, Ben TASKAR (2012). “**Determinantal Point Processes for Machine Learning**”. *Foundations and Trends in Machine Learning* (vol. 5, nb. 2–3, pp. 123–286), *Now Publishers, Inc.*
- Rithesh KUMAR, Prem SEETHARAMAN, Alejandro LUEBS, Ishaan KUMAR, Kundan KUMAR (2023). “**High-Fidelity Audio Compression with Improved RVQGAN**”. *NeurIPS* (vol. 36, pp. 27980–27993).
- Raphael LAFARGUE, Luke A. SMITH, Franck VERMET, Matthias LÖWE, Ian REID, Jack VALMADRE, Vincent GRIPON (2024). “**Oops, I Sampled It Again: Reinterpreting Confidence Intervals in Few-Shot Learning**”. *TMLR*.
- Brenden LAKE, Joshua TENENBAUM (2010). “**Discovering Structure by Learning Sparse Graphs**”. *Annual Meeting of the Cognitive Science Society* (vol. 32, nb. 32).
- Bonnie Yin Ka LAM, Yuan CAI, Huijing ZHENG, Jize WEI, Rufus AKINYEMI, Geert Jan BIESSELS, Hilde van den BRINK, Christopher CHEN, Marco DUERING, Deborah GUSTAFSON, Saima HILAL, Vincent Ming Ho HUI, Rajesh KALARIA, Sang Yun KIM, Maggie Li Man LAM, Frank Erik LEEUW, Ami Sin Man LI, Hugh S. MARKUS, Anna MARSEGLIA, John O’BRIEN, Leonardo PANTONI, Perminder S. SACHDEV, Eric E. SMITH, Joanna M. WARDLAW, Adrian WONG, Ho Ko, Vincent Chung Tong MOK (2026). “**Global Burden of Cerebral Small Vessel Disease Determined from Large MRI Studies: A Systematic Review and Meta-Analysis**”. *International Journal of Stroke*, *SAGE Publications*.

[LE MAGOAROU *et al.*, 2017]

 —  18

[LECUN *et al.*, 2015]

 —  33

[LEUS *et al.*, 2023]

 —  18

[LEWIS *et al.*, 2020]

 —  76

[LI *et al.*, 2024]

 —  6, 34

[LI *et al.*, 2023]

 —  40

[LI *et al.*, 2023]

 —  74

[LI *et al.*, 2025]

 —  40

[LIAQAT *et al.*, 2025]

 —  82

[LIOI *et al.*, 2021]

 —  32

[LIU *et al.*, 2024]

 —  34

[LIU *et al.*, 2023]

 —  82

Luc LE MAGOAROU, Rémi GRIBONVAL, Nicolas TREMBLAY (2017). “**Approximate Fast Graph Fourier Transforms via Multilayer Sparse Approximations**”. *TSIPN* (vol. 4, nb. 2, pp. 407–420), IEEE.

Yann LECUN, Yoshua BENGIO, Geoffrey E. HINTON (2015). “**Deep Learning**”. *Nature* (vol. 521, nb. 7553, pp. 436–444), Nature Publishing Group.

Geert LEUS, Antonio G. MARQUES, José M. F. MOURA, Antonio ORTEGA, David I. SHUMAN (2023). “**Graph Signal Processing: History, Development, Impact, and Outlook**”. *Signal Processing Magazine* (vol. 40, nb. 4, pp. 49–60), IEEE.

Patrick LEWIS, Ethan PEREZ, Aleksandra PIKTUS, Fabio PETRONI, Vladimir KARPUKHIN, Naman GOYAL, Heinrich KÜTTLER, Mike LEWIS, Wentau YIH, Tim ROCKTÄSCHEL, Sebastian RIEDEL, Kiela DOUWE (2020). “**Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks**”. *NeurIPS* (vol. 33, pp. 9459–9474).

Chunyuan LI, Zhe GAN, Zhengyuan YANG, Jianwei YANG, Linjie LI, Lijuan WANG, Jianfeng GAO (2024). “**Multimodal Foundation Models: From Specialists to General-Purpose Assistants**”. *Foundations and Trends in Computer Graphics and Vision* (vol. 16, nb. 1-2, pp. 1–214), Emerald Publishing Limited.

Junnan LI, Dongxu LI, Silvio SAVARESE, Steven HOI (2023). “**BLIP-2: Bootstrapping Language-Image Pre-Training with Frozen Image Encoders and Large Language Models**”. *ICML* (pp. 19730–19742), PMLR.

Kangming LI, Daniel PERSAUD, Kamal CHOUDHARY, Brian DECOST, Michael GREENWOOD, Jason HATTRICK-SIMPERS (2023). “**Exploiting Redundancy in Large Materials Datasets for Efficient Machine Learning with Less Data**”. *Nature Communications* (vol. 14, nb. 1, art. 7283), Nature Publishing Group.

Tianyi LI, Mingda CHEN, Bowei GUO, Zhiqiang SHEN (2025). “**A Survey on Diffusion Language Models**”. *arXiv preprint arXiv:2508.10875*.

Muhammad Irzam LIAQAT, Shah NAWAZ, Muhammad Zaigham ZAHEER, Muhammad Saad SAEED, Hassan SAJJAD, Tom DE SCHEPPER, Karthik NANDAKUMAR, Muhammad Haris KHAN, Ignazio GALLO, Markus SCHEDL (2025). “**Chameleon: A Multimodal Learning Framework Robust to Missing Modalities**”. *International Journal of Multimedia Information Retrieval* (vol. 14, nb. 2, art. 21), Springer.

Giulia LIOI, Vincent GRIPON, Abdelbasset BRAHIM, François ROUSSEAU, Nicolas FARRUGIA (2021). “**Gradients of Connectivity as Graph Fourier Bases of Brain Activity**”. *Network Neuroscience* (vol. 5, nb. 2, pp. 322–336), MIT Press Cambridge.

Fan LIU, Tianshu ZHANG, Wenwen DAI, Chuanyi ZHANG, Wenwen CAI, Xiaocong ZHOU, Delong CHEN (2024). “**Few-Shot Adaptation of Multimodal Foundation Models: A Survey**”. *AI Review* (vol. 57, nb. 10, art. 268), Springer.

Hong LIU, Dong WEI, Donghuan LU, Jinghan SUN, Liansheng WANG, Yefeng ZHENG (2023). “**M3AE: Multimodal Representation Learning for Brain Tumor Segmentation with Missing Modalities**”. *Conference on Artificial Intelligence* (vol. 37, nb. 2, pp. 1657–1665), AAAI.

- [LIU *et al.*, 2023]
 —  33
- [LOGOTHETIS, 2008]
 —  45
- [LOU *et al.*, 2024]
 —  40
- [LOWERRE, 1976]
 —  42
- [LU *et al.*, 2024]
 —  62
- [LU *et al.*, 2024]
 —  39
- [LUDWIG, 1991]
 —  12
- [MA *et al.*, 2024]
 —  80
- [MALLÉGOL, 2023]
 —  64
- [MARCHET *et al.*, 2021]
 —  84
- [MARQUES *et al.*, 2020]
 —  24
- [MASOUD *et al.*, 2024]
 —  62
- [MATEOS *et al.*, 2019]
 —  32
- Pengfei LIU, Weizhe YUAN, Jinlan FU, Zhengbao JIANG, Hiroaki HAYASHI, Graham NEUBIG (2023). “**Pre-Train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing**”. *Computing Surveys* (vol. 55, nb. 9, pp. 1–35), ACM New York.
- Nikos K. LOGOTHETIS (2008). “**What We Can Do and What We Cannot Do with fMRI**”. *Nature* (vol. 453, nb. 7197, pp. 869–878), Nature Publishing Group.
- Aaron LOU, Chenlin MENG, Stefano ERMON (2024). “**Discrete Diffusion Modeling by Estimating the Ratios of the Data Distribution**”. *ICML* (nb. 1333, pp. 32819–32848), PMLR.
- Bruce T. LOWERRE (1976). “**The Harpy Speech Recognition System**”. *Carnegie Mellon University*.
- Shaojun LU, Yongqi WANG, Min KONG, Weizhong WANG, Weimin TAN, Yingxin SONG (2024). “**A Double Deep Q-Network Framework for a Flexible Job Shop Scheduling Problem with Dynamic Job Arrivals and Urgent Job Insertions**”. *Engineering Applications of Artificial Intelligence* (vol. 133, art. 108487), Elsevier.
- Sheng LU, Irina BIGOULAIEVA, Rachneet SACHDEVA, Harish Tayyar MAD-ABUSHI, Iryna GUREVYCH (2024). “**Are Emergent Abilities in Large Language Models Just In-Context Learning?**”. *Annual Meeting of the ACL* (vol. 1, pp. 5098–5139).
- Mark A. LUDWIG (1991). “**The Little Black Book of Computer Viruses: The Basic Technology**”. *American Eagle Publications*.
- Jun MA, Yuting HE, Feifei LI, Lin HAN, Chenyu YOU, Bo WANG (2024). “**Segment Anything in Medical Images**”. *Nature Communications* (vol. 15, nb. 1, art. 654), Nature Publishing Group.
- Antoine MALLÉGOL (2023). “**Optimisation Multi-objectifs de Systèmes Énergétiques Couplés: Méthodes Exactes et Approchées**”. *École Nationale Supérieure Mines-Télécom Atlantique (phdthesis)*.
- Camille MARCHET, Christina BOUCHER, Simon J. PUGLISI, Paul MEDVEDEV, Mikaël SALSON, Rayan CHIKHI (2021). “**Data Structures Based on k -mers for Querying Large Collections of Sequencing Datasets**”. *Genome Research* (vol. 31, nb. 1, pp. 1–12), Cold Spring Harbor Laboratory Press.
- Antonio G. MARQUES, Santiago SEGARRA, Gonzalo MATEOS (2020). “**Signal Processing on Directed Graphs: The Role of Edge Directionality When Processing and Learning from Network Data**”. *Signal Processing Magazine* (vol. 37, nb. 6, pp. 99–116), IEEE.
- Mahmoud MASOUD, Ahmed ABDELHAY, Mohammed ELHENAWY (2024). “**Exploring Combinatorial Problem Solving with Large Language Models: A Case Study on the Travelling Salesman Problem Using GPT-3.5 Turbo**”. *arXiv preprint arXiv:2405.01997*.
- Gonzalo MATEOS, Santiago SEGARRA, Antonio G. MARQUES, Alejandro RIBEIRO (2019). “**Connecting the Dots: Identifying Network Structure via Graph Signal Processing**”. *Signal Processing Magazine* (vol. 36, nb. 3, pp. 16–43), IEEE.

[McMAHAN *et al.*, 2017]

 —  52, 85

[MOEHLING *et al.*, 2021]

 —  80

[MOHAN *et al.*, 2014]

 —  18

[MOHSEN *et al.*, 2022]

 —  6

[MONTELLANO *et al.*, 2021]

 —  79

[MOOR *et al.*, 2023]

 —  6

[MOORE, 1959]

 —  97

[MORTAHEB *et al.*, 2019]

 —  32

[MULLER *et al.*, 2024]

 —  81

[NARANG *et al.*, 2012]

 —  18

[NGUYEN *et al.*, 2022]

 —  52, 85

Brendan McMAHAN, Eider MOORE, Daniel RAMAGE, Seth HAMPSON, Blaise AGÜERA Y ARCAS (2017). “**Communication-Efficient Learning of Deep Networks from Decentralized Data**”. *AISTATS* (pp. 1273–1282), *PMLR*.

Taylor J. MOEHLING, Gihoon CHOI, Lawrence C. DUGAN, Marc SALIT, Robert J. MEAGHER (2021). “**LAMP Diagnostics at the Point-of-Care: Emerging Trends and Perspectives for the Developer Community**”. *Expert Review of Molecular Diagnostics* (vol. 21, nb. 1, pp. 43–61), Taylor & Francis.

Dhanya Menoth MOHAN, Muhammad Tayyab ASIF, Nikola MITROVIC, Justin DAUWELS, Patrick JAILLET (2014). “**Wavelets on Graphs with Application to Transportation Networks**”. *ITSC* (pp. 1707–1712), *IEEE*.

Farida MOHSEN, Hazrat ALI, Nady EL HAJJ, Zubair SHAH (2022). “**Artificial Intelligence-Based Methods for Fusion of Electronic Health Records and Imaging Data**”. *Scientific Reports* (vol. 12, nb. 1, art. 17981), Nature Publishing Group.

Felipe A. MONTELLANO, Kathrin UNGETHÜM, Laura RAMIRO, Aliona NACU, Simon HELLWIG, Felix FLURI, William N. WHITELEY, Alejandro BUSTAMANTE, Joan MONTANER, Peter U. HEUSCHMANN (2021). “**Role of Blood-Based Biomarkers in Ischemic Stroke Prognosis: A Systematic Review**”. *Stroke* (vol. 52, nb. 2, pp. 543–551), Lippincott Williams & Wilkins Hagerstown.

Michael MOOR, Oishi BANERJEE, Zahra Shakeri Hossein ABAD, Harlan M. KRUMHOLZ, Jure LESKOVEC, Eric J. TOPOL, Pranav RAJPURKAR (2023). “**Foundation Models for Generalist Medical Artificial Intelligence**”. *Nature* (vol. 616, nb. 7956, pp. 259–265), Nature Publishing Group.

Edward F. MOORE (1959). “**The Shortest Path Through a Maze**”. *International Symposium on the Theory of Switching* (pp. 285–292), Harvard University Press.

Sepehr MORTAHEB, Jitka ANNEN, Camille CHATELLE, Helena CASSOL, Geraldine MARTENS, Aurore THIBAUT, Olivia GOSSEY, Steven LAUREYS (2019). “**A Graph Signal Processing Approach to Study High Density EEG Signals in Patients with Disorders of Consciousness**”. *EMBC* (pp. 4549–4553), *IEEE*.

Thomas MULLER, Stéphane RAGOT, Laetitia GROS, Pierrick PHILIPPE, Pascal SCALART (2024). “**Speech Quality Evaluation of Neural Audio Codecs**”. *Interspeech* (pp. 4808–4812), *ISCA*.

Sunil K. NARANG, Antonio ORTEGA (2012). “**Perfect Reconstruction Two-Channel Wavelet Filter Banks for Graph Structured Data**”. *TSP* (vol. 60, nb. 6, pp. 2786–2799), *IEEE*.

Dinh C. NGUYEN, Quoc-Viet PHAM, Pubudu N. PATHIRANA, Ming DING, Aruna SENEVIRATNE, Zihuai LIN, Octavia DOBRE, Won-Joo HWANG (2022). “**Federated Learning for Smart Healthcare: A Survey**”. *Computing Surveys* (vol. 55, nb. 3, pp. 1–37), *ACM New York*.

[NGUYEN *et al.*, 2023]

 —  84

Eric NGUYEN, Michael POLI, Marjan FAIZI, Armin W. THOMAS, Michael WORNOW, Callum BIRCH-SYKES, Stefano MASSAROLI, Aman PATEL, Clayton M. RABIDEAU, Yoshua BENGIO, Stefano ERMON, Christopher RÉ, Stephen A. BACCUS, Christophe RÉ (2023). **“HyenaDNA: Long-Range Genomic Sequence Modeling at Single Nucleotide Resolution”**. *NeurIPS* (vol. 36, pp. 43177–43201).

[NGUYEN *et al.*, 2025]

 —  65

Son NGUYEN, Matthieu GADEL, Ke WANG, Jing LI, Xiaocai ZHANG, Siang-Ching KONG, Xiuju FU, Zheng QIN (2025). **“Maritime Decarbonization Through Machine Learning: A Critical Systematic Review of Fuel and Power Prediction Models”**. *Cleaner Logistics and Supply Chain* (vol. 14, art. 100210), Elsevier.

[NIE *et al.*, 2025]

 —  44

Shen NIE, Fengqi ZHU, Zebin YOU, Xiaolu ZHANG, Jingyang OU, Jun HU, Jun ZHOU, Yankai LIN, Ji-Rong WEN, Chongxuan LI (2025). **“Large Language Diffusion Models”**. *arXiv preprint arXiv:2502.09992*.

[NTAIOS *et al.*, 2012]

 —  79

George NTAIOS, Mohamed FAOUZI, Julia FERRARI, Wilfried LANG, Konstantinos VEMMOS, Patrik MICHEL (2012). **“An Integer-Based Score to Predict Functional Outcome in Acute Ischemic Stroke: The ASTRAL Score”**. *Neurology* (vol. 78, nb. 24, pp. 1916–1922), Lippincott Williams & Wilkins Hagerstown.

[OLAH *et al.*, 2020]

 —  44

Chris OLAH, Nick CAMMARATA, Ludwig SCHUBERT, Gabriel GOH, Michael PETROV, Shan CARTER (2020). **“Zoom In: An Introduction to Circuits”**. *Distill* (vol. 5, nb. 3, art. e00024–001).

[ORTEGA *et al.*, 2018]

 —  18, 32, 45

Antonio ORTEGA, Pascal FROSSARD, Jelena KOVAČEVIĆ, José M. F. MOURA, Pierre VANDERGHEYNST (2018). **“Graph Signal Processing: Overview, Challenges, and Applications”**. *Proceedings of the IEEE* (vol. 106, nb. 5, pp. 808–828).

[OUYANG *et al.*, 2022]

 —  34

Long OUYANG, Jeffrey WU, Xu JIANG, Diogo ALMEIDA, Carroll L. WAINWRIGHT, Pamela MISHKIN, Chong ZHANG, Sandhini AGARWAL, Katarina SLAMA, Alex RAY, John SCHULMAN, Jacob HILTON, Fraser KELTON, Luke MILLER, Maddie SIMENS, Amanda ASKELL, Peter WELINDER, Paul CHRISTIANO, Jan LEIKE, Ryan LOWE (2022). **“Training Language Models to Follow Instructions with Human Feedback”**. *NeurIPS* (vol. 35, pp. 27730–27744).

[PAL *et al.*, 2022]

 —  78

Ankit PAL, Logesh Kumar UMAPATHI, Malaikannan SANKARASUBBU (2022). **“MedMCQA: A Large-Scale Multi-Subject Multi-Choice Dataset for Medical Domain Question Answering”**. *Conference on Health, Inference, and Learning* (pp. 248–260), PMLR.

[PAN *et al.*, 2023]

 —  31

Erlin PAN, Zhao KANG (2023). **“High-Order Multi-View Clustering for Generic Data”**. *Information Fusion* (vol. 100, art. 101947), Elsevier.

[PARK *et al.*, 2023]

 —  84

Sangwoo PARK, Kfir M. COHEN, Osvaldo SIMEONE (2023). **“Few-Shot Calibration of Set Predictors via Meta-Learned Cross-Validation-Based Conformal Prediction”**. *TPAMI* (vol. 46, nb. 1, pp. 280–291), IEEE.

[PARKER *et al.*, 2025]

 —  81

Julian PARKER, Anton SMIRNOV, Jordi PONS, C.J. CARR, Zack ZUKOWSKI, Zach EVANS, Xubo LIU (2025). **“Scaling Transformers for Low-Bitrate High-Quality Speech Coding”**. *ICLR* (vol. 2025, pp. 51997–52021).

[PAZ-ARGAMAN *et al.*, 2020]

 —  82

Tzuf PAZ-ARGAMAN, Reut TSARFATY, Gal CHECHIK, Yuval ATZMON (2020). **“ZEST: Zero-shot Learning from Text Descriptions Using Textual Similarity and Visual Summarization”**. *EMNLP* (pp. 569–579), ACL.

- [PEREIRA *et al.*, 2025]
 —  62
- [PERRAUDIN *et al.*, 2017]
 —  18
- [POPPER, 1989]
 —  101
- [PRESS *et al.*, 2021]
 —  62
- [PRONIN *et al.*, 2002]
 —  103
- [RADFORD *et al.*, 2021]
 —  39, 40
- [RADFORD *et al.*, 2023]
 —  69
- [RADFORD *et al.*, 2018]
 —  40
- [RAI *et al.*, 2024]
 —  44
- [RAMSAY *et al.*, 2019]
 —  79
- [REBEDEA *et al.*, 2023]
 —  78
- [REN *et al.*, 2026]
 —  78
- [RETHO *et al.*, 2024]
 —  51
- Pierre PEREIRA, Frederic GIROIRE, Emanuele NATALE (2025). “Solving the Traveling Salesman Problem with Positional Encoding”.
- Nathanaël PERRAUDIN, Pierre VANDERGHEYNST (2017). “Stationary Signal Processing on Graphs”. *TSP* (vol. 65, nb. 13, pp. 3462–3477), *IEEE*.
- Karl Raimund POPPER (1989). “Logik der Forschung”. *Akademie Verlag*.
- Ofir PRESS, Noah A SMITH, Mike LEWIS (2021). “Train Short, Test Long: Attention with Linear Biases Enables Input Length Extrapolation”. *ICLR*.
- Emily PRONIN, Daniel Y. LIN, Lee ROSS (2002). “The Bias Blind Spot: Perceptions of Bias in Self Versus Others”. *Personality and Social Psychology Bulletin* (vol. 28, nb. 3, pp. 369–381), *Sage Publications*.
- Alec RADFORD, Jong Wook KIM, Chris HALLACY, Aditya RAMESH, Gabriel GOH, Sandhini AGARWAL, Girish SASTRY, Amanda ASKELL, Pamela MISHKIN, Jack CLARK, Gretchen KRUEGER, Ilya SUTSKEVER (2021). “Learning Transferable Visual Models from Natural Language Supervision”. *ICML* (pp. 8748–8763).
- Alec RADFORD, Jong Wook KIM, Tao XU, Greg BROCKMAN, Christine MCLEAVEY, Ilya SUTSKEVER (2023). “Robust Speech Recognition via Large-scale Weak Supervision”. *ICML* (pp. 28492–28518), *PMLR*.
- Alec RADFORD, Karthik NARASIMHAN, Tim SALIMANS, Ilya SUTSKEVER (2018). “Improving Language Understanding by Generative Pre-Training”. *OpenAI*.
- Daking RAI, Yilun ZHOU, Shi FENG, Abulhair SAPAROV, Ziyu YAO (2024). “A Practical Review of Mechanistic Interpretability for Transformer-Based Language Models”. *arXiv preprint arXiv:2407.02646*.
- LeeAnn RAMSAY, Marie-Lise QUILLÉ, Cyrille ORSET, Pierre LA GRANGE, Estelle ROUSSELET, Claude FÉREC, Gérald LE GAC, Emmanuelle GÉNIN, Serge TIMSIT (2019). “Blood Transcriptomic Biomarker as a Surrogate of Ischemic Brain Gene Expression”. *Annals of Clinical and Translational Neurology* (vol. 6, nb. 9, pp. 1681–1695), *Wiley Online Library*.
- Traian REBEDEA, Razvan DINU, Makesh Narsimhan SREEDHAR, Christopher PARISIEN, Jonathan COHEN (2023). “NeMo Guardrails: A Toolkit for Controllable and Safe LLM Applications with Programmable Rails”. *EMNLP* (pp. 431–445).
- Zhe REN, Yimeng CHEN, Dandan GUO, Guowei RONG, Tonghui LI, RB XIONG, Qingfeng LAN, Wenyi WANG, Li NANBO, Yibo YANG, Mingchen ZHUGE, Jürgen SCHMIDHUBER (2026). “Self-Improvements in Modern Agentic Systems: A Survey”. *arXiv preprint arXiv:2607.13104*.
- Elise RETHO, Yannick TASSENG, Maëlys CONSIGNY, Lisa LE BOURHIS, Amélie LEBLANC, Aurore JOURDAIN, François-Mathias MERRIEN, François ROUHART, Irina VIAKHIREVA-DOVGANYUK, Philippe GOAS, Caroline LAVENANT, Marie BRUGUET, Serge TIMSIT (2024). “Increased Incidence of Ischemic Stroke in Young: A Population-Based Stroke Registry Study from 2008 to 2018”. *Revue Neurologique* (vol. 180, nb. 3, pp. 182–194), *Elsevier*.

[RINCHI *et al.*, 2026]

 —  69

[RUSSELL *et al.*, 2021]

 —  33

[SAJADIEH *et al.*, 2026]

 —  34

[SANDRYHAILA *et al.*, 2014]

 —  18

[SAVER, 2006]

 —  51, 79

[SCHAUB *et al.*, 2022]

 —  24

[SEDI NZAKUNA *et al.*, 2025]

 —  48

[SEGARRA *et al.*, 2017]

 —  25

[SHAFIPOUR *et al.*, 2019]

 —  25

[SHAFIPOUR *et al.*, 2018]

 —  18

[SHAFIPOUR *et al.*, 2021]

 —  25

[SHAZEER *et al.*, 2017]

 —  44

[SHEN *et al.*, 2026]

 —  106

Bilal RINCHI, Sameer AL-DAHIDI, Mohammad ALRBAI, Osama AYADI (2026). “**Machine Learning for District Heating and Cooling: A Review of Open Problems and Deployment Barriers**”. *Energy and Buildings* (art. 117878), Elsevier.

Stuart RUSSELL, Peter NORVIG (2021). “**Artificial Intelligence: A Modern Approach, 4th Edition**”. *Pearson Education*.

Sha SAJADIEH, Loredana FATTORINI, Raymond PERRAULT, Yolanda GIL, Vanessa PARLI, Lapo SANTARLASCI, Juan PAVA, Nestor MASLEJ, Russ ALTMAN, Erik BRYNJOLFSSON, Carla BRODLEY, Jack CLARK, Virginia DIGNUM, Vipin KUMAR, James LANDAY, Terah LYONS, James MANYIKA, Juan Carlos NIEBLES, Yoav SHOHAM, Elham TABASSI, Russell WALD, Toby WALSH, Dan WELD (2026). “**Artificial Intelligence Index Report 2026**”. *arXiv preprint arXiv:2606.15708*.

Aliaksei SANDRYHAILA, José M. F. MOURA (2014). “**Big Data Analysis with Signal Processing on Graphs: Representation and Processing of Massive Data Sets with Irregular Structure**”. *Signal Processing Magazine* (vol. 31, nb. 5, pp. 80–90), IEEE.

Jeffrey L. SAVER (2006). “**Time Is Brain — Quantified**”. *Stroke* (vol. 37, nb. 1, pp. 263–266), Lippincott Williams & Wilkins.

Michael T. SCHAUB, Jean-Baptiste SEBY, Florian FRANTZEN, T. Mitchell RODDENBERRY, Yu ZHU, Santiago SEGARRA (2022). “**Signal Processing on Simplicial Complexes**”. *Higher-Order Systems* (pp. 301–328), Springer.

Pierre SEDI NZAKUNA, Emanuele D’AURIA, Vincenzo PACIELLO, Vincenzo GALLO, Ernest Nlandu KAMAVUAKO, Aimé LAY-EKUAKILLE, Kyandoghere KYAMAKYA (2025). “**Real-World Evaluation of Deep Learning Decoders for Motor Imagery EEG-Based BCIs**”. *Frontiers in Systems Neuroscience* (vol. 19, art. 1718390), Frontiers Media.

Santiago SEGARRA, Antonio G. MARQUES, Gonzalo MATEOS, Alejandro RIBEIRO (2017). “**Network Topology Inference from Spectral Templates**”. *TSIPN* (vol. 3, nb. 3, pp. 467–483), IEEE.

Rasoul SHAFIPOUR, Abolfazl HASHEMI, Gonzalo MATEOS, Haris VIKALO (2019). “**Online Topology Inference from Streaming Stationary Graph Signals**”. *DSW* (pp. 140–144), IEEE.

Rasoul SHAFIPOUR, Santiago SEGARRA, Antonio G. MARQUES, Gonzalo MATEOS (2018). “**Directed Network Topology Inference via Graph Filter Identification**”. *DSW* (pp. 210–214), IEEE.

Rasoul SHAFIPOUR, Santiago SEGARRA, Antonio G. MARQUES, Gonzalo MATEOS (2021). “**Identifying the Topology of Undirected Networks from Diffused Non-Stationary Graph Signals**”. *Open Journal of Signal Processing* (vol. 2, pp. 171–189), IEEE.

Noam SHAZEER, Azalia MIRHOSEINI, Krzysztof MAZIARZ, Andy DAVIS, Quoc LE, Geoffrey HINTON, Jeff DEAN (2017). “**Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer**”. *ICLR*.

Judy Hanwen SHEN, Alex TAMKIN (2026). “**How AI Impacts Skill Formation**”. *arXiv preprint arXiv:2601.20245*.

[SHIN *et al.*, 2026]

 —  41

[SHOEYBI *et al.*, 2019]

 —  33

[SHUMAN *et al.*, 2013]

 —  18

[SHWARTZ ZIV *et al.*, 2024]

 —  6

[SIMPSON, 1951]

 —  103

[SINGHAL *et al.*, 2023]

 —  7

[SNELL *et al.*, 2017]

 —  35

[SONG *et al.*, 2025]

 —  76

[SONG *et al.*, 2023]

 —  35

[SOUOP *et al.*, 2026]

 —  29

[SRIVASTAVA *et al.*, 2023]

 —  18

Suhyeong SHIN, Sujeong JI, Yeongwook YANG (2026). “**Frozen Inner Looping Is Brittle for Code Generation with BitNet b1.58**”. *Deep Learning for Code: Towards Human-Centered Coding Agents*.

Mohammad SHOEYBI, Mostofa PATWARY, Raul PURI, Patrick LEGRESLEY, Jared CASPER, Bryan CATANZARO (2019). “**Megatron-LM: Training Multi-Billion Parameter Language Models Using Model Parallelism**”. *arXiv preprint arXiv:1909.08053*.

David I. SHUMAN, Sunil K. NARANG, Pascal FROSSARD, Antonio ORTEGA, Pierre VANDERGHEYNST (2013). “**The Emerging Field of Signal Processing on Graphs: Extending High-Dimensional Data Analysis to Networks and Other Irregular Domains**”. *Signal Processing Magazine (vol. 30, nb. 3, pp. 83–98)*, IEEE.

Ravid SHWARTZ ZIV, Yann LECUN (2024). “**To Compress or Not to Compress—Self-Supervised Learning and Information Theory: A Review**”. *Entropy (vol. 26, nb. 3, art. 252)*, MDPI.

Edward H. SIMPSON (1951). “**The Interpretation of Interaction in Contingency Tables**”. *Journal of the Royal Statistical Society (vol. 13, nb. 2, pp. 238–241)*, Wiley Online Library.

Karan SINGHAL, Shekoofeh AZIZI, Tao TU, S. Sara MAHDAVI, Jason WEI, Hyung Won CHUNG, Nathan SCALES, Ajay TANWANI, Heather COLE-LEWIS, Stephen PFOHL, Perry PAYNE, Martin SENEVIRATNE, Paul GAMBLE, Chris KELLY, Nathaneal SCHARLI, Aakanksha CHOWDHERY, Philip MANSFIELD, Blaise AGUERA Y ARCAS, Dale WEBSTER, Greg S. CORRADO, Yossi MATIAS, Katherine CHOU, Juraj GOTTWEIS, Nenad TOMASEV, Yun LIU, Alvin RAJKOMAR, Joelle BARRAL, Christopher SEMTURS, Alan KARTHIKE-SALINGAM, Vivek NATARAJAN (2023). “**Large Language Models Encode Clinical Knowledge**”. *Nature (vol. 620, nb. 7972, pp. 172–180)*, Nature Publishing Group.

Jake SNELL, Kevin SWERSKY, Richard ZEMEL (2017). “**Prototypical Networks for Few-Shot Learning**”. *NeurIPS (vol. 30)*.

Xiangchen SONG, Aashiq MUHAMED, Yujia ZHENG, Lingjing KONG, Zeyu TANG, Mona T DIAB, Virginia SMITH, Kun ZHANG (2025). “**Position: Mechanistic Interpretability Should Prioritize Feature Consistency in SAEs**”. *arXiv preprint arXiv:2505.20254*.

Yisheng SONG, Ting WANG, Puyu CAI, Subrota K. MONDAL, Jyoti Prakash SAHOO (2023). “**A Comprehensive Survey of Few-Shot Learning: Evolution, Applications, Challenges, and Opportunities**”. *Computing Surveys (vol. 55, nb. 135, pp. 1–40)*, ACM New York.

Rigobert Fokam SOUOP, Laurent BITJOKA (2026). “**Harmonic Analysis on Graphs via Isometric Group Embedding: A Canonical Fourier Transform, Shift, and Convolution for Network Signals**”. *arXiv preprint arXiv:2607.13338*.

Divyanshu SRIVASTAVA, Ganesh BAGLER, Vibhor KUMAR (2023). “**Graph Signal Processing on Protein Residue Networks Helps in Studying Its Biophysical Properties**”. *Physica A: Statistical Mechanics and its Applications (vol. 615, art. 128603)*, Elsevier.

[STANKOVIĆ *et al.*, 2021]

 —  28

[SU *et al.*, 2024]

 —  62

[SUN *et al.*, 2024]

 —  6

[SUSHIL *et al.*, 2024]

 —  82

[TAÏGA *et al.*, 2026]

 —  44

[TANG *et al.*, 2024]

 —  6, 7, 78

[THANOU *et al.*, 2017]

 —  25

[THUMMERER, 2026]

 —  69

[TIBSHIRANI *et al.*, 2019]

 —  84

[TREMBLAY *et al.*, 2023]

 —  75

Ljubiša STANKOVIĆ, Jonatan LERGA, Danilo MANDIĆ, Miloš BRAJOVIĆ, Cédrick RICHARD, Miloš DAKOVIĆ (2021). “**From Time-Frequency to Vertex-Frequency and Back**”. *Mathematics* (vol. 9, nb. 12, art. 1407), MDPI.

Jianlin SU, Murtadha AHMED, Yu LU, Shengfeng PAN, Wen BO, Yunfeng LIU (2024). “**RoFormer: Enhanced Transformer with Rotary Position Embedding**”. *Neurocomputing* (vol. 568, art. 127063), Elsevier.

Kai SUN, Siyan XUE, Fuchun SUN, Haoran SUN, Yu LUO, Ling WANG, Siyuan WANG, Na GUO, Lei LIU, Tian ZHAO, Xinzhou WANG, Lei YANG, Shuo JIN, Jun YAN, Jiahong DONG (2024). “**Medical Multimodal Foundation Models in Clinical Diagnosis and Treatment: Applications, Challenges, and Future Directions**”. *arXiv preprint arXiv:2412.02621*.

Madhumita SUSHIL, Travis ZACK, Divneet MANDAIR, Zhiwei ZHENG, Ahmed WALI, Yan-Ning YU, Yuwei QUAN, Dmytro LITUIEV, Atul J. BUTTE (2024). “**A Comparative Study of Large Language Model-Based Zero-Shot Inference and Task-Specific Supervised Classification of Breast Cancer Pathology Reports**”. *Journal of the American Medical Informatics Association* (vol. 31, nb. 10, pp. 2315–2327), Oxford Academic.

Adrien Ali TAÏGA, James ASSIENE, Daniele CALANDRIELLO, Rahma CHAABOUNI, João GANTE, Tamara GLEHN, Nate KEATING, Chris KNUTSEN, Martin KUKLA, Tianlin LIU, Ivan LOBOV, Ofir NABATI, João Gabriel OLIVEIRA, Nicolas PEREZ-NIEVES, Nastasia PRUTIANOVA, Bobak SHAHRIARI, Jean TARBOURIECH, Pavel TYLETSKI, Çağlar ÜNLÜ, Cindy WU, Glenn CAMERON, Jerome CONNOR, Sertan GIRGIN, Maarten GROOTENDORST, Alon LEVKOVITCH, Eliya NACHMANI, Omar SANSEVIERO, Piotr STANCZYK, Quentin BERTHET, Andrew CAMPBELL, Clément CREPY, Valentin DE BORTOLI, Arnaud DOUCET, Romuald ELIE, Alexandre GALASHOV, Klaus GREFF, Alexis JACQ, David RUHE, Yu-Han WU, Sebastian FLENNERHAG, Brendan O'DONOGHUE, George SCRIVENER, Shantanu THAKOOR (2026). “**Dif-fusionGemma Technical Report**”. *arXiv preprint arXiv:2608.00146*.

Xiangru TANG, Anni ZOU, Zhuosheng ZHANG, Ziming LI, Yilun ZHAO, Xingyao ZHANG, Arman COHAN, Mark GERSTEIN (2024). “**MedAgents: Large Language Models as Collaborators for Zero-Shot Medical Reasoning**”. *Findings of the ACL* (pp. 599–621).

Dorina THANOU, Xiaowen DONG, Daniel KRESSNER, Pascal FROSSARD (2017). “**Learning Heat Diffusion Graphs**”. *TSIPN* (vol. 3, nb. 3, pp. 484–499), IEEE.

Tobias THUMMERER (2026). “**Scientific Machine Learning for Hybrid Modeling in Real World Engineering: Validated on Examples from Automotive and Medical Industry**”. *Universität Augsburg (phdthesis)*.

Ryan J. TIBSHIRANI, Rina FOYGE BARBER, Emmanuel CANDÈS, Aaditya RAMDAS (2019). “**Conformal Prediction Under Covariate Shift**”. *NeurIPS* (vol. 32).

Nicolas TREMBLAY, Simon BARTHELMÉ, Konstantin USEVICH, Pierre-Olivier AMBLARD (2023). “**Extended L-Ensembles: A New Representation for Determinantal Point Processes**”. *The Annals of Applied Probability, Institute of Mathematical Statistics*.

[TSATSARONIS *et al.*, 2015]

 —  78

George TSATSARONIS, Georgios BALIKAS, Prodromos MALAKASIOTIS, Ioannis PARTALAS, Matthias ZSCHUNKE, Michael R ALVERS, Dirk WEISSENBORN, Anastasia KRITHARA, Sergios PETRIDIS, Dimitris POLYCHRONOPOULOS, Yannis ALMIRANTIS, John PAVLOPOULOS, Nicolas BASKIOTIS, Patrick GALLINARI, Thierry ARTIÈRES, Axel-Cyrille NGONGA NGOMO, Norman HEINO, Eric GAUSSIER, Liliana BARRIO-ALVERS, Michael SCHROEDER, Ion ANDROUTSOPOULOS, Georgios PALIOURAS (2015). “**An Overview of the BIOASQ Large-Scale Biomedical Semantic Indexing and Question Answering Competition**”. *BMC Bioinformatics* (vol. 16, nb. 1, art. 138), Springer.

[TSUCHIDA *et al.*, 2024]

 —  80

Ami TSUCHIDA, Martin GOUBET, Philippe BOUTINAUD, Iana ASTAFEVA, Victor NOZAIS, Pierre-Yves HERVÉ, Thomas TOURDIAS, Stéphanie DEBETTE, Marc JOLIOT (2024). “**SHiVA-CMB: A Deep-Learning-Based Robust Cerebral Microbleed Segmentation Tool Trained on Multi-Source T2*GRE- and Susceptibility-Weighted MRI**”. *Scientific Reports* (vol. 14, nb. 1, art. 30901), Nature Publishing Group.

[TSUYUKI *et al.*, 2026]

 —  36

Shuhei TSUYUKI, Reda BENSaid, Jérémy MORLIER, Mathieu LÉONARDON, Naoya ONIZAWA, Vincent GRIPON, Takahiro HANYU (2026). “**Efficient Few-Shot Learning for Edge AI via Knowledge Distillation on MobileViT**”. *ICASSP* (pp. 19707–19711), IEEE.

[TURING, 1950]

 —  33

Alan M. TURING (1950). “**Computing Machinery and Intelligence**”. *Mind* (vol. 59, nb. 236, pp. 433–460).

[TURNER *et al.*, 2023]

 —  44

Alexander Matt TURNER, Lisa THIERGART, Gavin LEECH, David UDELL, Juan J. VAZQUEZ, Ulisse MINI, MacDiarmid MONTE (2023). “**Steering Language Models with Activation Engineering**”. *arXiv preprint arXiv:2308.10248*.

[UDANDARAO *et al.*, 2025]

 —  74

Vishaal UDANDARAO, Nikhil PARTHASARATHY, Muhammad Ferjad NAEEM, Talfan EVANS, Samuel ALBANIE, Federico TOMBARI, Yongqin XIAN, Alessio TONIONI, Olivier J. HÉNAFF (2025). “**Active Data Curation Effectively Distills Large-Scale Multimodal Models**”. *CVPR* (pp. 14422–14437).

[VAN ASCH *et al.*, 2010]

 —  51

Charlotte J.J. VAN ASCH, Merel J.A. LUITSE, Gabriël J.E. RINKEL, Ingeborg TWEEL, Ale ALGRA, Catharina J.M. KLIJN (2010). “**Incidence, Case Fatality, and Functional Outcome of Intracerebral Haemorrhage over Time, According to Age, Sex, and Ethnic Origin: A Systematic Review and Meta-Analysis**”. *The Lancet Neurology* (vol. 9, nb. 2, pp. 167–176), Elsevier.

[VASWANI *et al.*, 2017]

 —  41, 81

Ashish VASWANI, Noam SHAZEER, Niki PARMAR, Jakob USZKOREIT, Llion JONES, Aidan N. GOMEZ, Łukasz KAISER, Illia POLOSUKHIN (2017). “**Attention Is All You Need**”. *NeurIPS* (vol. 30).

[VELIČKOVIĆ *et al.*, 2018]

 —  28, 83

Petar VELIČKOVIĆ, Guillem CUCURULL, Arantxa CASANOVA, Adriana ROMERO, Pietro LIÒ, Yoshua BENGIO (2018). “**Graph Attention Networks**”. *ICLR*.

[VILLANI, 2008]

 —  57

Cédric VILLANI (2008). “**Optimal Transport: Old and New**”. *A Series of Comprehensive Studies in Mathematics* (vol. 338), Springer.

[VINYALS *et al.*, 2016]

 —  35

Oriol VINYALS, Charles BLUNDELL, Timothy LILLCRAP, Daan WIERSTRA (2016). “**Matching Networks for One Shot Learning**”. *NeurIPS* (vol. 29).

[VOVK, 2012]

 —  84

Vladimir VOVK (2012). “**Conditional Validity of Inductive Conformal Predictors**”. *ACML* (pp. 475–490), PMLR.

[WANG *et al.*, 2025]

 —  48

[WANG *et al.*, 2024]

 —  34

[WANG *et al.*, 2020]

 —  35

[WANG *et al.*, 2025]

 —  24

[WARDLAW *et al.*, 2013]

 —  80

[WEI *et al.*, 2022]

 —  39, 82

[WEI *et al.*, 2022]

 —  44, 78

[WISEMAN *et al.*, 2016]

 —  42

[WU *et al.*, 2025]

 —  78

[WU *et al.*, 2024]

 —  74

Jiquan WANG, Sha ZHAO, Zhiling LUO, Yangxuan ZHOU, Haiteng JIANG, Shijian LI, Tao LI, Gang PAN (2025). **“CBraMod: A Criss-Cross Brain Foundation Model for EEG Decoding”**. *ICLR* (vol. 2025, pp. 75310–75346).

Lei WANG, Chen MA, Xueyang FENG, Zeyu ZHANG, Hao YANG, Jingsen ZHANG, Zhiyuan CHEN, Jiakai TANG, Xu CHEN, Yankai LIN, Wayne Xin ZHAO, Zhewei WEI, Jirong WEN (2024). **“A Survey on Large Language Model Based Autonomous Agents”**. *Frontiers of Computer Science* (vol. 18, nb. 6, art. 186345), Springer.

Yaqing WANG, Quanming YAO, James T. KWOK, Lionel M. NI (2020). **“Generalizing from a Few Examples: A Survey on Few-Shot Learning”**. *Computing Surveys* (vol. 53, nb. 3, pp. 1–34), ACM New York.

Zehong WANG, Zheyuan LIU, Tianyi MA, Jiazheng LI, Zheyuan ZHANG, Xingbo FU, Yiyang LI, Zhengqing YUAN, Wei SONG, Yijun MA, Qingkai ZENG, Xiusi CHEN, Jianan ZHAO, Jundong LI, Meng JIANG, Pietro LIÒ, Nitesh CHAWLA, Chuxu ZHANG, Yanfang YE (2025). **“Graph Foundation Models: A Comprehensive Survey”**. *arXiv preprint arXiv:2505.15116*.

Joanna M. WARDLAW, Eric E. SMITH, Geert J. BIESSELS, Charlotte CORDONNIER, Franz FAZEKAS, Richard FRAYNE, Richard I. LINDLEY, John T. O'BRIEN, Frederik BARKHOF, Oscar R. BENAVENTE, Sandra E. BLACK, Carol BRAYNE, Monique BRETHER, Hugues CHABRIAT, Charles DECARLI, Frank-Erik LEEUW, Fergus DOUBAL, Marco DUEKING, Nick C. FOX, Steven GREENBERG, Vladimir HACHINSKI, Ingo KILIMANN, Vincent MOK, Robert OOSTENBRUGGE, Leonardo PANTONI, Oliver SPECK, Blossom C. M. STEPHAN, Stefan TEIPEL, Anand VISWANATHAN, David WERRING, Christopher CHEN, Colin SMITH, Mark BUCHEM, Bo NORRVING, Philip B. GORELICK, Martin DICHGANS (2013). **“Neuroimaging Standards for Research into Small Vessel Disease and Its Contribution to Ageing and Neurodegeneration”**. *The Lancet Neurology* (vol. 12, nb. 8, pp. 822–838), Elsevier.

Jason WEI, Yi TAY, Rishi BOMMASANI, Colin RAFFEL, Barret ZOPH, Sebastian BORGEAUD, Dani YOGATAMA, Maarten BOSMA, Denny ZHOU, Donald METZLER, Ed H. CHI, Tatsunori HASHIMOTO, Oriol VINYALS, Percy LIANG, Jeff DEAN, William FEDUS (2022). **“Emergent Abilities of Large Language Models”**. *TMLR*.

Jason WEI, Xuezhi WANG, Dale SCHUURMANS, Maarten BOSMA, Fei XIA, Ed CHI, Quoc V. LE, Denny ZHOU (2022). **“Chain-of-Thought Prompting Elicits Reasoning in Large Language Models”**. *NeurIPS* (vol. 35, pp. 24824–24837).

Sam WISEMAN, Alexander M. RUSH (2016). **“Sequence-to-Sequence Learning as Beam-Search Optimization”**. *Conference on Empirical Methods in Natural Language Processing* (pp. 1296–1306).

Chaoyi WU, Pengcheng QIU, Jinxin LIU, Hongfei GU, Na LI, Ya ZHANG, Yanfeng WANG, Weidi XIE (2025). **“Towards Evaluating and Building Versatile Large Language Models for Medicine”**. *NPJ Digital Medicine* (vol. 8, nb. 1, art. 58), Nature Publishing Group.

Minghao WU, Thuy-Trang VU, Lizhen QU, Gholamreza HAFFARI (2024). **“The Best of Both Worlds: Bridging Quality and Diversity in Data Selection with Bipartite Graph”**. *arXiv preprint arXiv:2410.12458*.

[XIE *et al.*, 2025]

 —  44

[XU *et al.*, 2024]

 —  81

[YAMADA *et al.*, 2020]

 —  24

[YAN *et al.*, 2024]

 —  24

[YANG *et al.*, 2025]

 —  74

[YAO *et al.*, 2023]

 —  77

[YE *et al.*, 2025]

 —  44

[ZEGHIDOUR *et al.*, 2021]

 —  81

[ZHANG *et al.*, 2025]

 —  74

[ZHANG *et al.*, 2019]

 —  24

[ZHAO *et al.*, 2024]

 —  62

[ZHAO *et al.*, 2026]

 —  40

Jinheng XIE, Weijia MAO, Zechen BAI, David Junhao ZHANG, Weihao WANG, Kevin Qinghong LIN, Yuchao GU, Zhijie CHEN, Zhenheng YANG, Mike Zheng SHOU (2025). **“Show-o: One Single Transformer to Unify Multi-modal Understanding and Generation”**. *ICLR* (vol. 2025, pp. 28240–28264).

Yaoxun XU, Hangting CHEN, Jianwei YU, Wei TAN, Rongzhi GU, Shun LEI, Zhiwei LIN, Zhiyong WU (2024). **“MuCodec: Ultra Low-Bitrate Music Codec”**. *arXiv preprint arXiv:2409.13216*.

Koki YAMADA, Yuichi TANAKA, Antonio ORTEGA (2020). **“Time-Varying Graph Learning with Constraints on Graph Temporal Variation”**. *arXiv e-prints*.

Yi YAN, Jiacheng HOU, Zhenjie SONG, Ercan Engin KURUOGLU (2024). **“Signal Processing over Time-Varying Graphs: A Systematic Review”**. *arXiv preprint arXiv:2412.00462*.

Yuming YANG, Yang NAN, Junjie YE, Shihan DOU, Xiao WANG, Shuo LI, Huijie LV, Tao GUI, Qi ZHANG, Xuan-Jing HUANG (2025). **“Measuring Data Diversity for Instruction Tuning: A Systematic Analysis and a Reliable Metric”**. *Annual Meeting of the ACL* (vol. 1, pp. 18530–18549).

Shunyu YAO, Jeffrey ZHAO, Dian YU, Nan DU, Izhak SHAFRAN, Karthik R. NARASIMHAN, Yuan CAO (2023). **“ReAct: Synergizing Reasoning and Acting in Language Models”**. *ICLR*.

Jiacheng YE, Zhihui XIE, Lin ZHENG, Jiahui GAO, Zirui WU, Xin JIANG, Zhenguo LI, Lingpeng KONG (2025). **“Dream 7B: Diffusion Large Language Models”**. *arXiv preprint arXiv:2508.15487*.

Neil ZEGHIDOUR, Alejandro LUEBS, Ahmed OMRAN, Jan SKOGLUND, Marco TAGLIASACCHI (2021). **“SoundStream: An End-to-End Neural Audio Codec”**. *TASLPRO* (vol. 30, pp. 495–507), *IEEE*.

Chi ZHANG, Huaping ZHONG, Kuan ZHANG, Chengliang CHAI, Rui WANG, Xinlin ZHUANG, Tianyi BAI, Qiu JIANTAO, Lei CAO, Ju FAN, Ye YUAN, Guoren WANG, Conghui HE (2025). **“Harnessing Diversity for Important Data Selection in Pretraining Large Language Models”**. *ICLR* (vol. 2025, pp. 72980–73003).

Songyang ZHANG, Zhi DING, Shuguang CUI (2019). **“Introducing Hypergraph Signal Processing: Theoretical Foundation and Practical Applications”**. *Internet of Things Journal* (vol. 7, nb. 1, pp. 639–660), *IEEE*.

Fuqing ZHAO, Changxue ZHUANG, Ling WANG, Chenxin DONG (2024). **“An Iterative Greedy Algorithm with Q-learning Mechanism for the Multiobjective Distributed No-Idle Permutation Flowshop Scheduling”**. *Transactions on Systems, Man, and Cybernetics: Systems* (vol. 54, nb. 5, pp. 3207–3219), *IEEE*.

Wayne Xin ZHAO, Kun ZHOU, Junyi LI, Tianyi TANG, Zican DONG, Yupeng HOU, Beichen ZHANG, Yingqian MIN, Junjie ZHANG, Peiyu LIU, Xiaolei WANG, Yifan DU, Chen YANG, Yushuo CHEN, Zhipeng CHEN, Jinhao JIANG, Ruiyang REN, Yifan LI, Xinyu TANG, Zikang LIU, Yiwen HU, Jian-Yun NIE, Ji-Rong WEN (2026). **“A Survey of Large Language Models”**. *Frontiers of Computer Science* (vol. 20, nb. 12, art. 2012627), *Springer*.

[ZHOU *et al.*, 2025]

 —  44

Chunting ZHOU, Lili YU, Arun BABU, Kushal TIRUMALA, Michihiro YASUNAGA, Leonid SHAMIS, Jacob KAHN, Xuezhe MA, Luke ZETTLEMOYER, Omer LEVY (2025). “**Transfusion: Predict the Next Token and Diffuse Images with One Multi-Modal Model**”. *ICLR (vol. 2025, pp. 6446–6469)*.

[ZHOU *et al.*, 2026]

 —  6, 83

Yucheng ZHOU, Huan ZHENG, Dubing CHEN, Hongji YANG, Wencheng HAN, Jianbing SHEN (2026). “**From Medical LLMs to Versatile Medical Agents: A Comprehensive Survey**”. *Authorea*.

Online Resources



[ACROBAT READER]

 August 22nd, 2026
 5

“**Adobe Acrobat Reader**”. *Adobe*.



[AGENTIC AI TRUST]

 August 21st, 2026
 77

Gabriel Morgan ASAFTEI, Roger ROBERTS, Abby STICHA, Cécile PRINSEN (2026). “**State of AI Trust in 2026: Shifting to the Agentic Era**”. *McKinsey & Company*.



[AI ACT]

 August 27th, 2026
 85

- (2024). “**Regulation 2024/1689 of the European Parliament (AI Act)**”. *European Union*.



[AI SAFETY REPORT]

 August 11th, 2026
 34

Yoshua BENGIO, Stephen CLARE, Carina PRUNKL, Malcolm MURRAY, Maksym ANDRIUSHCHENKO, Ben BUCKNALL, Rishi BOMMASANI, Stephen CASPER, Tom DAVIDSON, Raymond DOUGLAS, David DUVENAUD, Philip FOX, Usman GOHAR, Rose HADSHAR, Anson HO, Tiancheng HU, Cameron JONES, Sayash KAPOOR, Atoosa KASIRZADEH, Sam MANNING, Nestor MASLEJ, Vasilios MAVROUDIS, Conor MCGLYNN, Richard MOULANGE, Jessica NEWMAN, Kwan Yee NG, Patricia PASKOV, Shalaleh RISMANI, Girish SASTRY, Elizabeth SEGER, Scott SINGER, Charlotte STIX, Lucia VELASCO, Nicole WHEELER, Daron ACEMOGLU, Vincent CONITZER, Thomas G. DIETTERICH, Edward W. FELTEN, Fredrik HEINTZ, Geoffrey HINTON, Nick JENNINGS, Susan LEAVY, Teresa LUDERMIR, Vidushi MARDIA, Helen MARGETTS, John McDERMID, Jane MUNGA, Arvind NARAYANAN, Alondra NELSON, Clara NEPPEL, Sarvapali D. RAMCHURN, Stuart RUSSELL, Marietje SCHAAKE, Bernhard SCHÖLKOPF, Alvaro SOTO, Lee TIEDRICH, Gaël VAROQUAUX, Andrew YAO, Ya-Qin ZHANG, Leandro Angelo AGUIRRE, Olubunmi AJALA, Fahad ALBALAWI, Noora ALMALEK, Christian BUSCH, Jonathan COLLAS, André Carlos Ponce de Leon Ferreira CARVALHO, Aman-deep GILL, Ahmet Halit HATIP, Juha HEIKKILÄ, Chris JOHNSON, Gill JOLLY, Ziv KATZIR, Mary N. KEREMA, Hiroaki KITANO, Antonio KRÜGER, Kyoung Mu LEE, José Ramón LÓPEZ PORTILLO, Aoife McLYSAGHT, Olexii MOLCHANOVSKIY, Andrea MONTI, Mona NEMER, Nuria OLIVER, Raquel PEZOA, Audrey PLONK, Balaraman RAVINDRAN, Hammam RIZA, Crystal RUGEGE, Haroon SHEIKH, Denise WONG, Yi ZENG, Liming ZHU, Daniel PRIVITERA, Sören MINDERMAN (2026). “**International AI Safety Report**”.



[ALIVE ENGINE]

 September 4th, 2026
 77

“**Alive Engine Official Website**”. *Cerevial*.



[ARXIV (B.P.)]

👁 March 11th, 2026

📖 7, 15



[CHATGPT]

👁 March 12th, 2026

📖 7, 106



[CLAUDE CODE]

👁 March 12th, 2026

📖 7



[CNIL]

👁 August 18th, 2026

📖 51



[CPGE 2013 REFORM]

👁 March 9th, 2026

📖 92



[CPGE 2021 REFORM]

👁 April 4th, 2026

📖 100



[DARTMOUTH]

👁 August 10th, 2026

📖 33



[DEEPMIND BLOG]

👁 June 17th, 2026

📖 106



[DEFAKATOR]

👁 April 7th, 2026

📖 101



[ENDIVE]

👁 October 15th, 2025

📖 74, 75



[EVINCE]

👁 August 22nd, 2026

📖 5



[FEW-SHOT LEADERBOARD]

👁 August 11th, 2026

📖 37



[G MILGRAM]

👁 April 7th, 2026

📖 101

“ArXiv Search Results for Bastien PASDELOUP”. *ArXiv*.

“ChatGPT Online Chatbot”. *OpenAI*.

“Claude Code”. *Anthropic*.

“Le RGPD appliqué au secteur de la santé”. *CNIL*.

Dominique OBER, Robert CABANE, Frédéric CARLIER, Laurent CARROUÉ, Laurent CHÉNO, Brigitte FLAMAND, Raymond NICODÈME, Bertrand PAJOT, Norbert PERROT (2014). “**Mise en Place des Nouveaux Programmes de Première Année en CPGE**”. *Inspection Générale de l'Éducation Nationale*.

Édouard GEFFRAY, Sophie BROCAS, Brice LANNAUD (2021). “**Bulletin Officiel**”. *Ministère de l'Enseignement Supérieur, de la Recherche et de l'Innovation*.

John MCCARTHY, Marvin L. MINSKY, Nathaniel ROCHESTER, Claude E. SHANNON (1955). “**A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence**”.

Zoubin GHAHRAMANI (2026). “**Measuring the Impact of Learning with AI in Sierra Leone and Beyond**”. *Google DeepMind*.

“YouTube Channel for Defakator”. *YouTube*.

Bastien PASDELOUP (2024). “**Detailed Description of the ENDIVE Project**”.

“Evince Document Viewer”. *GNOME*.

Yaoyao LIU. “**Few-Shot Classification Leaderboard – Few Shot CIFAR-100 Leaderboard (5-class)**”.

“YouTube Channel for G Milgram”. *YouTube*.

-  [GDC]
 August 19th, 2026
 57
-  [GDPR]
 August 18th, 2026
 51
-  [GEMINI]
 March 12th, 2026
 7
-  [GITHUB (B.P.)]
 October 15th, 2025
 7, 15
-  [GITHUB COPILOT]
 June 11th, 2026
 106
-  [GITHUB EDUCATION]
 June 17th, 2026
 106
-  [GNU GPL v3]
 August 20th, 2026
 5
-  [GOOGLE SCHOLAR (B.P.)]
 August 31st, 2026
 49, 153
-  [GOOGLE SCHOLAR LABS]
 March 11th, 2026
 7
-  [GRAPHVIZ]
 June 15th, 2026
 154
-  [HAL (B.P.)]
 October 15th, 2025
 7, 15
-  [HAUTEURS UGA]
 April 7th, 2026
 101
-  [HUGGINGFACE]
 October 28th, 2025
 6
-  [HUGGINGFACE INCIDENT]
 September 2nd, 2026
 34, 78
- “Data Portal – Harmonized Cancer Datasets”. *National Cancer Institute, GDC*.
- (2016). “Regulation 2016/679 of the European Parliament (GDPR)”. *European Union*.
- “Gemini Online Chatbot”. *Google*.
- “GitHub Repositories by Bastien PASDELOUP”. *GitHub*.
- “GitHub Copilot Official Website”. *GitHub*.
- “Empowering the Next Generation of Developers”. *GitHub*.
- (2007). “GNU GPL v3”. *Free Software Foundation*.
- “List of Articles by Bastien PASDELOUP”. *Google*.
- “Google Scholar Labs Online Chatbot”. *Google*.
- “Graphviz NEATO Layout”. *AT&T*.
- “HAL Profile for Bastien PASDELOUP”. *CCSD*.
- Richard MONVOISIN (2014). “Zététique & Autodéfense Intellectuelle”. *Hauteurs UGA*.
- “Models, Any-to-Any Category”. *HuggingFace*.
- (2026). “The HuggingFace Incident and the Road Ahead”. *OpenAI*.

	[IBIS-CT] 👁 August 24 th , 2026 📖 79	Serge TIMSIT (2020). “ Identification of Biomarkers in Ischemic Stroke - Clinical Trial (IBIS-CT) ”. <i>CHU of Brest</i> .
	[MANUSCRIPT TEMPLATE] 👁 August 20 th , 2026 📖 5	Bastien PASDELOUP (2026). “ HDR / PhD L^AT_EX Template ”.
	[MICROSOFT AI] 👁 August 21 st , 2026 📖 77	Kees HERTOIGH (2026). “ Assessing Healthcare’s Agentic AI Readiness: New Research from Microsoft and The Health Management Academy ”. <i>Microsoft</i> .
	[MOONCAT STUDIO] 👁 April 3 rd , 2026 📖 95	“ Mooncat Studio Official Website ”. <i>Mooncat Studio</i> .
	[NERUALL] 👁 April 3 rd , 2026 📖 94	Lauren LEFUMEUR-PASDELOUP. “ Deviant Art Gallery for Neruall ”. <i>Deviant Art</i> .
	[NEYJOUR] 👁 April 3 rd , 2026 📖 94	Christina ROBERTS. “ Deviant Art Gallery for Neyjour ”. <i>Deviant Art</i> .
	[OPEN CODE PAPERS] 👁 August 11 th , 2026 📖 37	“ Open Code Papers – Few-Shot-Image-Classification-On-FC100-5-Way ”.
	[ORBI] 👁 August 26 th , 2026 📖 81	Vincent AUFFRET (2018). “ Observatoire Régional Breton sur l’Infarctus (ORBI) Risk Score ”.
	[OVERLEAF] 👁 August 22 nd , 2026 📖 5	John HAMMERSLEY, John LEES-MILLER. “ Overleaf, Online L^AT_EX Editor ”. <i>WriteL^AT_EX</i> .
	[PYRAT MOOC] 👁 April 3 rd , 2026 📖 95, 97	Vincent GRIPON, Patrick MEYER, Nicolas FARRUGIA, Carlos LASSANCE, Ghouthi BOUKLI-HACENE (2018). “ IMTx: Advanced Algorithmics and Graph Theory with Python ”. <i>edX</i> .
	[PYRAT PYPI] 👁 April 3 rd , 2026 📖 96	Bastien PASDELOUP (2025). “ PyRat ”. <i>PyPI</i> .
	[SCHOLAR INBOX] 👁 March 12 th , 2026 📖 7	Markus FLICKE, Glenn ANGRABEIT, Madhav IYENGAR, Vitalii PROTSENKO, Illia SHAKUN, Jovan CICVARIC, Bora KARGI, Haoyu HE, Lukas SCHULER, Lewin SCHOLZ, Kavyanjali AGNIHOTRI, Yong CAO, Andreas GEIGER. “ Scholar Inbox: Personalized Paper Recommendations for Scientists ”. <i>University of Tübingen</i> .
	[TAVI] 👁 August 19 th , 2026 📖 53	“ TAVI ”. <i>Parole & Réaction</i> .



[TGRS]
👁 August 10th, 2026
📖 34

“The Global Research Space”.



[VSCode PDF Viewer]
👁 August 22nd, 2026
📖 5

“VSCode PDF Viewer”. *Mathematics Inc.*



[WIKIPEDIA (STROKE)]
👁 August 19th, 2026
📖 52

“Stroke”. *Wikipedia.*



[WORLD WIDE WEB]
👁 November 3rd, 2025
📖 5

Tim BERNERS-LEE (1989). “World Wide Web”. *CERN.*

ACRONYMS



AAAI		Association for the A dvancement of AI
AAC		A dvanced A udio C oding
ACL		A ssociation for C omputational L inguistics
ACM		A ssociation for C omputing M achinery
ACML		A sian C onference on ML
ADASYN		AD Aptive SYN thetic Sampling
ADEME		A gence D e l' E nvironnement et de la M aitrise de l' E nergie
AFPy		A ssociation F rancophone P ython
AI		A rtificial I ntelligence
AI@IMT		AI@IMT pour une IA Responsable dans la Société et dans l'Industrie
AISTATS		I nternational Conference on AI and STAT istic S
ALF		A mdahl's L aw is F orever
ALiBi		A ttention With L inear B ias
ANR		A gence N ationale de la R echerche
ANSI		A merican N ational S tandards I nstitute
AP-HP		A ssistance P ublique - H ôpitaux de P aris
ARC		AI ₂ R easoning C hallenge
ARED		A llocations de R echerche D octorale
ARR		A ssociation for Computational Linguistics R olling R eviews
ARS		A gence R égionale de la S anté
ASCII		A merican S tandard C ode for I nformation I nterchange
ASR		A utomatic S peech R ecognition
ASTRAL		A cute ST roke R egistry and A nalysis of L ausanne
ASTRID		A ccompagnement S pécifique des T ravaux de R echerches d' I ntérêt D éfense
AT&T		A merican T elephone & T elegraph C ompany
AUC		A rea U nder the C urve
BCI		B rain- C omputer I nterface

BEETL	Benchmarks for EEG Transfer Learning
BERT	Bidirectional Encoder Representations from Transformers
BFS	Breadth First Search
BIOASQ	BIOmedical Semantic Indexing and Question Answering
BIOT	BIOsignal Transformer
BLIP	Bootstrapping Language-Image Pre-training
BMC	BioMed Central
BOLD	Blood-Oxygen-Level Dependent
BRAIN	BRoader Artificial INtelligence
BRAIn	Better Representations for Artificial Intelligence
BSR	Brest Stroke Registry
BTS	Brevet de Technicien Supérieur
CBraMod	Criss-Cross Brain Model
CCSD	Centre pour la Communication Scientifique Directe
CCTA	Coronary Computed Tomography Angiography
CDC	Complex Data Clustering
CDE	Contrat Doctoral Établissement
CDV	Coup De Venst
CEA-LIST	Commissariat à l'Énergie Atomique et aux Énergies Alternatives - Laboratoire d'Intégration des Systèmes et des Technologies
CERN	Conseil Européen pour la Recherche Nucléaire
CERV	Centre Européen de Réalité Virtuelle
CFD	Computational Fluid Dynamics
cHSR	computational Hyperemic Stenosis Resistance
CHU	Centre Hospitalier Universitaire
CHUV	CHU Vaudois
CICU	Cardiac Intensive Care Unit
CIFAR	Canadian Institute For Advanced Research
CIFRE	Conventions Industrielles de Formation par la REcherche
CISTER	Research Centre In Real-Time and Embedded Computing Systems
CLAP	Contrastive Language-Audio Pretraining
CLIP	Contrastive Language-Image Pre-Training
CM	Cours Moyen
CMB	Cerebral MicroBleeds
CNIL	Commission Nationale de l'Informatique et des Libertés
CNN	Convolutional Neural Network
CNRS	Centre National de la Recherche Scientifique
CNU	Conseil National des Universités

COCOLA		Coherence- O riented C ontrastive L earning of Musical A udio Representations
CODEVSI		C onception et D éveloppement de S ystème d'Ingénierie
CODHYPIA		C oDesign of H YPerfrequency Components Using A I
COMPSAC		C OMPuters, Software, and A pplications C onference
CORS		C anadian O perational R esearch S ociety
CORTICO		C ollectif pour la R echerche T ransdisciplinaire sur les I nterfaces C erveau- O rdinateur
CPGE		C lasses P réparatoires aux G randes É coles
CRIStAL		C entre de R echerche en I nformatique S ignal et A utomatique de L ille
CSBrain		C ross-Scale S patiotemporal B rain Foundation Model
cSVD		c erebral S mall V essel D isease
CT		C omputed T omography
CTEPD		C hronic T hrombo E mbolic P ulmonary D isease
CTEPH		C hronic T hrombo E mbolic P ulmonary H ypertension
CUB		C altech U CS D B irds
CV		C urriculum V itæ
CVPR		C omputer V ision and P attern R ecognition
D ₃ PM		D iscrete D enoising D iffusion P robabilistic M odel
D ₅ P ₄		P artition D PP for D iversity in P arallel D iscrete D iffusion D ecoding
DA ₂ PL		From Multiple Criteria D ecision A id T o P reference L earning
DAC		D escript A udio C odec
DARFI3D		D esign A vançé R adio- F réquence en I mpression 3D par I A
DECIDE		D ECision A id and I nformation D iscov E ry
DFPE		D iverse F inger P rint E nsemble
DFS		D epth F irst S earch
DGA/MI		D irection G énérale de l' A rmement / M aitrise de l' I nformation
DHN		D istrict H eating N etwork
DigIC		D igitalisation, I nnovation et C hangement
Dipper		D iversity in p rompts for p roducing LLM ensembles in r easoning T asks
DNA		D eoxyribo N ucleic A cid
DNN		D eep N eural N etwork
DPP		D eterminantal P oint P rocess
DS		D ata S cience
DSBM		D epressive S uicidal B lack M etal
DSW		D S W orkshop
EASY		E nsemble A ugmented- S hot Y - S haped L earning
ECAM		É cole C atholique des A rts et M étiers

ECE		E lectrical and C omputer E ngineering
ECG		E lectro C ardio G raphy
ECOS		International Conference on E fficiency, C ost, O ptimization, S imulation and E nvironmental I mpact of E nergy S ystems
EDF		É lectricité D e F rance
EEG		E lectro E ncephalo G raphy
EEG ₂ REP		EEG To REP resentation
EGC		E xtraction et G estion des C onnaissances
EJOR		E uropean J ournal of O perational R esearch
ELASTIQ		EEG-L anguage A lignment with S emantic T ask I nstruction and Q uerying
ELF		E mbedded L anguage F lows
EMBC		Annual International Conference of the E ngineering in M edicine and B iology Society
EMNLP		E mpirical M ethods in NLP
ENDIVE		EN couraging DIVER sity in FSL
ENIB		É cole N ationale d'Ingénieurs de B rest
ENS		É cole N ormale S upérieure
EPFL		É cole P olytechnique F édérale de L ausanne
ESC		E uropean S ociety of C ardiology
EURASIP		EU ropean A ssociation for S ignal P rocessing
EURO		EU ropean Conference on O perational R esearch
EUSIPCO		EU ropean S ignal P rocessing C onference
FabLab		F abrication L aboratory
FHIR		F ast H ealthcare I nteroperability R esources
FICUS		F ew-Shot I mage C lassification with U nsupervised S egmentation
FIP		F ormation d'Ingénieur sous Contrat de P rofessionnalisation
FISE		F ormation d'Ingénieur sous S tatut d'Étudiant
fMRI		f unctional M agnetic R esonance I maging
FSL		F ew-Shot L earning
FUSCHIA		FUS ion de C omposantes H étérogènes par IA
GAN		G enerative A dversarial N etwork
GAT		G raph AT tention network
GCN		G raph C onvolutional N etwork
GCS		G lasgow C oma S cale
GDC		G enomic D ata C ommons
GDPR		G eneral D ata P rotection R egulation
GdR		G roupement d e R echerche

GFT	Graph Fourier Transform
GFT^{-1}	Inverse GFT
GIPSA-Lab	Laboratoire de Grenoble Images Parole Signal Automatique
GlobalSIP	Global Conference on Signal and Information Processing
GNOME	GNU Network Object Model Environment
GNU	GNU's Not Unix
GPL	General Public License
GPT	Generative Pre-trained Transformer
GPU	Graphics Processing Unit
GRETSI	Groupeement de Recherche En Traitement du Signal et des Images
GRU	Gated Recurrent Unit
GSM	Grade School Math
GSP	Graph Signal Processing
GUI	Graphical User Interface
H&E	Hematoxylin & Eosin
HAL	Hyper Articles en Ligne
HDR	Habilitation à Diriger des Recherches
HellaSwag	Harder endings, longer Contexts and low-Shot activities for Situations with adversarial generations
HSR	Hyperemic Stenosis Resistance
<i>i.i.d.</i>	<i>i</i> ndependent and <i>i</i> dentically <i>d</i> istributed
IBIS-CT	Identification of Biomarkers in Ischemic Stroke - Clinical Trial
ICASSP	International Conference on Acoustics, Speech and Signal Processing
ICCV	International Conference on Computer Vision
ICH	IntraCerebral Hemorrhage
ICIP	International Conference on Image Processing
ICLR	International Conference on Learning Representations
ICML	International Conference on ML
ICTAI	International Conference on Tools with AI
IDE	Integrated Development Environment
IEEE	Institute of Electrical and Electronics Engineers
IET	Institution of Engineering and Technology
Ifremer	Institut français de recherche pour l'exploitation de la mer
IFSA	International Frequency Sensor Association
IG	Iterated Greedy
IHU	Institut Hospitalo-Universitaire
ILS	Iterated Local Search

IMT		Institut M ines T élécom ⁽²⁾
Inria		I nstitut n ational de r echerche en i nformatique et en a utomatique
Inserm		I nstitut n ational de la s anté et de la r echerche m édicale
IoT		I nternet of T hings
IoU		I ntersection-over- U nion
IRISA		I nstitut de R echerche en I nformatique et S ystèmes A léatoires
ISCA		I nternational S peech C ommunication A ssociation
ISEN		I nstitut S upérieur de l'Électronique et du N umérique
ISIS		I nformation, S ignal, I mage, V iSion
ITA		I nformation T heory and A pplications W orkshop
ITM®		I nstitut T rans- M éthodologique®
ITSC		I ntelligent T ransportation S ystems C onference
JCJC		J eunes C hercheuses et J eunes C hercheurs
JMLR		J ournal of M L R esearch
Lab-STICC		L aboratoire des S ciences et T echniques de l' I nformation, de la C ommunication et de la C onnaissance
LaBraM		L arge B rain M odel
LAMBADA		L anguage M odeling B roadened to A ccount for D iscourse A spects
LAMP		L oop-Mediated Isothermal A mplification
LaTIM		L aboratoire de T raitement de l' I nformation M édicale
LIGM		L aboratoire d' I nformatique G aspard- M onge
LLaDA		L arge L anguage D iffusion with M Aasking
LLaMA		L arge L anguage M odel M eta A I
LLM		L arge L anguage M odel
LM		L anguage M odel
LoRA		L ow- R ank A daptation
LSTM		L ong S hort- T erm M emory
LTS		L aboratoire de T raitement du S ignal
LUNA		L atent U nified N etwork A rchitecture
M ₃ AE		M ulti M odal M asked A uto E ncoder
MAA		M athematical A ssociation of A merica
MathSTIC		M athématiques et S ciences et T echnologies de l' I nformation et de la C ommunication en Bretagne Océane
MATLAB		M ATrix L ABoratory
MCDA		M ulti- C riteria D ecision A nalysis
MCE		M athematical and C omputational E ngineering

(2) Well... not technically always an acronym:  La marque IMT Atlantique : halte à la confusion !

MD		M edicinæ D octor
MDPI		M ultidisciplinary D igital P ublishing I nstitute
MedMCQA		M edical M ultiple- C hoice Q uestion A nswering
MedQA		M edical Q uestion A nswering
MedS		M edical S uper
MEG		M agneto E ncephalo G raphy
MERT		M usic U nd E Rstanding M odel with L arge- S cale S elf- S upervised T raining
MI		M otor I magery
MILP		M ixed- I nteger L inear P rogramming
MIMIC		M edical I nformation M art for I ntensive C are
MinLA		M inimum L inear A rrangement
MIT		M assachusetts I nstitute of T echnology
ML		M achine L earning
MLP		M ulti- L ayer P erceptron
MMLU		M assive M ultitask L anguage U nderstanding
MMORPG		M assively M ultiplayer O nline R ole P laying G ame
MNIST		M odified N ational I nstitute of S tandards and T echnology
MONICA		M ONItoring T rends and D eterminants in C Ardiovascular D isease
MOOC		M assive O pen O nline C ourse
MOVIOLIA		M ONtage de V idéOs d'Entreprise Assisté par L 'IA
MP		M etaphor- P araphrase
MP ₃		M oving P icture E xperts G roup A udio L ayer 3
MPC		M odel P redictive C ontrol
MRI		M agnetic R esonance I maging
MRMR		M inimum R edundancy M aximum R elevance
MtG		M agic: t he G athering
MUNCH		M etaphor U nderstanding C hallenge
MVA		M athématiques, V ision, A pprentissage
NeurIPS		A nnual C onference on N eural I nformation P rocessing S ystems
NeurIPT		N eural I nterfaces with a P re- T rained T ransformer
NeuroLM		N euroscience L anguage M odel
NIHSS		N ational I nstitutes of H ealth S troke S cale
NLP		N atural L anguage P rocessing
NP		N on-deterministic P olynomial
NPJ		N ature P artner J ournals
OLA		I nternational C onference on O ptimization and L eArning
OOP		O bject- O riented P rogramming

ORBI	Observatoire Régional Breton sur l'Infarctus
OSI	Oscillatory Shear Index
PascalVOC	Pascal Visual Object Classes
PDF	Portable Document Format
PEFT	Parameter-Efficient Fine-Tuning
PFSP	Permutation Flowshop Scheduling Problem
PH	Praticien Hospitalier
PhD	Philosophiæ Doctor
PI	Principal Investigator
PME	Petites et Moyennes Entreprises
PMLR	Proceedings of ML Research
PR	Precision-Recall
PRACom	Pôle de Recherche Avancée en Communications
PRAMeL	Projects on Recent Advances in ML
PRCE	Projet de Recherche Collaborative – Entreprise
PROCOM	Concevoir et Gérer un PROjet COMplexe
PRONTO	PROjet d'Ingénieur dans un MONde en TransitiONS
PU	Professeur des Universités
PubMedQA	PubMed Abstracts Question Answering
PyConFr	Convention Annuelle de la Communauté Python Francophone
PyPI	The Python Package Index
Q&A	Questions & Answers
QPES	Questions de Pédagogies dans l'Enseignement Supérieur
R0-FoMo	Robustness of Few-Shot and Zero-Shot Learning in Large Foundation Models
RAG	Retrieval-Augmented Generation
RCFE	Réseau Romand de Conseil Formation et Évaluation pour l'Enseignement Universitaire
ReAct	Reasoning and Acting
ReLU	Rectified Linear Unit
ResNet	Residual Neural Network
REVE	Representation for Electroencephalography with Versatile Embeddings
RF	Random Forest
RFE	Recursive Feature Elimination
RIA	Research and Innovation Actions
RL	Reinforcement Learning
RMP	Ranking with Multiple Profiles
RNA	RiboNucleic Acid

RNN	Recurrent Neural Network
RNTI	Revue des Nouvelles Technologies de l'Information
ROADEF	Société Française de Recherche Opérationnelle et d'Aide à la Décision
ROC	Receiver Operating Characteristic
RoFormer	Rotary TransFormer
RoPE	Rotary Positional Embedding
RVQGAN	Residual Vector Quantization GAN
SAE	Sparse AutoEncoder
SAGE	Sara A. GEorge
SAM	Segment Anything Model
SEDD	Score Entropy Discrete Diffusion
SequoIA	AI for Trust and Security
SHAP	SHapley Additive ExPlanation
SIGMAS	Cerebral Small Vessel Disease Imaging, Genomics and ML Analysis in Clinical Stroke
SIGMETRICS	Special Interest Group on Measurement and Evaluation (METRICS)
SimCLR	Simple Framework for Contrastive Learning of Visual Representations
SMOTE	Synthetic Minority Oversampling TEchnique
SNOMED CT	Systematized NOmenclature of MEDicine - Clinical Terms
SPL	Signal Processing Letters
SRMP	Simple Ranking Method Using Reference Profiles
SSL	Self-Supervised Learning
STEMI	ST-Segment Elevation Myocardial Infarction
STRATIF-AI	Continuous STRATIFication for Improved Prevention, Treatment, and Rehabilitation of Stroke Patients Using Digital Twins and AI
STRATUM	Stroke TRANscripTomic Utility Markers
SVM	Support-Vector Machine
T2*GRE	T2*-Weighted GRradient Echo Sequences
TabPFN	Tabular Prior-Data Fitted Network
TAF	Thématique d'Approfondissement de la Formation d'Ingénieur
TASLPRO	Transactions on Audio, Speech, and Language PROcessing
TaupIC	Concours d'Informatique des Taupins
TAVI	Transcatheter Aortic Valve Implantation
TAWSS	Time-Averaged Wall Shear Stress
TCGA	The Cancer Genome Atlas
TGRS	The Global Research Space
TIPE	Travaux d'Initiative Personnelle Encadrés
TMLR	Transactions on ML Research

TNNLS	Transactions on N eural N etworks and L earning S ystems
TPAMI	Transactions on P attern A nalysis and M achine I ntelligence
TSIPN	Transactions on S ignal and I nformation P rocessing over N etworks
TSN	Télécom et S ociété N umérique
TSP	T raveling S alesperson P roblem
TSP	Transactions on S ignal P rocessing
TSPLIB	TSP LIB rary
UBL	Université B retagne L oire
UBO	Université de B retagne O ccidentale
UCSD	University of California, S an D iago
UEB	Université E uropéenne de B retagne
UGA	Université G renoble A lpes
UML	U nified M odeling L anguage
UMR	U nité M ixte de R cherche
ViT	V ision T ransformer
VLM	V ision- L anguage M odel
VSCoDe	V isual S tudio C ode
wav2vec	w aveform t o v ector
WinoGrande	W inograd Schema Challenge at G rande Scale
ZEST	Z ero- S hot Learning from T ext Descriptions

Modern AI advances fastest where data is plentiful, shareable and benchmarked. Hospital data is none of those things. This manuscript carries advances from fundamental AI research to healthcare.

The work that underpins this manuscript revolves around one question: how can we take advantage of the structure or variability of the data, rather than of its amount? GSP was a first answer, exploiting the geometry of the domain signals live on; FSL a second, exploiting the geometry of representations; and diversity in the data or representations, a third. More recently, LLMs and foundation models shifted this question to how best to exploit powerful representations of data from many modalities.

This methodological evolution and the projects that arise from it are the main subjects of this manuscript. I develop it through past and ongoing projects, in collaboration with researchers and clinicians in neurology, cardiology and oncology, each project tackling the question from a different angle.

Leveraging these previous works, I am now convinced that the (mid-term) future of AI for healthcare is not likely to be one single very large model producing an opaque vector for a patient, but rather a diversity of specialized models, deployed where the data lives, exchanging information in natural language that an orchestrator can assemble and a practitioner can read, question and contest.

The second part of this manuscript is dedicated to the axes of research I plan to pursue in the future, in order to contribute to this vision. It lays out a research plan that gives a lot of space to multimodality, agentic systems, and LLMs.

Finally, this manuscript also includes a brief description of some chosen teaching activities, which I consider an important part of my interests and activities as an Associate Professor at IMT Atlantique.

