Transferring Advances from Modern AI to Healthcare

Habilitation to Supervise Research

Bastien Pasdeloup

November 12th, 2026

Université de Bretagne Occidentale
IMT Atlantique
Bretagne-Pays de la Loire
École Mines-Télécom

Outline



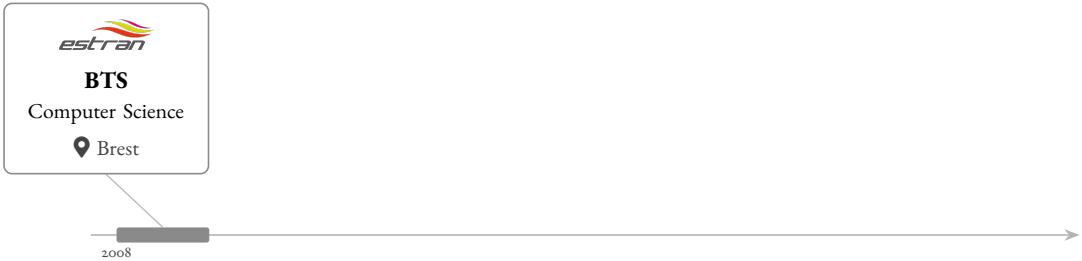




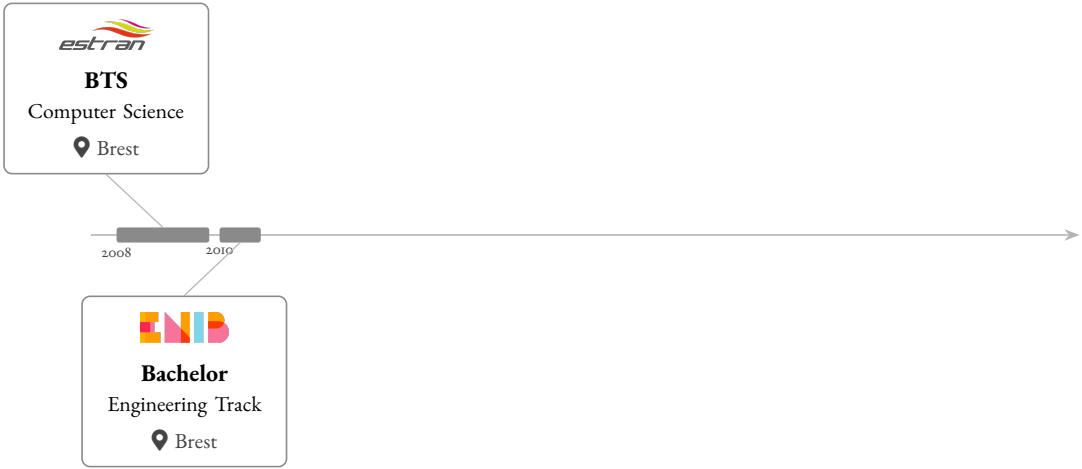


Introduction

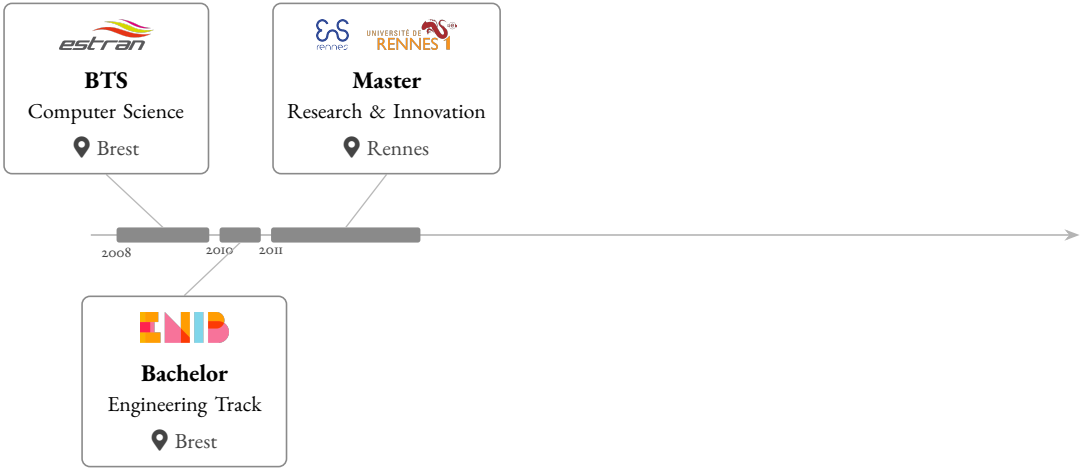
Academic path



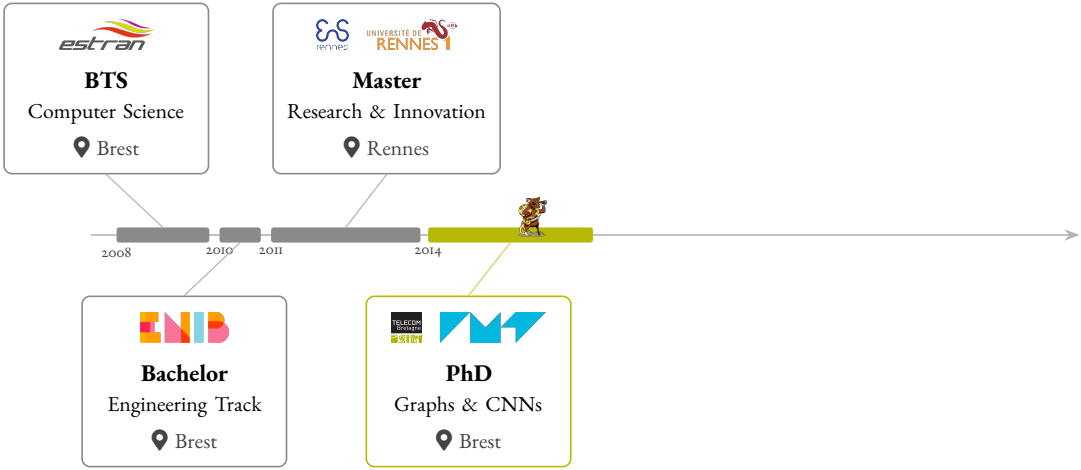
Academic path



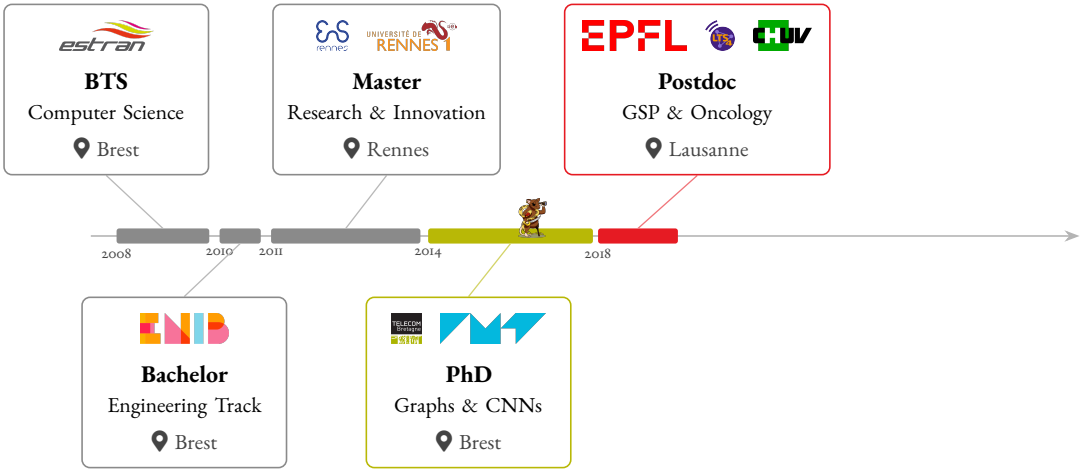
Academic path



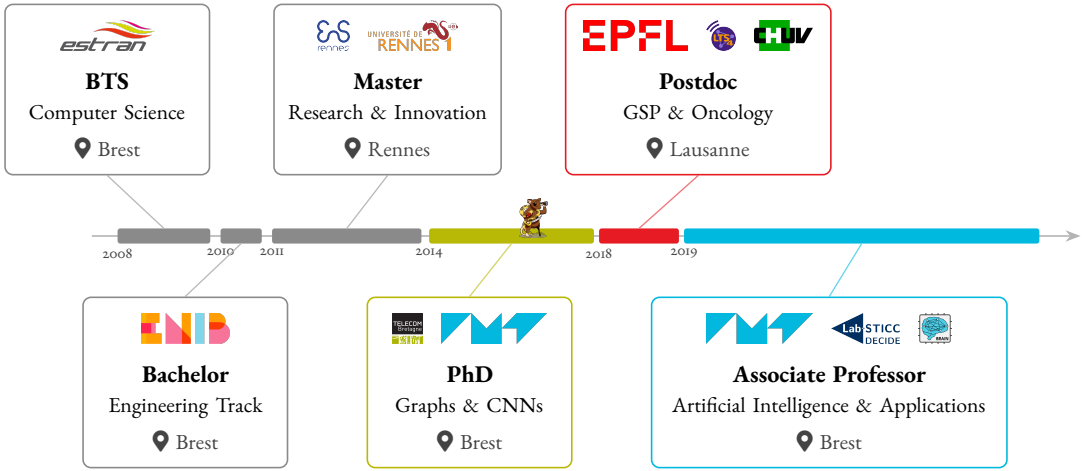
Academic path



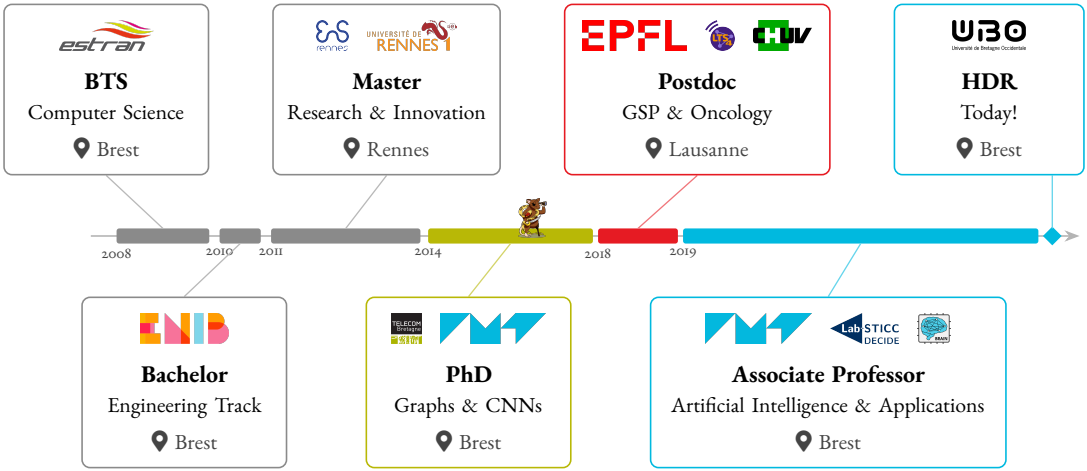
Academic path



Academic path



Academic path



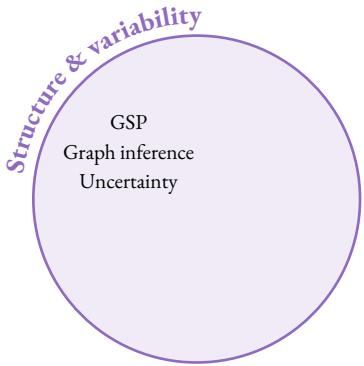
My research in one diagram

Central question

How to get the best out of data through its structure and variability, rather than its sheer amount?

- Follow the advances in core AI
- Transfer them to healthcare

My research in one diagram

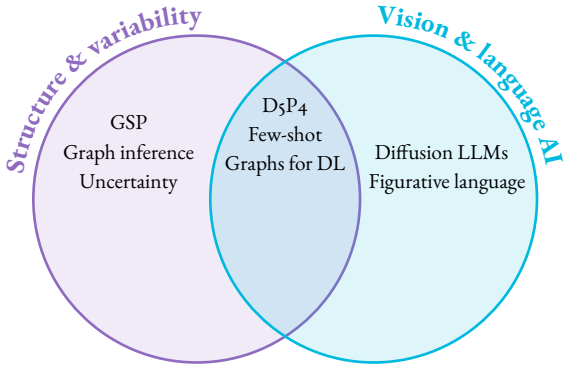


Central question

How to get the best out of data through its structure and variability, rather than its sheer amount?

- Follow the advances in core AI
- Transfer them to healthcare

My research in one diagram

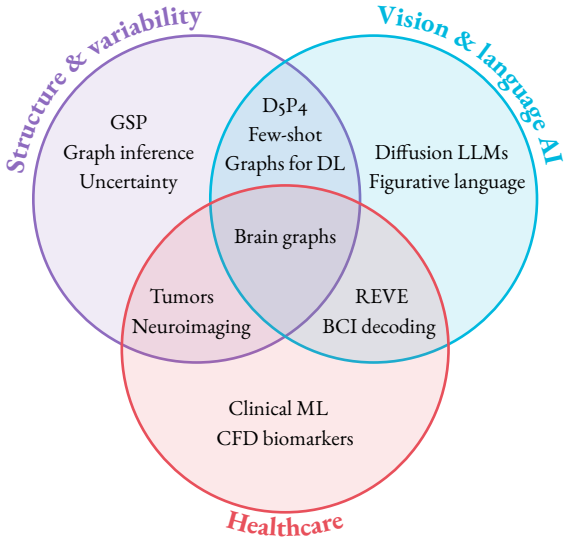


Central question

How to get the best out of data through its structure and variability, rather than its sheer amount?

- Follow the advances in core AI
- Transfer them to healthcare

My research in one diagram

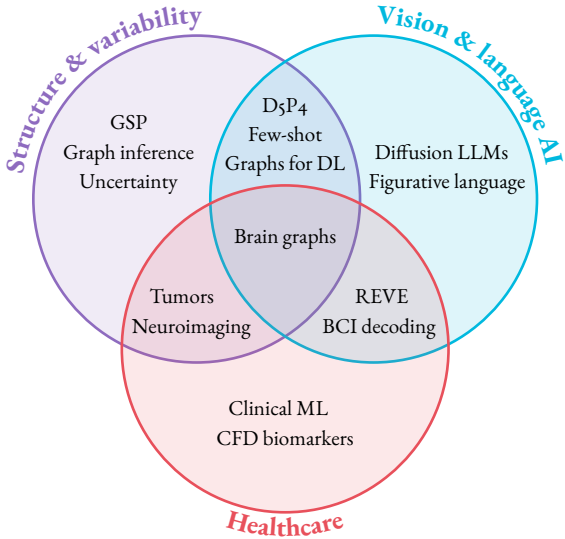


Central question

How to get the best out of data through its structure and variability, rather than its sheer amount?

- Follow the advances in core AI
- Transfer them to healthcare

My research in one diagram



Central question

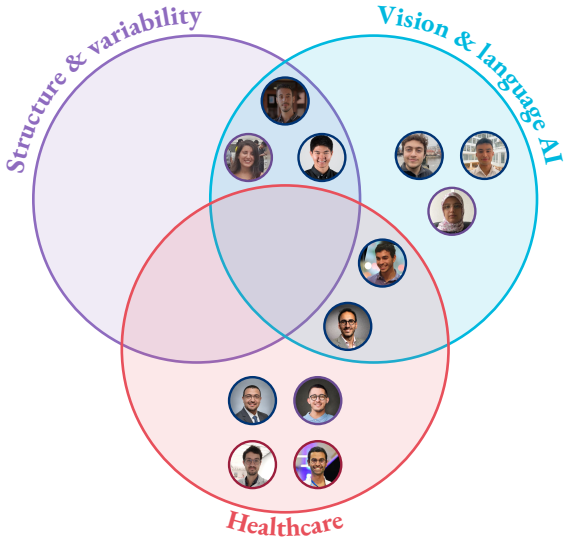
How to get the best out of data through its structure and variability, rather than its sheer amount?

- Follow the advances in core AI
- Transfer them to healthcare

Other collaborations

- Operational research
- Physics and time series
- Audio: ASR on out-of-distribution signals, neural codecs

My research in one diagram



Central question

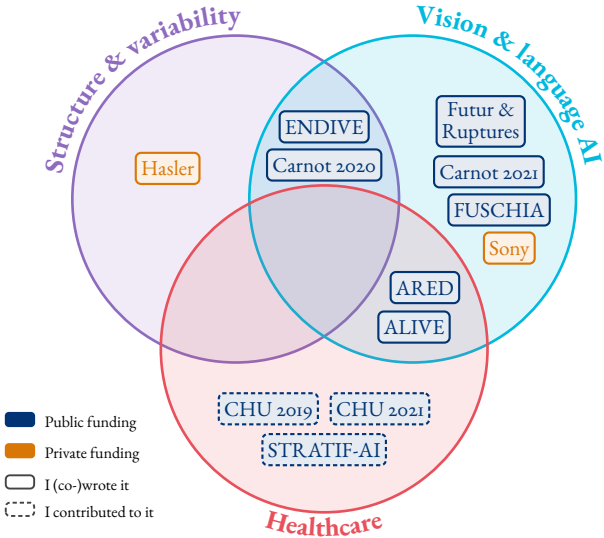
How to get the best out of data through its structure and variability, rather than its sheer amount?

- Follow the advances in core AI
- Transfer them to healthcare

Other collaborations

- Operational research
- Physics and time series
- Audio: ASR on out-of-distribution signals, neural codecs

My research in one diagram



Central question

How to get the best out of data through its structure and variability, rather than its sheer amount?

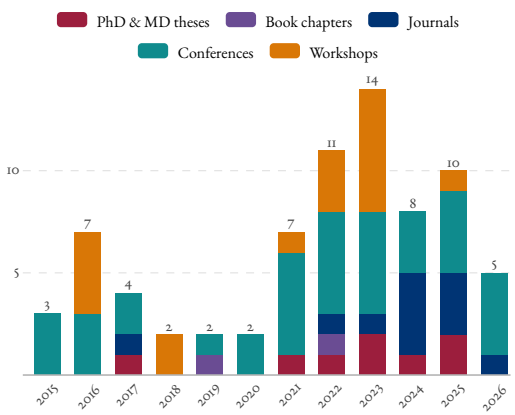
- Follow the advances in core AI
- Transfer them to healthcare

Other collaborations

- Operational research Chantiers
- Phys. & TS DARFI$_3$D CODHYPIA Pack
- Audio: ASR on out-of-distribution signals, neural codecs Orange

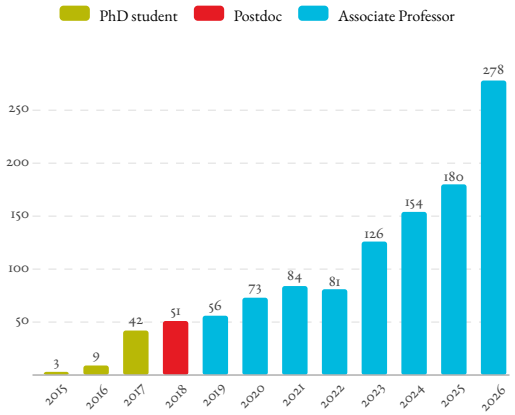
Research in numbers

75 peer-reviewed publications



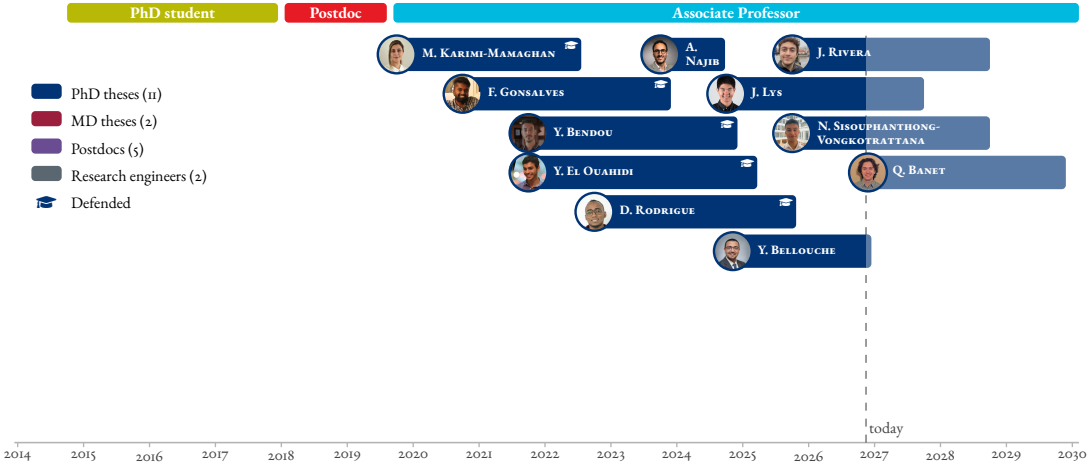
+6 preprints

1 142 citations, h-index 12

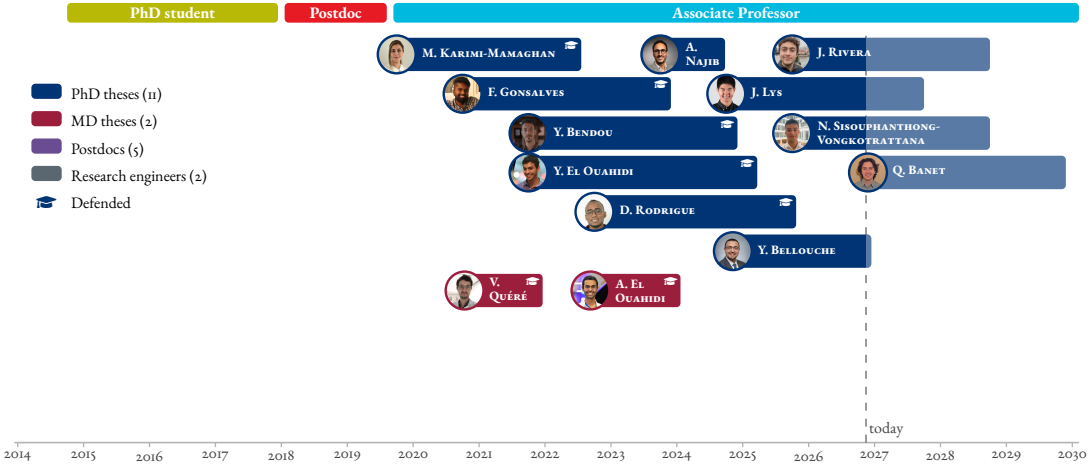


Google Scholar, as of November 3, 2026

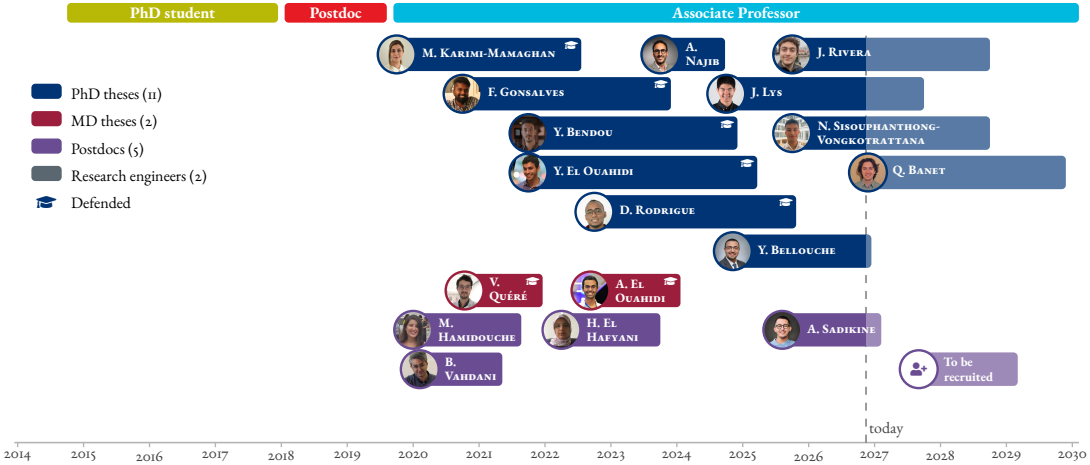
Supervisions



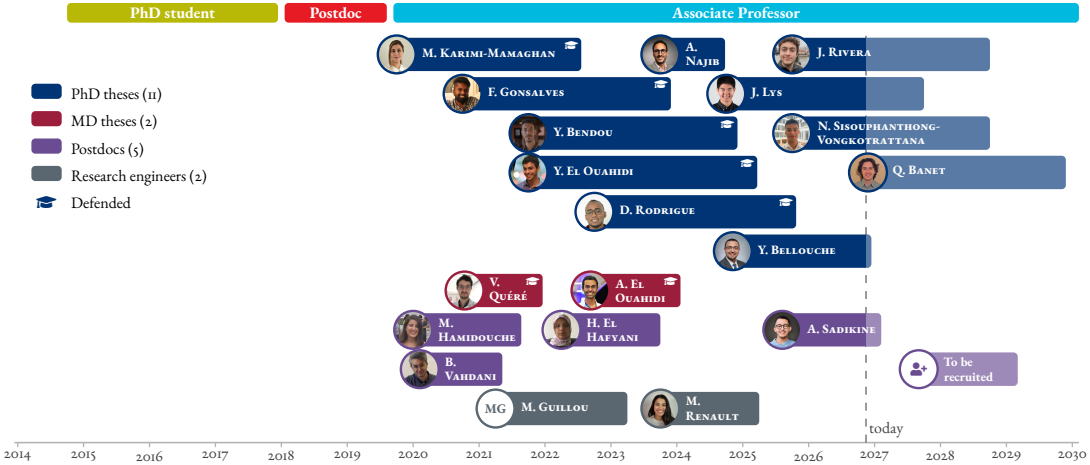
Supervisions



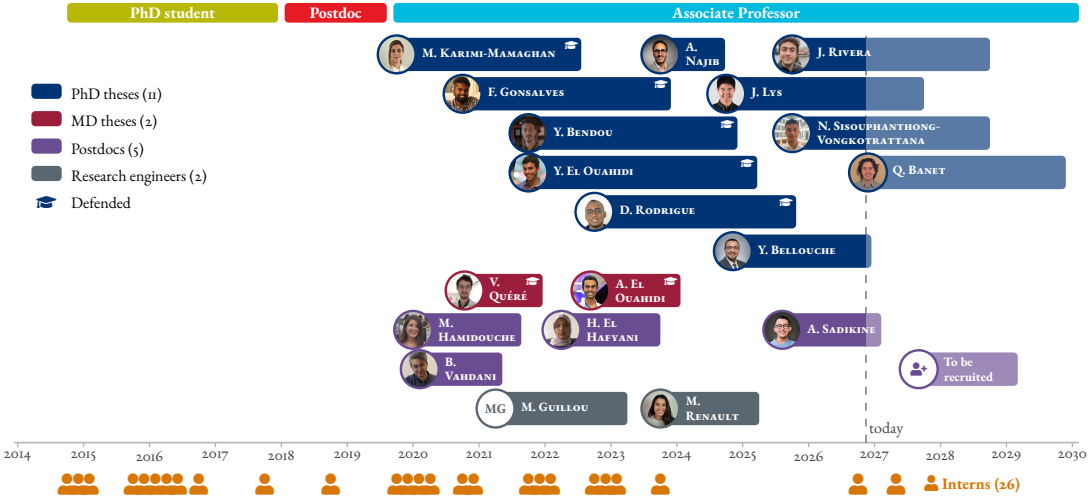
Supervisions



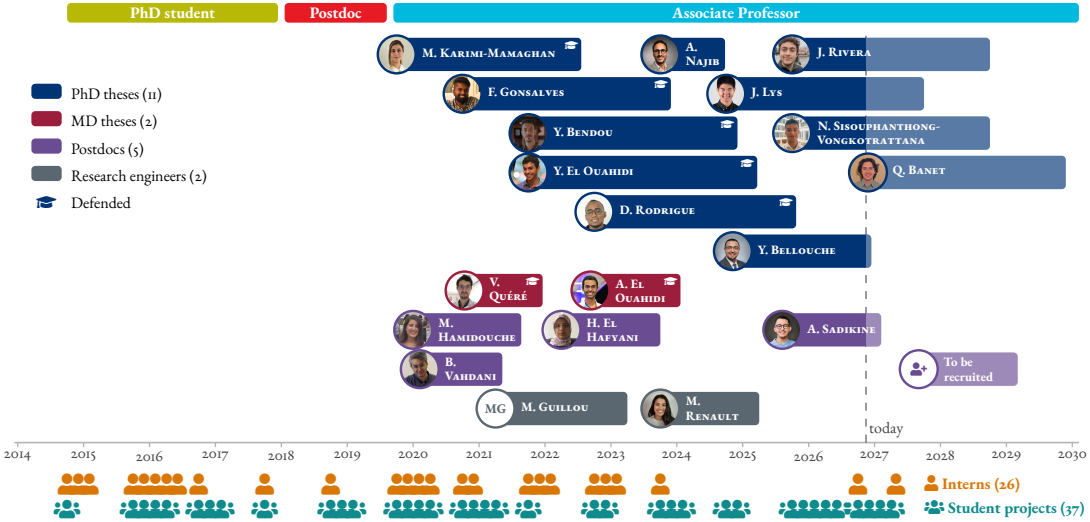
Supervisions



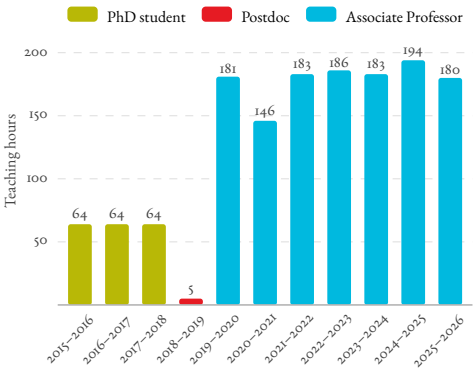
Supervisions



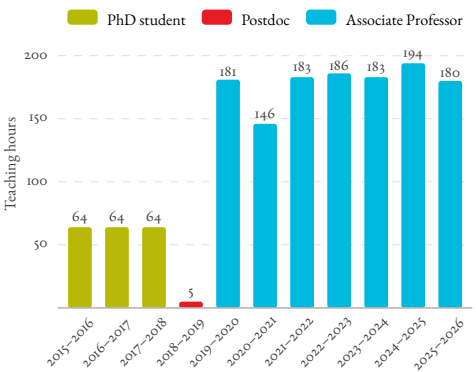
Supervisions



A glimpse of my teaching activities



A glimpse of my teaching activities



How to Lie with Statistics

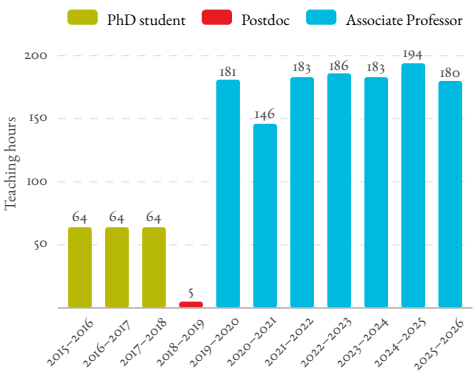
MASTER 1 & 2 · 2021–



Critical thinking

- Born from the misinformation on COVID-19
- Lie with data, then get debunked by peers

A glimpse of my teaching activities



How to Lie with Statistics

MASTER 1 & 2 · 2021–



Critical thinking

- Born from the misinformation on COVID-19
- Lie with data, then get debunked by peers

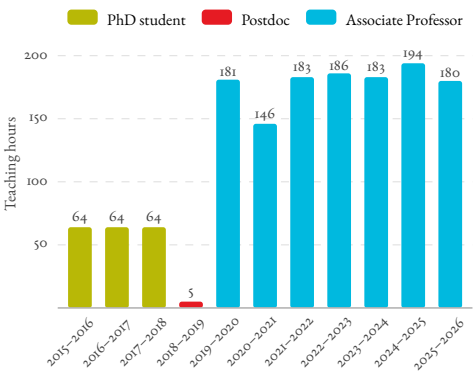
Data Science Track

MASTER 1 & 2 · 2024–

Introduction to data science

- Designed as demand for the AI elective tripled
- A proxy of each student's internship data

A glimpse of my teaching activities



How to Lie with Statistics

MASTER 1 & 2 · 2021–



Critical thinking

- Born from the misinformation on COVID-19
- Lie with data, then get debunked by peers

Data Science Track

MASTER 1 & 2 · 2024–

Introduction to data science

- Designed as demand for the AI elective tripled
- A proxy of each student's internship data

PyRat

BACHELOR 3 · 2015–



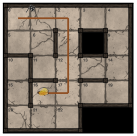
Algorithms and graphs

- A maze game to learn graphs and heuristics
- 400+ students a year, open source on PyPI

PyRat: algorithms through a maze game



Avoid walls
A graph



One cheese
BFS, DFS



With mud
Dijkstra



All cheese
TSP, backtracking



All, faster
Heuristics

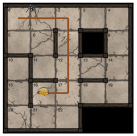


Tournament
Face an opponent

PyRat: algorithms through a maze game



Avoid walls
A graph



One cheese
BFS, DFS



With mud
Dijkstra



All cheese
TSP, backtracking



All, faster
Heuristics



Tournament
Face an opponent

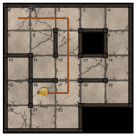
Handling heterogeneity

- Objectives of growing difficulty
- Basic and advanced content, for every level
- Teams of three of the same level

PyRat: algorithms through a maze game



Avoid walls
A graph



One cheese
BFS, DFS



With mud
Dijkstra



All cheese
TSP, backtracking



All, faster
Heuristics



Tournament
Face an opponent

Handling heterogeneity

- Objectives of growing difficulty
- Basic and advanced content, for every level
- Teams of three of the same level

Engaging students

- A game, not a calculator
- Short lectures, mostly practice
- Flipped classes

Science outreach and open science









In primary schools

2025

- Une Chercheure à l'École
- AI with Legos, in 9 classes of 7 schools

Science outreach and open science













In primary schools 2025

- Une Chercheure à l'École
- AI with Legos, in 9 classes of 7 schools

At Océanopolis 2019–2025

- Fête de la Science
- Nuit des Chercheures

Science outreach and open science















In primary schools 2025

- Une Chercheure à l'École
- AI with Legos, in 9 classes of 7 schools

At Océanopolis 2019–2025

- Fête de la Science
- Nuit des Chercheures

Open science

- All articles on arXiv or HAL
- Code and software under free licenses
- Reproducibility badge, GRETSI 2025



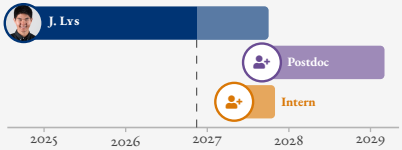
ENDIVE

ENcouraging DIVErsity in few-shot learning

ENDIVE

2024–2028

- ANR JCJC, as PI
- ∼294 k€

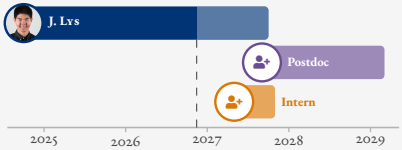


ENcouraging DIVErsity in few-shot learning

ENDIVE

2024–2028

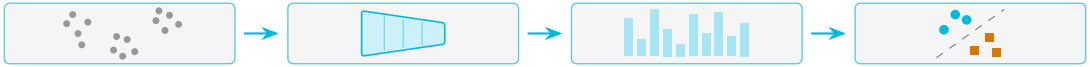
- ANR JCJC, as PI
- ∼294 k€



As proposed

2023

Learning from a few labeled examples, with a model pretrained on others, through the prism of diversity [ENDIVE 🌎]:



Examples to label
diverse, with a DPP

Pretrained model(s)
on other data

Features to keep
diverse and relevant

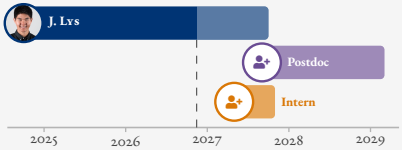
Classifier
from a few labels

ENcouraging DIVErsity in few-shot learning

ENDIVE

2024–2028

- ANR JCJC, as PI
- ~294 k€

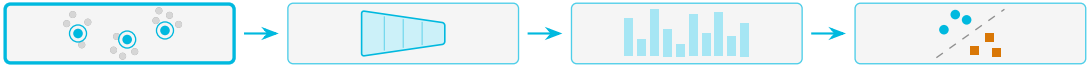


As proposed

2023

Learning from a few labeled examples, with a model pretrained on others, through the prism of diversity [ENDIVE]:

- **Which examples to label?** Draw them with DPPs, which repel similar items



Examples to label
diverse, with a DPP

Pretrained model(s)
on other data

Features to keep
diverse and relevant

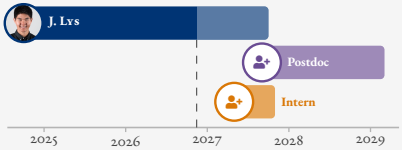
Classifier
from a few labels

ENcouraging DIVErsity in few-shot learning

ENDIVE

2024–2028

- ANR JCJC, as PI
- ∼294 k€

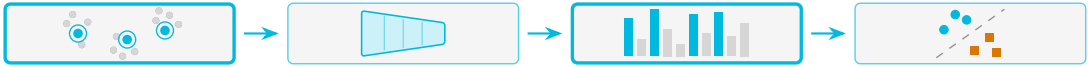


As proposed

2023

Learning from a few labeled examples, with a model pretrained on others, through the prism of diversity [ENDIVE]:

- **Which examples to label?** Draw them with DPPs, which repel similar items
- **Which features to keep?** Diverse and relevant dimensions of the representations



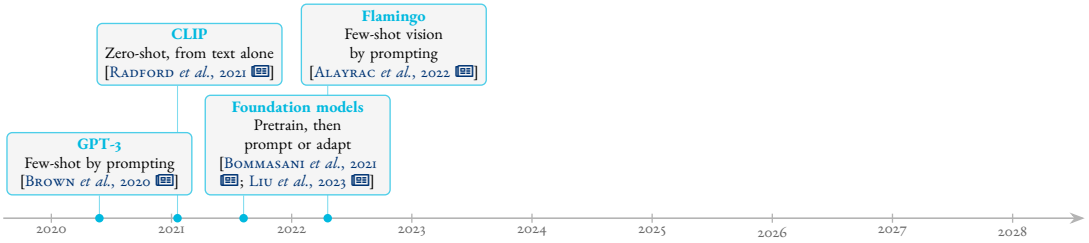
Examples to label
diverse, with a DPP

Pretrained model(s)
on other data

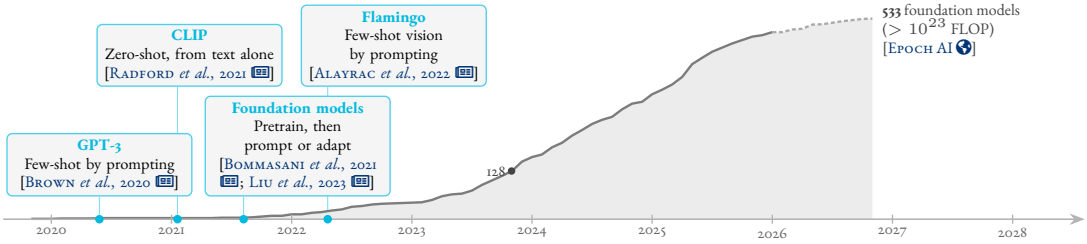
Features to keep
diverse and relevant

Classifier
from a few labels

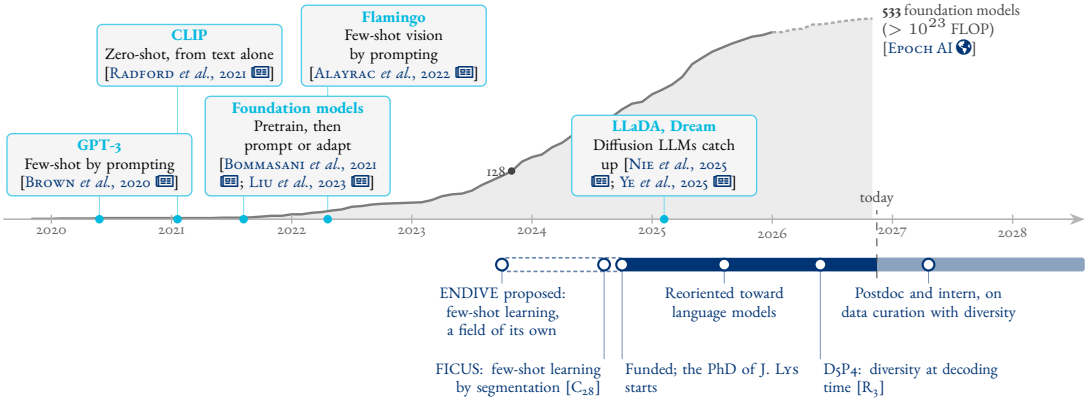
A field that moves faster than a project



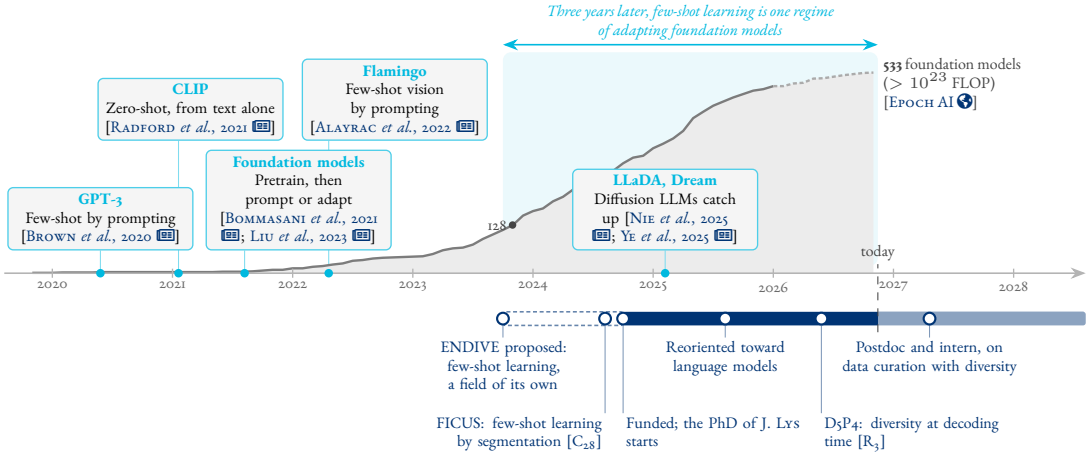
A field that moves faster than a project



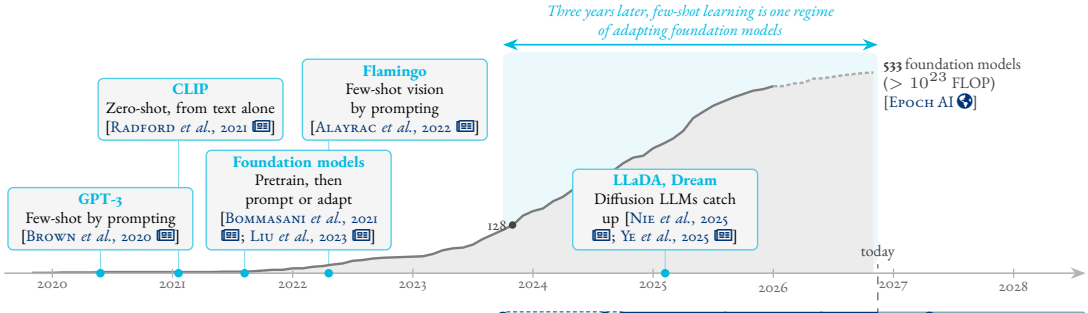
A field that moves faster than a project



A field that moves faster than a project



A field that moves faster than a project



- No race against the big actors: niches that are not on their roadmaps
- A broad objective, which survives the pace of the field

Two axes, reoriented

As proposed



Which examples to label?

Two axes, reoriented

As proposed



Which examples to label?

Diversity in the data

Now

- Curating what diffusion LLMs learn from: as good with less data?
- How the diversity of the data shapes the features they learn

Two axes, reoriented

As proposed



Which examples to label?

As proposed



Which features to keep?

Diversity in the data

Now

- Curating what diffusion LLMs learn from: as good with less data?
- How the diversity of the data shapes the features they learn

Two axes, reoriented

As proposed



Which examples to label?

Diversity in the data — Now

- Curating what diffusion LLMs learn from: as good with less data?
- How the diversity of the data shapes the features they learn

As proposed



Which features to keep?

Diversity in the representations — Now

- Which features: sparse autoencoders
- Which contexts: demonstrations, retrieval
- Which models: complementary ones

Two pieces of work

FICUS

Images

- Four Master students (student project)
- Drawn to research: a paper, not a report



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* [C28]

D5P4

Text

- A turn to discrete diffusion LLMs
- A collaboration with Sony Research



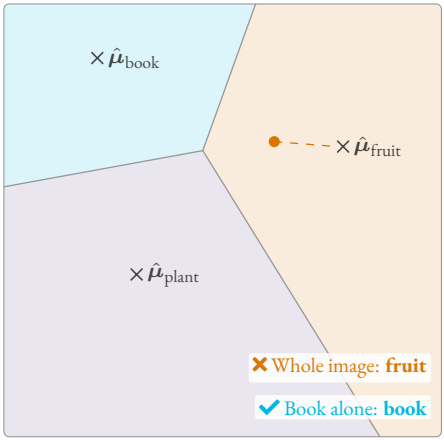
Jonathan Lys, Vincent Gripon, Axel Marmoret, Lukas Mauch, Fabien Cardinaux, Ghouthi Boukli-Hacene, Bastien Pasdeloup (2026). **"D5P4: Partition Determinantal Point Process for Diversity in Parallel Discrete Diffusion Decoding"**. *NAACL* (under review) [R3]

[FICUS] One label, several objects

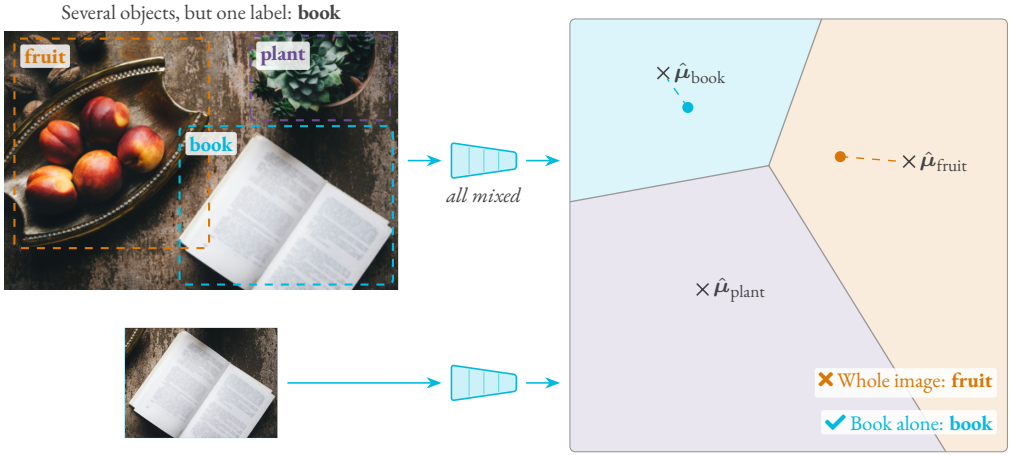
Several objects, but one label: **book**



all mixed



[FICUS] One label, several objects



[FICUS] Finding the objects, with no training

Input image



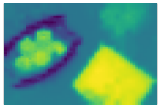
[FICUS] Finding the objects, with no training

Input image

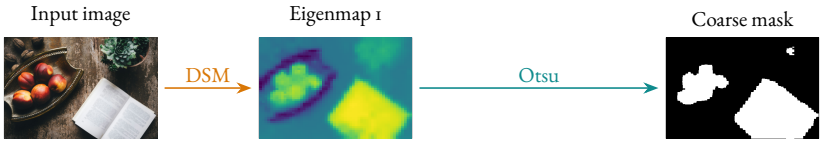


DSM

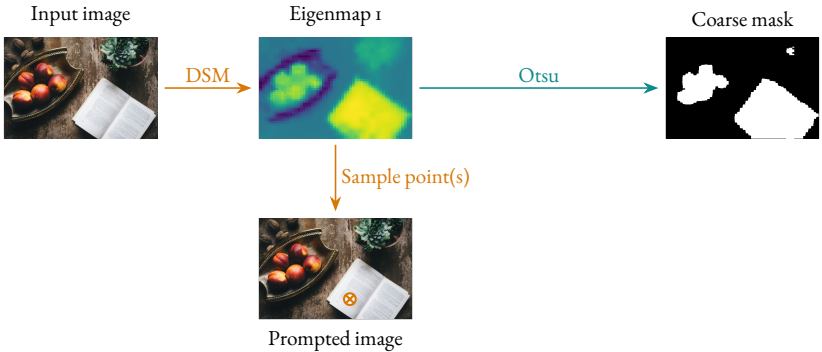
Eigenmap 1



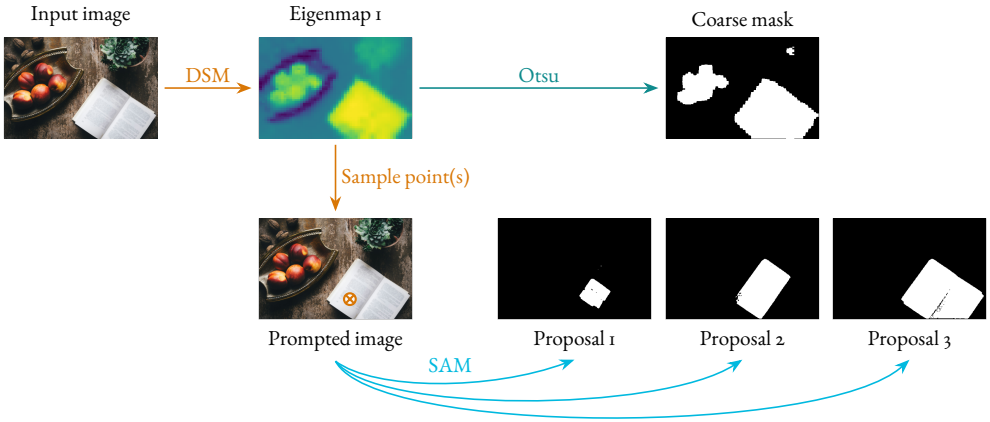
[FICUS] Finding the objects, with no training



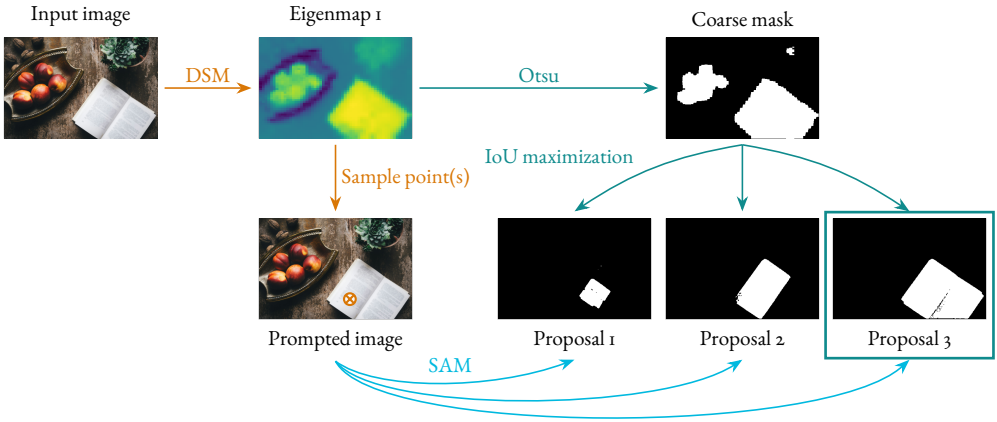
[FICUS] Finding the objects, with no training



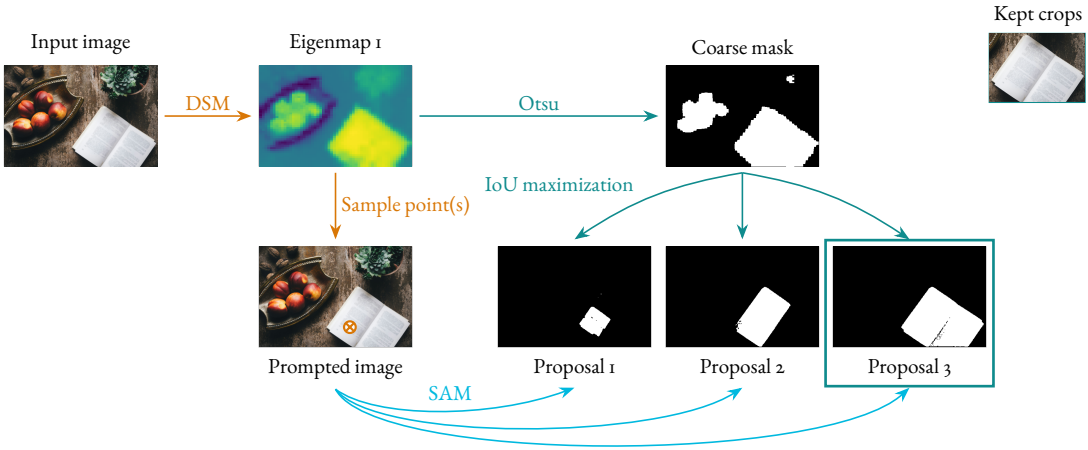
[FICUS] Finding the objects, with no training



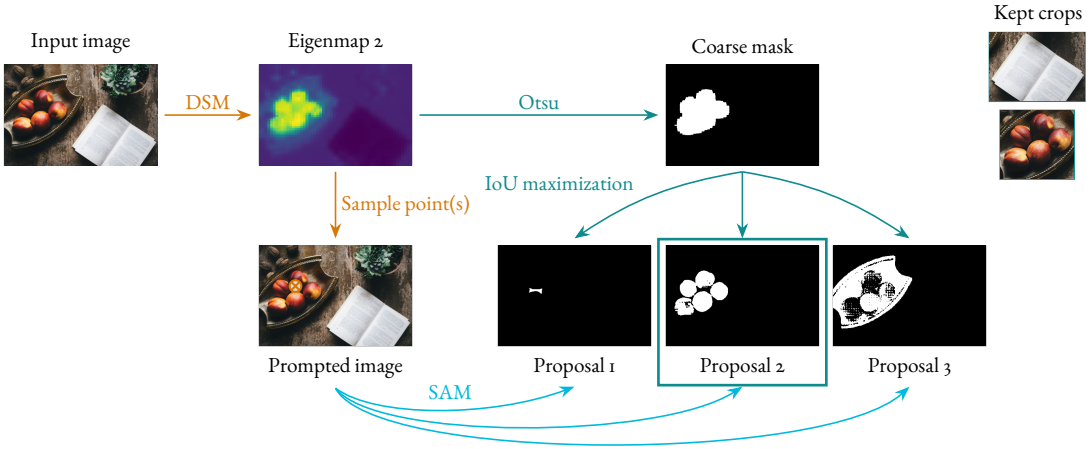
[FICUS] Finding the objects, with no training



[FICUS] Finding the objects, with no training



[FICUS] Finding the objects, with no training



[FICUS] Diversity, worth it before foundation models

Diversity in FICUS



- Distinct eigenvectors, prompts, and objects: one crop each
- Several diverse views, rather than one embedding
- The view closest to a class decides

Before foundation models

- Small backbones, small datasets: one embedding mixed the whole scene
- A few labels cannot teach where the object is: separating objects with no training helped

[FICUS] Diversity, worth it before foundation models

Diversity in FICUS



- Distinct eigenvectors, prompts, and objects: one crop each
- Several diverse views, rather than one embedding
- The view closest to a class decides

Before foundation models

- Small backbones, small datasets: one embedding mixed the whole scene
- A few labels cannot teach where the object is: separating objects with no training helped

With foundation models

- Zero-shot semantic segmentation now works: objects are found and named with no example
- Few-shot learning is no longer a problem in itself, but one way of adapting a foundation model

[D5P4] Why several different answers?

Three independent answers

A train leaves at 9:40 and arrives at 11:15. How long is the trip?

2 h 25 | 2 h 25 | 2 h 25

✖ All wrong; the vote too

Describe a sunset in one sentence.

The sun sets over the sea.

The sun sets over the sea.

The sun slowly sets over the sea.

✖ One sentence, three times

[D5P4] Why several different answers?

	Three independent answers	Three diverse answers
A train leaves at 9:40 and arrives at 11:15. How long is the trip?	2 h 25 · 2 h 25 · 2 h 25 ✗ All wrong; the vote too	2 h 25 · 1 h 35 · 1 h 75 ✓ One is right, and can be checked
Describe a sunset in one sentence.	The sun sets over the sea. The sun sets over the sea. The sun slowly sets over the sea. ✗ One sentence, three times	The sky turns orange as the day ends. Light fades slowly behind the hills. Evening paints the clouds pink. ✓ Three different formulations

[D5P4] Why several different answers?

	Three independent answers	Three diverse answers
A train leaves at 9:40 and arrives at 11:15. How long is the trip?	2 h 25 · 2 h 25 · 2 h 25 ✗ All wrong; the vote too	2 h 25 · 1 h 35 · 1 h 75 ✓ One is right, and can be checked
Describe a sunset in one sentence.	The sun sets over the sea. The sun sets over the sea. The sun slowly sets over the sea. ✗ One sentence, three times	The sky turns orange as the day ends. Light fades slowly behind the hills. Evening paints the clouds pink. ✓ Three different formulations

To cover

- One right answer is enough: pass@k

To select

- A vote only helps if answers differ

To learn

- RL learns from their differences

[D5P4] Determinantal point processes

A probability on subsets

Given N items and their similarities $\boldsymbol{L}$, a DPP draws a subset S with [Kulesza *et al.*, 2012]:

$$\mathbb{P}(S) \propto \det(\boldsymbol{L}_S)$$

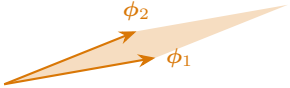
- $\det(\boldsymbol{L}_S)$ is the squared volume spanned by the features $\boldsymbol{\phi}_i$ of the items of S
- Similar items repel each other
- With $|S| = k$: a k-DPP [Kulesza *et al.*, 2011]

A probability on subsets

Given N items and their similarities $\boldsymbol{L}$, a DPP draws a subset S with [Kulesza *et al.*, 2012]:

$$\mathbb{P}(S) \propto \det(\boldsymbol{L}_S)$$

- $\det(\boldsymbol{L}_S)$ is the squared volume spanned by the features $\boldsymbol{\phi}_i$ of the items of S
- Similar items repel each other
- With $|S| = k$: a k-DPP [Kulesza *et al.*, 2011]



Similar items: small volume, rarely drawn together

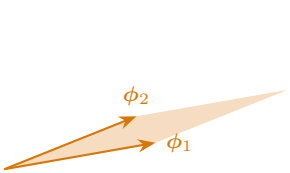
[D5P4] Determinantal point processes

A probability on subsets

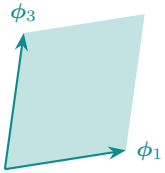
Given N items and their similarities $\boldsymbol{L}$, a DPP draws a subset S with [Kulesza *et al.*, 2012]:

$$\mathbb{P}(S) \propto \det(\boldsymbol{L}_S)$$

- $\det(\boldsymbol{L}_S)$ is the squared volume spanned by the features $\boldsymbol{\phi}_i$ of the items of S
- Similar items repel each other
- With $|S| = k$: a k-DPP [Kulesza *et al.*, 2011]



Similar items: small volume, rarely drawn together



Diverse items: large volume

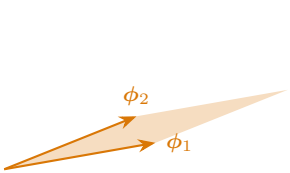
[D5P4] Determinantal point processes

A probability on subsets

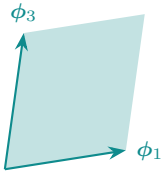
Given N items and their similarities $\boldsymbol{L}$, a DPP draws a subset S with [Kulesza *et al.*, 2012]:

$$\mathbb{P}(S) \propto \det(\boldsymbol{L}_S)$$

- $\det(\boldsymbol{L}_S)$ is the squared volume spanned by the features $\boldsymbol{\phi}_i$ of the items of S
- Similar items repel each other
- With $|S| = k$: a k-DPP [Kulesza *et al.*, 2011]

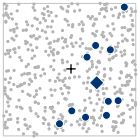


Similar items: small volume, rarely drawn together

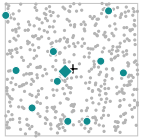


Diverse items: large volume

10 independent points



10 points of a DPP



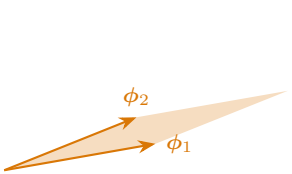
[D5P4] Determinantal point processes

A probability on subsets

Given N items and their similarities $\boldsymbol{L}$, a DPP draws a subset S with [Kulesza *et al.*, 2012]:

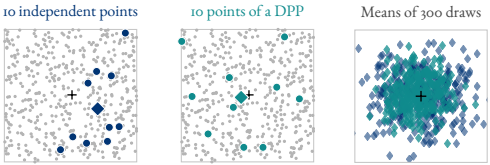
$$\mathbb{P}(S) \propto \det(\boldsymbol{L}_S)$$

- $\det(\boldsymbol{L}_S)$ is the squared volume spanned by the features $\boldsymbol{\phi}_i$ of the items of S
- Similar items repel each other
- With $|S| = k$: a k-DPP [Kulesza *et al.*, 2011]



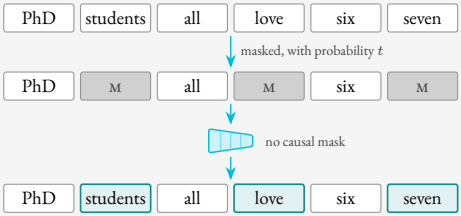
Similar items: small volume, rarely drawn together

Diverse items: large volume



[D5P4] Discrete diffusion for text

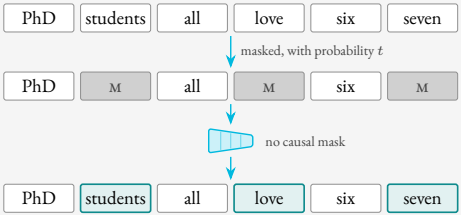
Training



- Each token masked with probability $t \sim \mathcal{U}(0, 1)$
- One transformer predicts all the masked tokens at once [Nie *et al.*, 2025]

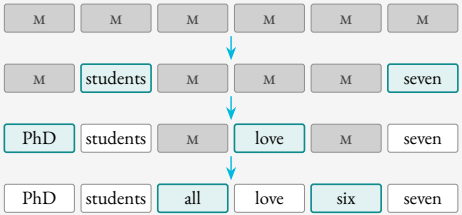
[D5P4] Discrete diffusion for text

Training



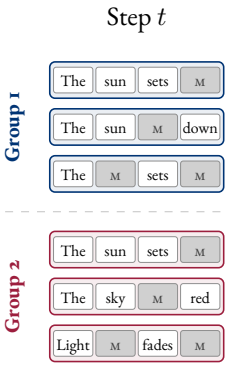
- Each token masked with probability $t \sim \mathcal{U}(0, 1)$
- One transformer predicts all the masked tokens at once [Nie *et al.*, 2025]

Inference

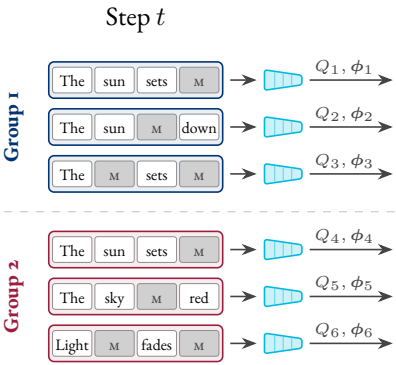


- From a fully masked answer, in a fixed number of steps
- Each step predicts every position, keeps the most confident, masks the rest again

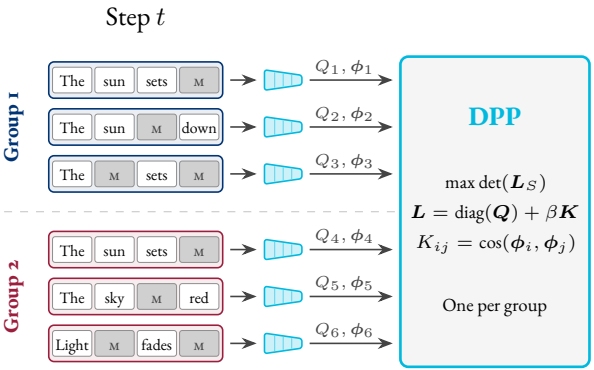
[D5P4] Diverse candidates at each step of inference



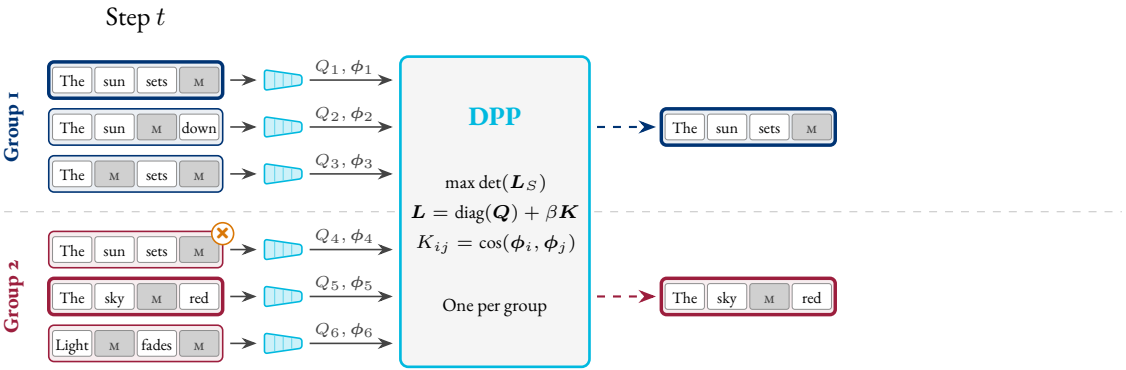
[D5P4] Diverse candidates at each step of inference



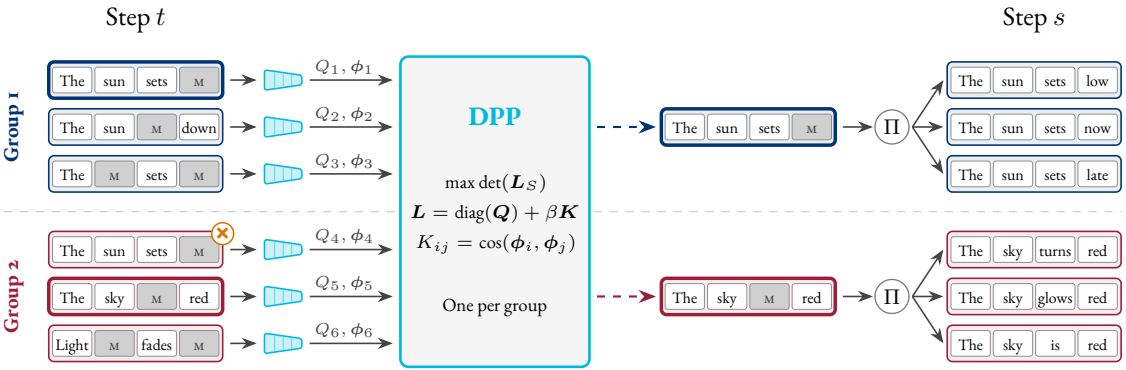
[D5P4] Diverse candidates at each step of inference



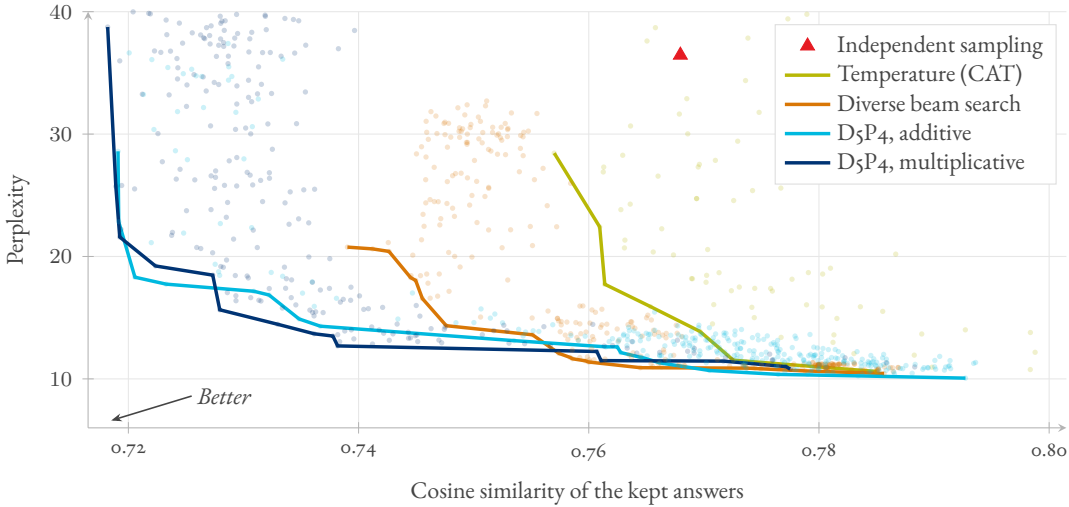
[D5P4] Diverse candidates at each step of inference



[D5P4] Diverse candidates at each step of inference



[D5P4] More diversity, same quality (open-ended generation)



Next steps

Diversity in the data

- **Curation:** training diffusion LLMs with less data, with DPPs (+ REVE 2)
- **Behavior:** how the diversity of the data shapes their trajectories and mechanisms

Next steps

Diversity in the data

- **Curation:** training diffusion LLMs with less data, with DPPs (+ REVE 2)
- **Behavior:** how the diversity of the data shapes their trajectories and mechanisms

Diversity in the representations

- **Features:** diverse, task-relevant ones, from sparse autoencoders
- **Contexts:** complementary demonstrations and retrieved documents
- **Models:** complementary foundation models, to select and combine

Next steps

Diversity in the data

- **Curation:** training diffusion LLMs with less data, with DPPs (+ REVE 2)
- **Behavior:** how the diversity of the data shapes their trajectories and mechanisms

Diversity in the representations

- **Features:** diverse, task-relevant ones, from sparse autoencoders
- **Contexts:** complementary demonstrations and retrieved documents
- **Models:** complementary foundation models, to select and combine

Recruitments

ENDIVE

- A postdoc for 18 months, from September 2027, on data curation and diversity
- A Master 2 intern for 6 months, from May 2027



AI & Healthcare

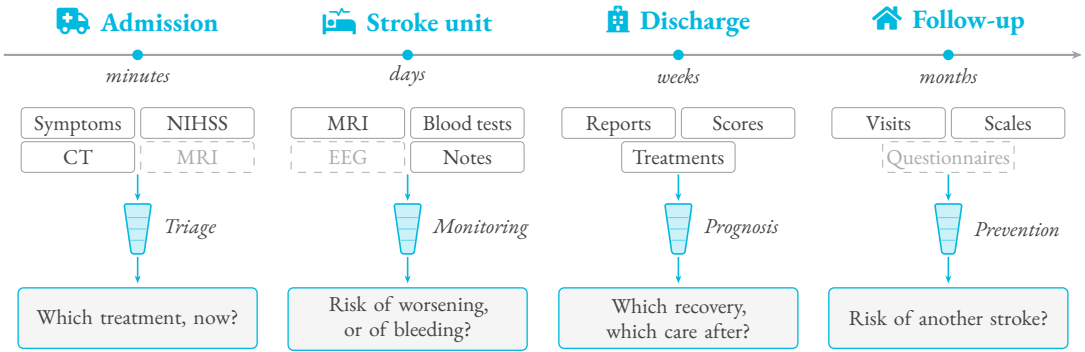
What healthcare asks of AI: the case of a stroke



What healthcare asks of AI: the case of a stroke



What healthcare asks of AI: the case of a stroke



From clinical scores to foundation models

Statistics

A few variables chosen by clinicians: interpretable, hard to beat on small cohorts

- Mortality after an ICH, CHU of Brest [T_2]
- Immunotherapy in melanoma, CHUV [R_4]

From clinical scores to foundation models

Statistics

A few variables chosen by clinicians: interpretable, hard to beat on small cohorts

- Mortality after an ICH, CHU of Brest [T_2]
- Immunotherapy in melanoma, CHUV [R_4]



Machine learning

Forests, SVM, boosting: nonlinear effects, and the variables that matter

- Pacemaker after a TAVI, CHU of Brest [T_4]
- Cardiac intensive care, AP-HP [J_{10}]
- Tumors as cell graphs, TCGA [C_{12}]

From clinical scores to foundation models

Statistics

A few variables chosen by clinicians: interpretable, hard to beat on small cohorts

- Mortality after an ICH, CHU of Brest [T_2]
- Immunotherapy in melanoma, CHUV [R_4]

Machine learning

Forests, SVM, boosting: nonlinear effects, and the variables that matter

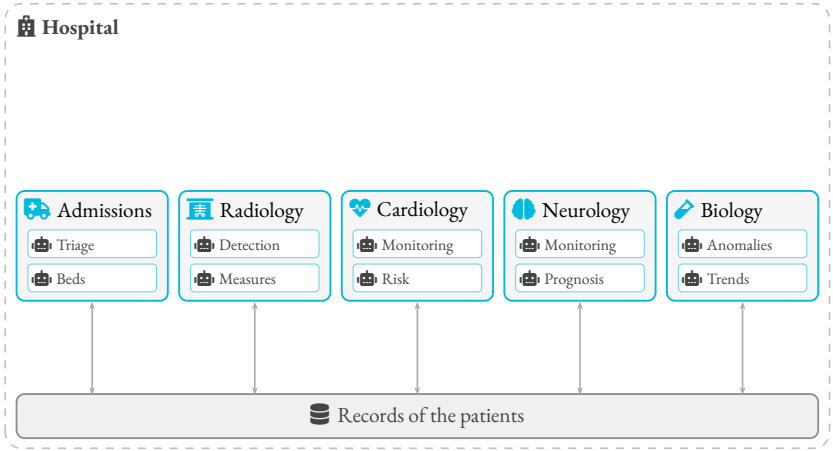
- Pacemaker after a TAVI, CHU of Brest [T_4]
- Cardiac intensive care, AP-HP [J_{10}]
- Tumors as cell graphs, TCGA [C_{12}]

Foundation models

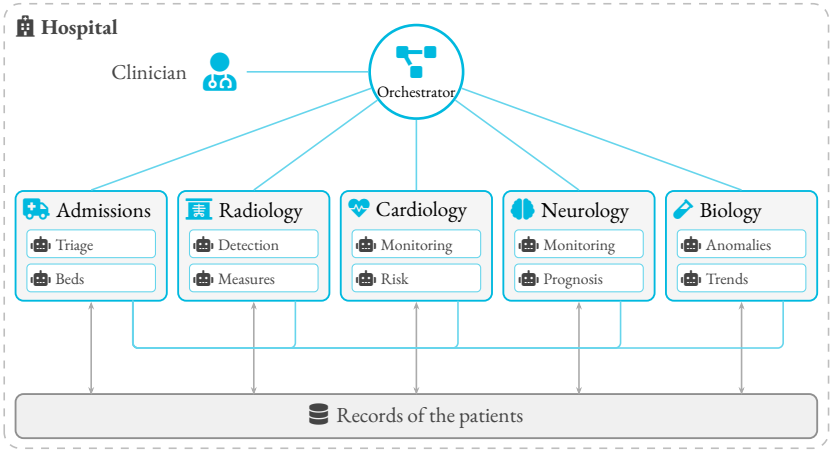
Pretrained on far more data than a cohort: useful even with few labels

- Multimodal ICH prognosis, STRATIF-AI [S_{13}]
- REVE, for EEG [C_{31}]

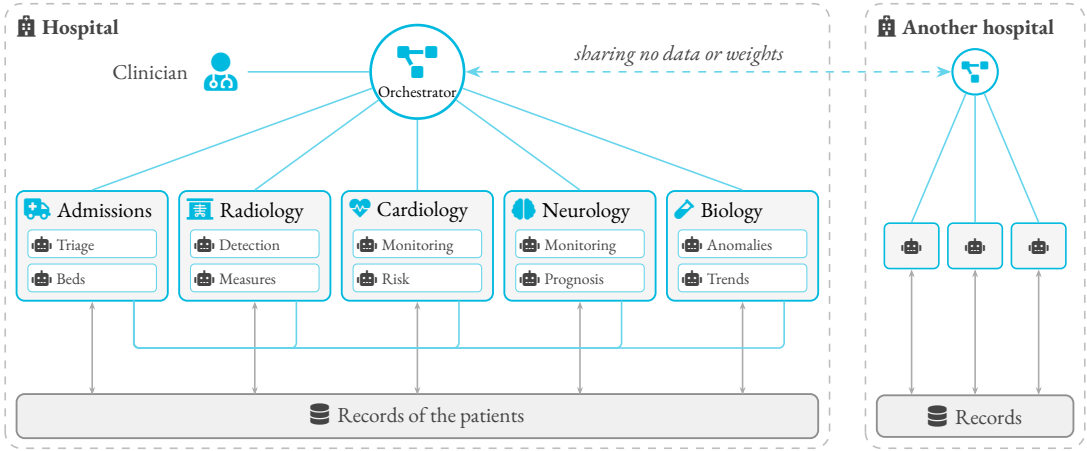
What an agentic system for the hospital could look like



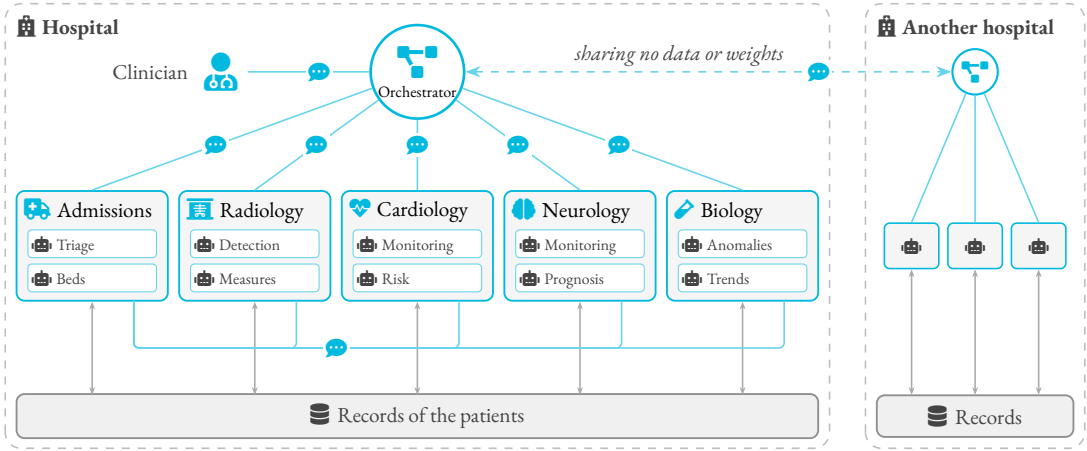
What an agentic system for the hospital could look like



What an agentic system for the hospital could look like



What an agentic system for the hospital could look like



Why agents that cooperate, rather than independent models

Independent models, one per task

- ✗ Each sees only its own inputs: triage ignores what follow-up learned
- ✗ A missing exam breaks a model built on fixed inputs
- ✗ Answers that may contradict each other, left to the clinician to reconcile
- ✗ A new task means a new model, and new labels

Why agents that cooperate, rather than independent models

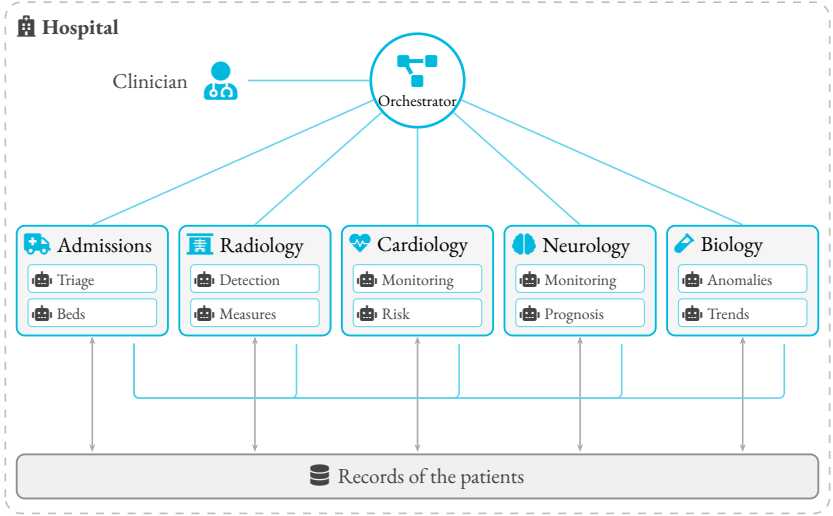
Independent models, one per task

- ✗ Each sees only its own inputs: triage ignores what follow-up learned
- ✗ A missing exam breaks a model built on fixed inputs
- ✗ Answers that may contradict each other, left to the clinician to reconcile
- ✗ A new task means a new model, and new labels

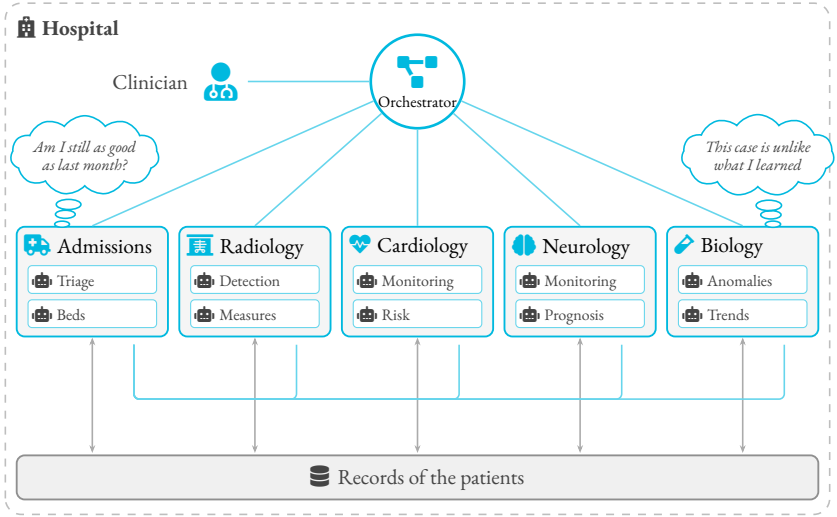
Specialized agents that cooperate

- ✓ An agent asks another for what it lacks, rather than failing without it
- ✓ Text between them: an orchestrator assembles it, a clinician reads it, and can disagree [Klang *et al.*, 2026; Zhou *et al.*, 2026]
- ✓ Each stays specialized, trained or adapted where its data lives
- ✓ Each can abstain, with a calibrated uncertainty [Angelopoulos *et al.*, 2023]

Trust in agents, and between agents



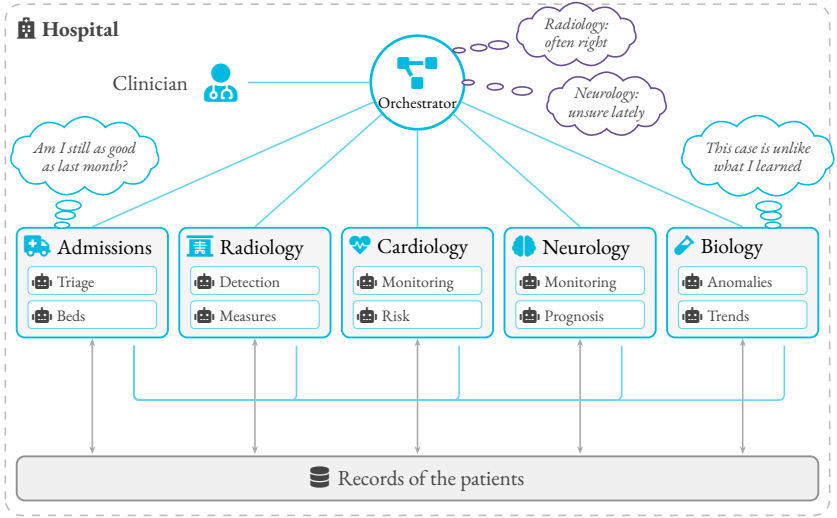
Trust in agents, and between agents



Self-perception

Regressions and drift: **ALIVE** (Cerevial) [S_{14}]

Trust in agents, and between agents



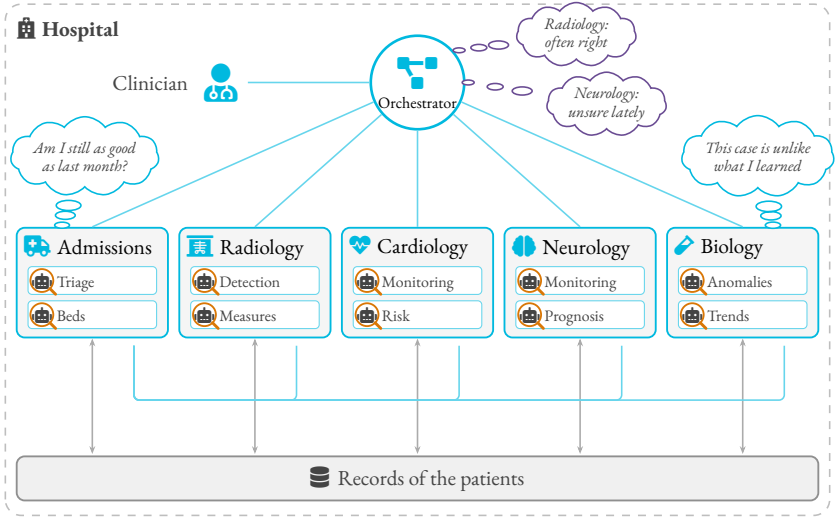
Self-perception

Regressions and drift: **ALIVE** (Cerevial) [S_{14}]

Perception of others

Trust from their history; abstain when unsure

Trust in agents, and between agents



Self-perception

Regressions and drift: **ALIVE** (Cerevial) [S_{14}]

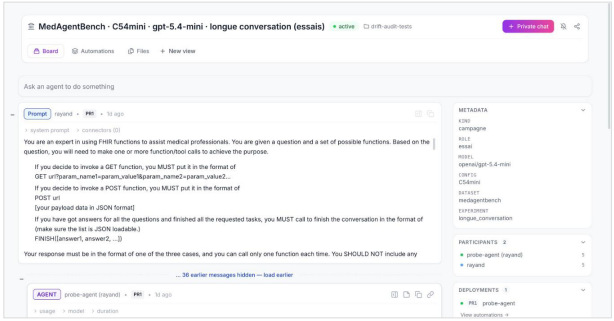
Perception of others

Trust from their history; abstain when unsure

Mechanistic interpretability

What an answer rests on, inside the model [Bricken *et al.*, 2023]

ALIVE: a project on the supervision of medical agents



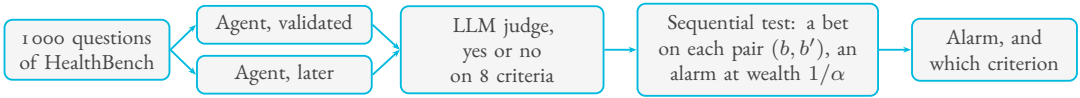
A maturation project

-  R.D
- France 2030, PUI *"In the Box"*
- Then a CIFRE PhD

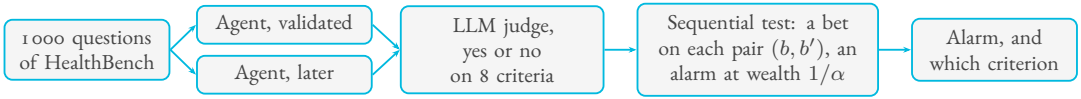
Objectives

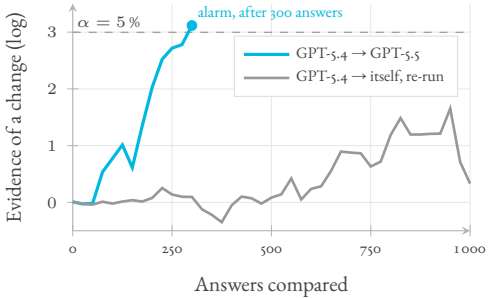
- Persistent agents, whose personality is revised by introspection
- A medical agent on intensive care records, to supervise: regressions, audit trail, drift

ALIVE so far: detecting the drift of an agent

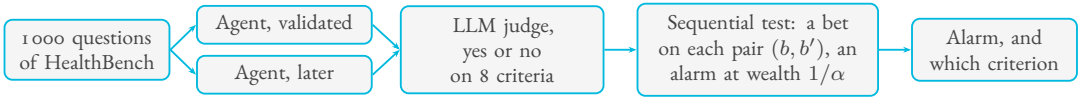


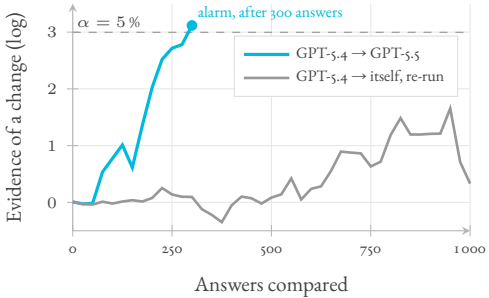
ALIVE so far: detecting the drift of an agent





ALIVE so far: detecting the drift of an agent





What we want

- An agent that changes: a model updated by its provider, a prompt rewritten by introspection
- See the drift from its answers alone, false alarms under 5 %, and name the criterion that moved
- Mechanistic interpretability to evaluate the impact of the drift



Conclusion

Where I am going

ENDIVE

2024–2028

Diffusion LLMs, from mainstream AI: diversity in their data and representations; REVE 2 curation

Where I am going

ENDIVE 2024–2028

Diffusion LLMs, from mainstream AI: diversity in their data and representations; REVE 2 curation

ALIVE 2026–2027

Agents, from mainstream AI, made trustworthy for medicine with Cerevial; then a CIFRE PhD

Where I am going

ENDIVE 2024–2028

Diffusion LLMs, from mainstream AI: diversity in their data and representations; REVE 2 curation

ALIVE 2026–2027

Agents, from mainstream AI, made trustworthy for medicine with Cerevial; then a CIFRE PhD

Applied projects 2027–

These advances on clinical questions: STRATUM, SIGMAS, weather and infarction; audio codecs

Where I am going

ENDIVE 2024–2028

Diffusion LLMs, from mainstream AI: diversity in their data and representations; REVE 2 curation

ALIVE 2026–2027

Agents, from mainstream AI, made trustworthy for medicine with Cerevial; then a CIFRE PhD

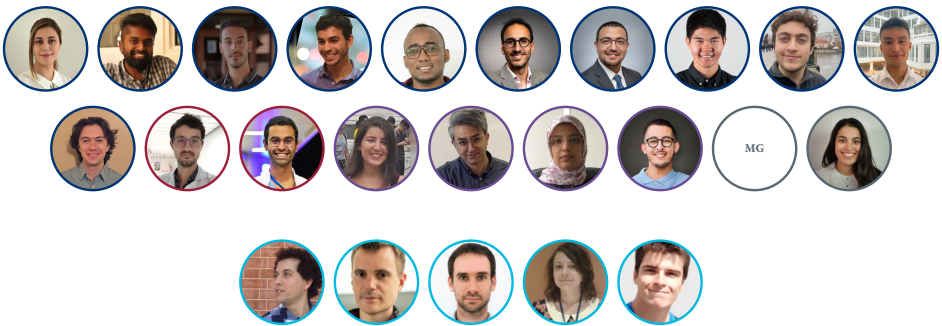
Applied projects 2027–

These advances on clinical questions: STRATUM, SIGMAS, weather and infarction; audio codecs

Specialized agents at the hospital LONG TERM

- Trained, deployed and audited by the hospital itself
- Text as the pivot, between agents and with clinicians
- Trust that circulates between agents, and calibrated abstention
- Diversity, to choose what each agent learns from
- Genomics, as the next modality

Thank You!



Supplementary slides

Supervisions 33

GSP 52

REVE 56

FSL 62

LLMs 64

Clinical works 69

Collaborations 72

Applied projects 66

Directions 68

Funding 73

Other activities 74

Teaching & AI 76



Supervisions

Maryam Karimi-Mamaghan



PhD

2019–2022

- Defended on July 20, 2022
- Funding: Lab-STICC, DECIDE team
- Now: Assistant Professor, Vrije Universiteit Amsterdam

Topics

Hybridization of metaheuristics and ML for combinatorial optimization

- A taxonomy of the ways to combine metaheuristics with machine learning
- Learning to select operators: Q-learning in iterated greedy and local search
- Applications to flowshop scheduling and to the traveling salesman problem

Publications

[T_3] Hybridizing Metaheuristics with Machine Learning for Combinatorial Optimization: A Taxonomy and Learning to Select Operators – *Manuscript, 2022*

[J_3] Learning to Select Operators in Meta-Heuristics: An Integration of Q-Learning into the Iterated Greedy Algorithm for the Permutation Flowshop Scheduling Problem – *EJOR, 2023*; also [W_9]

[C_{10}] An Online Learning-Based Metaheuristic for Solving Combinatorial Optimization Problems – *Congrès Annuel de la ROADEF, 2020*

[C_{15}] A Learning-Based Iterated Local Search Algorithm for Solving the Traveling Salesman Problem – *OLA, 2021*

[C_{16}] A Reinforcement Learning-Based Operator Selection in Iterated Local Search for Solving Scheduling Problems – *EURO, 2021*

[W_7] An Adaptive Iterated Local Search Using Reinforcement Learning for Solving Combinatorial Optimization Problems – *CORS Annual Conference, 2021*

Fred Michael Gonsalves



PhD

2020–2023

- Defended on November 29, 2023
- Funding: CIFRE with Chantiers de l'Atlantique
- Now: Data Scientist in Energy Efficiency and Environmental Performance, Chantiers de l'Atlantique

Topics

ML and operational research for cruise ship energy efficiency

- Predicting the propulsion power of cruise ships from operational data
- Designing and operating multi-energy systems on board, as a MILP

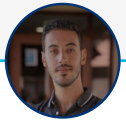
Publications

[T_5] Leveraging Machine Learning and Operations Research to Enhance Cruise Ship Energy Efficiency – *Manuscript, 2023*

[C_{13}] New Insights into the Propulsion Power Prediction of Cruise Ships – *ICTAI, 2021*

[P_3] Optimal Design and Operation of Cruise Ship Multi-Energy Systems: An MILP Formulation – *HAL preprint hal-04310676, 2023*

Yassir Bendou



PhD

2021–2024

- Defended on December 2, 2024
- Funding: Fondation Mines-Télécom and Institut Carnot
- Now: Applied Scientist, Mistral AI

Topics

Few-shot image classification

- EASY: a simple baseline among the best, best paper award
- Ambiguous one-shot tasks, and a statistical model that predicts generalization
- Text and vision-language models for few- and zero-shot classification

Publications

[T_6] A Data Modeling Perspective on Few-Shot Image Classification – *Manuscript, 2024*

[J_2] EASY: Ensemble Augmented-Shot Y-shaped Learning: State-of-the-Art Few-Shot Classification with Simple Components – *Journal of Imaging, 2022*

[C_{17}] Le Manchot, la Banane et la Bibliothèque... (de la Désambiguïsation d'une Tâche de Classification avec un Exemple) – *Colloque du GRETSI, 2022*

[C_{22}] A Statistical Model for Predicting Generalization in Few-Shot Classification – *EUSIPCO, 2023*; also [C_{23}], in French

[W_{11}] Inferring Latent Class Statistics from Text for Robust Visual Few-Shot Learning – *NeurIPS R0-FoMo Workshop, 2023*

[P_4] LLM Meets Vision-Language Models for Zero-Shot One-Class Classification – *arXiv preprint arXiv:2404.00675, 2024*

Yassine El Ouahidi



PhD

2021–2025

- Defended on March 20, 2025
- Funding: Région Bretagne (ARED), AI@IMT
- Now: Lead Applied Scientist, Mistral AI

Topics

Vision-inspired deep learning approaches for BCI applications

- Simple and strong deep learning baselines for motor imagery decoding
- Graphs for brain signals: pruning graph networks, interpolating electrodes
- REVE, a foundation model for EEG; collaborations in cardiology

Publications

[T_7] Deep Learning for BCI Applications: Leveraging Recent Advances in Computer Vision – *Manuscript, 2025*

[J_4] A Strong and Simple Deep Learning Baseline for BCI Motor Imagery Decoding – *Transactions on Neural Systems and Rehabilitation Engineering, 2024*

[J_5] Machine Learning for Pacemaker Implantation Prediction After TAVI Using Multimodal Imaging Data – *Scientific Reports, 2024*; also [W_8], in French

[J_{10}] Machine Learning Score to Predict In-Hospital Outcomes in Patients Hospitalized in Cardiac Intensive Care Unit – *European Heart Journal, Digital Health, 2025*; also [W_{16}]

[C_{19}] Pruning Graph Convolutional Networks to Select Meaningful Graph Frequencies for fMRI Decoding – *EUSIPCO, 2022*; also [C_{18}], in French

[C_{24}] Spatial Graph Signal Interpolation with an Application for Merging BCI Datasets with Various Dimensionalities – *ICASSP, 2023*; also [C_{25}], in French

[C_{27}] Unsupervised Adaptive Deep Learning Method for BCI Motor Imagery Decoding – *EUSIPCO, 2024*

[C_{31}] REVE: A Foundation Model for EEG: Adapting to Any Setup with Large-Scale Pretraining on 25,000 Subjects – *NeurIPS, 2025*

[W_{12}] Cross-Dataset Few-Shot Learning for Motor Imagery BCI Classification – *International BCI Meeting, 2023*; also [W_{13}]

Dubon Rodrigue



PhD

2022–2025

- Defended on October 23, 2025
- Funding: IMT Atlantique (CDE), IMT and ANR (AI@IMT)
- Now: Postdoc, Idiap Research Institute

Topics

Simplification of energy network models through AI

- Reducing the topology of district heating networks by aggregating nodes
- Clustering substations, and learning the dynamics of the clusters
- Hybrid simulation and model predictive control with neural networks

Publications

[T_8] Topology Simplification of District Heating Networks via Artificial Intelligence-Based Node Aggregation – *Manuscript, 2025*

[J_7] Topology Reduction Through Machine Learning to Accelerate Dynamic Simulation of District Heating – *Energy and AI, 2024*

[C_{29}] Leveraging Neural Networks in a Hybrid Model Predictive Control Framework for District Heating Networks – *ECOS, 2024*

[C_{33}] Validation of Two Distinct Simulation Models of District Heating Networks: Application to Efficient Looping Analysis – *Journal of Physics: Conference Series, 2025*

[W_{15}] Learning the Dynamic of Clusters of Vertices in a Graph for District Heating Network Simplification – *GSP Workshop, 2023*

[J_{11}] Flexible Clustering of Substations for Accurate and Rapid Hybrid Simulation of District Heating – *Energy Informatics, 2026*

Arsalan Najib



PhD

2023–2024

- Stopped after 11 months
- Funding: CHU of Rouen (salary)

Topics

Classification of ocular diseases using deep learning approaches

- Deep learning on ophthalmic images
- Stopped to focus on his work at the hospital

Publications

No publication, the PhD having stopped early.

Yahia Bellouche



PhD

2024–2026

- Ongoing
- Funding: CHU of Brest (salary)

Topics

Computational fluid dynamics in cardiovascular medicine: from application to a clinical perspective

- A hyperemic resistance index computed from cardiac CT, for coronary syndromes
- Wall shear stress of the pulmonary artery, a non-invasive biomarker for CTEPH

Publications

[J8] Performance of a Novel Computational Hyperemic Resistance Index Derived from Cardiac CT in Coronary Chronic Syndromes – *Journal of Clinical Medicine*, *2025*

[J9] Proximal Pulmonary Artery Wall Shear Stress Derived from Computational Fluid Dynamics: A Non-Invasive Biomarker for CTEPH and Perfusion Mismatch – *Physiological Reports*, *2025*

Jonathan Lys



PhD

2024–2027

- Ongoing
- Funding: ANR JCJC ENDIVE

Topics

Exploring diversity in FSL and diffusion LLMs

- Diversity at decoding time for diffusion language models: D5P4
- Inference of pretrained transformers: inner loops, residual connections
- Foundation models for EEG, with REVE

Publications

[C_{31}] REVE: A Foundation Model for EEG: Adapting to Any Setup with Large-Scale Pretraining on 25,000 Subjects – *NeurIPS, 2025*

[C_{34}] Key Ingredients for EEG Foundation Models: What Works - and What Doesn't – *EUSIPCO, 2026*

[C_{35}] Inner Loop Inference for Pretrained Transformers: Unlocking Latent Capabilities Without Training – *EUSIPCO, 2026*

[C_{36}] Residual Connections and the Causal Shift: Uncovering a Structural Misalignment in Transformers – *EUSIPCO, 2026*

[R_3] D5P4: Partition Determinantal Point Process for Diversity in Parallel Discrete Diffusion Decoding – *NAACL, 2026*

Joshua Chaz Rivera



PhD

2025–2028

- Ongoing
- Funding: IMT Atlantique (CDE), SequoIA

Topics

Diffusion LLMs for unusual data

- Figurative language: metaphors and idioms
- The blind spots of discrete diffusion models

Publications

[R_1] Evaluating LLMs on Sentence-Level Metaphors – *Findings of the ACL, 2026*
[R_2] Idiomatic Blind Spots in Discrete Diffusion Models – *EMNLP, 2026 (accepted)*

Nolan Sisouphanthong-Vongkotrattana



PhD

2025–2028

- Ongoing
- Funding: IMT Atlantique (CDE), PRACom

Topics

Diffusion LLMs and mechanistic interpretability

- Diffusion language models
- Mechanistic interpretability of their representations

Publications

No publication yet, the PhD having started in 2025.

Quentin Banet



PhD

2026–2029

- Starting soon
- Funding: CIFRE with Orange

Topics

Neural coding for audio signals

- Neural codecs for audio signals

Publications

No publication yet, the PhD starting.

Victor Quéré



MD thesis

2020–2021

- Defended on October 13, 2021
- Funding: CHU of Brest (salary)

Topics

Comparison of ML approaches for prediction of ICH early mortality

- Logistic regression against random forests
- On the population-based Brest stroke registry

Publications

[T_2] Comparison of Logistic Regression Analysis with Random Forest for Prediction of ICH Early Mortality in the Population-Based Brest Stroke Registry – , *2021*

Amine El Ouahidi



MD thesis

2022–2023

- Defended on September 9, 2023
- Funding: CHU of Brest (salary)

Topics

Use of ML to predict pacemaker implantation following a TAVI procedure

- Predicting pacemaker implantation after TAVI
- From multimodal imaging data

Publications

[T_4] Utilisation de Méthodes d'Apprentissage Automatique pour Prédire l'Implantation d'un Pacemaker dans les Suites d'une Procédure de TAVI – , *2023*

[J_5] Machine Learning for Pacemaker Implantation Prediction After TAVI Using Multimodal Imaging Data – *Scientific Reports, 2024*; also [W_8], in French

Mounia Hamidouche



Postdoc 2020–2021

- 10 months, then 6 months
- Funding: Carnot TSN

Topics

Using GSP to improve few-shot classifiers

- Graph filtering of features to improve classification accuracy
- Graphs as tools to improve deep learning methods
- Then, theoretical foundations of deep learning

Publications

[B_2] Graphs as Tools to Improve Deep Learning Methods – *Advances in AI: Reviews, 2022*

[C_{14}] Improving Classification Accuracy with Graph Filtering – *ICIP, 2021*

Behnam Vahdani



Postdoc 2020

- 10 months
- Funding: Carnot TSN

Topics

Radio-frequency design via 3D printing

- AI for the design of radio-frequency components

Publications

No publication from this work.

Hafsa El Hafyani



Postdoc

2022–2023

- 18 months
- Funding: ANR ASTRID, with InPixal

Topics

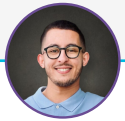
Multimodal approaches for image segmentation

- Multimodal object detection
- Integrating probabilistic a priori information

Publications

[P_5] Multimodal Object Detection via Probabilistic a priori Information Integration
– arXiv preprint arXiv:2405.15596, 2024

Amine Sadikine



Postdoc 2025–2027

- 18 months, ongoing
- Funding: Horizon Europe STRATIF-AI

Topics

AI for stroke recovery analysis

- Fusing clinical and imaging data after a stroke
- Predicting 30-day mortality after a hemorrhagic stroke

Publications

[S_{13}] Multimodal Data Integration for 30-Day Mortality Prediction in Hemorrhagic Stroke – *STRATIF-AI Consortium, 2026*

MG

Research engineer

2021–2023

- 24 months
- Funding: ANR Plan Relance, with Elliptika

Topics

Codesign of hyperfrequency components using AI

- AI for the design of hyperfrequency components

Publications

No publication from this work.

Manon Renault



Research engineer 2023–2025

- 18 months
- Funding: ADEME Logistique 4.0, with Goodfloow

Topics

Prediction from sensor data using embedded deep learning models

- Classifying events from accelerometer data
- Embedded deep learning, for monitoring industrial packages

Publications

[W$_{17}$] Event Classification of Accelerometer Data for Industrial Package Monitoring with Embedded Deep Learning – *COMPSAC, 2025*



Graph Signal Processing

Graph signal processing in a nutshell

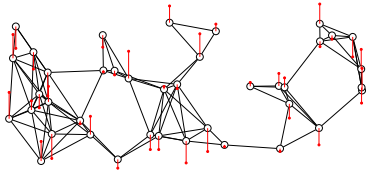
Fourier on a graph

A signal $\boldsymbol{x}$ gives a value to each vertex of a graph of weights $\boldsymbol{W}$ and degrees $\boldsymbol{d}$ [Shuman *et al.*, 2013]:

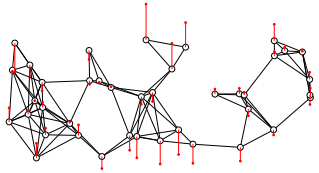
$$\begin{aligned}\boldsymbol{L} &= \operatorname{diag}(\boldsymbol{d}) - \boldsymbol{W} \\ &= \boldsymbol{V} \operatorname{diag}(\boldsymbol{\lambda}) \boldsymbol{V}^\top \\ \hat{\boldsymbol{x}} &= \boldsymbol{V}^\top \boldsymbol{x}\end{aligned}$$

Eigenvectors play the role of Fourier modes, and eigenvalues of frequencies. A filter $\hat{\boldsymbol{h}}$ acts in the spectral domain:

$$\boldsymbol{y} = \boldsymbol{V}\left(\hat{\boldsymbol{h}} \odot \hat{\boldsymbol{x}}\right)$$



A random signal on a random geometric graph



After a low-pass filter $\hat{h}[i] = e^{-0.1i}$: neighbors take similar values

Inferring a graph from diffused signals

Graph inference under stationarity

Signals are observed after white noise has diffused on an unknown graph, through an operator $\boldsymbol{S}$. Their covariance is then a function of $\boldsymbol{S}$, and both share their eigenvectors [J1]:

$$\boldsymbol{\Sigma} = \boldsymbol{V}\,\mathrm{diag}(\boldsymbol{\sigma})\,\boldsymbol{V}^\top$$
$$\Rightarrow\ \boldsymbol{S} = \boldsymbol{V}\,\mathrm{diag}(\boldsymbol{\mu})\,\boldsymbol{V}^\top$$

- Only the eigenvalues $\boldsymbol{\mu}$ remain unknown
- Admissible operators, symmetric with a null diagonal and non-negative entries elsewhere, form a convex polytope
- A sparse one is chosen in it, with a linear program

What came after

- Graph shift operators from spectral templates [Segarra *et al.*, 2017]
- Graphs from heat diffusion [Thanou *et al.*, 2017]
- Non-stationary signals, through graph filters [Shafipour *et al.*, 2021]

The covariance of signals as a key to their graph: my most cited article, 217 citations.

Uncertainty and translations on graphs

An uncertainty principle

A signal cannot be localized both on the vertices and in the spectrum [Agaskar *et al.*, 2013]. With spreads in the two domains:

- Weighted graphs, with distances that follow the weights [C_2]
- Signals that reach the uncertainty curve [C_5]
- A bound on rings, none on complete or star graphs [B_1]

From translations to graph CNNs

Translations first, as maps that preserve neighborhoods, then convolutions built from them [C_4]:

- On cyclic grids, only the usual translations
- Infer a graph, find its translations, move a kernel with them: a CNN without a given graph [P_2]
- An NP-complete search, overtaken by spectral and message-passing GCNs [Defferrard *et al.*, 2016; Kipf *et al.*, 2017]

GSP for models, and for the brain

Denoised features

With M. Hamidouche : a graph of neighbours per class in feature space, low-pass filtered before the classifier is retrained [C_{14}]; reused by others [Pan *et al.*, 2023].

Latent geometry graphs

Label smoothness on graphs of intermediate representations tells a well-trained model from an overfitted one [B_2]; a batch of views as a graph gives a contrastive loss [W_{14}].

Decoding fMRI

Graphs over the regions of an atlas, from geometry or functional connectivity; sampling or Fourier projection before classifying: one of the first uses of GSP in neuroscience [C_8].

Stationarity of an EEG

Sliding windows of a recording, their covariances against a common eigenbasis: the residual measures a departure from stationarity, a trace of the class [W_5].



REVE

REVE, a foundation model for EEG

REVE

NeurIPS 2025

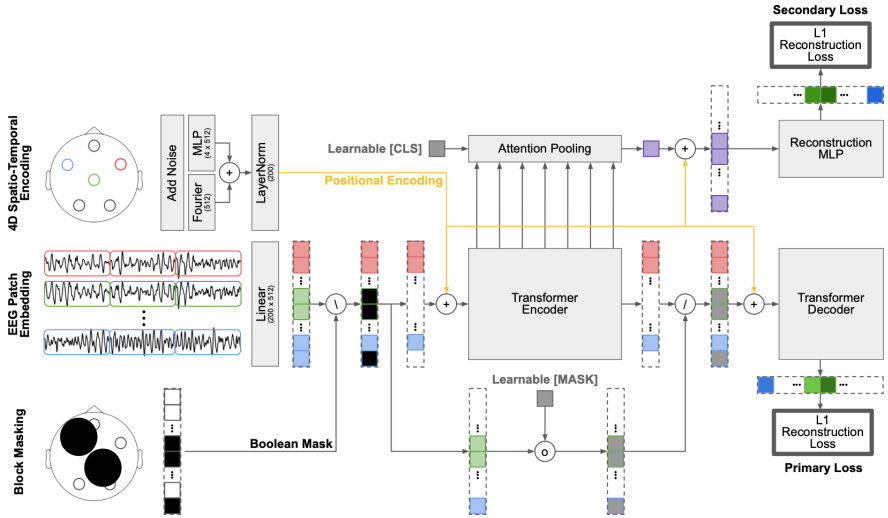
Y. El Ouahidi , with G. Lioi, helped by J. Lys  [C31]

- More than 60 000 hours of EEG, from 92 public datasets and about 25 000 subjects
- A transformer, pretrained by masked autoencoding
- A 4D positional encoding, the 3D position of each electrode and time: any montage, any duration
- State of the art on 10 downstream tasks, even with linear probing

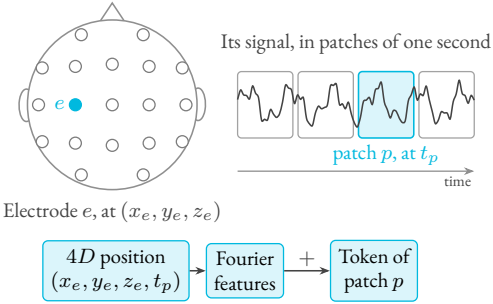
Its reception

- First overall in NeuralBench [Banville *et al.*, 2026]
- 117 citations within a year
- Weights, code and tutorials released
- The first NeurIPS article of the team, written by PhD students together

The architecture of REVE



Any montage, any duration: the $4D$ positional encoding



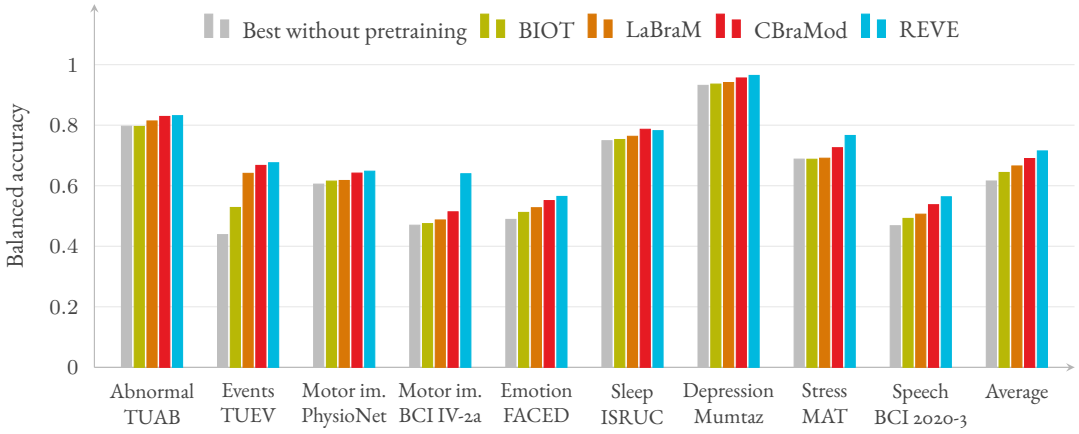
One montage per dataset — The obstacle

Each dataset has its own electrodes, device and duration: 92 of them in pretraining, dozens of montages. Earlier models fix one montage, onto which recordings are projected: channels selected, or interpolated.

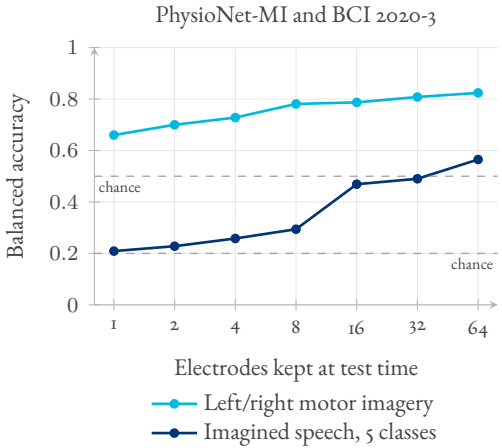
The layout as part of the input — $4D$ encoding

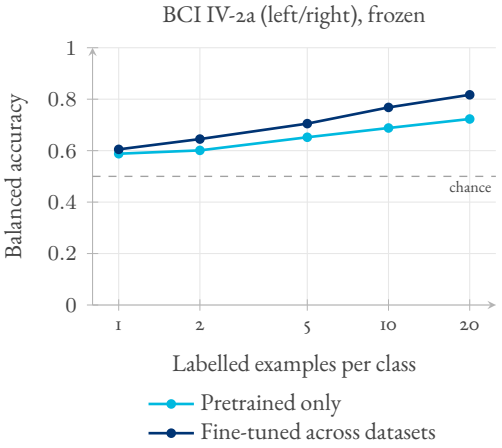
Every patch carries the $3D$ position of its electrode and its time, as Fourier features added to its token. Any number and arrangement of electrodes, any duration: a bipolar montage never seen in pretraining (TUEV), inputs of 30 s after pretraining on 10 s (sleep staging).

REVE against the other models, after fine-tuning



Beyond fine-tuning: fewer electrodes, fewer labels





Before and after REVE

EEG-SimpleConv — Baseline

A plain $1D$ CNN reaches the state of the art in motor imagery; normalization is what matters most [J_4].

Montages, unified — GSP

Missing electrodes are interpolated on a learned graph, to train across datasets [C_{24}].

New users, no labels — Adaptation

Alignment and normalization statistics updated online bring cross-subject decoding close to offline accuracy [C_{27}].

REVE 2 — Ongoing

MEG, other positional encodings and losses, new datasets, and curation: five current or former PhD students, with G. Lioi.

A review of EEG foundation models: scale alone does not beat a small model trained for the task [C_{34}]



Few-Shot Learning

Few-shot learning

EASY

Best paper award

With Y. Bendou : ensembles, multi-crop inference and a nearest-class-mean classifier, a simple baseline still among the best [J_2].

Predicted accuracy

EUSIPCO 2023

Gaussian classes and an unbiased distance between their means predict the error from the shots alone [C_{22}].

Less can be more

Pretraining

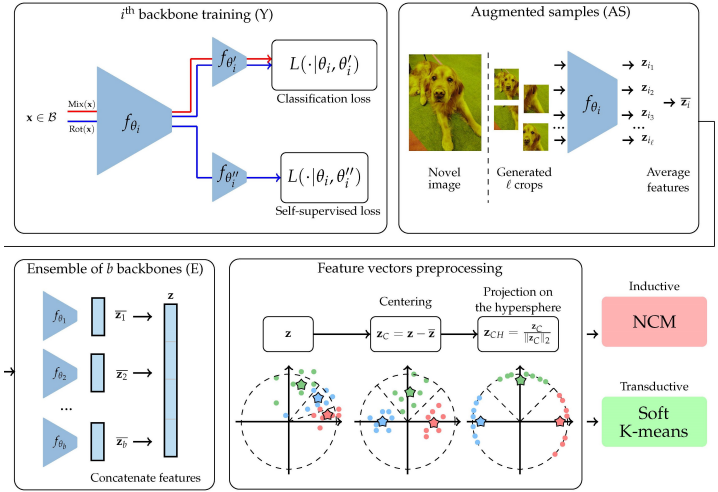
Removing one class from pretraining can improve few-shot accuracy; a signal-to-noise ratio tells when [C_{21}].

Wider intervals

Benchmarks

Benchmarks underestimate confidence intervals, by sampling episodes but not data [Lafargue *et al.*, 2024].

The pipeline of EASY





Language Models

Autoregressive and diffusion language models

Autoregressive

One token at a time, each from the previous ones, as in GPT:

$$p(\boldsymbol{x}) = \prod_{t=1}^{T} p(x_t \mid \boldsymbol{x}_{<t})$$

Left to right, and each choice is final.

Diffusion

Tokens are masked with a probability t, at positions $\mathcal{M}_t$, and the model learns to unmask them [Austin *et al.*, 2021; Nie *et al.*, 2025]:

$$\mathcal{L} = -\mathbb{E}\Big[\tfrac{1}{t} \sum_{i \in \mathcal{M}_t} \log p(x_0^i \mid \boldsymbol{x}_t)\Big]$$

All positions are denoised together, over several steps.

Overlooked by the major actors when we started, now LLaDA, Dream, and DiffusionGemma [Nie *et al.*, 2025; Ye *et al.*, 2025; Taïga *et al.*, 2026]

Other works on language models

Inner loops
EUSIPCO 2026

With J. Lys : re-applying middle blocks of a frozen transformer, with a regularization, gives small but consistent gains [C_{35}].

Causal shift
EUSIPCO 2026

Hidden states turn from the current token to the next one late in the network; a learned gate on the residual helps [C_{36}].

Novel metaphors
In revision

With J. Rivera : models rank a metaphor against its literal version, but fail to label it alone [R_1].

Idioms
EMNLP 2026

Diffusion models know idioms, but tokens committed early keep them from producing them [R_2].



Projects Ahead

Applied projects

STRATUM

Under review

Blood markers of stroke in the first hours: 720 patients, and 3 to 5 genes against stroke mimics [Ramsay *et al.*, 2019]; my part, multimodal models with missing data.

SIGMAS

CHU of Brest

Small vessel disease on clinical MRI: segmenting its markers, measuring what they add to prognosis, then their links to genomics.

Weather and infarction

How many cardiologists on duty? Daily admissions for myocardial infarction, predicted from the weather forecast [ORBI].

Audio codecs

CIFRE Orange

Q. Banet, with A. Marmoret: one representation for mono and stereo, small models, an interpretable latent space.

ALIVE: detecting the drift of an agent, in detail

The sequential test

$\alpha = 5\,\%$

- Two answers per question, before and after the change, each scored: a pair (b, b'). A wealth W, 1 at first, is multiplied by $1 + \varphi(b) - \varphi(b')$ at each pair; the alarm sounds when $W \geq 1/\alpha = 20$ [Richter *et al.*, 2025]
- φ, a small network, is refitted every 25 pairs on those seen so far, to bet on the right side. If nothing changed, every bet is fair, and Ville's inequality keeps the false alarms under α, however often one looks

The score of an answer

- 1 000 questions of HealthBench, the same before and after, answers only [Arora *et al.*, 2025]
- One LLM judge votes yes, no or not applicable on 8 fixed criteria of HealthBench Consensus; the score is the share of yes
- Same judge, same provider on both sides: a change of provider alone set off the alarm

Which criterion moved

- One test per criterion, each at $\alpha/8$: a little slower, but it names the criterion; the fused evidence of the 8 tests [Bharti *et al.*, 2026] sounds earlier
- Four criteria up with GPT-5.5, emergencies first; none flagged on the agent against itself, over 150 orders



Directions

Four directions, and where they would start

Text as the pivot

Specialized models describe each modality in text, for an orchestrator, and for a clinician who can disagree [Klang *et al.*, 2026]; first test: the ICH prognosis, its tabular branch replaced by text.

Trust that circulates

Agents as a graph: the attention of a GAT as trust, revised over exchanges [Veličković *et al.*, 2018]; abstention by conformal prediction, calibrated on a set drawn by a DPP: does validity survive?

Genomics as the next modality

Four letters, close to text [Ji *et al.*, 2021]; k-mers and de Bruijn graphs [Marchet *et al.*, 2021]: a sequence and a graph at once, with far more variables than patients. Partners in Montpellier, Lille and Paris-Saclay.

A cross-service environment

One agent per service, where its data lives, an orchestrator in natural language; text circulates, not weights: auditable, unlike a gradient [McMahan *et al.*, 2017]; a proof of concept on open data first [Johnson *et al.*, 2023].



CLINICAL WORKS

Neurology

Mortality after an ICH

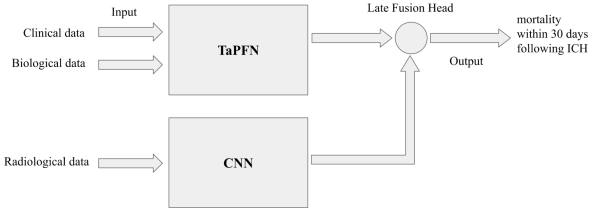
MD thesis

Brest stroke registry: dropping incomplete records removes 65 % of the patients, the most severe; imputing them raises the AUC from 0.78 to 0.83 [T_2].

STRATIF-AI

Horizon Europe

Stratification after a stroke, in 15 centers; with A. Sadikine , clinical and imaging data fused for 30-day mortality [S_{13}].



Tabular data through a tabular foundation model, images through a CNN, fused for the 30-day mortality after an ICH [S_{13}]

Oncology

Cells as a graph

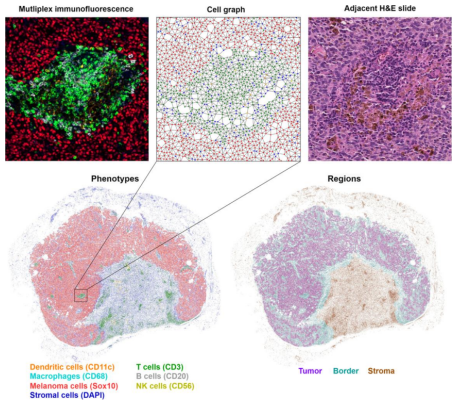
LUNG CANCER

94 patients: distances between lymphocytes and tumor, compared by optimal transport, separate survival groups without any label [C_{12}].

Immunotherapy

MELANOMA

77 patients: two spatial features, contacts between dendritic and T cells and B-cell clusters, reach an AUC of 0.74 on validation [R_4].



From a stained slide to a graph of cells, by Delaunay triangulation [R_4]

Cardiology

Pacemaker after a TAVI

MD THESIS

520 patients, 67 variables: an SVM reaches an AUC of 0.92, the membranous septum on CT being key; a clinical research award [T_4; J_5].

Intensive cardiac care

AP-HP

1 499 patients in 39 units: 7 variables at admission reach an AUC of 0.90, held on an external cohort [J_{10}].

Coronary lesions

CFD

With Y. BELLOUCHE : 64 patients, 106 lesions; an index computed from CCTA and fluid dynamics reaches 96 % accuracy [J_8].

After a pulmonary embolism

CFD

90 patients: a low wall shear stress in the proximal arteries flags thromboembolic disease, with an AUC of 0.92 [J_9].



Other Collaborations

AI for other domains

Choosing operators

Q-learning picks the perturbation of a metaheuristic, for flow-shop scheduling; 167 citations [J3].

Diverse parents

k-DPPs choose the parents of a genetic algorithm that learns ranking models [C26; J6].

Cruise ships

Predicting propulsion power: random splits of time series leak, and make models look better than they are [C13].

District heating

Clusters of substations replaced by neural surrogates, chosen by physics, to speed up simulations [J7; J11].

Reusable packages

A $1D$ CNN on accelerometers, for a sensor that has to last for years [W17].

Metal vocals

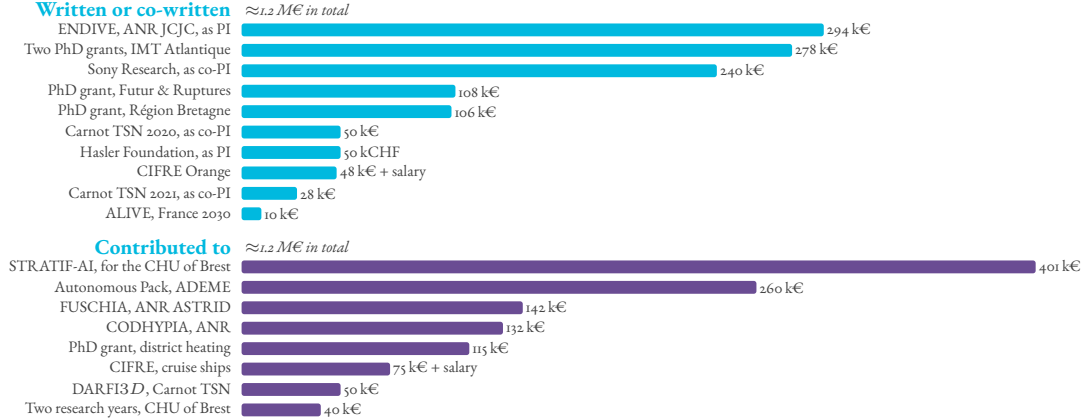
Speech recognition on 14 subgenres of metal, out of distribution; a badge for reproducibility [C32].

A frequent request of physicists: fast models to replace costly simulations



Funding

Funded projects





Other Activities

Research stays, invited talks

Research stays 4

- LTS2, EPFL, 2014: P. Vandergheynst
- McGill, two months, 2015: M. Rabbat
- Sony Research, Stuttgart, 2022: F. Cardinaux
- CRIStAL, Lille, 2025: R. Bardenet

Invited talks 13

- PyRat: UBL (2016), ECAM (2017)
- Graph CNNs: GIPSA-Lab, Nantes, Inria Rennes, LTS4 (2017)
- GSP for ML and healthcare: ENS Lyon, CRIStAL (2018)
- Foundation models for text: IMT Atlantique, Lab-STICC (2024)
- Diversity in FSL: JCJC committee (2024), CRIStAL (2025)
- Multimodal and agentic AI for health: STRATIF-AI consortium (2026)

Events, reviews, challenges, qualifications

Events organized 4

- GSP Workshop, EPFL, 2018
- Signal processing over graphs, neuroscience data: GdR ISIS, 2019
- Tutorial day on foundation models, IMT Atlantique, 2024
- Workshop on FSL, IMT Atlantique, 2024

Reviewing 29 REVIEWS

- IEEE TSP: 10, since 2016; IEEE TSIPN: 4, since 2019
- GRETSI: 7, 2022–2025; ACL ARR: 3, 2026
- Signal Processing, IEEE SPL, EURASIP, ACM SIGMETRICS, Bioinformatics: 1 each
- ANR JCJC projects: 2, 2025–2026

Programming challenges 3

- Enigmatic, Télécom Bretagne, 2014
- TaupIC, Télécom Bretagne and ENS Rennes, 2015 and 2016
- As a co-organizer, each time

Qualifications, training 4

- CNU qualifications, sections 27 and 61, 2018
- Flipping a class, EPFL, 2019
- Course design, RCFE and EPFL, 2019



Teaching & AI

Teaching at the time of generative AI

The question

How to teach, when generative AI solves the exercises faster than students can learn what they need to solve them?



An engineer must use AI tools, but needs a systemic view of a problem to use them well.

What I observed in PyRat

- Documentation, code and tests generated separately, never checked together
- Prompts out of context: *"Give me a Dijkstra implementation for PyRat"*
- Copilot taken for the language model itself [GitHub Copilot]
- Grades push to pass rather than to understand

What I changed in my courses

Agents in the editor

A practical session on agentic AI in the IDE, in first year and in the data science track, with free student accounts.

High-level skills

Agents write code, not software: structure the program, test each part, and deliver what a client asked for.

Two steps

Agents off, to understand the concepts; then on, to be efficient, and to find better solutions than by hand.

What teachers still bring: a high-level understanding, in short lectures,
which students now value more than flipped classes, agents tutoring them outside

Building one's own tools

The aim

Step back from one's own needs, to use generative AI well, in a project that is actually useful.

How

- Start from a need of one's own, not from what the tool can do
- Write the rules: what it must do, and what must never happen
- The agent writes the code; one decides the architecture, and checks
- Used every day: errors show up fast

Example: Marvin

Since October 2026

- A personal assistant: mails, calendars, research projects, supervisions, courses
- Explicit rules first, AI only at given points: it proposes, I decide
- Every action logged and undoable; what is read is data, never instructions
- Corrections become rules, applied without AI

Generative AI in this manuscript

The template

A LaTeX template built with ChatGPT, Gemini, then Claude [Claude Code]: code and advice, all reviewed and tested.

References

Models as search engines, Scholar Labs and Scholar Inbox among them; the articles read afterwards.

Summaries

First drafts of the summaries of my past works, from the articles; then reviewed and adapted by hand.

Scripting

The graph of co-authors, partly generated, then checked and adapted by hand.

Syntax

Written in English directly, with final touches by a model: I am still French.

These slides

All graphical elements and layouts, with Claude Opus 5.5 and Fable 5.1 [Claude Code].